The Abbott Guide to Rheology

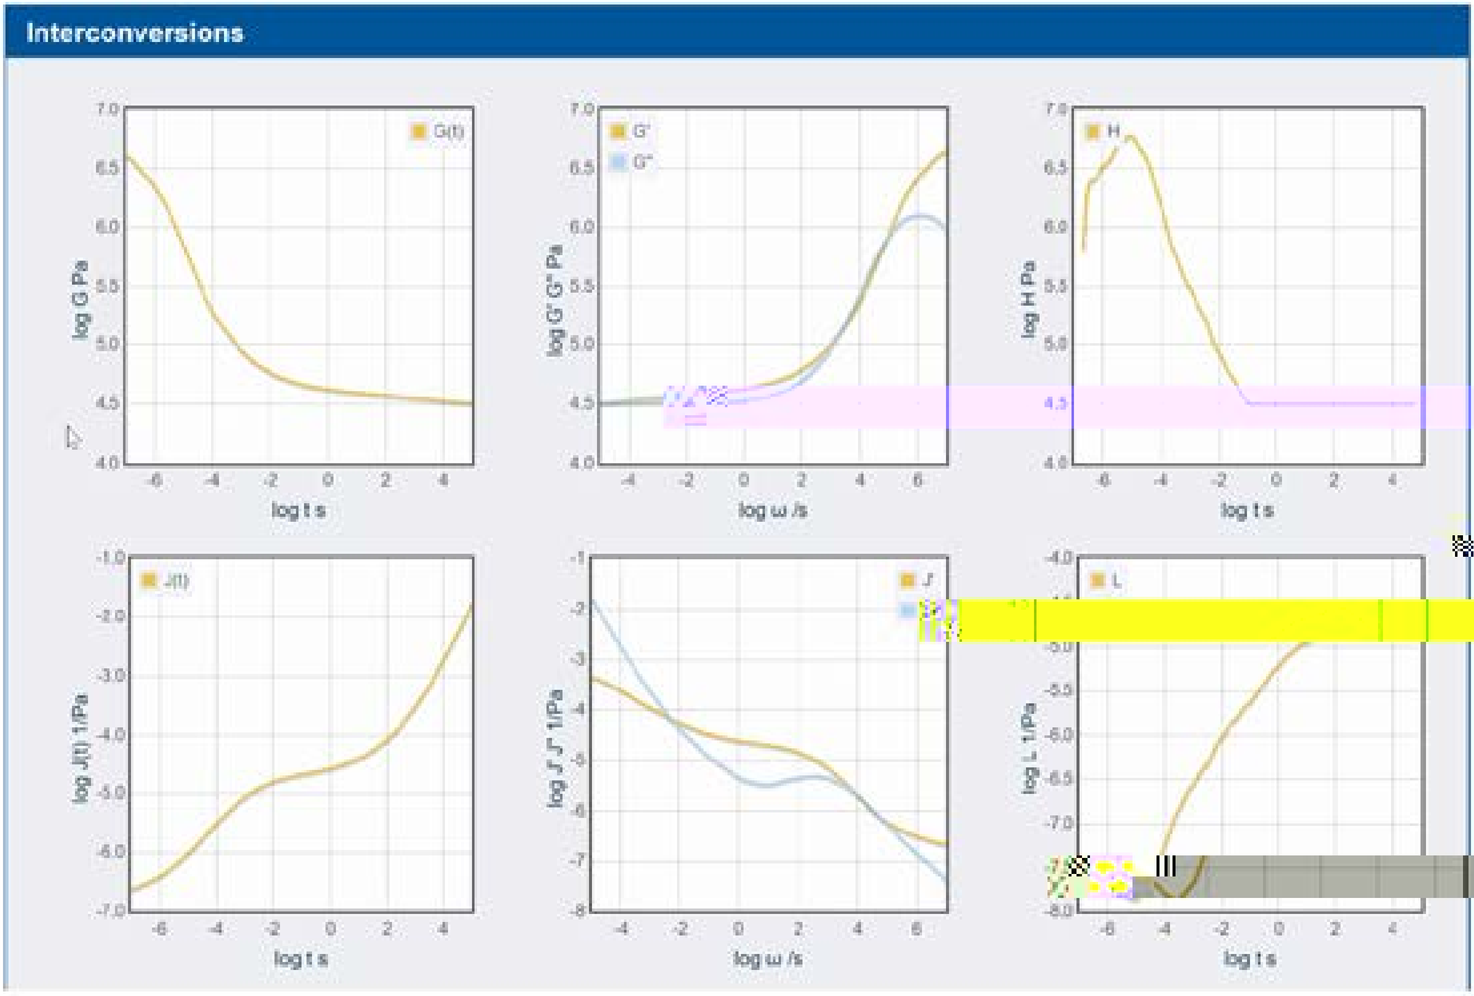

Prof Steven Abbott

The Abbott Guide to Rheology

Prof Steven Abbott

Steven Abbott TCNF Ltd, Ipswich, UK
and Visiting Professor, University of Leeds, UK

steven@stevenabbott.co.uk
www.stevenabbott.co.uk

Version history:

First words written

Contents

Preface 4

1 Setting the scene 6
- 1.1 Never measure ***a*** viscosity 8

2 Shear-rate dependent viscosity 10
    - 2.1.1 Shear thinning and your process 12
    - 2.1.2 What causes shear thinning? 13
- 2.1 Thixotropy 17
    - 2.1.1 What to do with the relaxation times 19
    - 2.1.2 What to do with thixotropy measurements 21

3 Yield Stress 22
- 3.1 Yield Strain 25
- 3.2 Using yield stress or strain values 26

4 Semi-Solids 27
- 4.1 G', G" and tanŭ 28
    - 4.1.1 Storage and Loss 30
- 4.2 TTE/TTS/WLF 30
- 4.3 How do we use G':G" and WLF information? 32
    - 4.3.1 Your own G':G", WLF insights 34
- 4.4 Are G':G" all we need? 35

5 Interconversions 36
- 5.1

Preface

"Maybe we should check out the rheology of this system" is the sort of sentence that can create panic and alarm in many people. Rheology has a reputation as being obscure and difficult, so people avoid it as much as possible. This is deeply unfortunate because, despite its bad reputation, rheology can deliver key insights to formulators and trouble-shooters and isn't, in fact, so difficult.

There are some excellent books on general rheology. I have especially valued Thomas Mezger's *Rheology Handbook* as he writes as an expert who knows his subject, his equipment and his audience well. This guide in no way wishes to compete with the 400pp of Mezger's book. There are also many excellent books on specific aspects of rheology such as Christopher Macosko's *Rheology Principles, Measurements and Applications* or Jan Mewis and Norman Wagner's *Colloidal Suspension Rheology,* from each of which I have learned a great deal. And the TA Instruments Webinar Series on YouTube is a great way to learn the details of rheological techniques from experts.

What is lacking in all these (precisely because they contain so much valuable information) is a stepping back from the detail to see rheology from the point of view of the harassed formulator who has an idea that rheology might be helpful but cannot see which measurements will help to solve a specific problem.

So this guide attempts to describe the problems we need to solve, along with the appropriate bit of rheology that helps provide the answers. In order to describe the problem we have to introduce some key ideas such as entanglement and relaxation times that are not at all hard but which are generally unknown or under-appreciated.

To make life easier for the reader and for myself (after all, I need to understand this stuff as well), I have written a set of apps on my Practical Rheology website, where you can see things live when they are discussed in this guide. Just click on the link to immediately start exploring. The apps are standard HTML5/Javascrip/CSS3 so they run on phones, tablets and laptops, are safe on corporate networks, and are free and free of ads.

Along the way, the guide will point out why rheologists have created so much confusion for the rest of us, and, via apps, bring some order to that confusion. If any real rheologist reads this guide, maybe they will see the world through our non-expert eyes and try, in future, to be less confusing. I will also point out some huge flaws in the world of rheology. This is not to make fun of rheologists. The flaws are there because it is a genuinely difficult scientific challenge to fix them. If the flaws could be fixed then both rheology and the practical formulator would benefit strongly, for reasons I hope to make clear.

Because I am not a rheologist and have always struggled with rheology, I have had to rely on the wise input from a number of world-class rheology experts. I am especially grateful to one expert who, after glancing at an early prototype of Practical Rheology said "It's fine, except for the fact that you have omitted almost everything of importance". That assessment was entirely accurate and the resulting site is much improved from that crude prototype.

I would like to acknowledge my debt to, in alphabetical order, the experts:

- Prof Paul Bowen, Neil Cunningham, Seth Lindberg, Roelof Luth, Dr Hans-Martin Sauer, Dr Saeid Savarmand.

They have each helped me considerably, though I stress that all opinions and errors in the Guide are my own.

It was an automotive engineer from Germany, Sebastian Abbott, who mentioned how "Maybe we should check out the rheology of the system" seemed to induce panic. Rheology should be a routine methodology for those who need it as part of their day job, just as they need other measurements, such as spectra, in which they do not have to be great experts. It should not be a source of alarm. That quote was the inspiration for putting this guide together. The act of writing it has reduced my own sense of panic and alarm. I hope that it will do the same for you, the reader. Sebastian also provided a fine critique of a late draft of the guide. This prompted a significant rewrite and, I hope, a more user-friendly book.

Steven Abbott, Ipswich, 2018

1 Setting the scene

What problems are we trying to solve?

Your formulations can flow. At the very minimum you need some objective measure of the flow characteristics, and viscosity is the most basic method. Before worrying about what our viscosity should be, we need to be confident that we know what viscosity is.

A rather different problem is that talk of viscosity tends to include terms like stress and strain, and units like /s or Pa.s. We need to be confident about these too.

Rheology got off to a bad start for me thanks to two 6-letter words which in normal English mean the same and which start with the same 3 letters: stress and strain. As we cannot make the words go away we have to learn not to get confused.

1. **Stress** is a force per unit area applied to an object. Because force is in Newtons, stress is N/m^2 or Pa.
2. **Strain** is the % increase in length of an object when it is stressed; i.e. it is increase of length divided by the original length. It is a pure number, without units.

You can't get a strain without a stress and you don't have a stress without a strain. Rheologists seem to outsiders to make random decisions about the

As we shall see, most materials have a mixture of elastic and viscous properties so we need that equation later. For now we want to focus on viscosity, where I will try to use ɖ consistently though you often see ɛ and I'm guilty of sometimes swapping accidentally between them.

For viscosity, strain is not useful because viscosity appears only when motion is involved. That is why we need strain *rate*, which is shown as ɔ with a dot above it, ɔ,–though my attempts to produce it always place the dot a bit lopsided. I find it very hard to spot whether people are discussing ɔ or ɔ,–another example of how rheologists have made our lives unnecessarily hard. Those who created the nomenclature knew that dot means "derivative" - without thinking through the implications for ordinary people who have to use their nomenclature.

Let us see how the idea of viscosity, strain rate and the required stress are inter-related.

The equation also tells us that viscosity must have the units of stress/shear_rate, which is Pa/(1/s) giving us our viscosity unit of Pa.s. A viscosity of 1 Pa.s is quite high, so for convenience, lower viscosities are usually cited as cP, where 1 cP=1 mPa.s. It is annoying that we have P for Poise and Pa.s for Pascal seconds and many of us have messed up by mis-reading the unit and therefore being out by a factor of 10. So, look out for P (which is 0.1 Pa.s) and Pa.s and don't feel too bad if you occasionally get them confused.

