

MITOCW | Lec 22 | MIT 2.830J Control of Manufacturing Processes, S08

SPEAKER: The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or to view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: Well, today I'm going to give the fourth and final case study of the course. And everything I'm going to talk about is my thesis research in one form or another. So this is work that's definitely in progress. And in a lot of cases, we haven't really drawn conclusions yet. So it'll be quite rough and ready. But I hope it'll be informative and raise some questions for you.

The focus of my thesis research is-- well, broadly speaking, finding approaches to manufacturing microfluidic devices using polymers. And for many of you, I think microfluidics will be something you know at least a little about. Essentially, microchips that manipulate very small volumes of fluids to perform various experimental, or diagnostic, or other engineering tasks.

And we were interested in developing processes and choosing materials that will help us make these things very cheaply so that they can be used at the point of care in the case of diagnostic devices, and so forth. Now, one particular process that's of great interest for this is the imprinting of thermoplastic polymeric layers. Hot embossing is what we're going to call it. And for a start, let me say why polymers are interesting for making these sorts of devices.

I think I see three main reasons. Firstly, they're cheap compared to silicone or glass, which are other materials that you might think of as being obvious candidates to fabricate these microscopic fluidic channels, and valves, and pumps, and so forth.

So polymers are cheap. Second thing is, they're often transparent. And that's really important. When you have a biological sample inside this device and you're trying to look at how it behaves in response to certain stimuli, either you're exposing it to a dye that will make cells with certain properties fluoresce. You need to see that fluorescence. Or you're prodding a cell somehow to see whether its stiffness gives you information about its disease state, something along these lines.

Anyway, but transparency is an important property. And thirdly, a lot of plastics are mechanically tough, which makes them highly suitable to use in the field at the point of care, in doctors' surgeries and so forth. So these are the three reasons why polymers are interesting. And imprinting them is interesting as a process, because that itself has the potential to be quick and cheap.

You can imagine a cycle imprinting a thermoplastic polymer occurring in around a minute. And you can choose what kind of size workpiece you machine. You could machine anything from the size of a single microfluidic chip, all the way up to continuous reels of polymeric film.

And indeed, the reel to reel printing approach is something that's been used in industry already for quite rudimentary devices. 3M has a patent on hot embossing reels of thermoplastic polymers. And so it's a pretty flexible process. And microfluidics aren't the only sorts of things you might think of making with it.

A lot of optical devices could be made this way. You could think of imprinting refractive elements. Or actually 3M's patent is so that they can imprint the reflective films that are laminated onto the front of road signs. Imprinted into those films are many parallel, V-shaped grooves that act as corner [? key ?] reflectors so that the street signs reflect light incident from any angle. And so those sorts of optical applications are also pretty crucial.

But anyway, let me just describe how the process works. What I'm illustrating here is not a role to role printing process, but something where the workpiece is about the size of a wafer. So tens of millimeters in diameter.

And first thing we need is a hard stamp, a tool that's been microfabricated using some other process. This tool might have been etched from silicon. It might have been electroplated with nickel. It might itself be a polymer, but one that doesn't soften at the processing temperatures. Elastomers like PDMS, polydimethylsiloxane are actually quite attractive as stamp materials.

So we have this stamp. We have two heated plates to heat the stamp and the workpiece, and a means of applying a load normal to the surface of the workpiece. Usually what happens is the material is heated above what's called its glass transition temperature.

And for anyone who isn't familiar with this, thermoplastics essentially are composed of an entangled network of polymer molecules. And there are one or more temperatures for a given polymer at which those entanglements loosen. And the chains can slide past one another. And the stiffness and the effective viscosity of the material around that temperature falls by several orders of magnitude.

So above that temperature, it's highly flexible, easily formed. Below that temperature, it's much stiffer and useable as a device. So we go above the glass transition temperature to soften it. How far we go above is an important question to answer. That's a crucial parameter to choose.

We then apply a load. That load is going to be ramped up over some time that we have a choice about. That might be an important parameter. We hold that load for some time to allow the material to flow, to fill cavities in the stamp. And while the load is still applied, we cool down usually, cool the material to below its glass transition temperature, and finally, remove the load. And often, that removal step has a lot of technical challenges involved in it, because of differential thermal contraction of the workpiece in the tool.

So that gives you some idea that there are at least four parameters that we need to decide upon, temperature, a load, and two times associated with that load. There are probably more parameters that we're interested in.

Now, at this point, I should also mention a rather similar process that has quite different applications. And that is thermal nano imprint lithography. This is, essentially, hot embossing, but done on very thin layers of polymer, thinner than a micron. And these layers have been spun onto a much harder surface, usually a silicon wafer.

And the beauty of this process is that it allows sub micron features to be transferred to the wafer. This is, I think, a really exciting process. It has the potential to revolutionize lithography and semiconductor manufacturing, because the resolution with which features can be transferred is not limited by the wavelength of light.

Usually you would pattern a photo resist layer with some kind of projection, optical projection system. And we've really got to the stage where critical dimensions of transistors are making that a big challenge, making the optical lithography systems cost tens or hundreds of millions of dollars.

So in this case, we have this spun on polymer layer. It's dissolved in an organic solvent, spun on the solvent. It evaporates. And then the spun on layer is a thermoplastic. It can be softened. You press the the nano fabricated stamp into the wafer. And ideally, squeeze all the material from below the individual features of the stamp.

Now, in reality, that's not possible. And there is a small residual polymer layer that's left there. So you can imagine, once we've removed the stamp or the mold, we have some extremely thin regions of polymer. And then imagine that we want to pattern, to etch away material in the wafer underneath. The plasma or the etching solution to which the wafer is exposed will quite readily break through that residual polymer layer and remove the material.

