EE412 Fall 2014 Final Report

Surface Micromachining Method for Releasing a Range of Micron-Scale Membranes

Martin Winterkorn, Anup Dadlani, Yongmin Kim

Mentors: J Provine, Michelle Rincon

1. Motivation

For many MEMS devices that contain released structures, the mechanical properties and therefore overall performance depend strongly on the positioning of the anchors. Achieving the desired locations is particularly challenging when the dimensions of the released membrane or structure are very small, since absolute errors have a much larger relative impact. For this project, our aim was to release membranes in sizes ranging from 0.5-100 by 4-100 um. In this regime, the classical ways of achieving a release all have major shortcomings:

- Using DRIE from the backside would most likely not work at all for the smallest membrane sizes such as 0.5x4um, or take very long for the larger ones, since the etch rate decreases dramatically with the opening size. Typically, the open area dimensions for Bosch process backside releases are several 100 um which is not the case here.
- While backside KOH or TMAH etching itself is a very well-controlled process, the uncertainty in the thickness specifications of typical 500 um thick Si wafers, which is usually around 25 um, creates a lateral uncertainty on the same order (see Figure 1a). One could account for that by measuring each wafer's thickness individually and adjusting the mask size (possibly even by changing the overlap of partial images on the stepper). However this is very time-consuming and becomes even more challenging when also trying to account for the thickness variation within each wafer. Another source of error in this process would be the alignment uncertainty of the flat to the wafer's crystal axes.
- By using an SOI wafer and first removing a large, non-size-critical area of the handle layer, then doing lithography again on the bottom of the exposed device layer and using KOH etching, one could achieve higher accuracy since the thickness specifications of the device layer are usually much tighter. Disadvantages of this approach are the higher cost of SOIs and the complexity of doing the second litho step on the bottom of the trenches.
- An approach fundamentally different from the previous ones would be to passivate the wafer surface, e.g. with aluminum oxide, then create holes in that passivation layer and use xenon difluoride vapor-phase etching to remove the silicon underneath. By timing the etch, it is possible to confine it to fairly small areas. However, unless the target release

area is circular and one is willing to create a hole at the center of it, it also fundamentally always results in a larger than desired area to be released (see Figure 1b).

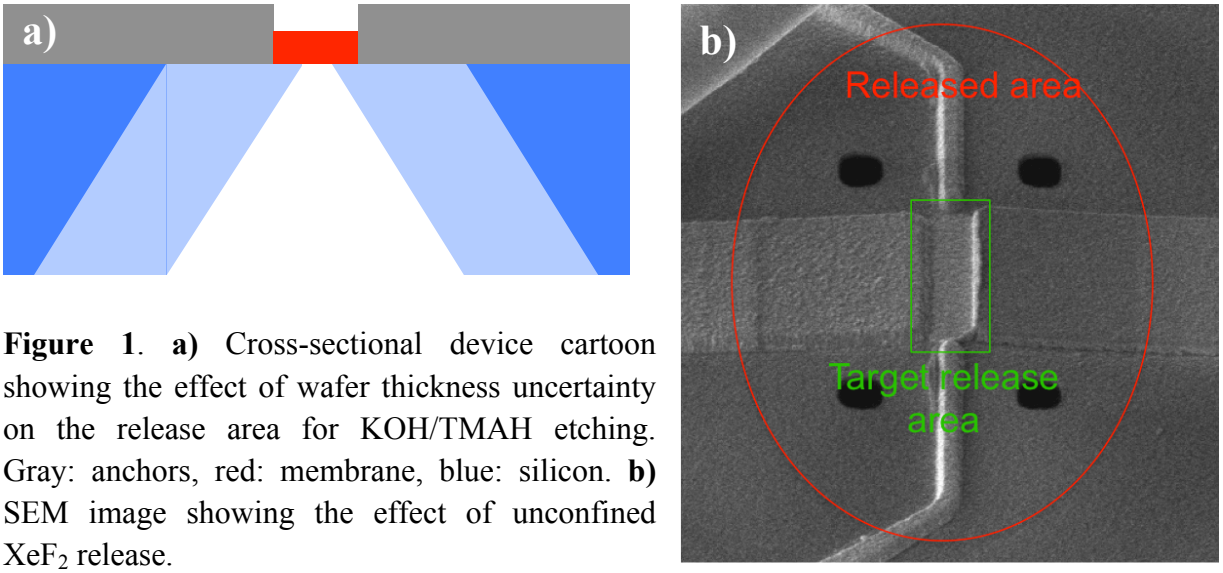

Figure 1. **a)** Cross-sectional device cartoon showing the effect of wafer thickness uncertainty on the release area for KOH/TMAH etching. Gray: anchors, red: membrane, blue: silicon. **b)** SEM image showing the effect of unconfined XeF_2 release.

Therefore, the goal of this project was to employ and characterize a novel method that does not have any of these shortcomings. It is described in detail in the following section.

2. Fabrication Process Overview

A cross-section cartoon of the process is shown in Figure 2. The overall idea is to use a passivation layer and xenon difluoride etching as described in the last example of the previous section, but to pre-define the release area by first creating a trench structure with passivated sidewalls, which we call “etch sandbox”.

To create the sandbox, we simply etch into the single-crystal silicon wafer a few microns deep using RIE (1). The sidewalls and bottom of the trenches as well as the wafer surface are then passivated with Al_2O_3 deposited by ALD (2). To be able to make a membrane over the sandboxes, they are refilled with a poly-silicon layer from LPCVD that is at least as thick as the depth of the trenches (3). In order for the membranes to be flat, the whole wafer is polished using CMP (4). At this point, the wafer is ready for the deposition and patterning of the actual membrane or MEMS structure. For the case that its materials also get attacked by XeF_2, it needs to be sandwiched between two additional Al_2O_3 layers that are deposited by ALD before (5) and afterwards (6). When the structure is ready to be released, holes are etched into the Al_2O_3 layer (7). Differently from what is shown in the figure, these do not have to be inside the membrane or structural area, but can be immediately adjacent – to the front or back – as long as they are included in the inside of the box area created in step (1). Finally, XeF_2 etching is used to remove the polysilicon from the inside of the sandbox (8). Other areas will not get attacked since the

sidewalls and bottom of the boxes, the bottom and top of the membrane as well as the rest of the wafer are all covered in Al_2O_3.

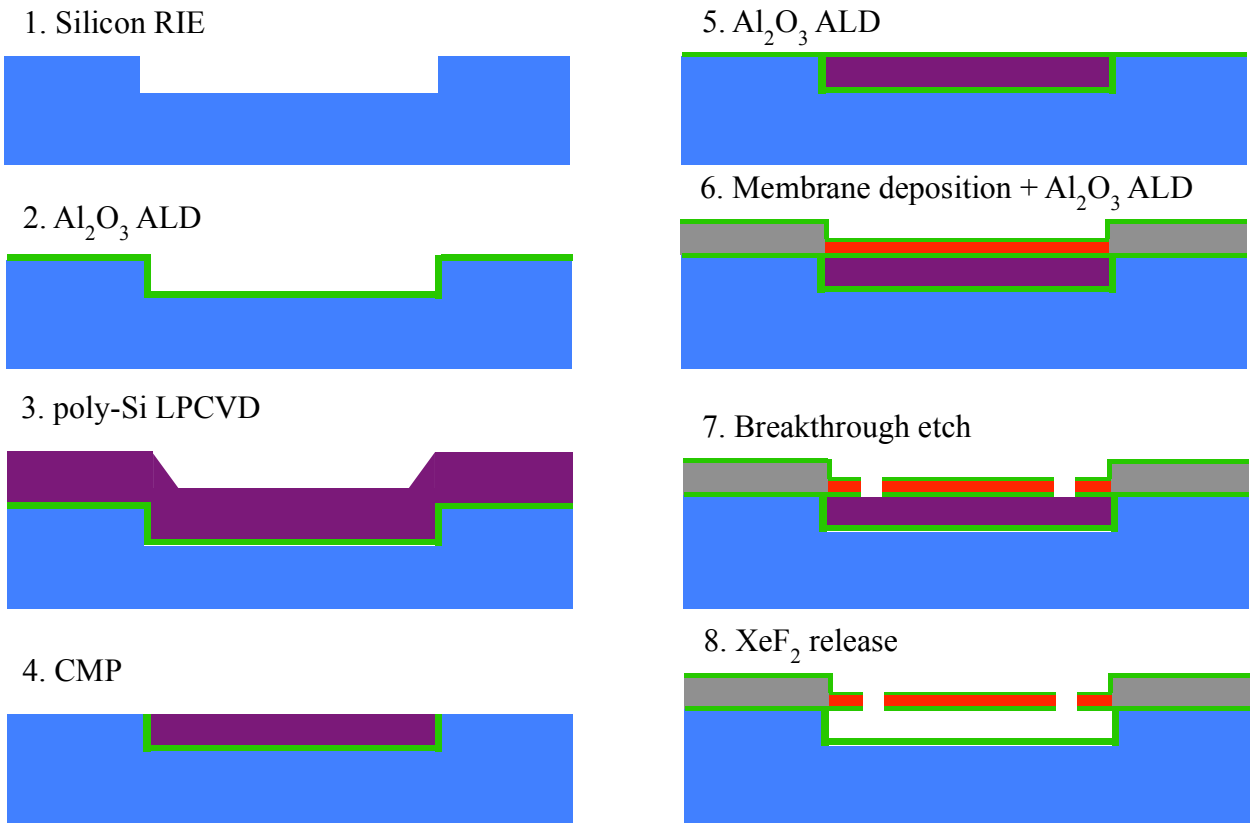

Figure 2. Process flow cartoon showing the major processing steps in cross-section. Blue: silicon, green: aluminum oxide, purple: polysilicon, gray: anchors, red: membrane.

3. First Steps: Box Etching, Passivation and Polysilicon Refill

As the tool for the etching of the boxes we used lampoly in order to get the smoothest sidewalls possible. The result prior to passivation can be seen in Figure 3. The etches were timed, after using CCI-HD to determine the etch rate, to give depths of 250 nm, 1 um and 4 um. This was limited by the maximum thickness that the LPCVD furnace can deposit in one run.

For each of the depths, 10 nm and 50 nm thick Al_2O_3 layers were deposited on fiji1, 2 or 3 using the standard plasma Al_2O_3 recipes at 200C for 100 and 500 cycles, respectively. On dummy wafers without boxes, intended for CMP characterization, the thicknesses chosen were 5 nm, 20 nm and 50 nm.