And while we are talking about confusion we also have the rheology world usually talking in terms of to control their product. With modest effort they could have found whether the viscosity was shear thinning (which would have helped) and how the viscosity changed with % solids. It turned out that, as is frequently the case, the solids had been pushed to a practical limit and that a small increase (via evaporation on the machine) leads to a significant increase in viscosity which in turn could create the coating defects I was there to troubleshoot. The common fact that a small change in polymer concentration can make a large difference to the viscosity is discussed later.

The first advice, therefore, is to never measure *a* viscosity and certainly to never use a flow cup. This is the 21st century and it is neither hard nor expensive to measure viscosities over a range of shear rates and over a relevant range of concentrations. Part of the reluctance to make such measurements is the perception that rheology is hard. I will make the case that rheologists make rheology unnecessarily hard, but rheology itself is not that difficult.

2 Shear-rate dependent viscosity

What problems are we trying to solve?

Giving our fluid a single "viscosity" usually is not good enough, so we ne is not

you are in terms of shear rate and, therefore, viscosity. I like to use a simple adhesive as an example.

The user opens the cap of the tube and wants to squeeze out the adhesive. You don't want the viscosity to be too low, otherwise it will flood out. Equally, it cannot be too high, otherwise it will be too hard to squeeze. You have your Cross model, now all you need is your shear rate. The tube radius is 0.25cm and you want the adhesive to come out at 2.5cm/s, so V/Y = 2.5/0.25 = 10/s. Now the user needs to get a thin, even layer of adhesive by sliding together two surfaces. We need a really low viscosity for the adhesive to flow nicely. So what is the shear rate? We can estimate 0.5m/s for the act of sliding, and we need 50εm of adhesive, so V/Y=0.5/0.00005 = 10,000/s. For manufacturing who need to mix the adhesive with a stirrer there is no easy way to estimate shear rate; it is well-known that a "high speed stirrer" is somewhere in the range of 100-1000/s.

This sort of quick analysis on your own formulation is very powerful. It combines your knowledge of how the formulation *should* behave with how it *will* behave (as judged by the Cross model and your estimates of shear rate). If your rough estimates tell you that each process is comfortably in the right domain, then no further work is required. If they tell you that one part of the process is problematic, then you have more work to do. If your estimate of shear rate tells you that things are borderline then you need to refine your estimate with some more experiments. For example, if others say that the flow rate from the tube will be 1, not 5 cm/s and your estimate tells you that between 2 and 10 /s viscosity will plummet, then you need to do some serious testing of glue squeeze rates and, perhaps, a re-design of the nozzle. If you know that the viscosity hardly changes across this velocity range then you don't have to do the experiments.

2.1.2 What causes shear thinning?

What is far more interesting than the details of the model is the reason for shear thinning itself. This is important to grasp as it will fit in with much else that rheology has to offer us. The reason for shear *thickening* (walking on cornstarch) is very different and is discussed in the particles chapter.

We start with a discussion of shear thinning in *polymeric* systems which is important in its own right and leads us onto a key theme that deserves our attention because we should be spending more of our rheological time trying to understand it. The shear thinning of *particle* systems does not happen below a specific volume fraction, showing that the flow behaviour of particles is more complex. Particle systems (which may be solids, emulsions or foams) have a chapter to themselves once we have discussed the broad range of rheological behaviours.

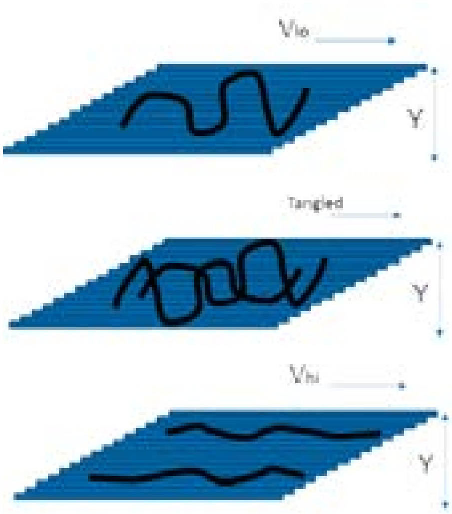

We have our streamlines. At low shear values, a polymer in the liquid spans a number of streamlines. Those in the fast stream are being held back by those in the slow stream, so there is a general increase in resistance to flow, i.e. a high viscosity. If the polymers are at a sufficient concentration to tangle then the effect on the viscosity is much higher. At high shear, however, the polymer chains are stretched along the streamlines and are untangled, so resistance to flow is low and therefore the viscosity is low. The reason we need a Cross model with four parameters is that the whole process is complex. The ability for a polymer to tangle depends on its concentration (of course), on its MWt[3] (bigger MWt will tangle more, of course). It also depends on how "entangleable" the polymer is which is more puzzling. For some polymers, a low MWt is enough to tangle, for others it needs a large MWt. If we know the "critical entanglement MWt" then we can predict the low shear viscosity.

Entanglement is important across large areas of formulation space. It is what gives us strong adhesion, it gives us strength and toughness, and it is what (often) stops us from adding as much polymer as we might like because the formulation gets too viscous to handle. For many formulated products and systems, entanglement is good, yet it is entanglement which makes it hard for us to formulate them.

Given that polymers have a critical entanglement MWt, M_c, we should all know their values so we can work out whether we want to be below or above it. The problem is that although values are known for a handful of simple polymers, they are unknown for most of the real-world polymers we tend to use. Until suppliers get into the habit of providing M_c values, we should get into the habit of estimating them ourselves. Two apps help you to do this.

---

[3] It gets tedious to keep typing molecular weight, so I use MWt.

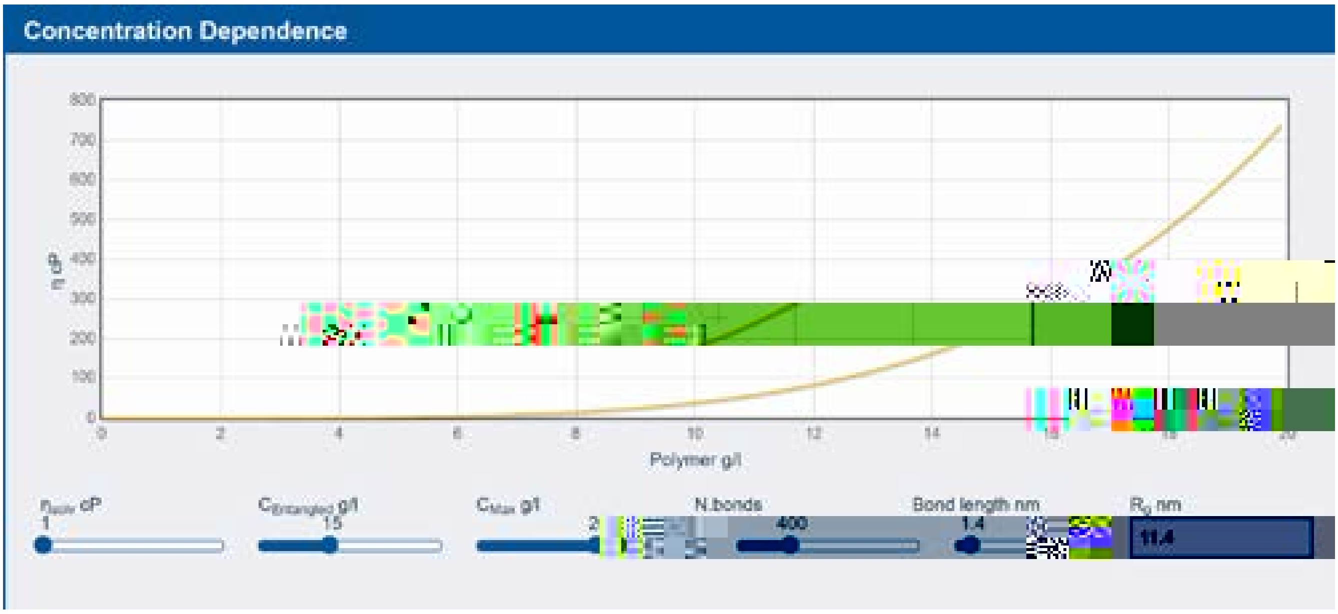

η_{solv} cP 1

$C_{Entangled}$ g/l 15

C_{Max} g/l 2

N.bonds 400

Bond length nm 1.4

R_g nm 11.4

App 2-2 https://www.stevenabbott.co.uk/practical-rheology/Low-Shear-Polymers.php

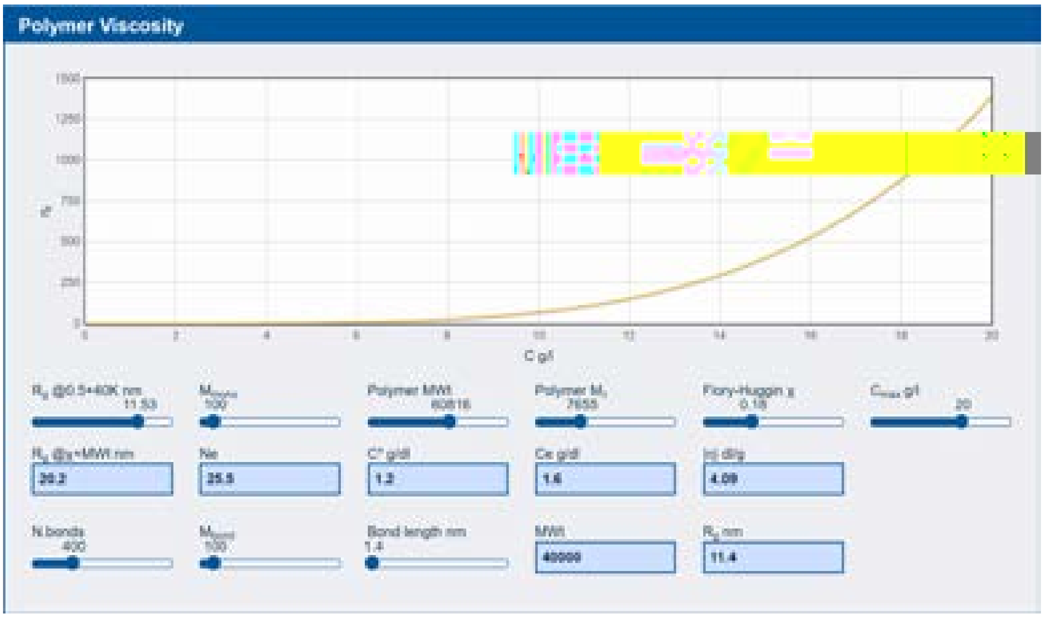

R_g @0.5+40K nm 11.53

M_{mono} 100

Polymer MWt 60816

Polymer M_c 7655

Flory-Huggin χ 0.18

C_{max} g/l 20

R_g @χ+MWt nm 20.2

Ne 25.5

C* g/dl 1.2

Ce g/dl 1.6

[η] dl/g 4.09

N bonds 400

M_{bond} 100

Bond length nm 1.4

MWt 40000

R_g nm 11.4

App 2-3 https://www.stevenabbott.co.uk/practical-solubility/polymer-viscosity.php