So these processes are similar. In many ways, they actually often use the same polymers. But the boundary conditions, the mechanics are rather different.

And so I'll talk a little bit about nano imprint, but mostly the results that I'm going to show you are to do with micro hot embossing, where the workpiece is substantially thicker than the diameters or the depths of the features that are being patterned.

AUDIENCE: [INAUDIBLE] about creating the mold, during the mold removal process, do you have the mold walls getting [? dilated? ?] Because the pattern is getting taken off and it might have a [? dimension ?] [INAUDIBLE]?

PROFESSOR: Yes, that was an excellent question. You've hit on one of the most important challenges of [INAUDIBLE] embossing. And [INAUDIBLE] was asking, are there problems with dimensional stability of the walls, the side walls during de-molding? And I'll show you some results that relate to that. So yes, good point.

Anyhow, as I already mentioned, there are several parameters to do with the process that we need to pick, temperatures, load times, loading durations. The range that those parameters can take will be limited by the capabilities of the machine we have, the physical properties of the polymer we're trying to pattern, what temperature does it start to burn, and so forth.

What about the material underneath? If we're doing nano imprint lithography, are there temperature constraints that the wafer underneath has to observe? So that's what constrains the process. And, of course, we might also be making engineering decisions about what type of polymer we're going to use to achieve a certain end.

And indeed, we might have some flexibility in the pattern that we are going to emboss. We might think about designing a pattern with an eye to manufacturability, something that will be easier to imprint with less variability. And that is one of the questions that I'm trying to answer. Mohammed.

AUDIENCE: [INAUDIBLE] how do you [INAUDIBLE] the stamp [? control? ?]

PROFESSOR: Right, well, the stamp can be made in a variety of ways that are usually traditional microfabrication. So one particularly easy way of doing it would be to use a process called deep reactive ion etching. And here I'm talking about micro embossing, where the features are several microns to many microns.

And essentially, this is just etching very deep trenches into a silicon wafer using a plasma process. I could go into great detail about how that works, but essentially, it's a fluorine based chemistry. The precursor gas is SF₆. And you build up these very deep, narrow trenches in the silicon wafer by a series of steps.

So in fact, you would start with the silicon with a photoresist mask or something. You'd do a brief fluorine based plasma etch that would remove small amounts of material. That's roughly isotropic. So if you want vertical sidewalls, you've got to do something else. And the something else is depositing a polymeric passivation layer.

So you change the gas in the chamber to C4F8, which creates what's essentially a teflon coating that passivates the sidewall of the trench that's developing. Then you go back to the SF6, build up a new notch, and do this many times, alternate many times. So whatever way you pattern the stamp.

Silicon actually, as we will see later, for de-molding reasons, is actually a lousy material to use for the stamp. So then you might go one step further, electroplate nickel into that etched wafer, peel off the nickel, and use that as the stamp. Much tougher.

Or you might use more novel materials, things like metallic glasses, which are alloys with an amorphous structure that are both very hard and very tough. They would be great materials for a stamp.

AUDIENCE: So what you mentioned-- you said that you're not limited to wavelengths of the light in this process. But however to make the stamp, then you are limited to--

PROFESSOR: Yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Well, yes.

AUDIENCE: [INAUDIBLE] limitation exists.

PROFESSOR: So what I just described here was for micro embossing. I'm sort of describing two processes in parallel that have great similarities, but not in terms of scale. So this would be for features that are a micron or larger.

If you want to make sub micron features, then you're quite right. You need a process that will let you make a stamp with sub micron features, and electron beam lithography is a great candidate for that where you actually would-- you would start with a silicon wafer with a 100 nanometer thick layer of a radiation sensitive resist. Then you would just actually steer a focused electron beam across the surface, scan it across the surface in the pattern that you wanted. And that will give you 5 nanometer resolution or better. The trouble is, it takes ages because it's a serial process. You're scanning the electrons over the surface.

So once you've made the stamp, you've paid however many \$1,000 for the stamp. You can stamp it in in a minute into as many substrates as you like. And that's really the reason for nano imprint.

OK, so keep both of those processes in mind. But a lot of the assertions I'm going to make from now on are all about the micron scale embossing. So as I said, there are lots of different parameters we can choose to do with the process, to do with the material we're embossing, to do with the pattern that we're embossing. And the overall mission that I suppose we have is to try to provide tools for people who are using this process to achieve the desired microstructure with as little time, energy, cost as possible. So if you're processing a certain material, can you pick a pattern that will take 30 seconds to replicate instead of a minute at the available processing conditions?

So that presupposes, of course, that we have some specification on what satisfactory replication is. The most obvious specification would be that every cavity in the stamp has been filled with polymer, even the narrowest ones. That would be a good specification to use.

Another good specification that's really important for nano imprint lithography is to make the thickness of those residual polymer layers very uniform. And that's important because it does take some time for the etching process to break through that residual layer.

And if the residual layer varies in thickness, the time taken to break through it will be different. And then you'll get different properties of the underlying structure. So let's say you've imprinted this 100 nanometer thick polymer layer. And on this side, the residual layer is really thin. And on this side it's twice as thick.

You then expose this with a fluorine plasma. You're trying to etch the silicon underneath here. You're going to end up with a deeper trench here than you are here because of the time it took to get through that.

So coming up with ways of designing the stamp so that the residual layer ends up being as uniform as possible is a really pressing challenge actually for this lithography technique to be adopted.