To ensure a complete refill of the boxes, the target thicknesses for the polysilicon depositions were chosen around 25% above the box depths. Due to being semi-clean or gold-contaminated from ALD, the only polysilicon furnace available to use was thermcopoly2. For the wafers with

250 nm and 1 um boxes we used the P580POLY recipe, which resulted in very smooth and shiny layers, almost indistinguishable from blank silicon wafers. To avoid extremely long deposition times we went to higher temperatures for the 4 um boxes with the P620POLY recipe, which gave a visibly rougher surface. The thickness uniformity was characterized using nanospec (for dummy wafers with SiO_2 below the polysilicon) and woollam (for wafers with Al_2O_3). While their results for the absolute thickness differed by around 10%, both showed that there was less than 1% variation within each wafer, making it reasonable to assume perfect uniformity as the starting point for CMP characterization later.

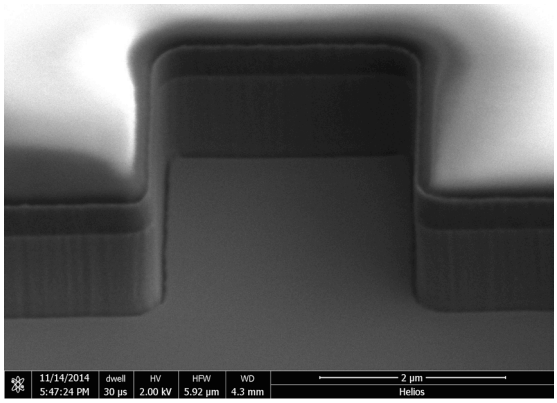

Figure 3. SEM image of 1 um deep boxes after lampoly etch, before resist removal.

4. CMP

Chemical Mechanical Planarization (CMP) is a common technique used in the semiconductor industry to flatten out features or smoothen surfaces. It was used in this process to remove the excess polysilicon and create a flat surface for the membrane patterning. In addition we were interested to see if alumina, our etch protection layer from XeF_2 would act as an adequate polish stop for the CMP process. We were also interested in process repeatability, a common issue previous users had conveyed from our interactions. Uniformity was also another cause for concern as we had learned that the wafer etches much faster on the perimeter as compared to the center, and our mentor Michelle had told us about donut shapes forming from her past experience. Another huge issue we would likely encounter would be dishing, a very common issue that often results from the CMP process.

4.1 First Tests

From our initial CMP runs it seemed that the technique would not be so valuable because there were issues in repeatability and big gradients in the depth of the material at different locations on the wafer. The process consisted of running the wafer for a certain time and checking by eye what had occurred from the color change of the polysilicon layer, and iterating this process until the desired depth was reached. Due to non-uniformity seen with the naked eye and later confirmed by woollam the process looked like it would come up short from what we were aiming for. However this was before endpoint detection was used and before we fine-tuned some of the parameters to mitigate some of these concerning issues.

4.2 Endpoint Detection

The CMP available at SNF has the capability of monitoring the friction force and the temperature, which can in principle be used as indicators when a different layer has been reached. However none of the previous users had tested that and we were the first to attempt using this for endpoint detection. It turned out that the instrument is more than capable of at least detecting that another material had been reached on the wafer. While the temperature does not yield useful information, the friction force over time shows clear peaks that can be used as the indication to manually end the polishing. A typical endpoint curve is shown in Figure 4.

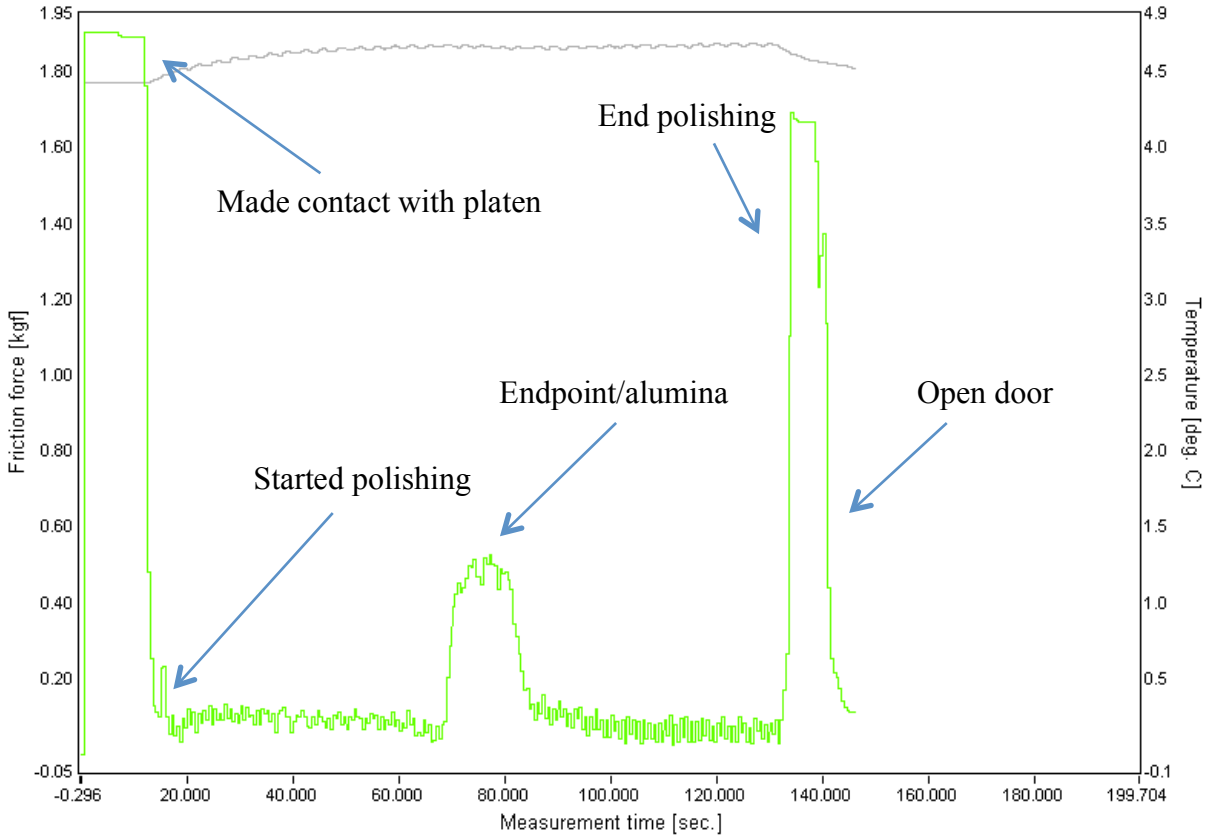

Figure 4. Endpoint detection in action showing friction force and temperature (not used) versus time, with added labels when each event occurred.

This worked even in extreme cases, where we polished through a little over 1 um of polysilicon and had only a 5 nm thick layer of alumina, as can be seen in Figure 5. All the different combinations of polysilicon and alumina thicknesses that we tested and whether they could be detected is shown in Table 1.

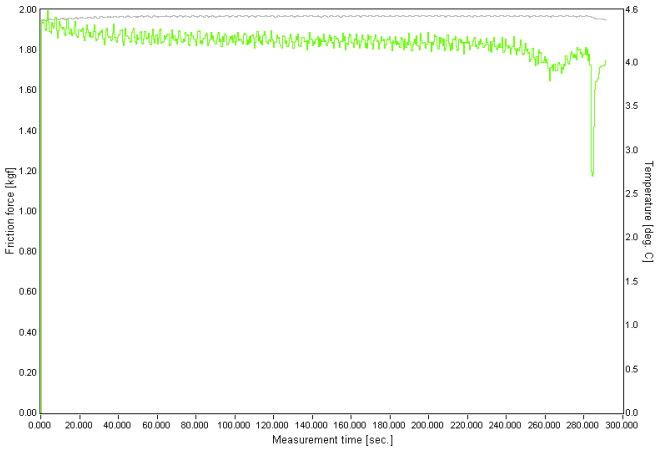

Figure 5. Friction cure used for endpoint detection for the case of having only 5 nm of Al_2O_3 below 1 um of polysilicon. The smaller bump marks reaching the Al_2O_3 layer.

Thicknesses Al_2O_3 \ poly	250 nm	1 um	4 um
5 nm	yes	yes	no
10 nm	yes	yes	
20 nm	yes	yes	
50 nm	yes	yes	yes

Table 1. Visibility of an endpoint peak for different combinations of alumina verus polysilicon thickness. Empty boxes mark cases that were not tested.

4.3 Stopping Capability

This endpoint detection results were encouraging but we wanted to see if we were actually able to stop the polish on the alumina and not penetrate further into the silicon lying underneath. We define "stopping" in this context as having some alumina left on every part of the wafer, while having cleared the polysilicon everywhere except possibly edge areas without devices. The presence of alumina was tested using 55-point measurements on wollam. We found that stopping was possible for a number of combinations of trench depth and alumina thicknesses, mostly those having at least 20 nm of alumina or more, which is shown in Table 2.

We would like to emphasize that it is by no means required to stop on the alumina in order for the sandboxing process to work. (In fact, the process overview in Figure 2 shows the alumina getting removed). The only requirement is to clear the polysilicon while not polishing through and completely removing the boxes, which is why having endpoint detection is reassuring. The main advantage of stopping would be to leave the box depths unaffected, i.e. to not add any depth variation from CMP non-uniformity. For potential applications where the box depth is critical, such as capacitive membranes or microfluidic channels, this would be highly desirable.

Thicknesses Al_2O_3 \ poly	250 nm	1 um	4 um
5 nm	no	no	no
10 nm	yes	no	no
20 nm	yes	yes	no
50 nm	yes	yes	no

Table 2. Capability to stop on different thicknesses of alumina versus the polysilicon thickness, using the endpoint detection system.

For the application we are aiming for, however, we decided that having a smooth wafer surface and a uniform alumina thickness before adding the core membrane layer would be more important, so we deliberately polished a few seconds past the endpoint to remove all the alumina.

4.4 Repeatability

Since using endpoint detection, which made it possible to polish each wafer continuously in one go instead of stopping in between and checking it multiple times, the polish uniformity as determined by the naked eye was much better. To further improve the repeatability, we also started to condition the pad after each wafer as opposed to the recommended suggestion listed in the manual, which is to recondition only after every 10 minutes of polishing. This led to the variation in polish rates – as determined from the time until hitting the alumina layer for wafers from the same batch – to be greatly reduced to less than 10%.

4.5 Uniformity

Despite the improvements, the cause of in some cases not being able to stop on alumina even with the endpoint being detected, was likely a remaining degree of non-uniformity, so we decided to characterize this further. This was done by deliberately stopping the polish well short of the endpoint and then using woollam to measure a 55-point thickness profile of the remaining polysilicon. The same was done for a wafer that only had around 50 nm of alumina on it; it was also polished using the slurry intended for silicon, as it would happen in the sandbox application. For the polysilicon measurement, several new multi-point analysis models had to be created on wollam. Since the tool is not really intended for measuring layers as thick as several hundred nm, the starting guess for the thickness has to be fairly close, usually within less than 50 nm of the actual value in order for the fit to converge on all 55 points. (Manually adjusting the fit for select points is unfortunately not possible with the current software version.)