This app takes in M_c as a parameter and includes a measure of solvent-polymer compatibility, the G parameter which goes from 0 (totally compatible, polymer fully expanded) to 0.5 (the polymer is bordering on unhappy). You can read the detailed text of the app to learn more. The point is that although polymer viscosity behaviour is more complicated than we would like it to be, it *is* possible

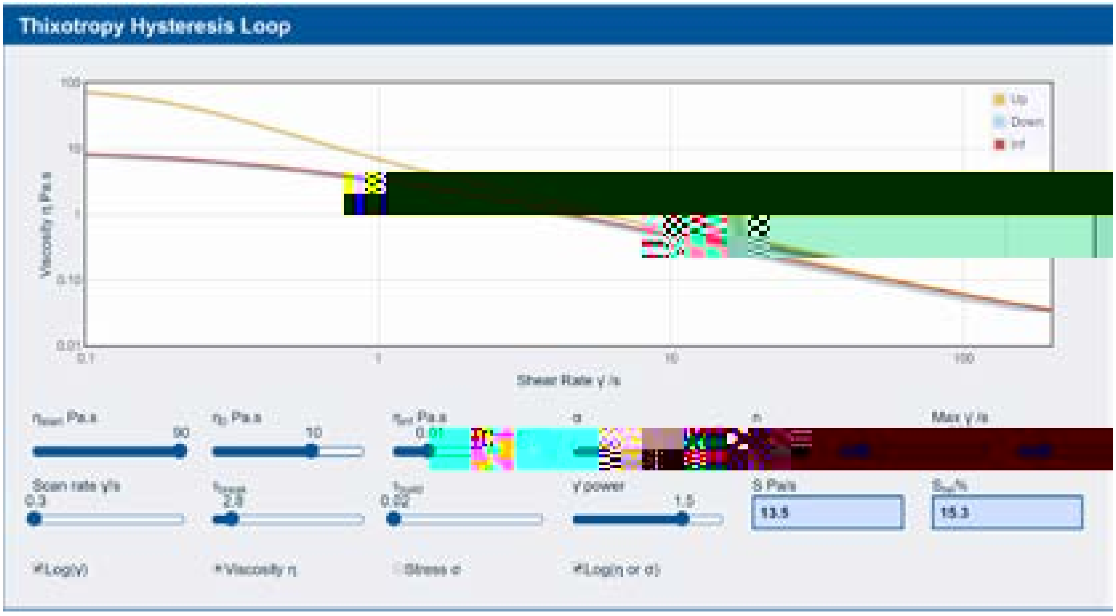

Thixotropy Hysteresis Loop

η_{start} Pa.s 90

η_0 Pa.s 10

η_{inf} Pa.s 0.01

α

n

Max γ̇ /s

Scan rate γ̇/s 0.3

τ_{break} 2.8

τ_{build} 0.02

γ̇ power 1.5

S Pa/s 13.5

S_{rel}% 15.3

☑ Log(γ̇)

● Viscosity η

○ Stress σ

☑ Log(η or σ)

App 2-4 ____ some sort of measure, though it is arguably better as the "relative area" obtained by dividing by the area under the up curve.

Out of personal interest, I wrote an app that describes an oscillatory method, giving the G' value with time after stopping a fast shear.

Equu. 2-7 $$G'_t = G'_0 + A.t^n$$

I do not see how it can be of much use because it just gives a power law, n, for the timescale and a value, A which sort of describes the difference in G' after shear and at infinite time. What can I do with such values?

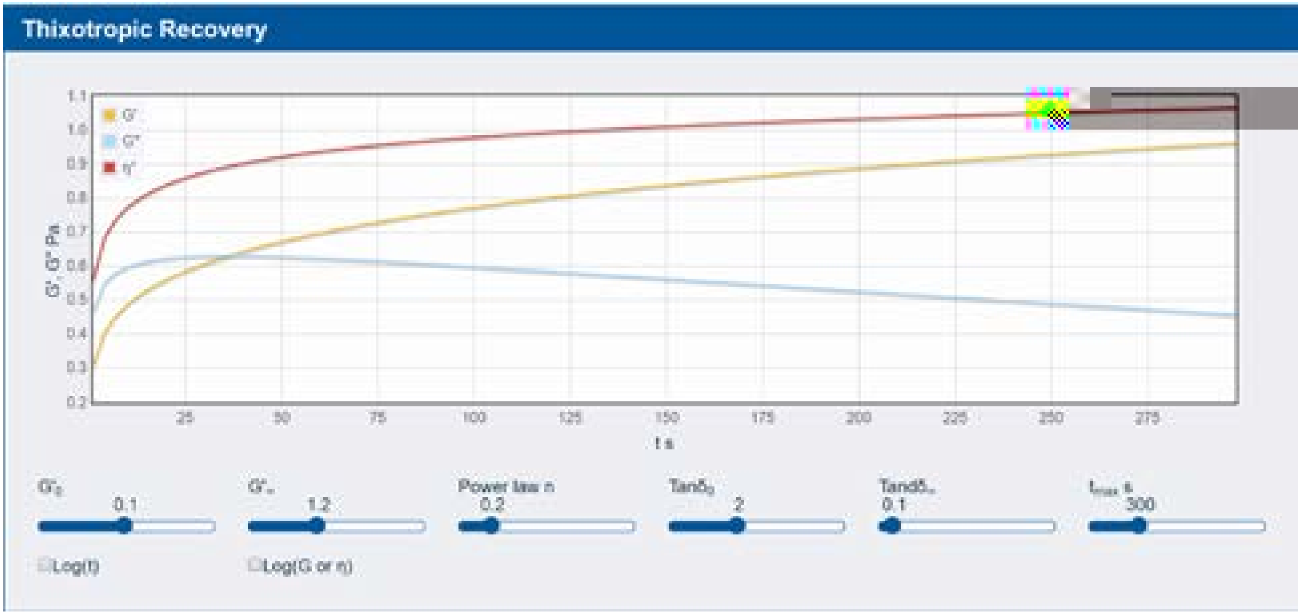

G'_0 0.1

G'_∞ 1.2

Power law n 0.2

$\tan\delta_0$ 2

$\tan\delta_\infty$ 0.1

t_{max} s 300

☐ Log(t)

☐ Log(G or η)

With hindsight I can see that my ignorance of entanglement was behind my puzzlement. I had never heard of critical entanglement MWt and how it varies wildly between different types of polymers, so I had no intellectual framework to work through the issue of choosing the right polymer. A specific example of this was in screen printing, I was disappointed that recovery of viscosity after shear was so slow. A particle-based ink printed side-by-side with a polymer-based ink behaved much better because the particle-based ink recovered its viscosity near instantly.

I then "discovered" associative thickeners. These were a revelation because their recovery is so much faster. The difference is entanglement. For normal polymers, the tangles that give them their viscosity also slow down their return to viscosity as the chains' ability to move back to their tangled state is increasingly limited by the tangles (this is reptation physics). Associative thickeners are tangled via a network of weak local bonds (such as hydrogen bonds) that can re-form with very little main chain motion. Particle systems recover quickly because there is little need for rearrangement or motion to return to the low-shear state.

All this is obvious with the right language of timescales and entanglement and rather obscure without it.

The implications for the formulator are clear. If you need an entangled system for other purposes (such as adhesion), don't expect fast relaxation times and low thixotropy. Because the amount of entanglement is concentration-dependent, expect thixotropy to change dramatically over a relatively small range around a critical concentration.

If you require viscosity coupled with low thixotropy, then associative or particle thickeners are effective. There is a subtlety with those particle thickeners that rely on strong self-association, i.e. flocculation. An app allows us to explore this subtlety:

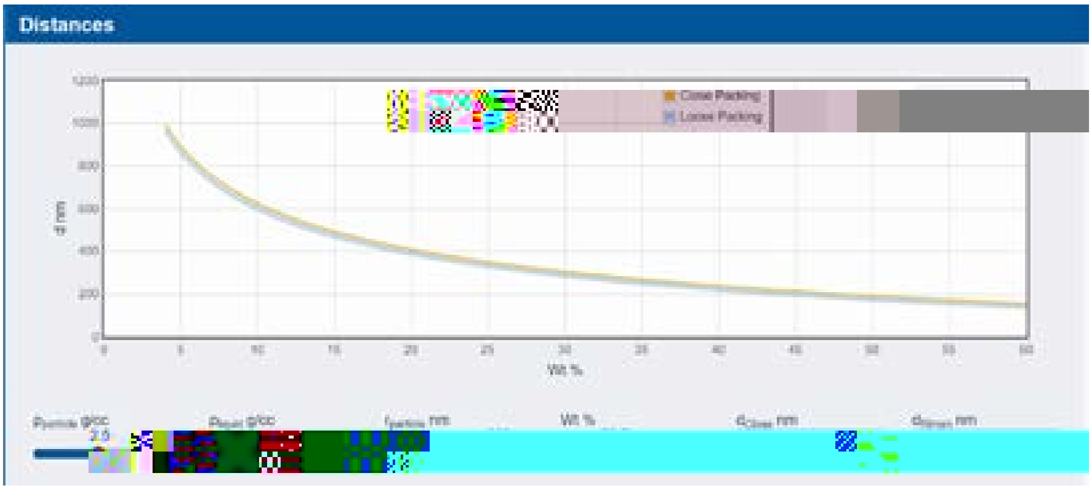

App 2-6 https://www.stevenabbott.co.uk/practical-coatings/distances.php

The app shows the average distance, d, between particles as a function of wt% of particles. The rate of flocculation recovery depends on the probability of particles touching. If you have a given volume of particles as large radius spheres, d is, on average, large, so it will take a long time for them to travel the

on yield stress and you will find measurements being made in a bewildering variety of ways, each stated with confidence as being "the" way of doing it, with no indication of why that method was chosen over others that might be equally good or bad. The apparent variety is even larger because of the rheologists' delight in plotting the same data in different manner, so it is even harder to work out what is going on.

Fortunately a wonderful paper[5] from Daniel Bonn's group at U Amsterdam describes many of the different ways, making it easy for me to write an app that shows what you would see if you measured the same yield stress, σ_y, (as an input to the app) in six different ways. Whenever I have to read a yield stress paper I go back to my app to find the nearesm an to go near ¿ 6

In a real experiment, K and n are fitting parameters. In the app they are set to plausible values of 1 and 0.5. Other curves, such as Casson, can be used - they are each approximations, with different industries deciding that one type is better than another. The world of chocolate, for example, use the Casson method, even though there is a detailed paper from the industry explaining that Casson doesn't really work well for chocolate. The n=0.5 I have chosen for Herschel-Bulkley is the same dependence on shear rate as Casson (which also uses the square root of $ů_y$).

The second method (bottom left) takes the same experiment but somehow extracts a viscosity from the data and plots that against ů. I confess to not knowing how this is done, so the values in the app are plausible (I checked with some papers) without being calculated. Maybe it is so obvious how to extract ɖ that the authors did not have to explain. I will happily update the app and this paragraph if someone would enlighten me.