Anyway, so what we're all about is simplified modeling tools that can be used as part of the design loop that are computationally efficient enough to give results in a reasonable amount of time, and that represent reality well. So that implies that we need some way of characterizing the tool that we're using, the materials that we're using efficiently so that we can get information about the physical properties of these things with a minimum of fuss.

Now, I'm just going to split the problem up into three notional length scales so that we can have a clearer picture of what's going on. And at the largest length scale you can think of variations in the quality of the embossed features from one side of the substrate to the other. And that might occur because the two heated plates that are compressing the stamp and the substrate are not perfectly parallel. So the pressure's greater on one side. It might be because one of the plates is bowed or wavy. Or it might be to do simply with the mechanics of the substrate itself.

If that substrate is behaving as a viscous fluid, say, then as you compress that substrate, that flat, polymer substrate, you're going to expect to see, in fact, a parabolic distribution of pressure across it. So there are various things to consider there.

And then at what I'm going to call the device scale, there are these pattern dependencies where the arrangement of features on the stamp can have a profound effect on how well defined the features end up being. And what I've sketched here, which is really just a sketch, it suggests if there are features that are slightly more closely packed together, then they may fill earlier, or maybe sometimes later than regions that are less densely packed. And if that's the case, then maybe you want to put some design rules in place that will constrain the variation of density of features on the stamp.

Finally, we've got at the smallest scale effects to do with individual features. So if an individual gap is too small, say, that's going to be very hard to force polymer into it. And that could be a showstopper. So we've got these three different length scales.

I've alluded to these sorts of problems, bowing of the patterns, non parallelism of the patterns. They relate to the substrate scale effects. This shows one particular substrate that we embossed. This is a square of PMMA, which is acrylic, which is about 100 millimeters square. We embossed with a wafer that had been uniformly patterned with some silicon posts, roughly 100 microns in diameter. And we just use this as a probe to look at our embossing machine and see how uniform a pattern could be produced.

Now, aside from the fact that huge chunks of the silicon wafer broke off during de-molding-- and that's a good illustration of some of the problems that we face-- we were able to look at the topography of the pattern at different locations on the wafer. We did that using an interferometer, scanning from above, measuring the topography. And so we were able to see various things, that the emboss depth was higher on this side than on this side. Then near the edge of the wafer the embossed features tailed off in height.

And so these are the sorts of things that we need to come up with ways of predicting so that we can counteract them. Anyway, what I'm going to concentrate on for the rest of this talk, in fact, is pattern dependent non uniformity, which is, in many ways, the most challenging to deal with. Effects to do with the edge of the wafer become important when you're trying to eke out one or two extra devices. You're trying to save a little bit of material. But if the actual design of the device-- there may be many devices across a wafer. If the design of the individual device is faulty and it will never replicate properly, then you're in real trouble.

So trying to get an understanding of what patterns will form well and what patterns will not is crucial. And so although we are working towards a unified way of dealing with this, I'm just going to talk mostly about the pattern dependencies now.

So I've mentioned all of these different questions that need to be answered, factors that need to be chosen when you're designing a new process. And some of them are going to be continuous variables, and some of them will be discrete choices. And it's not absolutely clear which will be which.

I've really split the decisions up into three categories, the decisions you have to make about the pattern you're going to emboss, what are the shapes of the features? Rectangles, circles, triangles, how big they are, and how they are oriented on the stamp?

And that can be important when you're starting to think about radial thermal contraction of the parts, or about anisotropic material properties. So there you've got-- you know, size is a continuous variable, and the feature shape you might think of as being a discrete choice.

When it comes to the substrate itself, then you've got to decide on the material. Various materials soften at different temperatures. Some of them exhibit a range of temperatures over which they behave rather like a rubber. Others don't really have much of a rubbery region. They just flow very easily as soon as you're above the glass transition temperature.

So that's a decision that has to be made. And in micro embossing, often you're picking between a variety of off the shelf substrates, PMMA, plexiglass, polycarbonate, materials like this, where these are technical grade materials. They're molecular weight. In other words, the average length of the polymer chains is not necessarily well documented or well defined.

So in that sense, your material choice might be a discrete decision. On the other hand, if you're spinning the polymer onto the wafer, and you're dissolving it, then you do have the option of choosing the molecular weight, and therefore, defining the viscosity as essentially a continuous variable. So those are things that might need to be thought about.

And, of course, there's the thickness. Either how thickly do you spin this resistor? And that, a lot of research in the field demonstrates the optimal thickness is a strong function of what pattern you're trying to emboss, how many voids there are in the stamp. So the thickness of that might be important, and also the thickness in micro embossing where the substrate is a millimeter thick, say. That's also a decision that has to be made. So and the third category, of course, is the process parameters, the temperature, pressure, hold time, and so forth.

Now, there's an awful lot of variables there. And if one were to just take a purely empirical approach and try doing a full factorial or a fractional factorial exploration of the space, it would start to look pretty much like a nightmare. And luckily, we don't have to stumble around in the dark, because we do already have some reasonably good physical intuition about how these materials behave.

And, of course, there's a huge research field to do with characterizing these bulk materials, building theoretical models of how they behave, fitting data to those models. And I'm sure mechanical engineers here will be familiar with this sort of way of describing polymers in a simplified manner, where you build up a model. In this case, this is a one dimensional model that sums up the viscoelastic behavior of the polymer using a combination of springs, elastic, nondissipative elements. And these elements here are dashpots, which are essentially viscous components and damp. Yes, exactly, they dissipate energy.