The resulting profiles are shown in Figures 6 and 8. The uniformities calculated from them, defined as thickness standard deviation divided by polish depth, are ~2.6% for polysilicon and ~12% for alumina. (Note that the uniformity percentages printed by woollam are different since they are calculated for a deposition starting at zero thickness, which is not the case here.) For our calculation we implicitly assume that the starting layer is perfectly uniform, which is reasonable since we had earlier verified the uniformity of polysilicon LPCVD and Al_2O_3 ALD to be better than 1% within each wafer. The average polish rates for the default recipe can also be determined from the profiles, they are 3.2-3.8nm/s for polysilicon and 1.3nm/s for alumina. The key numbers are summarized in Table 3.

Material	poly-Si	Al_2O_3
Polish rate [nm/s]	3.2 – 3.8	1.3
Uniformity	2.6%	12%

Table 3. Summary of CMP characterization results for Si slurry and default recipe.

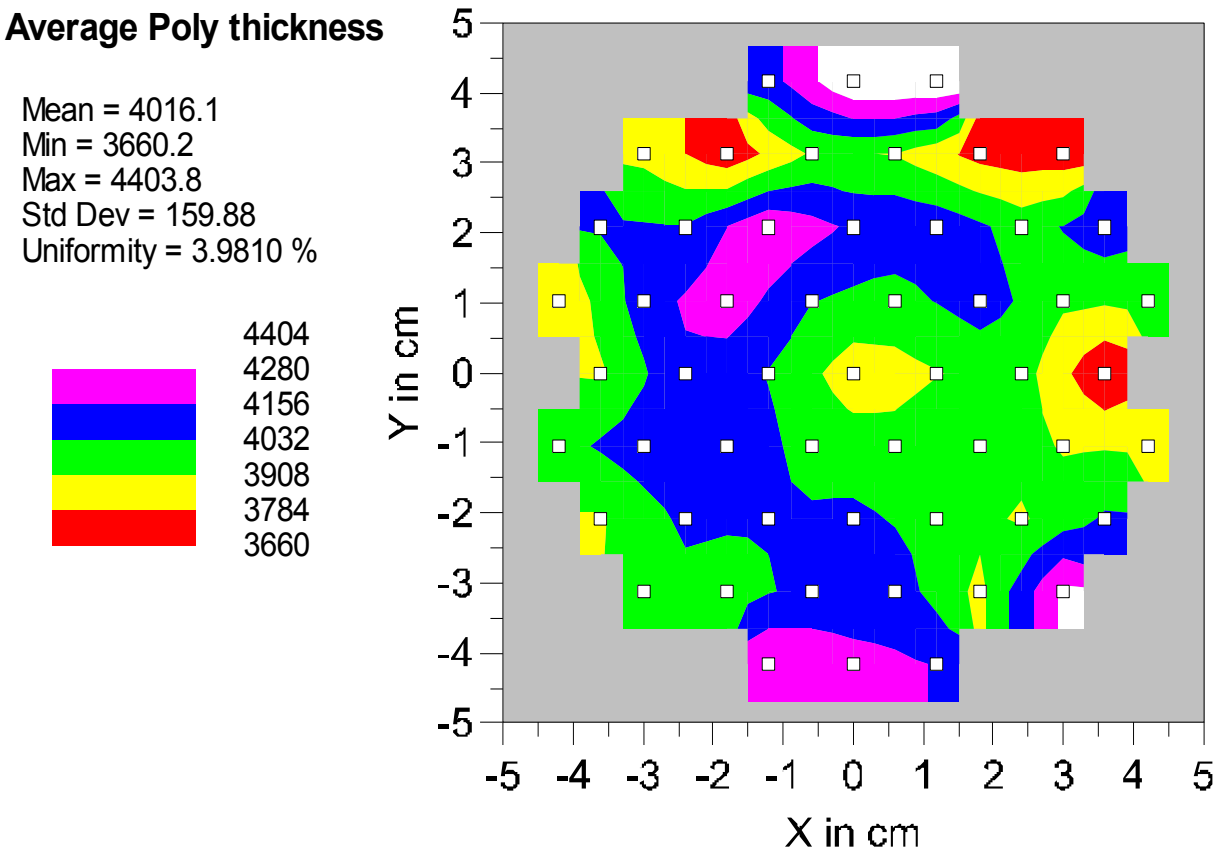

Figure 6. 55 point measurement on woollam of the remaining polysilicon thickness after polishing a 1 um thick layer for 180 seconds. All thicknesses in A. The resulting CMP uniformity turns out to be ~2.6%.

4.6 Dishing

One of the anticipated issues with CMP was dishing – the effect of some features receding below or protruding above the wafer surface after polishing, as illustrated in Figure 7. To address this, we designed several mask variations to employ strategies commonly used in industry to combat dishing, such as having dummy boxes around the actual features, or arraying the features with varying spacing. In addition, we also wanted to characterize the effects of different feature sizes and shapes (circular vs. square vs. rectangular), leading to a total of over 50 device variations on the mask.

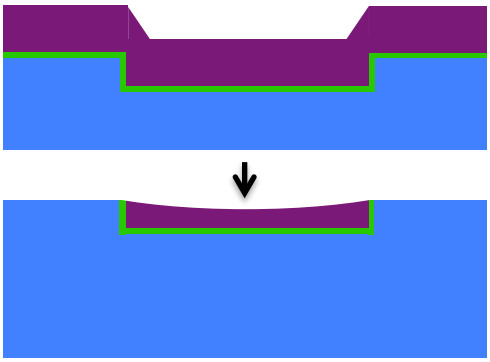

Figure 7. Cross-sectional device cartoon illustrating dishing.

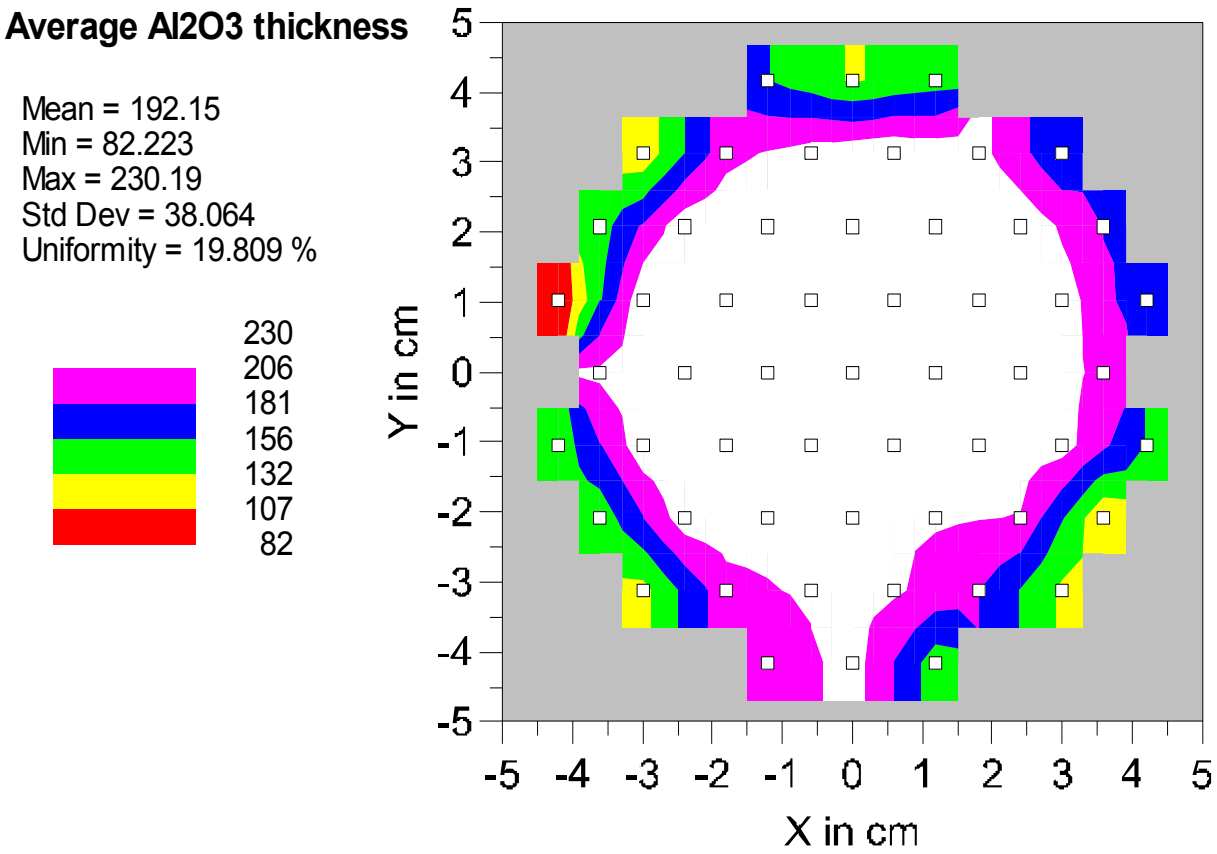

Figure 8. 55 point measurement on woollam of an Al_2O_3 layer that was polished for 20 seconds, starting with a thickness of 48.5 nm. All thicknesses in A. The resulting CMP uniformity turns out to be ~12%.

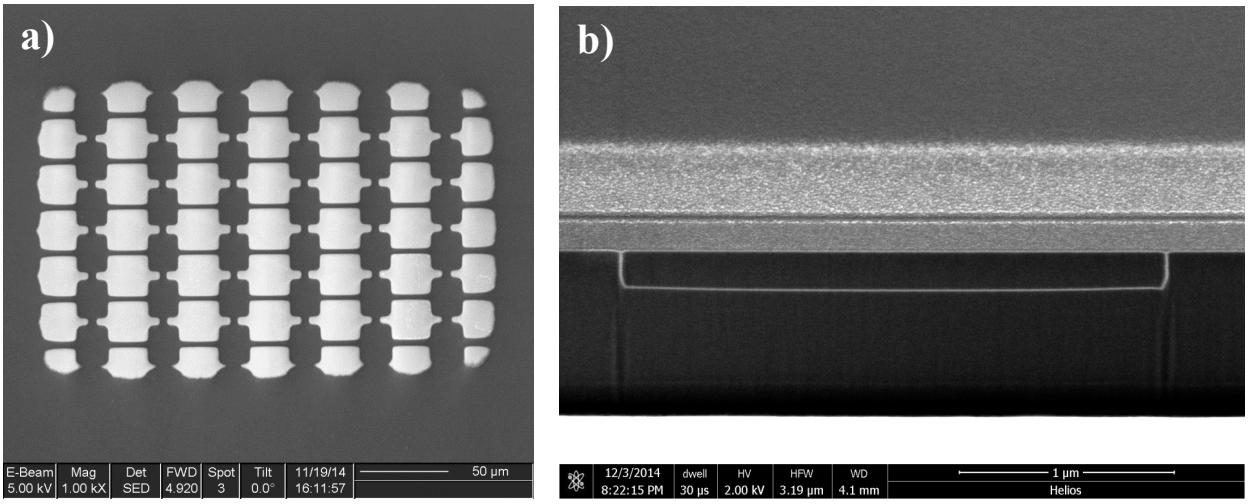

Figure 9. SEM characterization of dishing. **a)** Top-down view of a sandbox array polished after a failed polysilicon deposition. **b)** Cross-section of a 250 nm deep sandbox with regular poly. Bright layers are protective platinum deposited by electron and ion beam before FIB milling.