The next technique provides three ways of extracting a value from the data. The basic experiment to measure G' and G'' will be described in the Semi-Solids chapter but hopefully readers are at least aware that these are standard values obtained from oscillatory rheology. The stress is steadily increased (so the best machine is a controlled stress machine rather than controlled strain) until there is an obvious change in the measured parameters. Now we can extract the values by whichever technique we happen to prefer. The first two methods use the data from the top-middle, the third uses data from the lower-middle plot:

1. Very popular, because it is simple for humans and computers to spot it, is to take ů at the cross-over point when G'' becomes larger than G', i.e. the system becomes more viscous than elastic. This value is usually an overestimate and is sometimes called the Flow stress rather than yield stress.
2. More true to $ů_y$ is to identify when the G' value starts to decrease which, by definition, means that it's starting to lose its elastic nature. This is tricky to do by computer and rather obvious by the human eye.
3. Because rheologists love to plot data in other manners, the complex strain, ɔ*, is calculated as Stress/Modulus, i.e. ů/G* where G* is the complex modulus given by G*ǀ=G'ǀ+G''ǀ. It is said that you can see $ů_y$ from intersection of the two relatively straight lines.

Next (top right), you set up a (very) small shear rate and watch how stress grows with strain. The yield stress is the plateau value where the strain just keeps going.

Finally (bottom right), you can do a set of creep experiments where you set a fixed stress ů and measure how strain ɔ evolves. You plot compliance, J = ɔ/ů versus time, and for a purely elastic material this is a constant, though for real-world materials there is a slight slope with increasing time. At the yield stress, the material can flow so J increases more significantly over time. You can therefore work out from a set of plots a value for $\mathring{u}_y$. If you guess badly about the range ofof a a aof f Ä ƶ

Whether you measure yield stress or strain depends on what question you are trying to answer.

3.2 Using yield stress or strain values

For many formulators, yield stress is a non-issue and this chapter can largely be ignored. For some, it is a vital parameter for a product such as the hand creams or foodstuffs discussed at the start.

The science of predicting/controlling yield stress is difficult, and I have, sadly, had little to say about it, though when we come to particulate systems the situation is better. My concern here is with two issues.

1. Ensuring that your measurement really relates to the end product performance. This often involves consumer test panels who poke, rub, lick or otherwise interact with a range of different formulations and say which they prefer. With luck you can relate that to an objective measure of yield stress.
2. Ensuring that your objective measure of yield stress is comparable between labs, between operators and between samples - in other words that it has a reasonable Gauge R&R[6], showing that the noise in the measurement is less than the signal. Trying to hit a target of 43 Pa when measurements vary randomly between 30 and 60 is not a great way to formulate.

The fact that we have at least six methods of measuring yield stress is a clue that measurement is difficult. The only encouragement I can offer is that *you are not alone*. I have not met anyone in the yield stress world who is happy with those two issues. A famous paper on chocolate yield stress[7] involving a world-wide team and some high-powered labs found many ways in which yield stress is frustratingly hard to measure (the paper does not cover the impact of yield stress on consumers).

One final tip I learned from the experts. If you spend a large amount of time worrying about the precise details of Herschel-Bulkley or Casson etc. the result will be that it really makes no difference compared to getting a good-enough, robust method that works day-to-day on real-world samples. Combine that observation with the fact that you are not alone in your frustrations with yield stress and, with luck, you'll be able to say that good enough is good enough, and get on with your formulating.

---

[6] Gauge Repeatability and Reproducibility - found in any reputable process control statistics package. It is amazing, and depressing, how many "industry standard" tests fail elementary Gauge R&R analyses and are, therefore, unfit for purpose.

[7] J.-M. Aeschlimann and S.T. Beckett, International Inter-Laboratory Trials to Determine the Factors Affecting the Measurement of Chocolate Viscosity, J. Textural Studies, 31, 541-576, 2000

4 Semi-Solids

What problems are we trying to solve?

Many of our formulations are neither pure solids nor pure liquids. So we need to be able to understand how much they are of one or the other. A straightforward measurement of G' and G'' will do this for you. But it is missing a key point. *Everything* is a solid at a low-enough temperature or a short-enough timescale and *everything* is a liquid at a high-enough temperature or a long-enough timescale.

Because most formulations have to perform under different timescales (e.g. a pressure sensitive adhesive must stick quickly yet not creep over long times), we need to be able to both measure *how* things change over temperature and time, and understand *why* they change so we can formulate intelligently.

There are no such things as solids. At a high-enough temperature (obvious) or a long-enough timescale (less obvious) "everything flows", *παντα ρει, panta rhei* from which we obtain rheology.[8] The prophetess Deborah in Judges 5:5 stated that "The mountains flowed before the Lord", meaning that in a long-enough timescale everything flows. Not explicitly pointed out by Deborah is that mountains also flow if they get hot enough.

We therefore end up with the Deborah number De = $t_{relax}/t_{observed}$, i.e. the ratio of the time over which the system relaxes to the time over which the system is observed. When De is large, the object is a solid, when it is small, the object is a liquid. Water observed on a picosecond timescale is a formidable solid[9] and a mountain on a billion year timescale is a liquid. Reduce the temperature and water becomes a solid with a glacial relaxation time, and at high temperatures a mountain becomes lava which can be a liquid even at a ms timescale, though chemistry and temperature may require measurements on timescales of seconds, minutes or hours[10].

If, at normal temperatures and timescales we have a liquid, our default is to measure its viscosity. If we have a solid, our default is to do tensile testing for its modulus. If you are reading this guide then you probably have a semi-solid

---

[8] The delightful note on παντα ρει and the origins of the Deborah number can be found in Markus Reiner's account in *The Deborah Number*, Physics Today, 17, 62 (1964). There is a problem faced back then which has re-surfaced in a different form for the 21st century. Reiner used to get letters addressed to the Theology Department. The sure sign that a rheology pdf has been created by OCR is when you find the word "theology" within it.

[9] There are many stories about the lab in Wales that measured the high-speed properties of water, using a gun to fire bullets into the water.

[10] Those lavas that seem not to be moving but over a couple of days can spread over a small village.

with a mixture of properties that will vary strongly with temperature and with timescale.

The classic way to get at this mixture of properties is via G' and G''. As these are easily described,we can spend most of this chapter discussing why other ways of looking at the properties of semi-solids should be a normal part of our rheology repertoire and of ways of thinking about materials in general.

4.1 G', G'' and tanδ

The mixture of fear and ignorance about G' and G'' is unnecessary and unfortunate. They are simple to understand and are powerful indicators of key aspects of many formulations. Most of us are happy with E, the tensile modulus, which is simply stress/strain, where the strain is *along* the material sample. G is the shear modulus, stress/strain, with the strain being *across* the material. So there is nothing mysterious about G values in general. The reason we have G' and G'' is that we need to describe the stress response to the strain (or strain response to the stress, which gives us equivalent terms J' and J'' discussed later) in terms of a purely elastic term, i.e. as if it were a pure solid, and a purely plastic (or viscous) term as if it were a pure liquid. We can best describe how this comes about with an app:

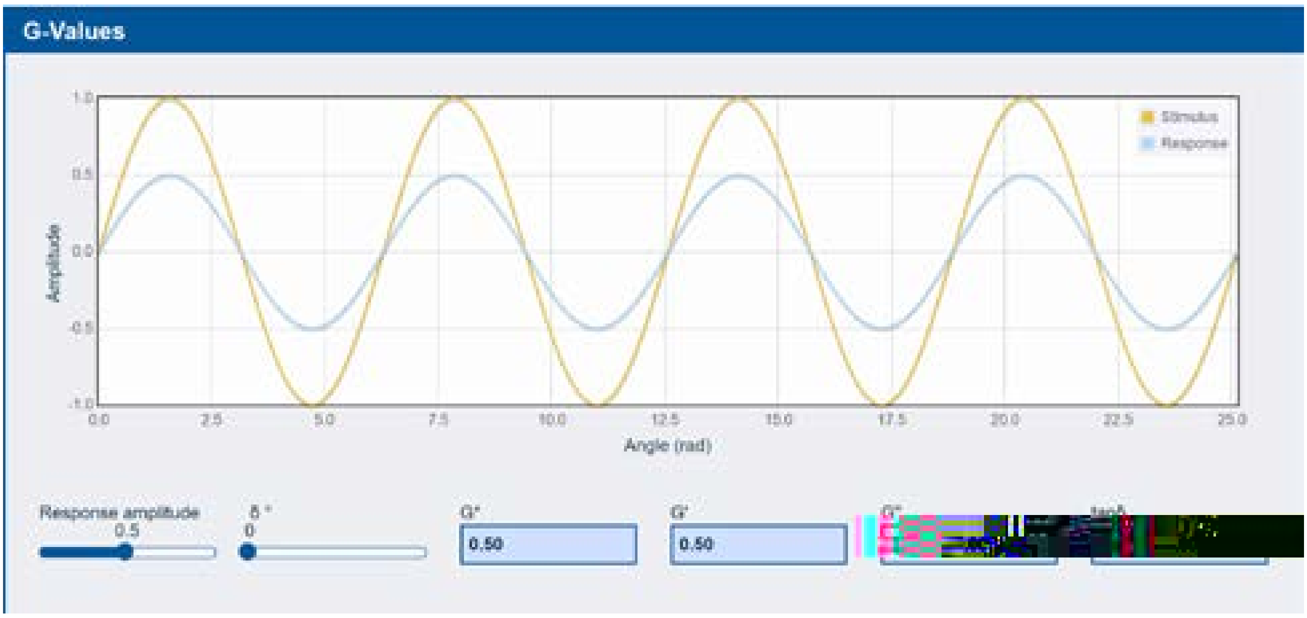

Response amplitude 0.5

δ ° 0

G* 0.50

G' 0.50

App 4-1 https://www.stevenabbott.co.uk/practical-rheology/G-Values.php

We have the sample squeezed between two plates, one of which can oscillate. We might oscillate it with a known maximum stress or maximum strain, creating a stimulus. We measure the response either as a stress on the lower plate (for oscillating strain) or as a rotational speed (strain rate) on the plate to which a known stress is applied. These are, respectively, the controlled strain and controlled stress rheometers, with endless discussions about which is better or better value for money.

For the measurement of G' and G'' the official tradition has been to oscillate the strain and measure the stress; doing it the other way round, oscillating the stress, is equivalent to measuring J' and J'' discussed briefly below. Because precise conversion between G and J is possible, we can feel free to choose whichever type of machine best suits all the other things we need from our rheometer. In any case we either have whatever rheometer we have, or modern technology is making the differences in price and performance between controlled stress and controlled strain machines ever smaller.

In the image above, the stimulus creates a response of amplitude 0.5 (again, we don't care about the number or units at this stage). The important point is that the response exactly follows the stimulus - it is "in phase", with a phase angle, ŭ

The other number shown in the app is tanŭ, the tangent of the phase angle which is G''/G'. When ŭ = 45°, tanŭ = 1 and we have equal components from the elastic and viscous elements.

Armed with this straightforward information you can now tell at a glance from a typical oscillatory rheometry plot how the system behaves as parameters (such as temperature) are systematically changed. If tanŭ < 1 then the system is dominated by elastic behaviour and when tanŭ > 1, the system is dominated by viscous behaviour.