So the various stiffnesses and dissipation factors for the springs and dashpots may be non-linear functions of temperature. They may be non-linear functions of strain rate, and, in fact, almost always are. And so there are many theses that have been written trying to characterize these materials.

And building on those experimental approaches are simulation models. There are finite element models of these polymers that have been lovingly built up. And so with enough information about the polymer with a detailed model of the stamp that you were thinking of embossing, and knowing something about the dimensional tolerances of the machine you have, you could do a simulation and perhaps predict what was going to happen.

The trouble is, it would probably take years. Because when you start having microscale patterns that you're trying to emboss, the computational burden of using these full nonlinear, finite, deformation polymer models becomes immense. And that's great if you have one particular device that absolutely must be perfect.

However, what we're interested in doing in our research is finding some sort of middle ground between the empirical approach and the really rigorous, thorough, theoretical approach that's based on a highly controlled set of experiments. And so the idea is to find approximate descriptions for the material that are efficient to run, efficient to simulate patterns, wafers that have thousands of features across them.

So let's start to think about how we can build out the bare bones of a model that are going to inform the experiments that we do about this embossing process. Now, what I've done here-- and I won't go into laborious detail about, but essentially, this is just building intuition about one particular polymer. This is a model for polymethylmethacrylate, plexiglass, acrylic, call it what you will, which a student at MIT has recently written a thesis about, doing compression experiments of it, building up a highly faithful model of it.

But there's a one dimensional implementation of that model that we can run. We can look at what happens when you give it various load profiles over time and look at the stresses and the strains in each of those components of the model over time. So what I've done here is I've just done two what are essentially thought experiments. The left hand column of graphs shows what happens if the material is heated above its glass transition temperature and then given a 1 and 1/2 megapascal compressive load which is held for 10 seconds. That load is maintained while the temperature is reduced below the glass transition.

The glass transition here is about 105 degrees C. So then we looked at how the model responded to that. There's a compressive strain, of course. There are stresses built up in the left hand branch of that model, this spring and this dashpot. And there are stresses in the right hand side as well.

And essentially what this shows is that most of the deformation is frozen in place, frozen by this cooling below glass transition. So when the load is removed, this amount of strain remains, this P amount here.

However, if you don't cool down before you remove the load, and that's what is shown on the right, then we see this recovery. We see the material springing back, almost to its original shape. And so there's a very small residual compressive strain in that case. And just looking at a distance, from a distance at this simulation, which is done at a particular temperature, 140 degrees C, essentially the material is behaving in a springlike way predominantly. Most of the deformation is recovered without cooling it to freeze that deformation in place.

And you could run similar thought experiments at different temperatures, different strain rates with the model and get an impression of how much of that deformation is recovered when the temperature is kept at its elevated level when you remove the load. And so loosely speaking, if most of the deformation is recovered at high temperature, we would describe the material as being rubbery over here.

This is a plot essentially of the ratio of that strain to that strain, P over Q, as a function of temperature and peak compressive stress. If that ratio is large, we call it rubbery. If the ratio is small, in other words, if there's been a lot of plastic deformation that doesn't spring back upon unloading, then we would sort of classify the material as being in a glassy state.

And sometimes glassy is useful to you. Sometimes rubbery is preferable. So in any case, this sort of intuition building is helpful to us in that it begins to give us a starting point for a physical model that will let us choose what experiments we're going to do, and will let us build quick, efficient simulations.

And so for that particular temperature that I looked at in detail, if you wanted the simplest possible model of the polymer, you might just think of making the model a linear, elastic model, where the Young's modulus was a function of temperature. And there are many shortcomings of that.

If you hold the load for long enough, then there will of course be plastic flow that's permanent. And the material is not linear elastic. It changes its stiffness as the strain increases. But we might as well start simple and see if it works. It gives us information about what types of features replicate and what do not.

So that's what we initially started with. And the idea is that increasing the temperature reduces that modulus by three or four orders of magnitude. We form the shape, and then we cool down to freeze that topography in place.

And what this is going to let us do is predict, for an arbitrary stamp design, what the shape of the embossed part will be, what its surface topography will be for a chosen temperature and pressure for embossing.

And here's the computational approach that we decided upon. And this is really quick to run. The idea is that we want to describe how the surface of the material responds to a point load. And if the material were linear elastic, and if it were substantially thicker than the dimensions of the features of interest, then the topography in response to a point load would effectively go as one over the radial distance from the point where the load is applied. And that's a standard contact mechanics result.

But we're talking about computing approximate topographies. And this is going to be done in a discretized way, where we split the surface of the polymer into a series of square elements so we can translate this point load response into the response of the surface to unit pressure, applied over one element of our discretized surface. So this is the plane of the surface of the wafer, the surface of the polymer. And we are trying to express how far that part of the surface goes down when you apply unit pressure here. And that's the answer.

So we tried this with some very simple geometries, lots of parallel channels that we had etched into silicon. They were about 20 microns deep, these silicon channels, and on the order of 120 microns in pitch.

There are many channels going across the screen. But what we show here is on the left, scanning electron micrographs of cross sections through those embossed structures, where a variety of different loads have been applied. And the loads were held for less than a minute. So it's our working assumption here that all the deformation was to do with the rubbery behavior of the material, and there wasn't really enough time for plastic flow to occur.

So we're just using our zero [INAUDIBLE] model to figure out what the Young's modulus of the material was essentially. There's one parameter that we fit in doing this simulation. And that is the Young's modulus.

Here are simulations done using that model. So we're roughly capturing the shape. And it seems to all stack up if we choose this Young's modulus, 5 mega pascals at 130 degrees C.