We did observe significant dishing on the first batch of wafers, which hand gone through the polysilicon LPCVD furnace while it had a huge leak, as it later turned out. The resulting film was transparent despite being over 1 um thick and therefore obviously not polysilicon, but likely silicon oxide, nitride or a combination of the two. The result of polishing it for testing purposes is shown in Figure 9a. After the furnace's tube was replaced to address the leak, all subsequent polysilicon depositions went as planned. On those wafers, we were unable to observe dishing using FIB/SEM, neither from the shapes viewed top-down nor in cross-section. An example of this can be found in Figure 9b. Note the upper part of the sandbox forming a seemingly perfectly straight line.

Therefore an effort was made to use a characterization method with higher resolution in the z-direction. This is offered by the CCI-HD optical 3D Profiler in SNF, claiming a vertical resolution of 0.1 A. Its disadvantage compared to SEM is a much lower lateral resolution of 1 um, so that the smallest feature we were able to image well were 7x7 um boxes. Since it requires a reflective surface, we used it to characterize wafers before the release of the sandboxes, however this should not affect the degree of dishing.

Figure 10 shows a profile of a 7 by 7 array of 100x100 wide um boxes that are 4 um deep; the two marked ones are examined in more detail in Figure 11. Note that all 2D and 1D profiles have the same vertical range for comparison purposes. The black squares are the release boxes which are out of range below the surface due to the overetch done when breaking through the alumina.

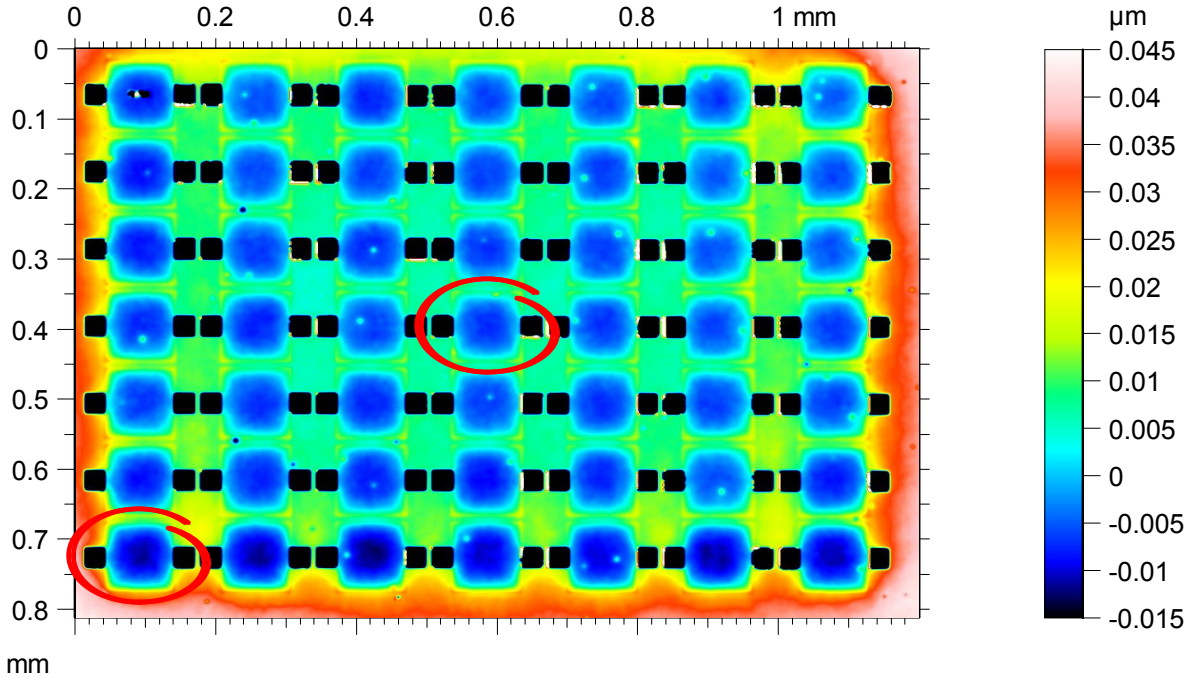

Figure 10. Surface profile measured with CCI-HD of an array of 100x100 um wide, 4 um deep boxes, after CMP and before xenon difluoride release. Circled boxes shown in detail in Fig 11.

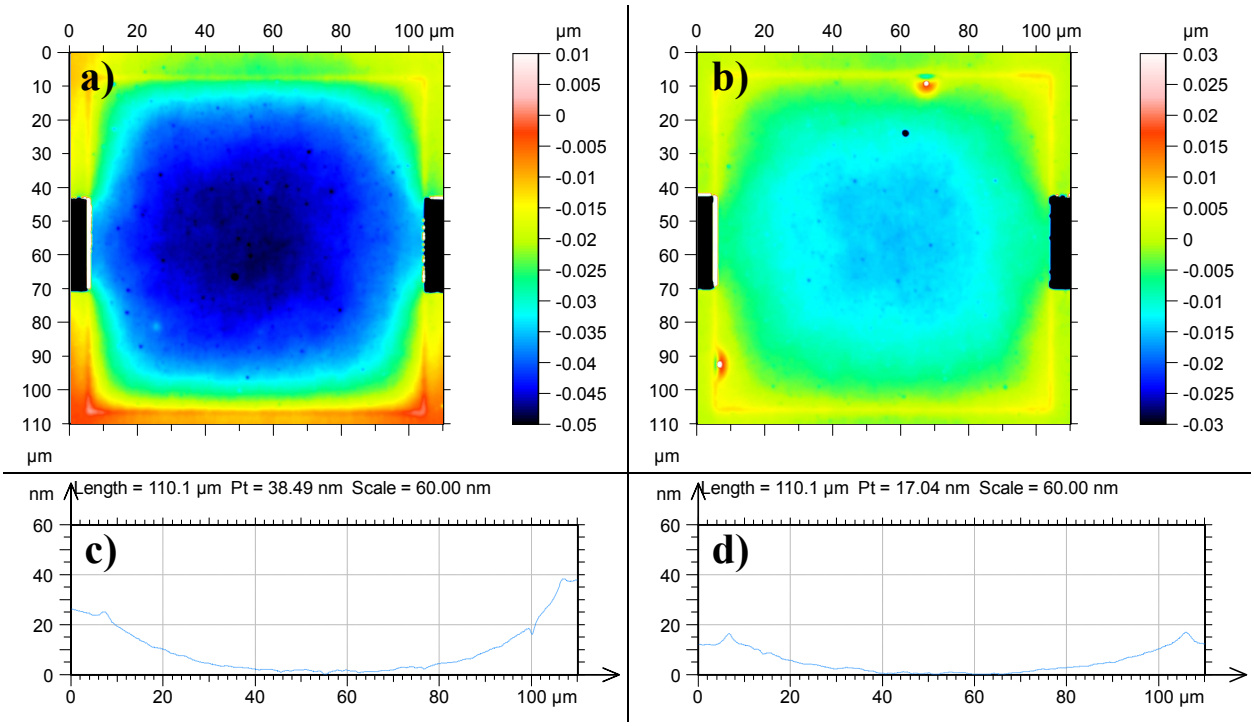

Figure 11. Higher resolution surface profiles from CCI-HD of the bottom left **(a)** and center **(b)** box of the array shown in Figure 10, with 1D profiles of their vertical center line in **(c)** and **(d)**.

The profiles show that there is a small but noticeable amount of dishing, leading to the membrane center being 38.5 nm below one of the anchors for the extreme case of the corner box (Fig. 11c). The edge boxes also show significant asymmetry with the height difference to the inner anchor only being 25 nm in this particular case. The profile of the center box is approximately symmetric as expected, and shows decidedly less dishing at only 17 nm.

To determine the dependence on the box sizes, location within the array and array spacing, the same measurements were done overall for 10x10, 25x25 and 100x100 um wide boxes, in each case for 7 by 7 arrays with large and small array spacing and for boxes in the corner and in the center of the arrays. The resulting trends are plotted in Figure 12.

As to be expected, the absolute surface deformation goes down for smaller box sizes. The spacing of the boxes in the array does seem to have an effect, with closer spacing resulting in less dishing, but this effect diminishes as the boxes get smaller, as does the difference between the corner and center boxes. In fact, for the 10x10 um boxes the height is very close to 15 nm for all cases. For the two larger sizes, however, the array center does have a significant advantage. To examine whether the same can be achieved with a smaller array, a measurement was performed 100x100 um boxes arranged in a 3 by 3 pattern, which is shown in Figure 13a. The corner boxes of this array give a similar result to the larger array at around 39 nm, while the center, at 20 nm, is noticeably worse than previously. This shows that the array size does matter.

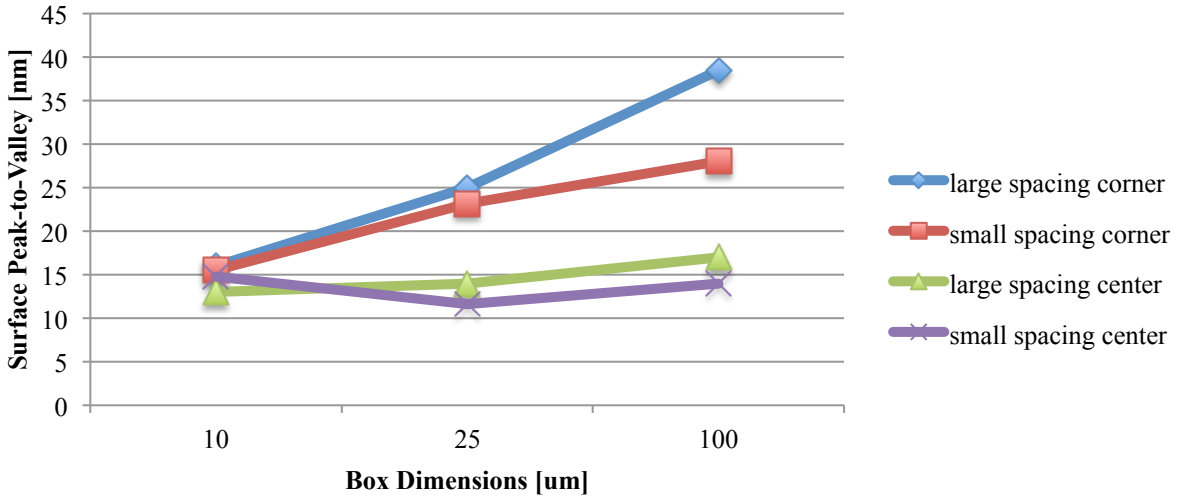

Figure 12. Dependence of dishing on box dimensions, array spacing and position in array. All data from measurements by CCI-HD with profile extraction done analogously to Figure 11.