When, for example, the G' line dips below G'', that is an indication of a significant change. We have seen this in one form of yield stress measurements. For classic polymers the temperature at which this happens is the glass transition temperature, Tg. When you find the conditions on the graph that apply to your formulation you will know whether it is mostly elastic or mostly viscous.

4.1.1 Storage and Loss

So far I have described G' as elastic and G'' as viscous. I could also call G'' plastic, which allows us to talk about the "elastic and plastic" components. Plastic in popular language means something that doesn't easily flow (a plastic bag returns to its original shape once the shopping is removed from it), but in technical language it means irreversible flow as in "plastic deformation".

The other common terms are "storage" for elastic and "loss" for viscous. A purely elastic material stores the energy under strain and returns 100% of it on release. A purely lossy, viscous material has no recovery after being strained and all the energy has been lost as heat.

You can feel free to choose whichever pair of terms pleases you; you just need to be aware that others might choose a different pair.

There is, however, a catch. To say that your formulation is "mostly elastic" or "mostly viscous" is true only for a specific set of conditions. Your formulation is likely to pass through a wide range of conditions so you need to know how the relative elastic and viscous proportions will change. For that we need to return to the Deborah number.

4.2 TTE/TTS/WLF

whichever way you prefer to think about it. Similarly, at normal frequencies, water has a respectable G' below 0°C and rock has a nice G'' above 1000°C, so we have a temperature dependence as well.

Returning to reality, in a typical rheometer we can measure only a narrow range of frequency and temperature. Let us say that our specific rheometer can only cover 0.01 to 10Hz and -10 to 80°C. Fortunately, we can create measured values over a much larger *virtual* range. For example, a measurement made at 0.01Hz at 80°C might be like a measurement made at 0.0001Hz at room temperature. Think about it; at high temperatures things flow and at low frequencies things flow, so the 80°C measurement is rather like a very slow measurement at room temperature. Going in the other direction, a measurement at -10°C and 10Hz might be like a measurement at 100Hz at room temperature, because things are more rigid at low temperature and at high frequency. So our limited rheometer can now, effectively, measure from 0.0001 to 100Hz.

The fact that time is equivalent to temperature or that you can superimpose time-related data with temperature-related data is well-established across a remarkable range of practical materials and you can generally assume that it applies to your specific formulation. We therefore invoke TTE (Time Temperature Equivalence) or, according to local preferences, TTS (Time Temperature Superposition). If you want to appear super-smart you can call these tTE and tTS because time is usually a small t and temperature is usually a large T. To do

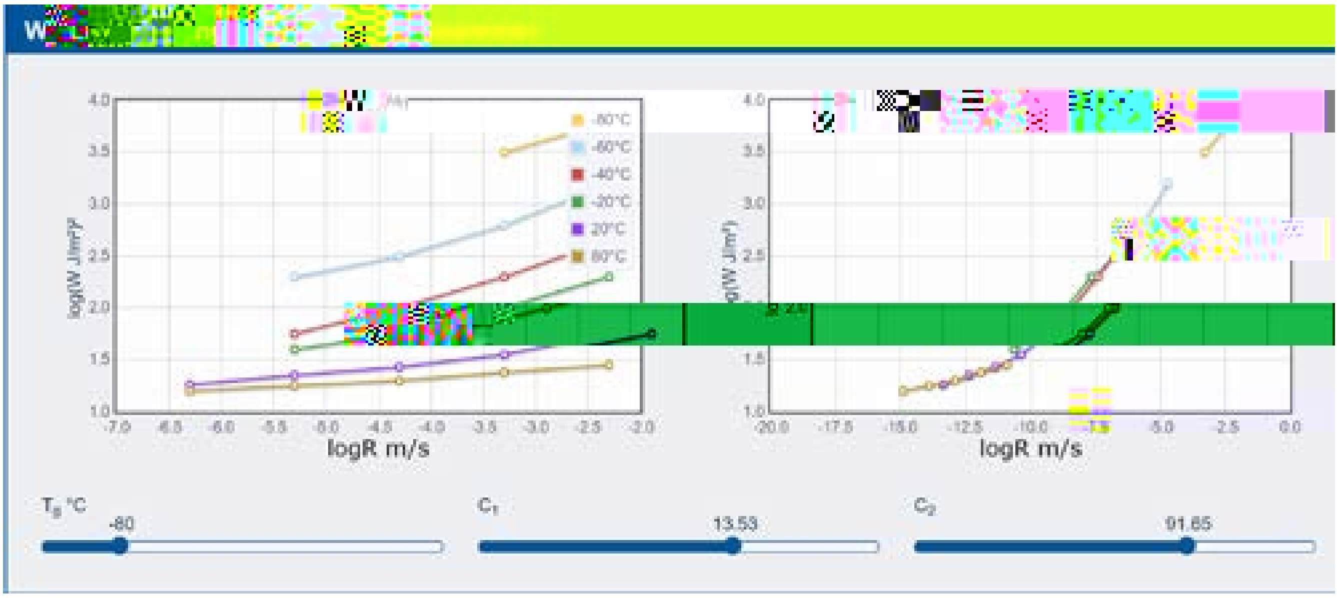

App 4-2 https://www.stevenabbott.co.uk/practical-rheology/WLF.php

The data on the left aren't rheology data, though they could be - WLF is a universal phenomenon. They happen to be a set of adhesion peel energy curves versus the rate, R, of measurement done from -80 to 80°C. At high temperatures the adhesion is low because the polymer is soft, at low temperatures it is high because the polymer is rigid. Adhesion increases with peel speed. So we have 6 curves which are very nice, but hard to fit into a big picture of what is going on. By knowing that in this case T_r = -80°C and by playing with the C_1 and C_2 sliders it is possible to create a single WLF curve, where the colours of the different portions (which can overlap) relate to the colours in the original graph.

The WLF curve tells us what the peel would be at -80°C at a rate of 10^{-14} m/s. And that is a nice give-away. You often see impressive G':G" plots spanning a vast frequency range and you can imagine that they must have an amazingly good rheometer - until you think about the timescale on the graph.tualo3FEFF0020>BDC

[illegible]adĐsYcsYdCV secondsrsas whrdd tha yearLW7q–þpLW%aAi0LW"à–Þ•™[illegible]7ers

I had been talking to a team working on a pressure sensitive adhesive, PSA. I had pointed out in theoretical discussions how the WLF version of the results would help understand the adhesive properties.

I

If we all had full plots such as these for our formulations we would know how they behaved with sudden shocks (high frequency) or over large timescales, or low or high temperatures. These plots are a super-compact way of providing answers to a lot of questions - within the limits of the assumption behind the measurement which is of small strains.

For those who remain cynical about all this and regard WLF plots as some sort of physicist's game unrelated to reality, I have one more thing to say about PSA. You can create Burgers-style G':G" plots via WLF. But what do they have to do with the important issue which is the adhesion of the PSA? The answer is "a lot". You can measure the same PSA's adhesion over a large range of temperatures and times and create a single plot via WLF. When you do so, you find that the three parameters, C_1, C_2 and T_r are identical for the rheology and adhesion datasets. So at a very fundamental level, adhesion, WLF and the Burgers model are inter-related.

4.3.1 Your own G':G", WLF insights

You may recall the idea from an earlier chapter of becoming a superstar by being able to do rough calculations of shear rates and, therefore, knowing what your viscosities might be at different parts of the product's lifecycle. You can (and should!) do a similar analysis for your semi-solid formulations. For example, for PSAs it turns out (and we will discuss it more in the next chapter) that even such an approximate analysis is of great help to the formulators. Here is the logic:

- What are the timescales of PSA processes? The key process of sticking has to be around 1/s, you need it to stick "immediately". Now imagine peeling a piece of adhesive tape for an adhesion tape test at 1 cm/s. The thickness of the adhesive is (in round numbers) 0.01 cm, so we are ripping it apart at 1/0.01 ~ 100 /s. If we apply a weight to the tape so that it is being sheared, we are worried about velocities of 10 cm/day, or around 10^{-4} cm/s, giving us $10_{-4}/0.01$ ~ 0.01 /s.

This rough analysis tells us that at the very least we should understand the G':G" behaviour over a 0.01 to 100 /s timescale. A large G" at the 0.01 /s timescale would be as catastrophic (easy shear) as a small G' at the 100 /s timescale (insufficient peel strength). It also turns out (the so-called Dahlquist criterion, https://www.stevenabbott.co.uk/practical-adhesion/dahlquist.php) that a G' greater than 0.3 MPa at 1 /s renders a typical PSA too slow to stick quickly.

As has been shown many times in the PSA world, getting the right G':G" at all three timescales does *not* guarantee success - PSAs are far too complex. On the other hand, having any of those three wrong *does* guarantee failure.

I can therefore say with certainty that if your product has to perform at different timescales (and few products don't), you will gain a large benefit from a quick analysis of what those timescales will be and what general G':G'' values are needed. By meeting those general requirements, success is *not* guaranteed; failing to meet them *is* guaranteed to render your formulation sub-optimal.

It is highly likely that to get the values at all relevant timescales will require some WLF work. Trust me; if I can do it with zero notice late at night with no app to help me, you will have no problem doing it if you can create the quality time away from daily pressures of formulation to become comfortable with a WLF analysis of your own system. Your rheometer probably has some presets that will

5 Interconversions

What problems are we trying to solve?

Semi-solids can do many things such as flow, creep and relax. The first problem is that if you think of each behaviour as something separate (e.g. you do G':G'' measurements and, on a different machine, creep measurements) you are missing out on a lot of information. The more important problem is that each type of measurement has a natural, convenient timescale so it might not be possible to get, with one technique, the data you want concerning a timescale that is important in your formulations.

It turns out that all these behaviours are mathematically equivalent so, in principle, a measurement of, say, creep over one timescale, can tell you about G':G'' behaviour in cases where that timescale is not readily accessible. It also turns out that we can get a grand view of these processes via "spectra" with respect to time. We therefore have the potential to combine and interconvert our separate measurements into whichever views provide us with the fullest insights into our formulation challenges.

Everyone, except me, knew that it was possible to convert between rheological measurements. Once I had found this out for myself I read many papers saying why interconversion method A is no good and how method B is better, without actually telling me how to do it. This is a reflection of the academic ethos which accepts that it is OK to publish a complicated paper on a topic that is said to be of real-world importance (which justifies doing the research) without having to include any consideration of whether anyone will ever use that research. All of us who formulate could really benefit from being able to grasp interconversions and use them to our advantage. Yet the academic community seems not to have cared, other than (and I applaud this) providing some algorithms to the few companies large enough to sell rheometers and the associated software.

I do not have the brain power or mathematical training to do the job that the specialist academics should have done. But at least I have made an interconversion app that shows the possibilities and discusses the principles. My own grasp of the key inter-related concepts in rheology has improved greatly during the creation of the app and my aim in this chapter is to persuade you why you should be interested in the ideas and, ideally, influence the rheology community to make interconversions more usable for us all.

First we must understand about relaxation and creep.