So that's all very well. But having these parallel channels doesn't give us too much information for a given experiment. So we have to think quite carefully about what type of characterization patterns to use in these experiments. We're trying to think of ways of doing a minimal number of experiments that will give us as much physical information as possible. And what we came up with was this.

Here is a plan of the surface of a silicon stamp that we etched. And it's a square array of various patches of features. Some of them are parallel channels. Some of them are square holes. And we have parallel channels running in both directions.

So for every line to space ratio of channels, there's one set running horizontally somewhere, one set running vertically. And they were arranged in a random order across the stamp.

Now, there's one twist to this, which is that because we do our simulation using a discrete Fourier transform, a fast Fourier transform, the representation of the space that we're simulating is assumed to be periodic in space. So we accommodate that by actually making the pattern periodic.

What I show here is one replicate of the pattern. And this is about 4 millimeters in diameter. But, in fact, what we do is we have-- I suppose on the stamp we have a 3x3 array where what I've shown on the screen is contained in this region. So it's pseudo periodic from the point of view of the material surrounding it.

The material's a millimeter thick. This is 4 millimeters. And so what's going on at the edge here looks exactly the same to the material as this material sees over here. And that makes the simulation match up with the experiments nicely.

So we etch this stamp. We choose some embossing conditions, press it into the substrate, and use white light scanning interferometry to measure the surface topography. The smallest features in this particular stamp are 5 micron in diameter.

What we're also doing now actually is a set of experiments where we've scaled down this pattern 100-fold to make a nano imprint lithography stamp of the same pattern. And so we're trying to tie all the mechanics together and see if similar effects occur at the nanoscale.

Anyhow, in the bottom left here is a composite map of the surface, where the red colors show that the higher topography is where more material has penetrated further into the stamp cavities. And on the right here are eight cross-sections through that topography, labeled one, two, eight.

Experimental data is shown in black, underneath. And the prediction of the linear elastic model, with only one fitted parameter, the Young's modulus, is shown in red. And you can see that there's remarkable correspondence between the two. We get both a prediction of how far on average material penetrates cavities of a given diameter. And also, we do capture these interactions of patterns of different densities.

You can see, for example, let's say-- well, I guess this is a good example. Here you're seeing many instances of a given trench diameter, but trenches that are closer to different patterns feel less far than those that are in the center of the patch.

And this is all captured reasonably well by this rubbery model of the polymer. So looks like we've got the basis here for doing efficient simulations. You can run that simulation in 20 seconds in Matlab, and it's giving you, I think, probably 95% of the information a full finite element simulation would give you, as far as topography is concerned.

So we can also abstract our experimental results and view them in simpler ways. And one way of doing that is to look within each region of features and measure the peak penetration of material into the cavity. So we just measure the range of heights within the central 80% of the area of each patch of features. And then plot that peak penetration against some parameter. In this case, I plotted the pattern density, which is the ratio of the cavity width to the cavity pitch.

And the different colors in this graph represent different feature pitches. So the pitch being the distance from one cavity to the next. And what I show is the elastic model matching up pretty well with the experimental results.

Of course, you might think that this is leading towards the possibility of non dimensional groups to describe features and how they fill, and putting down design rules that would tell you how to scale the embossing pressure as your pattern's scaled down, and these sorts of things. So this is a great way of viewing the results and getting intuition about them.

Of course, as you increase the pressure and material starts to touch the tops of the stamp cavities, this straight line breaks down, of course. And the red symbols show what happens when the pressure is increased to a stage where the larger cavities fill.

What you can also do is use a series of these experiments to get material properties as a function of temperature or of some other parameter. And that's exactly what we did here, by doing a series of embossing tests at different temperatures. And the gradient of that line penetration to pressure will be inversely proportional to the effective elastic modulus of the material.

So we did that. We got out these values for the Young's modulus against temperature. And luckily, they match up really nicely with a model for Young's modulus that was derived from bulk compression experiments, that was just done in a big and strong load frame using bulk materials.

So you could start to see how, with careful choice of test embossing patterns, we can get a lot of information out of a new material quite quickly, scale information, strain rate information perhaps, temperature information. And really, one of the big challenges here is deciding what are the embossing tests you're going to do, what combination of temperatures and pressures are you going to go for first that will give you the most information?

Now, I said before that this rubbery model is really only useful if the whole time is comparatively short. And often that's the case. You want the cycle time to be as short as possible. So if you can get the material into a state where you press hard enough, you get the deformation almost instantaneously, then that's great. But in a lot of cases, you can't do that, and you do get this plastic flow.

So what I show here is the evolution of an embossed topography over time as we hold the load at high temperature. The black symbols down here show what happens if you start cooling the material down as soon as the load reaches its peak. So the material is bowing into the cavities, but hasn't gone very far.

And then right up here, you've waited for 10 minutes. And all but the smallest features have at least touched the tops of the cavities. Of course, this y-axis is showing the peak penetration. So it's the distance from the top of the feature to the bottom. And it tells you nothing about whether the corners of the features have been filled. But this is a reasonably useful measurement.

Now, the question is, can we adapt our simple model to capture this viscoelastic behavior without increasing the computational burden very much? And well, the next thing you might think of doing is adding a linear dashpot into the system. And maybe it's linear. That would be the first thing to try. Maybe you do need to capture the idea of a yield stress as well. And maybe you need to capture the idea of a strain rate dependence, which in PMMA is such that the yield stress increases substantially with strain rates above about 10 to the minus 2 per second.