Another comparison made was between arrays of circular boxes with 25 um diameter (Figure 13b) and square ones of the same side length. The results were very close, e.g. 13 nm for the center circular box, suggesting that the shape is largely irrelevant.

Since the dishing does decrease with array spacing, one could imagine it would go away almost completely if one had a large enough area that was almost entirely polysilicon. Such an arrangement not possible with the square or circular boxes investigated above because of their adjacent release boxes, but it does exist for bar-shaped rectangular boxes. The corresponding profile is shown in Figure 14. The structure does exhibit unprecedentedly low dishing at around 9 nm across each bar's width of 10 um, but it is still present, despite the single crystal silicon walls being only 0.5 um wide. One possible explanation would be that the 50 nm thick alumina on the sidewalls does play a role since it is mechanically harder than silicon.

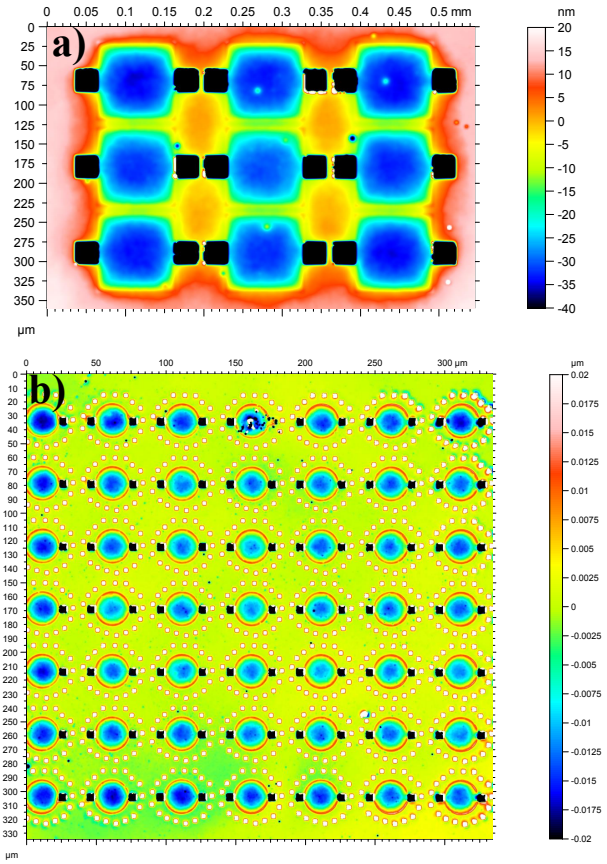

Figure 13. CCI-HD dishing profiles to compare smaller arrays **(a)** and circular boxes **(b)** to the previous results.

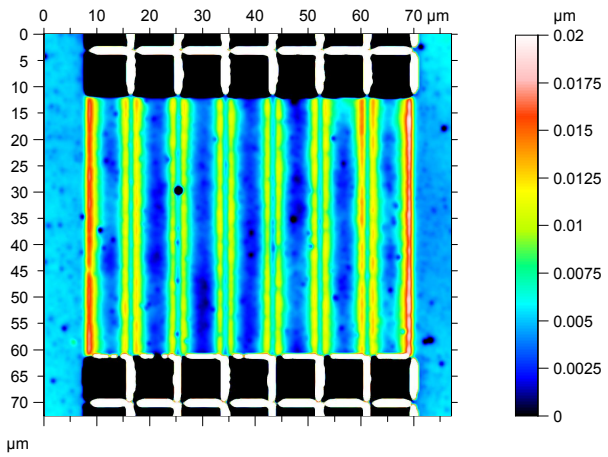

Figure 14. CCI-HD profile of bar-shaped rectangular boxes. Note the lower-range scale.

Overall, dishing does seem to be always present, but with the height differences generally being less than 0.1% of the box dimensions and it not having any effects on the shape (as to be seen in Fig. 11a for the case for the failed polisilicon deposition), it might not be an issue for most applications. In particular if the membrane material has tensile stress it seems likely that it would simply straighten itself out, eliminating the issue completely.

Besides dishing, another possible concern is a small elevation around the edges of the membranes, which can be seen very well at the corners of Figure 15, which shows the same profile as Fig. 11a in a 3D reconstruction.

Due to time constraints and limited availability of unreleased wafers, all the above measurements were performed on 4 um deep refilled boxes with 50 nm alumina sidewalls. Further investigations might look into whether shallower boxes or thinner alumina sidewalls lead to less dishing or a lower bead around the membrane.

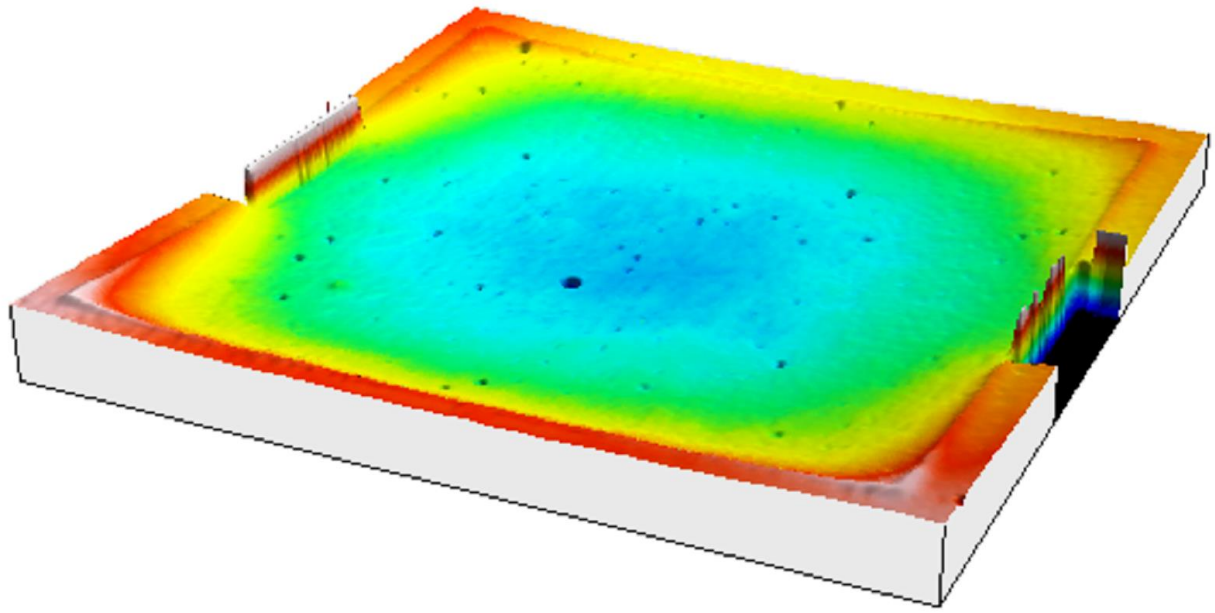

Figure 15. 3D surface reconstruction (not to scale in z-direction) from CCI-HD measurement of a 100x100 um box at the corner of a 7 by 7 array. Note the bead around the edge of the square, particularly visible in the lower left corner.

4.7 Nanoparticle Contamination

Another issue that we ran into was contamination of our wafers from the silica nanoparticles that are contained in the slurry. Although copious amounts of water and a sponge were used to wipe down the wafers, some residual nanoparticles always managed to remain on the wafers, as shown in Figure 16. This did not affect the release of the sandboxes but is likely to become an issue for potential future applications, so we needed to consider how best to remove all the nanoparticles. One attempt that we made was to etch them away in the uetch HF vapor etcher, taking advantage of the fact that anhydrous HF does not attack alumina. We used recipe 2 for 120 seconds which should etch a 100 nm thick SiO_2 layer and therefore be more than sufficient for the 50-100 nm wide nanoparticles. However, in SEM we always saw a few nanoparticles remaining. Another possibility of addressing this problem would be to flow in some DI water along with slurry (1 drop/s), which may be sufficient to prevent the slurry from solidifying on the wafer before we get to rinse it.

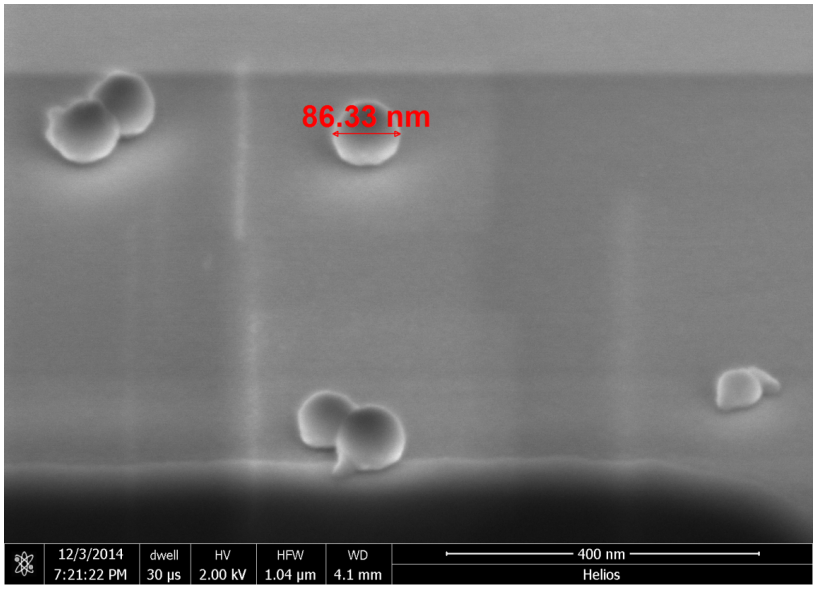

Figure 16. SEM image of residual silica nanoparticles from the CMP slurry, close to the sandboxes. We were unable to remove them using various methods.

4. Ultrathin Pinhole-Free Layers with Plasma-Enhanced ALD

The next steps after CMP are to deposit another alumina layer with ALD, deposit and pattern the membrane and then cap it again in alumina. The requirement for this second and third alumina layer is to hold up to the xenon difluoride release etch in order to protect the membrane. While this could be easily achieved by making them fairly thick, this time it is actually beneficial to go as low in the thickness as possible since both layers will end up being part of the membrane and it is undesirable that its mechanical properties are dominated or affected by the passivation layers.