5.1 Relaxation and creep

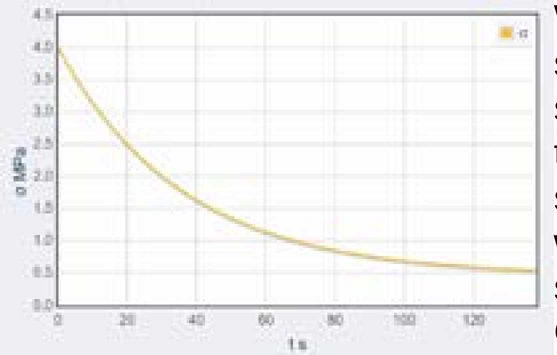

When you apply a sudden, fixed strain to a sample it will start to relax and the measured stress will decrease. The image shows a typical plot of stress ů over time. If the initial strain had been doubled, the initial stress would also have been doubled. So you also see equivalent plots of Relaxation Modulus G=ů/strain. Such plots are Relaxation curves.

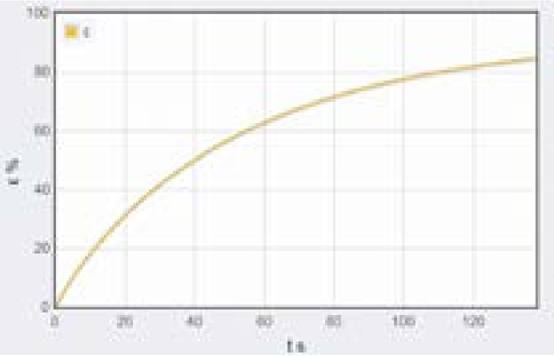

If you apply a fixed stress to a sample, it will suddenly start to creep, i.e. the strain increases. The image shows a typical plot of strain (here shown as Ů for historical reasons) over time. If the initial stress had been doubled the resulting strains would double. So you also see equivalent plots of Compliance J=strain/ů. Such plots are Compliance curves.

You would think that looking at stress for constant strain and strain for constant stress would be highly inter-related - and they are. However, in just about every reference I could find, they were discussed via entirely different equations, implying that they are totally unrelated.

Relaxation is typically modelled using a Maxwell model:

Equ. 5-1

$$\frac{\delta\gamma}{\delta t} = \frac{\sigma}{\eta} + \frac{1}{E}\frac{\delta\sigma}{\delta t}$$

What this tells us is that the rate of change of stress with time (which is what interests us) depends on the modulus (here we use the tensile modulus E but this could equally be the shear modulus G) and on the ratio of the current stress ů to the "viscosity" ɖ. It also depends on the change of strain, ɔ, with time.

Creep is typically shown via the Kelvin-Voigt model:

Equu. 5-2 $$\frac{\delta\gamma}{\delta t} = \frac{1}{\eta}\left(\sigma - E\gamma\right)$$

As with Maxwell, we are interested in the response to a sudden increase to a constant

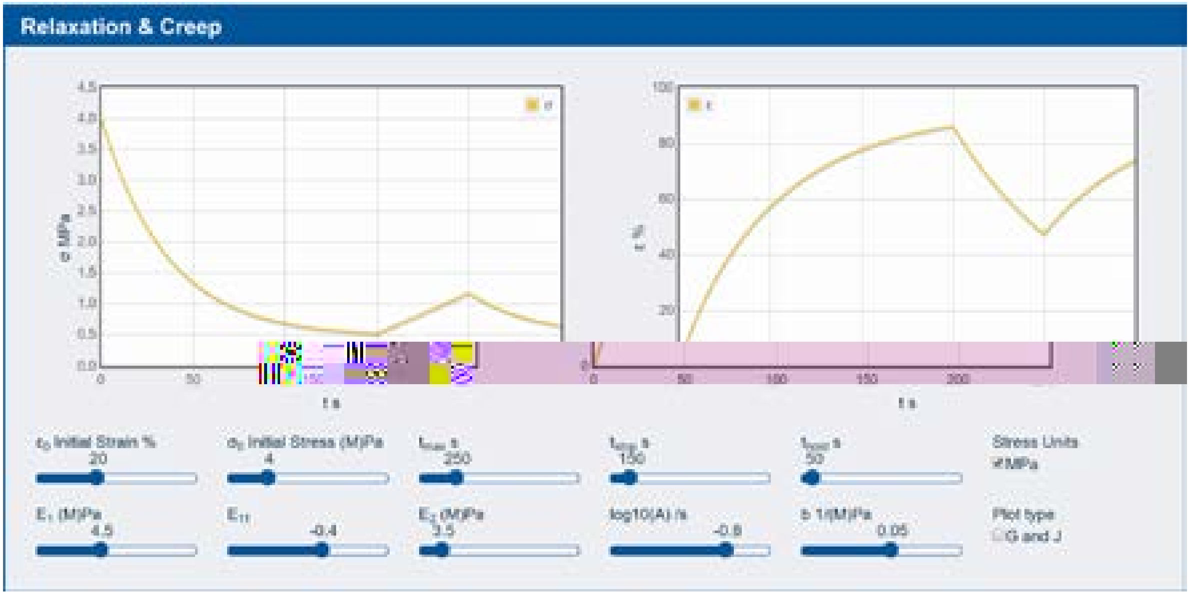

ε_0 Initial Strain % 20

σ_0 Initial Stress (M)Pa 4

t_{max} s 250

t_{stop} s 150

t_{hold} s 50

Stress Units ☑ MPa

E_1 (M)Pa 4.5

E_{11} -0.4

E_2 (M)Pa 3.5

log10(A) /s -0.8

b 1/(M)Pa 0.05

Plot type ☐ G and J

App 5-1 https://www.stevenabbott.co.uk/practical-rheology/Relaxation-Creep.php

We can now play with the various stresses and strains, timescales and spring and dashpot parameters. Because such experiments are carried out on relatively fluid materials using low stresses in the Pa range and on semi-solids with stresses in the MPa range, you can choose the units - these merely change the graph labels.

As mentioned at the start of the book, I am being inconsistent: strains are shown as ε rather than γ. In this sort of plot it would be odd to use γ. As I said at the start, rheology is full of inconsistent units.

The curious shape of the curves comes because after 150s (see the t_{stop} slider) the strain or stress is reduced to zero and the system starts to return to its original value via its elastic response though clearly it will never fully return because the viscous component is irreversible. After 50s, (t_{hold}) the stress is re-applied and the strain or stress increases further. Such experiments are very much part of the repertoire of relaxation and creep measurements, though I currently have no idea how you use the data from multiple starts and stops.

Earlier I said that all real-world relaxation and compliance curves can be described with a series of Maxwell or Kelvin-Voigt G:ɖ pairs. The app has bypassed this by using more complex springs and dashpots. In reality the stress/strain-dependent elements could have been emulated via a series of pairs. This fact will be key to the discussions in the next section.

As mentioned above, if the relaxation experiment is done with a different strain, then the recorded stress will be different. For example, doubling the strain will double the stress. Similarly, if the creep experiment starts with double the stress then the strains will double. If you need to compare samples measured at different strains or stresses, just looking at the plots will be unhelpful. We can solve this problem. Instead of plotting ů we can plot ů/Ů, as G(Ű) the Relaxation Modulus at any given timescale Ű in units of Pa which will give Relaxation curves that are identical over a reasonable range of strains. And instead of Ů we can plot Ů/ů, as J(Ű), the Compliance in units of 1/Pa and the Compliance curves, again, will be identical over a reasonable range of stresses. Using G(Ű) or J(Ű) curves we can therefore make meaningful comparison between samples which might have been measured at different strains or stresses.

What happens when linearity of response breaks down is discussed in a very brief chapter.

These two plots are based on elastic and viscous components. So they must be related somehow to G' and G''. The timescales in the screenshot are in the 100s range, so 0.01HZ. A typical oscillatory measurement of G' and G'' takes place in the 0.01 to 10Hz range, So these data should have some links to low frequency values. If we did a much slower creep test (and why not?) maybe we could get G':G'' data in the super-low frequency range. If we wished to know about relaxation and creep behaviour in the 0.1s timescale (which is tricky because the assumption of "instant" stress or strain becomes difficult to justify) we could, perhaps, use the 10Hz G':G'' data.

5.2 The power of interconversions

As mentioned above, all real-world relaxation curves are constructed from G:ɖ pairs that work over different time-scales. If we happen to know all such pairs across the whole time domain then we can plot the relaxation modulus, G(Ű) , for all times from microseconds to macroseconds. We also know that all G':G'' measurements are frequency dependent. It would be rather helpful to see which are the frequencies where "interesting" things happen and which are those where there is not too much of interest.

If we had such a plot of "interest" versus frequency we could then think through the timescales of the processes that are important to our formulation and know with confidence that our system has the chance of giving the correct response. What does "interesting" mean, and what would be the correct response? Only you can answer those questions - it is your product, not mine.

My shock at learning that such a plot was possible and (in theory at least) routine, was mostly to do with the fact that because I had been unaware that such a thing could be done, I had never asked myself the sorts of questions that the technique would have answered. If I had known, for example, that a rapid change was taking place at a timescale important to my process, I would have been alerted to the process being on a knife-edge and could have devised methods to shift the peak to a different timescale.

The plot of H (representing the degree of interest) versus t is called the *relaxation spectrum*. I hope, therefore, that you will take the trouble to learn about H (and its equivalent in terms of creep is the compliance spectrum, L), and find out how to obtain them.

Before looking at relaxation spectra we need to think more about creep-style measurements. The compliance curve, the creep strain experiment normalised by the applied strain, can, just like the relaxation curve, able to assemble a good-enough collection of algorithms by combining Ferry, Tschoegl and Park.[13]

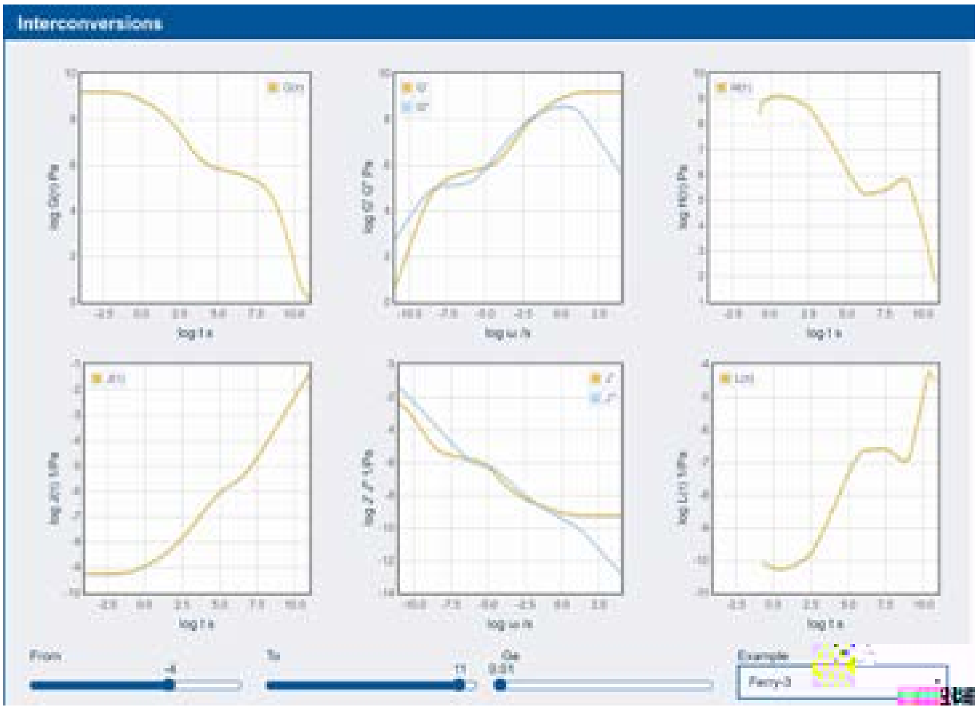

From -4

To 11

Ge 0.01

Example Ferry-3

App 5-2 https://www.stevenabbott.co.uk/practical-rheology/Interconversions.php

What the screen shot doesn't show is the rather messy method for creating the inputs - a series of G(Ű)

Because compliance is, at heart, just 1/relaxation, it is not surprising that the curves look somewhat like mirror images. If it were that simple then life would be much easier. In fact the conversions are not at all straightforward and the curves, at the detail level, are not mirror images.