So, in fact, we just tried adding in a linear dashpot. So what this does, in effect, is takes the point load response associated with purely rubbery, purely linear elastic behavior, which would be like this, and scales it by a factor that's related to the time the load is applied. And all that's saying is imagine this was viscoelastic. You applied a point load you get some instantaneous deformation. But over time, I'm saying this surface is just scaling down. All the points are staying in the same ratio of height, but it's just scaling down.

And we assume that that scale factor was proportional to the hold time, but you might not assume that. You might say, well, what if there's actually a sort of limiting strain where the polymer network is as stretched out as it is it can be for a particular load? And in that case, you would have I guess strain against time would asymptotically approach some limiting value that was related to that spring.

And that would also be quite easy to compute. You compute a topography at the start of the whole thing, a topography that would be associated with infinite time, and then maybe interpolate using an exponential function of hold time. But anyway, we just went with a linear scaling. And it seems to work remarkably well.

For these particular conditions, 110 degrees C, which is a few degrees above glass transition. So we're definitely in a region where plastic flow is significant, is important. And what we see, these are the top experimental data, 3D plot of the measurements taken using optical interferometry on the left for less than a minute holding time, loading duration, and on the right when we left that load in place for 10 minutes.

So you can see that obviously over time the narrower features start to fill. And the simulation, to a reasonable extent, has captured that. And in fact, the shortcoming of the simulation is that it underestimates how quickly the narrower features fill for the parameters that we fit. And so that is definitely an imperfect model, but it's pretty useful, I think, for getting a first cut of a simulation.

So, of course, now that we've got this characterization idea, we can apply that to different materials. And we did it with this cyclic olefin polymer, ZEONOR, which is a thermoplastic. This one softens around 135 degrees C. It has quite a few advantages over PMMA. Although, it's more expensive, it doesn't absorb water as readily, which is pretty relevant to microfluidic devices.

It transmits light to shorter wavelengths, which is relevant when you're exciting fluorescent tags inside the chip with UV light. And so it's useful to be able to characterize these new materials as they come along, and when you don't necessarily have an unlimited supply of the material. So again, this is a plot of peak penetration against cavity diameter. And although most of the data sit on a roughly straight line, when you have cavity diameter being a large proportion of the pitch, in other words, there are very narrow-- hello. There are very narrow walls of the stamp material pressing into the polymer, that's when the penetration does seem to deviate from this straight line.

So anyway, now, I've shown you results from a few carefully chosen operating points. And I showed a set of results where only temperature was varied. And I showed some for ZEONOR that was only one temperature, one hold time, one load. And we need to start thinking about what combination of operating parameters would we want to deploy if we knew nothing about a material except for its approximate glass transition temperature? What would we want to vary first?

And so you might think about doing some sort of fractional factorial experiments where your variables of interest were those connected with the process, the temperature, the load, the hold time, the time over which the load is applied. And that, indeed, is what I did with a third material, another brand of cyclic olefin polymer called TOPAS.

And this is a material that softens around 85 degrees C. And so we did a 2 to the 4 minus 1 fractional factorial where the four parameters were embossing temperature, peak embossing force, hold time, and the time over which the load was ramped up to its peak. And the astute among you will notice that these variables are not going to be independent, because the average load over the loading time will contribute to the hold time. So it's a rough and ready set of experiments. But nevertheless, it gives you some idea of what we might do.

And here we have, again, cavity penetration results as a function of cavity width for different feature pitches. The top plot is for features of a pitch of 100 microns. And this is for 50 and 25.

And you see that, again, we have this nice trend where the peak penetration increases up to a maximum value. That maximum value is the height of the cavities in the stamp. And you'll also notice that I've put these little black dots on the graph as well. And those are the measured heights of the stamp cavities.

So you can see that actually the heights of the stamp cavities fall off for the narrower cavities. And that's associated with the etching process that's used. The actual plasma etch that makes the trenches in the silicon stamp has a harder time etching the narrower trenches. But it's a good sanity check that we're actually measuring the penetration of polymer into the cavities.

Now, here are some 3D plots of results. This in the top left, this is essentially our standard, our center point run set of parameters where for a 4 minute hold time you get a reasonably good filling of most of the feature sizes, but it's not perfect. So we can get an idea of process variability.

This is a set of parameters that was really no use at all, 100 newtons. We haven't even leveled out the non-parallelism in the machine. And one side of the substrate was contacted with the stamp, and the other was not. So there's a lot of missing data here.

And at 100 C, 900 Newtons, 8 minutes, we pretty much got complete filling of most of the features. So what we also did, I should have mentioned it randomized the order of the samples and interspersed these center point runs. Mohammed.

AUDIENCE: [INAUDIBLE] with feature sizes [INAUDIBLE] small, that forces [INAUDIBLE] [? 100 ?] Newtons, does it damage the stamp?

PROFESSOR: In these experiments we haven't damaged the stamp. Actually, the compressive stresses experienced by the silicon are nowhere near enough. I mean, to break the tensile strength of silicon is over 100 mega pascals. But actually where the danger arises when you're cooling the substrate. And we'll get onto that.

But it's really the lateral forces applied to the protruding stamp features that are dangerous, because they create moments on the features, and can cause cracks at the base of the feature as it propagates.

So good question. All right, so anyway, interspersed center point runs. And then what we've done to provide some basis for a test of significance of the effects is to take the average of these peak penetrations. So we've just taken the mean over all feature sizes, and orientations, and everything, and put it down here.

And so you can do an ANOVA. And you can look at the significance of various effects. And actually we get the impression-- I mean, certainly temperature, force, and hold time is significant. That is exactly what we expected. The question down here is, well, at the 5% level, loading rate, the time over which that load is ramped up is not appearing quite to be significant.