So we tried to measure the etch rate of alumina in xenon difluoride in order to find the selectivity towards polysilicon. We deposited alumina films with different thicknesses on blank silicon wafers and measured the film with woollam before and after etches of varying numbers of cycles. In all cases, the thickness reported by woollam ended up being higher after the etch than before, which might be explained by added surface roughness or the deposition of some polymer.

Either way, it did not seem like the alumina was getting attacked, resulting in a selectivity towards polysilicon that is essentially infinite.

Therefore the only limiting factor in how thin the alumina layers can be made is the minimum number of cycles that give a pinhole-free film. Investigating this for various ALD films was the main topic of a previous EE412 project. Their testing process consisted of depositing the films on blank silicon wafers using thermal ALD in savannah and then subjecting them to several XeF_2 etch cycles in xactix to see if the underlying silicon was getting attacked. Their result for alumina was that 20 cycles did form a pinhole-free film while 10 did not.

Using this as a starting point, we decided to repeat this characterization for plasma-enhanced ALD in fiji2 and fiji3, since this might give different results from thermal, and to use closer spacings for the cycle numbers. The results are shown in Table 4. Interestingly, the limit is much lower for fiji3, lying between 8 and 9 cycles instead of 14 to 15 cycles for fiji2, as further shown in Figure 17.

All depositions were carried out at 200C using the standard recipes. Using exposure mode on fiji2 did not improve the results. The XeF_2 test etches were done at 3.0 Torr XeF_2 pressure, 0.0 Torr N_2, pressure 30 s exposure time per cycle and 30 cycles.

fiji2						10	11	12	13	14	15
fiji3	5	6	7	8	9	10	11	12	13	14	15

Table 4. Pinhole testing results for PEALD Al_2O_3 layers deposited in fiji2 and fiji3 with various cycle numbers, at 1A/cycle growth rate. Green: Pinhole-free, red: pinholes, empty: not tested.

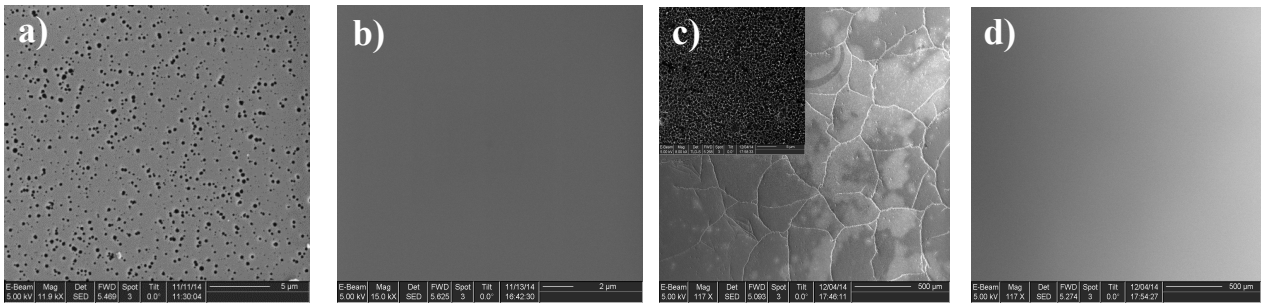

Figure 17. SEM images used as confirmation for pinhole testing results. **a)** 14 cycles Al_2O_3 fiji2, **b)** 15 cycles Al_2O_3 fiji2, **c)** 8 cycles Al_2O_3 fiji3, **d)** 9 cycles Al_2O_3 fiji3.

The above data shows that it is possible to synthesize pinhole-free layers with thicknesses of only 9 A, which is very attractive as an encapsulation for thin membranes in this process. To ensure that films as thin as this would hold up to the entire sandbox release etch, samples with 10 and 15 cycles of alumina ALD were subjected to a 10 times longer XeF_2 etch of now 300 cycles

with otherwise identical parameters as above. Again, there was no sign of the alumina or the underlying silicon getting attacked.

The significant difference between fiji2 and fiji3 can probably be explained with the history of the tools. Firstly, fiji3 has a higher general cleanliness standard as it only allows oxide depositions and restricts certain high vapor-pressure materials, despite also being gold-contaminated. Possibly even more important, it is much less heavily used than fiji2 and therefore its chamber walls should be in a better condition. In fact, at the time of this testing the built-up deposition in fiji2 had just reached a critical limit so that the chamber was scheduled to be sent out for cleaning just a few weeks after. Repeating the pinhole test in the refurbished state could verify this hypothesis and give further insights.

5. XeF_2 Etching

Creating the openings in the alumina layer was achieved by dry etching in PT-MTL using a BCl_3-based recipe. After resist stripping in PRS-3000 – despite not having metal on the wafer, piranha cannot be used since it etches alumina – the release is done in xactix using the same parameters as for the pinhole testing, i.e. 3.0 Torr XeF_2 pressure, 0.0 Torr N_2, 30 seconds cycle time and 0 seconds cooling time between cycles. To get an idea for the required cycle number for the release, a test was done on a wafer without etch sandboxes that was just coated in alumina and then had the same openings etched into as the device wafer. The lateral extent of an etch using the above recipe for 30 cycles for various sizes of the square openings was measured using SEM, the results and an example image are shown in Table 5 and Figure 18, respectively.

Opening size [um]	Lateral etch distance [um]
1 x 1	7
2.5 x 2.5	10
4 x 4	12.5
7 x 7	15.5
27 x 27	20

Table 5. Etch distances for 30 cycles of XeF_2 etching using the described xactix recipe.

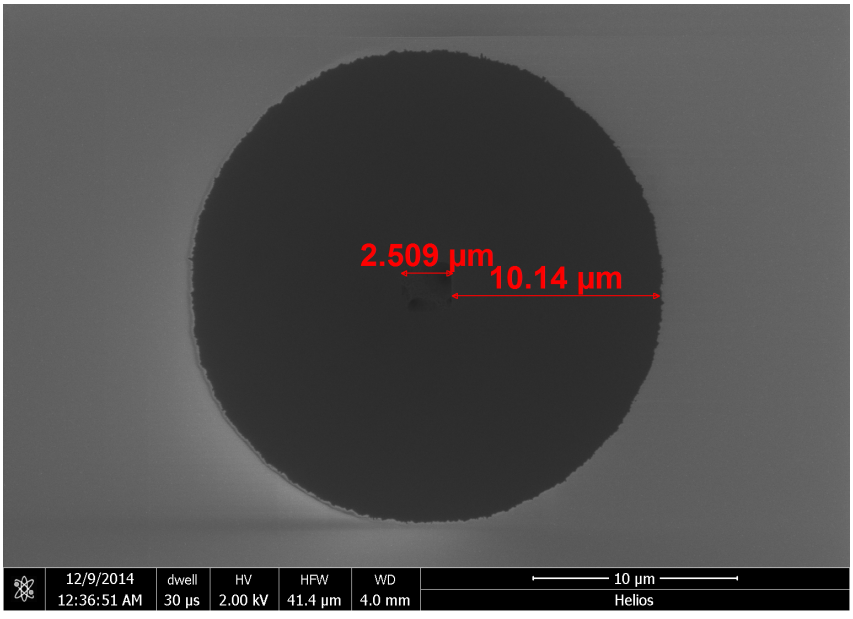

Figure 18. SEM characterization of the lateral etch distances for 30 cycles of XeF_2 etching using the described xactix recipe.

From these figures, we calculated that we would need approximately 85 cycles to release the 100 x 100 um squares. As a safety margin, the actual number used was 100. In addition, several partial releases with smaller cycle numbers such as 30 were done on other wafers. Overall, the release went well with a typical result shown in Figure 19. Several other SEM images of released membranes can be found in the appendix. Usually, all membranes survived, as can be seen in the image with the array. From the fact that the alumina of the membrane, having a starting thickness of 5 nm, remained after the etch while 50 um were undercut into the silicon, we can calculate that the Al_2O_3 to Si selectivity of the XeF_2 etch must be at least 10,000:1.

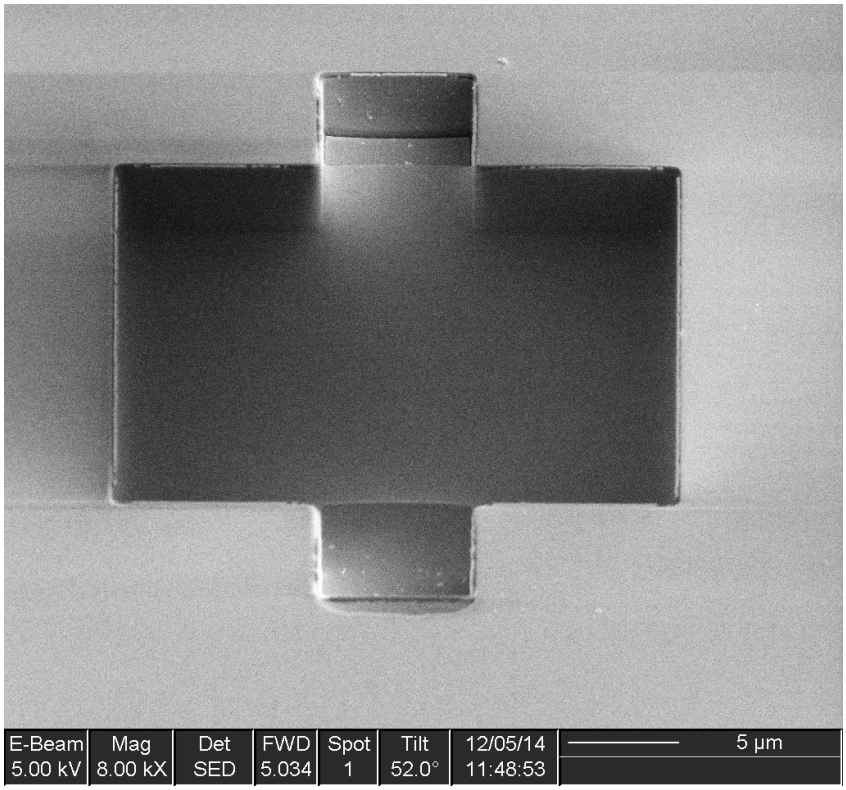

Figure 19. SEM image of a released sandbox, 10x10 um wide and 4 um deep with a 5 nm alumina membrane.

6. Other Process Issues

Several concerns that came up during processing were related to the ASML alignment marks. The first was whether they would even survive the CMP process. Since the XPA marks are by default etched only 120 nm deep into the wafer, it seemed distinctly possible that they would get completely removed just from not hitting the polishing endpoint perfectly, particularly on wafers with 4 um deep boxes, where a mere 3% overpolish would already be problematic. To alleviate this, we tested whether deliberately making the marks deeper than recommended would still work, and found that the tool was able to detect them for every case, i.e. 250 nm, 1 um and 4 um depth. Therefore for the second batch of wafers we combined the zero and first layers. This ensured that the alignment marks would always survive if the boxes survived.