Some of the curves have calculation artefacts. That is partly a limitation of the algorithms I used and also a characteristic of the problem of these conversions. I have the advantage of data from Ferry, covering many orders of magnitude of timescale, which makes it possible to show the full conversion capabilities. In real life you will have only a limited set of data unless you do extensive WLF conversions. The fewer the number of datapoints, the more likely the conversions will contain artefacts. As many papers stress, some of these conversions are "ill-posed problems" which means that there is no unique solution.

Out of personal interest I attempted to create the interconversions for the Burgers model shown in the previous chapter and discussed further below. You can judge my attempt by selecting the option from the combo box. Although I can convert from a series of G(Ű) values to G':G" I don't have a method to go in the other direction so that we could take any G':G" dataset and create the full interconversions. To get these G(Ű) values I used trial and error. Automating the G':G" conversion may be a future challenge.

The difficulties of interconversion might make it hard to realise my dream scenario which is the ability to combine data from whichever techniques happen to be convenient for different timescalf

realise in ! نم ҳ ّ

why interconversion is so important, simply because I know a lot about them and already have a good model for what they should be like. If I had more experience in other areas I could equally have used them as examples because the principles are general. Macosko's *Rheology Principles* book is especially good for those who want to explore other implications of entanglement.

The common idea of a strong adhesive (such as an epoxy) is that it gets its strength from having a high elastic modulus. A polymer with lots of entanglement (such as a cross-linked epoxy) will behave as a pure G' over the normal timescales that we typically measure. Such a polymer would make a useless PSA because these have to be soft and accommodating to the surface - they must have a G' no larger than 0.3MPa, compared to a typical epoxy's 4GPa.

Any polymer without entanglement is going to behave with a large G'' element as there is nothing much to stop the polymer chains from sliding past each other. A PSA without tangles might still achieve the 0.3MPa G', while having a G'' significantly higher. As an adhesive it will be useles-

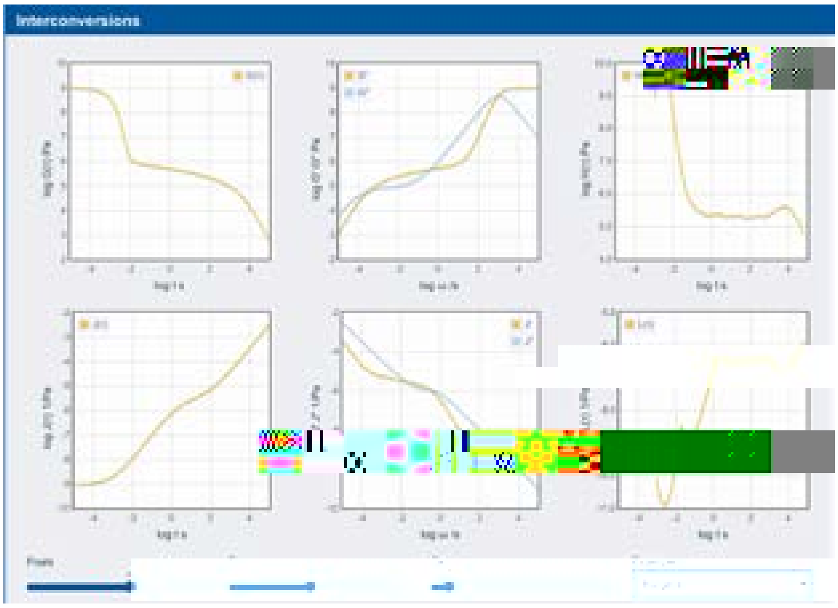

Figure 5-1 The Burgers Viscoelastic model for a PSA fully interconverted

What I especially like about this interconversion is the ability to think better about creep behaviour. It is rather easy to design an excellent PSA which fails in the application because of creep. The standard "creep" test is simply "how long does it take for this sample to fall from the test rig", providing very little information about what is going on. The interconverted creep curves allow me to think more about the timescales over which creep will be a liability. If the PSA world routinely looked at creep this way, it would be much more knowledgeable than it currently is.

5.4 M_c from Likhtman-McLeish theory

Suppose you have a full G', G" curve covering many orders of magnitude. There is one more thing you can do with it - work out the critical entanglement molecular weight M_c. The theory was developed by Likhtman and McLeish[15] some years ago, but most of us were totally unaware of it because it's the sort of thing known in the specialist M_c community, not among those of us who need it. An appified version, https://www.stevenabbott.co.uk/practical-rheology/L-M_Mc.php, was made possible by the kind assistance (acknowledged on the app page) of those who are supporting and developing this, and other theories for the future. Indeed, they have provided L-M theory as part of the impressive RepTate software [illegible] https://reptate.readthedocs.io/ [illegible]

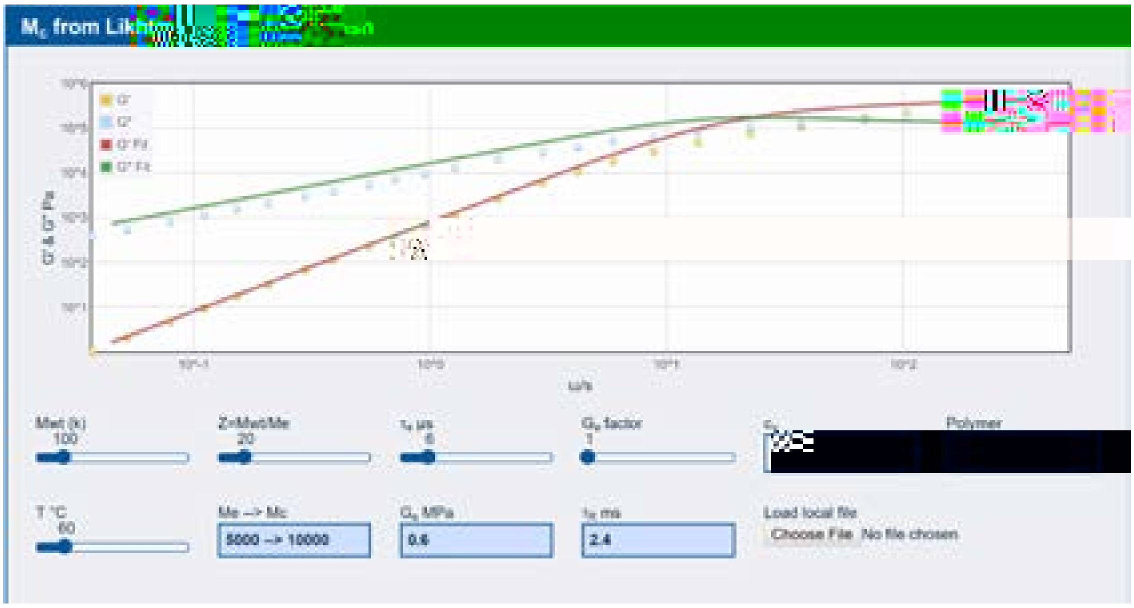

Mwt (k) 100

Z=Mwt/Me 20

τ_e μs 6

G_e factor 1

T °C 60

Me --> Mc: 5000 --> 10000

G_e MPa: 0.6

τ_R ms: 2.4

Figure 5-2 The Likhtman-McLeish route to finding M_c values.

The app is mentioned here because L-M theory provides us with two key outputs - the M_c itself and a Rouse relaxation time, τ_R. That relaxation time is based on on input value, τ_e, which is the entanglement timescale. The other inputs are MWt of the polymer being measured and Z, the number of entanglements per that MWt. What the app does is take those 3 inputs and calculates the G':G'' curves. What you do is move the sliders till the calculated and experimental values are a good-enough match. Although it would be possible to add an automated fitting, I find it more instructive for the user to see the trade-offs between the different possibilities. There is an extra fudge factor (it shouldn't be needed, but the experts recommended it) to adjust the maximum G' value, G_e. This makes the fitting easier.

The paper itself isn't so hard to read and the formulae not so hard to understand. Unfortunately one of the formulae is near-impossible to calculate. So in the app I make use of (as does RepTate) some heroic pre-calculations of curves depending on all (relevant) possible input values. The app, therefore, is not doing any calculations, just showing the curves expected from the current parameters. This means that we have a very powerful app without the need for vast computational power.

The science behind the app can be described very briefly. The high-frequency G' depends only on the number of entanglements - the more entangled, the higher the modulus. The low-frequency effects depend on τ_e, because, as we know from the interconversions discussion, relaxation times are related intimately to the other possible curves.

If (big if) you have a polymer with a low polydispersity index (<1.1) and if (big if) you have a full WLF set of G':G" data then the fitting is easy and if (big if) you have the WLF data corrected for a temperature that is a reasonable balance between T_g and T_m the results are sound. Most of us, most of the time, have polydisperse polymers so what should we do? The purist answer is to wait till some newer, better theories are available (and added to RepTate). My answer is to be pragmatic and get at least a working idea of what's going on with your polymer via a pragmatic M_c.

5.5 Impossible interconversions?

The book started in the familiar world of rotational viscosity, where we can go from shear rates of 0.01 to 10,000 with no great difficulty. It seemed obvious to me that there was one more interconversion we needed - taking the rotational viscosity data and using it to fill out tricky parts of the six interconversions.

I realised that I had never come across such a conversion methodology and was rather mad at myself for having missed something so important. I then did a literature search and came up blank. Then I asked a rheologist guru how to do it. He laughed and told me that not only was it impossible but that a further conversion, from extensional viscosity was also impossible - indeed, it often wasn't possible to convert extensional viscosity results between different ways of measuring the same material. This point is nicely described in the opening of the abstract of a famous paper[16]: *"The issue of whether extensional viscosity is a concept that causes more confusion than enlightenment is addressed".* The conclusion of the paper is that extensional viscosity causes plenty of confusion and that reducing the confusion looks to be a difficult task.

I will say no more about extensional viscosity other than to express my regret that a technique which should be so relevant to so many formulation issues remains unfamiliar to most of us and, apparently, so difficult to master.

I am genuinely astonished that after decades of rheological studies, there is (apparently) no methodology for combining rotational and oscillatory rheology. Readers will say "What about the Cox-Merz rule?" The answer is that it is not a rule and if it applies at all it is to a rather narrow domain (and does *not* apply at all to particle systems[17]). Cox-Merz says that the viscosity, η, at a given shear rate is equal to the complex viscosity, η^*, when measured at the equivalent frequency, ω. As a reminder, G^* is the complex modulus and the complex viscosity is given by $\eta^* = G^*.\omega$. There are at least two reasons why Cox-Merz is not a rule. First, there is no theoretical reason why it should be the case (i.e.