But you know, we don't know, because we didn't make a very judicious choice of aliasing arrangements. We don't know whether that's really at the 7% level significant that loading time is a relevant factor, or whether it's the interaction of temperature, force, and hold time.

Just thinking about this intuitively, you do expect the interactions to be important. You expect the product of temperature and force to be relevant, or some non-linear combination of the two parameters, just because a temperature imparts a strain rate for a given force. So increasing the force will increase the strain rate and so forth.

So anyway. And also, as we probably would have expected, there's significant curvature in the results. And harking back to that graph of modulus against temperature, that's exactly what we would expect.

Actually, I should have highlighted when I showed you this graph that there are trade offs to be made in picking these operating parameters. On the one hand, you might say I don't want to heat the material any hotter than is necessary. I'd rather just press as hard as I need to, and maybe operate at 120, 115 degrees C where I can deform the material. And then because I'm not at a very high temperature, when I cool down, there isn't going to be very much differential thermal contraction of the stamp on the substrate. So I won't have big residual stresses, and I won't risk breaking the stamp.

Now, that would be a really logical thing to say. And from a process control perspective, however, you can see that this may not be a very good place to operate, where modulus is highly sensitive to temperature. And our personal experience with lab apparatus-- and this is probably different in industry. But controlling the temperature is fairly difficult. I mean, this is quite an optimistic error bar, in fact, to put on the temperature.

So that's, I think, a really relevant consideration for choosing hot embossing parameters, how sensitive we are to temperature.

AUDIENCE: [? Given ?] the temperature readouts, plus minus 1 degree?

PROFESSOR: That's what I estimated for that particular machine, yeah. And I think that there are also big challenges to do with temperature uniformity across [INAUDIBLE]. Yeah, that's really something that people are looking at.

So, in fact, the side of the industry that's more advanced is really the nano imprint lithography side, more so than the micro embossing side. And in that case, it's more attractive to go to very high temperatures, to make the material behave more like a viscous fluid than like a rubbery material or a viscoelastic. So then temperature variation isn't so much of an issue. So you're right down in the 4 megapascal range.

So but yes, good point, absolutely. Now, yes, there's one really interesting thing about these results. And I mentioned that I had two sets of each feature size. And one was oriented vertically on the stamp, and one horizontally.

And I just put those on there on the off chance that there was anisotropy in the materials. And it turns out, with this TOPAS sample, which is a gift from the supplier, there is really strong anisotropy.

And so you can see that the filled in symbols are for an arbitrarily defined orientation, 90 degrees, and the open symbols for 0 degrees. So turning the feature through 90 degrees gives you this really big difference in penetration. And these results are based on five replicates. You can see the error bars there at one standard deviation of the five replicate results.

And there's definitely a significant anisotropy. Whether that's to do with residual stress in the cast sheet, alignment of the polymer chains because of some feature of its processing, I mean, you can start to see that if materials like this have such a big orientation dependence, and that is so dependent on how they've been processed, you can't necessarily know the whole processing history of what you're supplied. It will be useful to have these quick checks of the properties of materials that are given to you. And you wouldn't be able to get that information just by a bulk compression experiment.

So these carefully designed test patterns are really important. I should say that this is anisotropy in the material, not in the stamp. Because if we rotate the stamp relative to the sample, the results reverse. So it's definitely in the material.

Anyhow, that was for TOPAS. And I showed you a few slides ago what happens when hold time is relevant, when you've got plastic flow. And we saw that the introduction of this linear dashpot, the scaling of the point load response function was a reasonably good way of capturing that. What we're working on now is this ability to capture yield stress to iron out some of the shortcomings of the model, the fact that those narrower features actually filled more than the model predicted.

And the other thing that we're trying to do is extend this simulation approach to thin substrates for the nano imprint lithography case, where this type of point load response function isn't correct. The features are the same size roughly as the thickness of the substrate. You have to start thinking about lateral transported material. And that's quite an exciting topic.

Now, we had a couple of good questions about damage to the stamp and de-molding issues. And this is perhaps the biggest impediment to the use of imprinting, a microscale structure. There's not that much of a problem for nano imprint. But microscale structures, it's crucial.

And here are some examples of problems that have occurred during cooling and de-molding. This is an SEM picture of a PMMA component that has been embossed with a hexagonal post. This was made from silicon, in fact. And if you imagine the substrate being mostly on this side of the feature, cooled under load, the polymer having at least 10 times the thermal expansion coefficient as silicon, it has pushed itself against the sides of the features. And when we de-mold, we get these really substantial ridges.

I think this was embossed at 130 degree C, 20 degrees above the glass transition. And so what we haven't yet worked out is whether these contact stresses are large enough actually to push the stamp off the substrate, or whether it's just elastic potential energy stored in the polymer that causes this ridge to spring up when you peel the stamp off the substrate.

Here are some optical micrographs showing a similar problem, a triangular hole embossed into a polymer substrate. The material is contracted from right to left relative to the stamp feature. And in fact, sometimes we see these shards of polymer being sheared off, and again, a hole where there's been a few tens of microns relative contraction.

AUDIENCE: I'm curious, is there any sort of like lubricant that you can use to increase the load?

PROFESSOR: There are products you can get that you can spray onto the mold, or, indeed, treatments that can be given to the mold to reduce the coefficient of friction between the mold and the substrate. That won't reduce the magnitude of the lateral force. That's to do with Young's modulus and the coefficient of thermal expansion.