The other question related to the XPA marks was whether the contrast between polysilicon and silicon would be sufficient, since the marks also get flattened out during CMP. Again the testing was successful in that ASML was able to find the marks without manual intervention, no matter

the trench depth or alumina thickness. However, this is likely to change when potentially depositing a reflective layer in future fabrication steps and might require a flood exposure to recover them.

The other smaller issue with the process as originally proposed is related to the protection of the membrane material from the xenon difluoride etch. In case the membrane is not patterned lithographically such that it is not present in the areas where the release holes are, then, as shown in Figure 2 step 8, parts of it will actually be exposed to the etchant on the sidewalls of the breakthrough. This might not be a huge issue in many cases, unless the membrane is fairly thick or made from a material that etches very fast in xenon difluoride such as silicon. Nevertheless, there is a way to address this without adding another lithography step: After the breakthrough etch and before the release, one can coat the wafer again in a few nm of alumina and then do a timed anisotropic etch aiming to remove that same thickness. This will open up the release holes again at the bottom but due to the anisotropic nature of the etch, the sidewalls will remain passivated.

Appendix A – SEM Images of Released Sandboxes

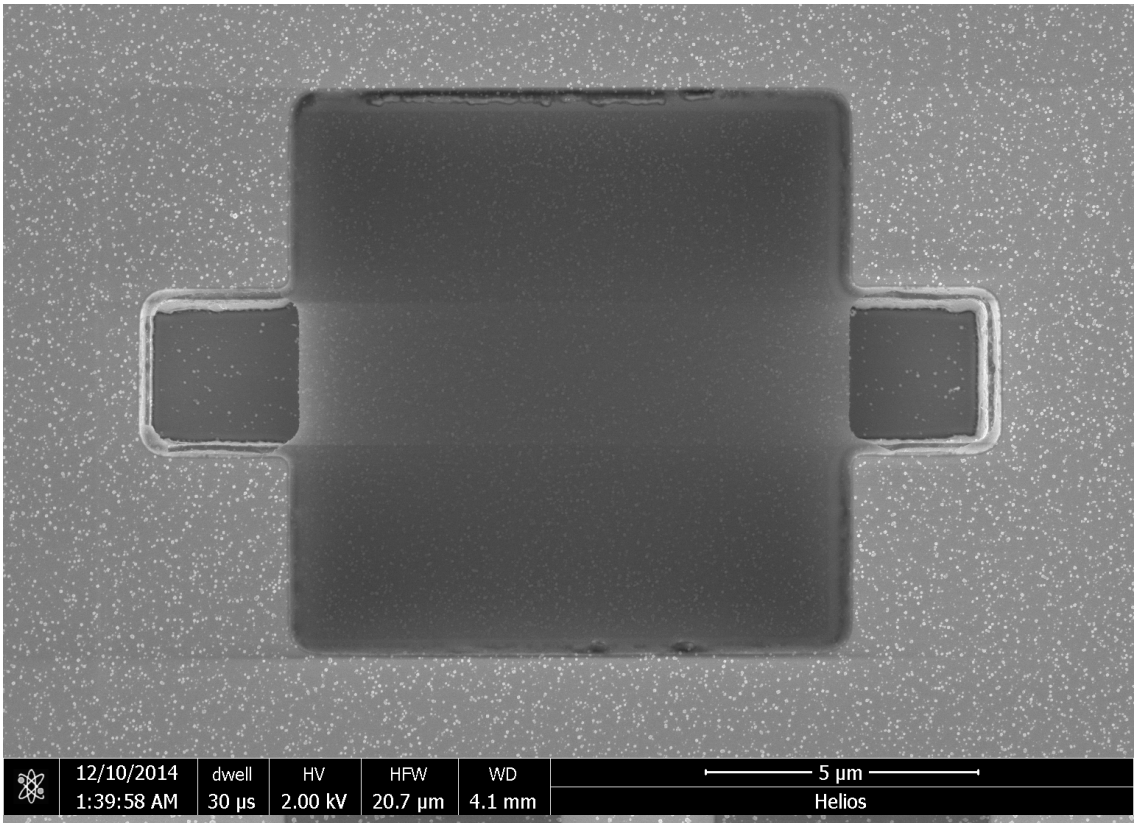

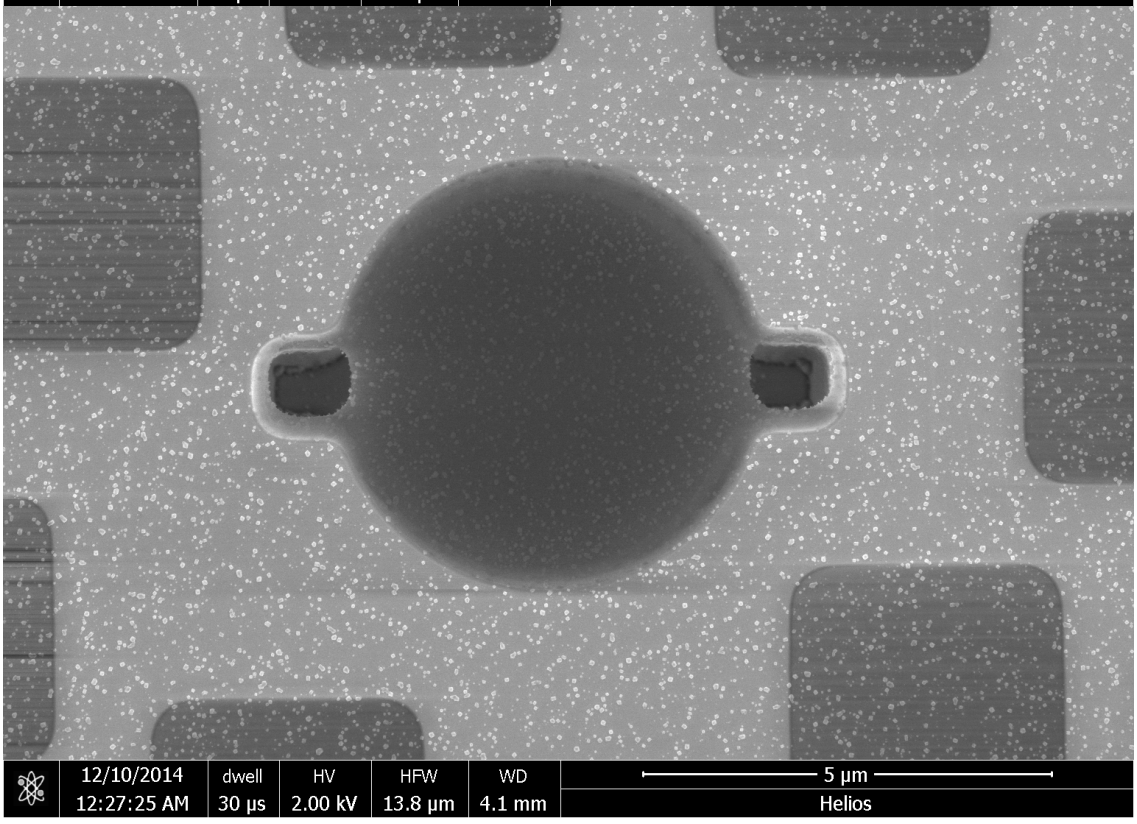

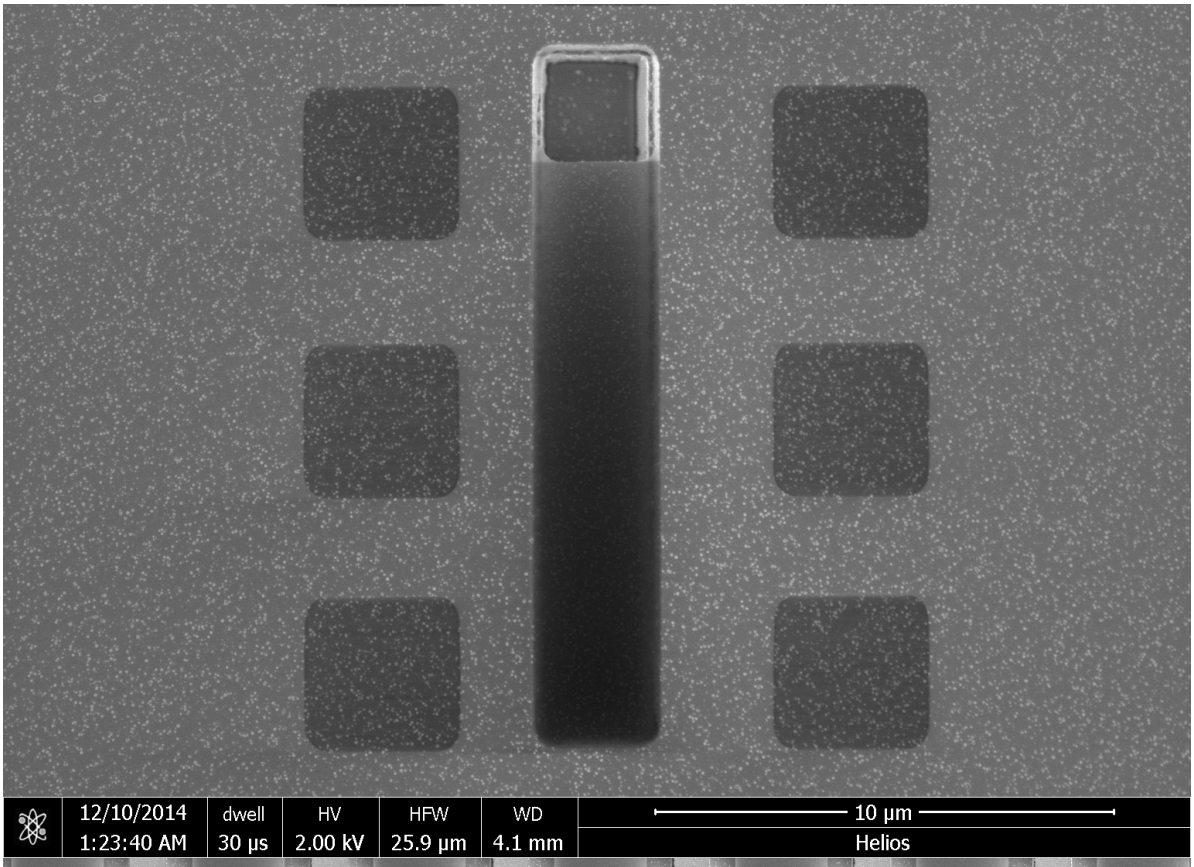
12/10/2014
1:23:40 AM
dwell
30 µs
HV
2.00 kV
HFW
25.9 µm
WD
4.1 mm
10 µm
Helios