But yes, it will potentially make it easier once that lateral force is exerted to pull the stamp out. However, it's not - the problem isn't all to do with friction. Sometimes you have-- either your stamp is imperfectly made, and there's a negative draft.

I mean, I exaggerate, but if that were the stamp and that were the polymer, that's a problem that people run into quite a lot. And then if the stamp has been made by deep, reactive ion etching where there's this cyclic etching process, and there are these little, sub micron scallops, you're susceptible to mechanical locking. And there's plastic deformation of the polymer as you pull it off.

So it's a good point. Yes, and actually, one thing that does work quite well is to leave the Teflon polymeric coating on the silicon after etching. And that does do quite a lot. It has a very limited lifetime, though.

Anyway, the point of saying all this is, as I alluded to a few minutes ago, you might think of reducing the temperature swing to a small a value as you can so that amount of thermal contraction is restricted. You might even think of removing the load slightly above the glass transition temperature. And if you're in a region where some of the deformation, or most of the deformation is actually plastic, which is true in a region about 10 degrees above glass transition for PMMA-- if you can exploit the material properties to avoid having a big temperature swing, then that could potentially be very exciting. And together with Matt Dirks, who's in the lab for manufacturing and productivity, we looked into this a couple of years ago, and found that if we control the temperature really carefully, then it was promising.

If you de-mold at 50 degrees C, so well below glass transition, and then pull the stamp off-- this is a cross-section through the part-- then you get this horrendous ridge on the side of the part nearer the edge, the side of the feature nearer the edge of the part.

If you de-mold, say, well above glass transition, 120, you get this blue line. So you get some bowing of the base of the feature where material has sprung back. And you get splaying of the sidewalls. And this was repeatable over many samples.

However, if we can hit 110 degrees C on the nose, we get a nice, flat feature. We get nice, flat sidewalls. And we get no ridge at all. And that would be great.

But again, that's the thing. Can you, in a production setting, make it 110, and not 115 or 120? And that, I think, is probably possible to achieve, but would increase the cost of the apparatus.

So this is definitely something to be aware of. Another thing that is of importance is that the risk of the whole part bowing if the de-molding temperature is too high. We've observed problems with that. So that needs to be thought of as well.

Here we can see also that if you're going to remove the load at or above the glass transition temperature, you have to pull out the stamp, cool it down jolly quickly. Because otherwise the part will start to contract back to its original flat surface. And so you leave it for 10 minutes, things start to look pretty poor.

OK, that is an overview of about half of my thesis. And since we're at half past 9:00, I think we should stop. But are there any questions from anyone?

If not, then I guess the floor is open for questions about the projects.

PROFESSOR: [INAUDIBLE] mention one thing [INAUDIBLE]. I think later today, hopefully, or by tomorrow at the latest, I'll post the tentative schedule for what groups will be presenting on each of Tuesday and Thursday. It will be in the normal class period. So we'll plan for that. And I'll have some other instructions, like it'll be best, if possible, to send those presentations to me a little bit early.

I see we do have a USB port. So it's possible we can add up the presentations on the fly. I don't know quite how well that works out in Singapore. We'll have to figure that out as well.

But be alert. Some of you, based on random draw, will be doing presentations on Tuesday. And some groups will be doing that on Thursday. Are there questions?

AUDIENCE: Is there a rough guideline for how long the IEEE letter should be?

PROFESSOR: Oh, the IEEE paper?

AUDIENCE: Yeah.

PROFESSOR: I think typically-- whatever it takes to do a reasonable job. I think something typically on the order of six, seven, eight pages is, I think, what we often saw in the past. Sometimes it might be a little bit shorter, depending on how much background you have.

So I don't have a strict page guideline. Good. Any questions in Singapore on that? OK, great.

AUDIENCE: Professor.

PROFESSOR: Yes.

AUDIENCE: We have a question for our project. We actually have two questions.

So one question is that we have two outputs in our-- can you hear me?

PROFESSOR: Yeah, speak up a little bit.

AUDIENCE: OK, so one question is that we have two outputs, and then we are measuring the data. But these two outputs, they are negatively correlated. And our objective is to maximize one output and minimize the other output.

So we understand that we can do it separately in our optimization. But when we want to sort of create a combined an objective function and optimize these two together, as they come from the same set of inputs, how can we do that?

PROFESSOR: Yeah, so the classic approach there is it's a trade off. Right? One or the other. And what you typically do is apply a weight to the two alternatives. So you have a different numeric weight based on your engineering importance, or the customer importance, or whatever of the two effects.

So you have one function that combines, say, squared deviations in the two outputs, but with a different weight on the two outputs.

AUDIENCE: OK, I see.

PROFESSOR: OK? You're all set?

AUDIENCE: We have another question actually.

PROFESSOR: Oh, OK.

AUDIENCE: I have one more question. It's regarding the email you sent us. So we were told that we can explore one factor [INAUDIBLE] strategy. We are not quite sure what strategy you are referring to. Is it studying the output and one input at a time? Or is it more like stepwise regression?

PROFESSOR: No, I'm referring to the optimization approach, the One Factor At a Time, OFAT, that Professor Dan [? Frey ?] talked about in his lecture. So where you explore the corners. You look and see if you've improved it or not, to decide if you keep that move. And then once you do make a move, then you randomly decide which next corner to experimentally try.

So you've actually got the data already done from a DOE. But now you can pretend you were actually running the experiment just one experimental point at a time using the OFAT approach. And I think that would be interesting for online optimization to compare also to response surface modeling optimization.

AUDIENCE: OK. All right, thanks.

PROFESSOR: Great, all right. See you guys later.