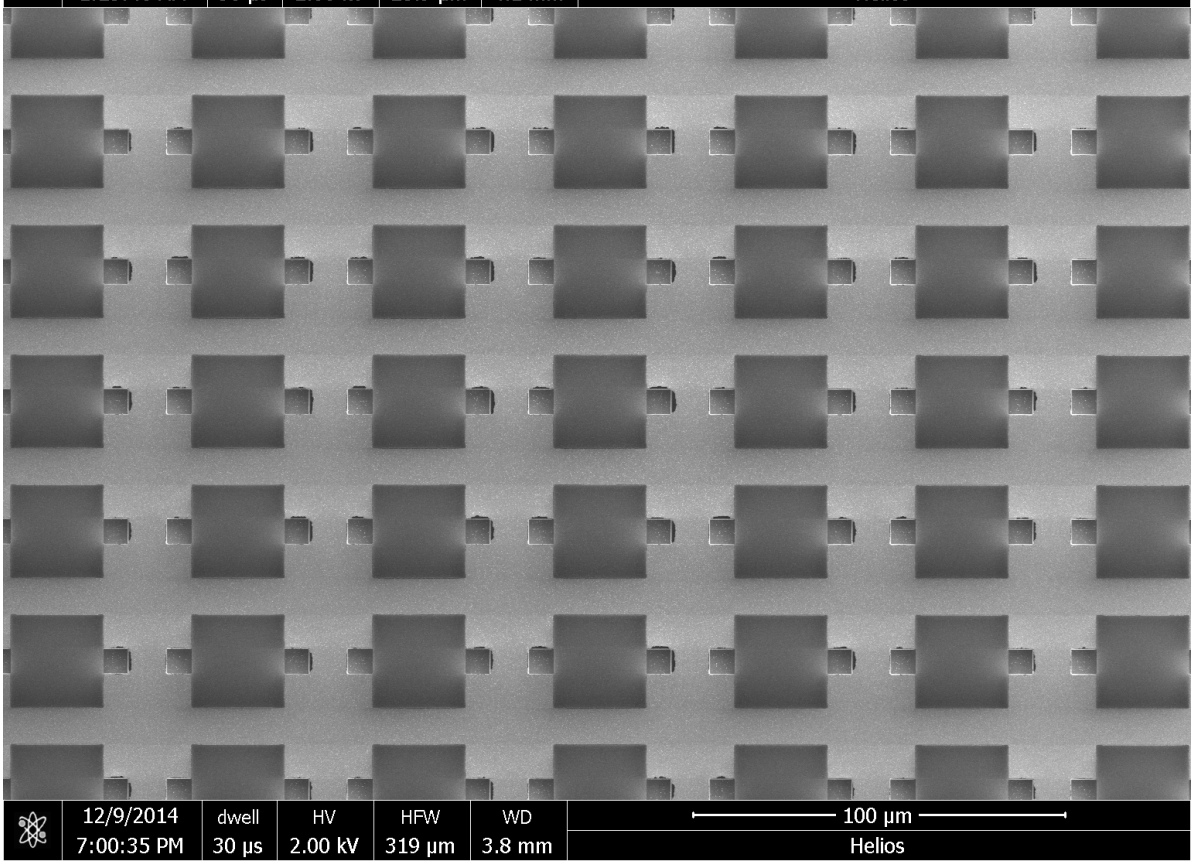
12/9/2014
7:00:35 PM
dwell
30 µs
HV
2.00 kV
HFW
319 µm
WD
3.8 mm
100 µm
Helios

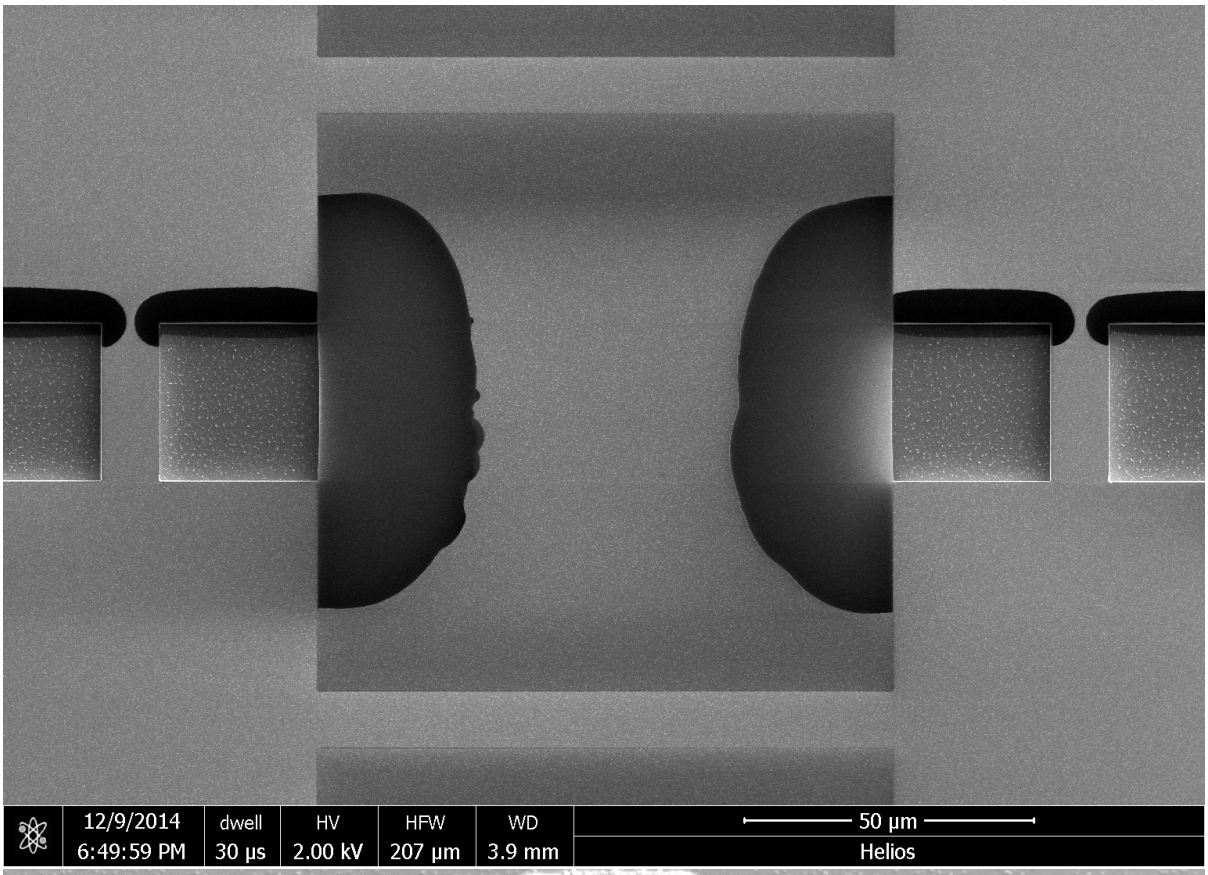
12/9/2014
6:49:59 PM
dwell
30 µs
HV
2.00 kV
HFW
207 µm
WD
3.9 mm
50 µm
Helios

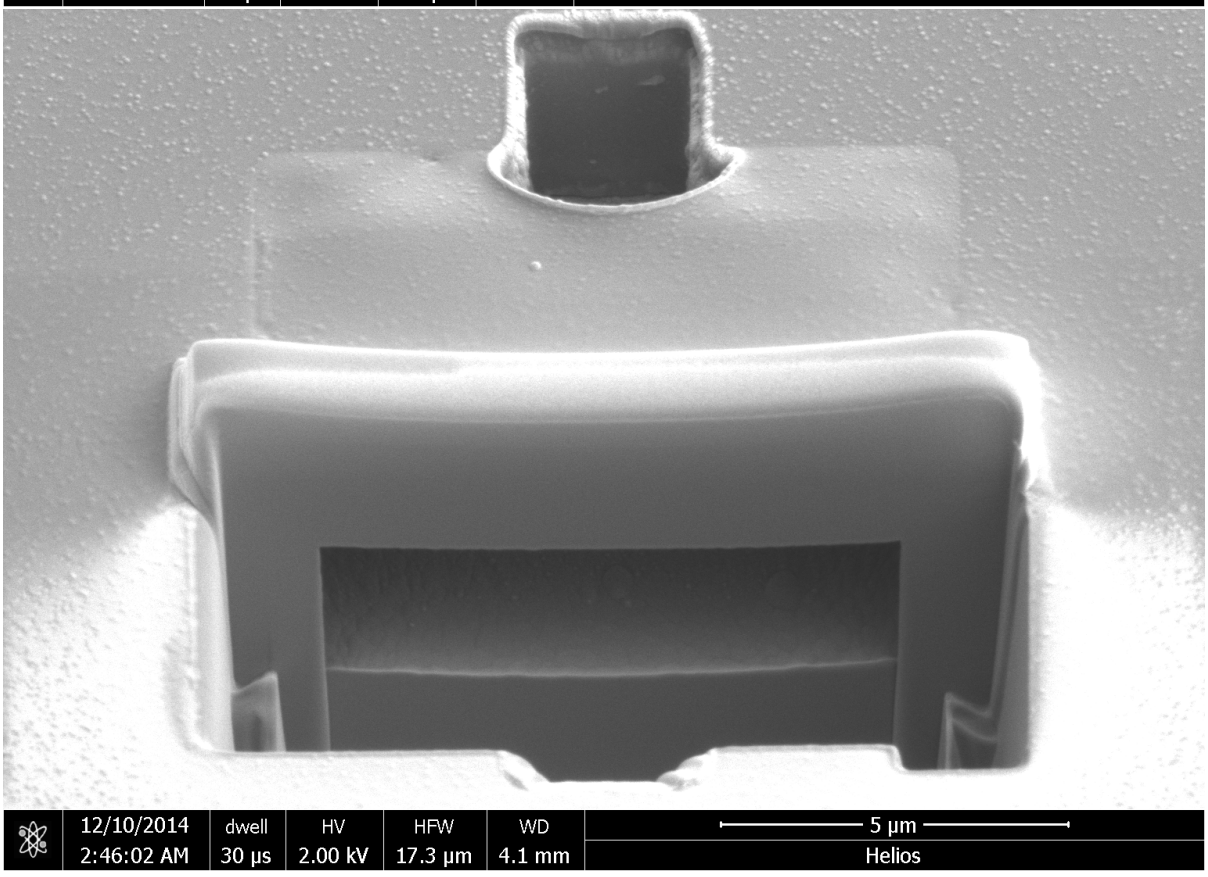
12/10/2014
2:46:02 AM
dwell
30 µs
HV
2.00 kV
HFW
17.3 µm
WD
4.1 mm
5 µm
Helios