---

[16] Christopher J.S. Petrie, *Extensional viscosity: A critical discussion*, J. Non-Newtonian Fluid Mech., 2006, 137, 15–23

[17] For those who are comfortable with "shear stress equivalent inner shear rate", Cox-Merz can be applied to particulate systems. I confess to not having the slightest idea what that phrase means. Those who are interested can find the ideas in the work of Prof Gleissle in Karlsruhe.

you cannot derive it from first principles). Second, ɖ is a real number and ɖ* is a complex number - you cannot just declare them as being equal.

5.6 What can we do with interconversions?

My view is that the ability to convert between (at least) the six views of the same data is of huge importance, that it is just about possible and that, despite numerous academic papers, this ability is very little understood and used. Also the potential for combining values from different techniques is enormous yet has been, as far as I can tell, hardly used. In addition we have extensional viscosity which isn't even self-consistent within itself. And we have rotational and oscillatory techniques, the twin pillars of rheology, that have no way to communicate with each other.

This means that the current answer to the question of what to do with interconversions is "not much".

If I were a rheologist I would see this as the grandest of grand challenges. Whoever could bring all those things together into a workable tool would transform our ability to better understand our systems via a series of systematic measurements.

So despite being at times rather critical of rheologists, I shall end this chapter on that positive note: maybe someone will take up the challenge, gain justified accolades, get the approach appified, and everyone will be much better off. Or, at the very least, the chapter might have encouraged some readers to go to the little-used Interconversions portion of their rheometer software to see what it can do for them.

6 Beyond linearity

What problems are we trying to solve?

Our rheology experiments are performed rather gently so we do not end up in "non-linear" domains where the standard analysis tools fail. Our real world formulations often operate within non-linear domains so ideally we would routinely be experimenting in these wilder areas.

This is a very, very short chapter.

Our real-world formulations are likely to be subjected to large stresses and strains. Yet our rheological measurements generally stay in a comfortable range of small stresses and strains - for good reason. The theories make sense within small stresses and strains where the elastic response is linear. The G' and G" analyses that we use most often are very clearly limited to small strains so any standard analysis falls apart when the amplitudes get large enough to represent real-world stresses and strains.

6.1 LAOS

Most modern rheometers can now carry out Large Amplitude Oscillatory Shear (LAOS) measurements. With a few Lisajou figures and Fourier transforms, a lot of interesting information becomes available. In planning this guide I asked a rheologist guru how much effort I should put in to understanding and explaining LAOS. His reply was clear: "Don't bother. It is early days, it is very tricky, and you and your readers already have enough problems with standard rheology". I am happy to accept that advice.

Because it is clearly a technique of great potential importance, I urge readers to get the friendliest-possible introduction via Prof Gareth McKinley's brilliant TA Instruments webinar at https://www.youtube.com/watch?v=LtoRnBt70wM. I had to use the pause and rewind buttons many times to keep up with the flow of exciting ideas, so make sure you have some quality time to spare if you want to dive in.

If a LAOS expert thinks that this chapter should be greatly expanded, to allow the technique to be explained to a wider audience, I would be delighted to work with them to make it happen.

7 Particle rheology

What problems are we trying to solve?

We already have the rheological methods that apply just as much to particulate

The key is that for particles there is a maximum volume fraction, ű$_m$ called the packing fraction at which point the system is effectively a solid. The frequent big debates about the precise value of ű$_m$ are missing the point. Few of us[19] know our real ű$_m$ values as our formulations are not perfect spheres. The point is that all formulations have a ű$_m$ and we can either stay far away from it so that small changes in ű$_m$ make little difference, or we adopt a few tricks to change it to a higher value to keep viscosities and yield stresses as low as possible.

The most common equation for describing the effect of ű is Dougherty-Krieger. Others, such as the one from Pal, may or may not be superior in some ways, but again it hardly matters for our real-world formulations of imperfect particles:

Equ. 7-1 $$\eta = \eta_0 \left(1 - \frac{\varphi}{\varphi_m}\right)^{-2.5\varphi_m}$$

The important point is that as ű app" M PRVWD

systems in the 2-3 range give lower yield stresses than systems with bad dispersants in the 8-10 range. The particle radius is r, with smaller particles, other things being equal, giving a higher yield strength as our instincts tell us. The constant k is my fudge value to provide plausible results. A screenshot of the relevant app follows shortly because it is incorporated into the high-shear app thanks to the commonality of many of the key inputs and calculations.

Although the effects of the interaction energy and radius are important (proportional to E and 1/r respectively) the biggest effect is the rapid rise with ű as it starts to approach ű$_m$. So paying attention to small changes in ű can be more rewarding than worrying too much about E or the meaning of r in a system with a range of particle sizes. If ű

7.2 High Shear

We know that viscosity goes to high levels at low shear and high ű. This would mean that high-solids formulations would be impossible to handle. Fortunately, they shear-thin strongly, though caution is required because at very high shear rates some formulations can shear thicken alarmingly.

The theory I've chosen[25] has, not surprisingly, many elements that are common to the many alternative theories found in the literature. It is especially useful for an app because it relies mainly on just two inputs which we already know: $ű_m$ and $ű_c$. The key equations are:

Equ. 7-3
$$\eta = \gamma^{n-1} \text{ and } n = \frac{\left(\frac{\varphi}{\varphi_m}\right)^{2/3} - 1}{\left(\frac{\varphi_c}{\varphi_m}\right)^{2/3} - 1}$$

Other elements of the equation are described in the app - which we need because it is hard to grasp what such an equation will produce. The results turn out to be rather straightforward, at least when viewed as a log plot:

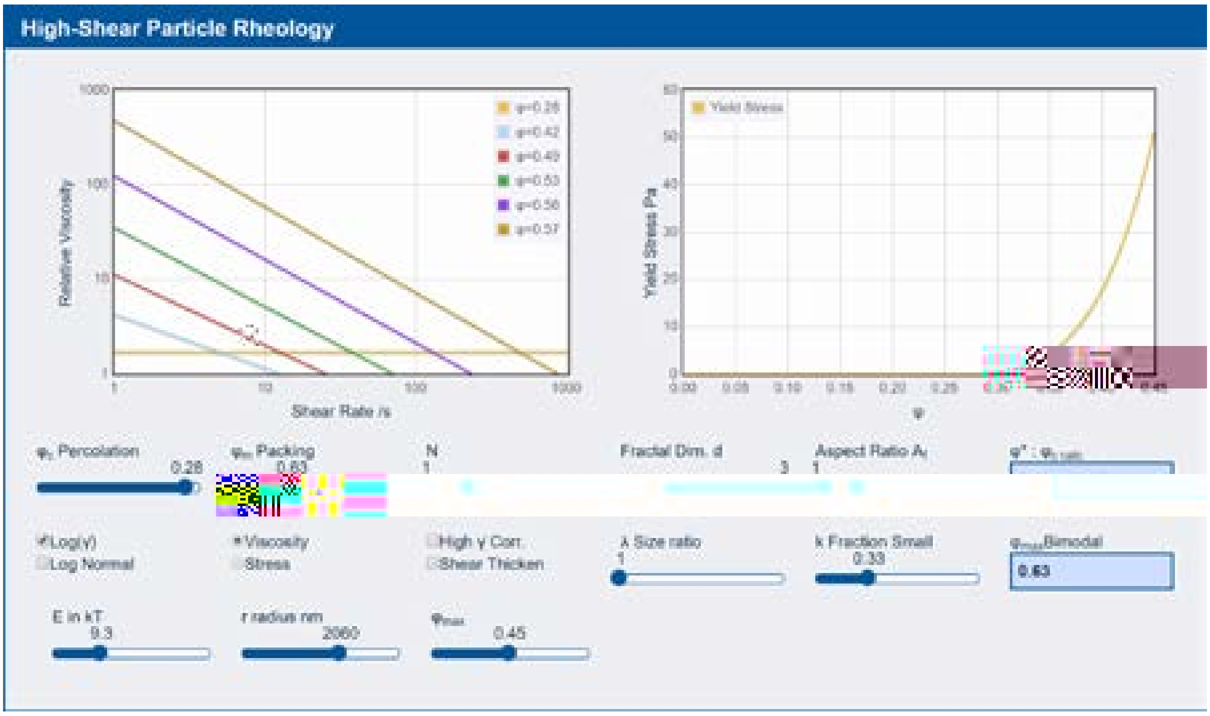

φc Percolation 0.28

φm Packing 0.63

N 1

Fractal Dim. d 3

Aspect Ratio Ar 1

φ* : φc calc

☑ Log(y)

☐ Log Normal

● Viscosity

○ Stress

☐ High γ Corr.

☐ Shear Thicken

λ Size ratio 1

k Fraction Small 0.33

φmax Bimodal 0.63

E in kT 9.3

r radius nm 2060

φmax 0.45

App 7-2 https://www.stevenabbott.co.uk/practical-rheology/High-Shear-Particles.php

The st

slide past each other. The effect is modelled by increasing φ_m with shear rate, an effect well-validated in the literature.

If you decrease the Fractal Dimension from 3 (i.e. the system is no longer pure particles) and increase N, the number of particles in a floc, the viscosity increases sharply. How do you know the fractal dimension and the number of particles in a floc? In general we don't. The Bicerano paper[26] from which I extracted the relationship tells us that a platelet formulation has a dimension of 2, so any normally flocculated system will be something like 2.5, which is the default used in some of the graphs in that paper. In terms of N, you have to decide whether the shear stresses at a high N value are likely to rip the flocculated particles apart or, as confusingly happens, makes further flocculation more likely as we will discuss shortly. To explore the effects, values of N from 2-5 seem to be reasonable.

Why do particulate systems shear thin so much? It is because the shear drives the particles into neat lines like vehicles flowing down the lanes of a highway rather than annoying lane-hopper motorists moving between lanes and causing everyone (including themselves) to slow down.

Why do some systems, after shear thinning as normal, shear thicken at very high shear rates. Think of this as a sort of Bernouilli effect where the air flow between the vehicles sucks them together, creating a massive pile-up: particles are structured across the rheometer gap, causing high viscosity. This is the transition from "hydrodynamic flow" to "lubrication dynamics" or "hydroclustering". There seems to be no convenient formula for this (and the effect is very formulation-dependent), so the Shear Thicken option merely illustrates the effect. This hydroclustering is discussed further in the shear thickening section.

7.2.1 Increasing φ_m

Even a small increase in φ_m can help with a high solids formulation. Increasing the aspect ratio makes this happen at high shear, with problems down at low shear thanks to the decrease in φ_c, so this is not a strategy to be encouraged, even if you could change the aspect ratio which, generally, you cannot.

Instead, the standard trick for our mostly-spherical formulations is to make sure that smaller particles in the formulation can fit in the holes created when the larger particles pack.

---

[26] Jozef Bicerano, Jack F. Douglas, and Douglas A. Brune, *Model for the Viscosity of Particle Dispersions*, Journal of Macromolecular Science, Part C, 39:4, 561-642

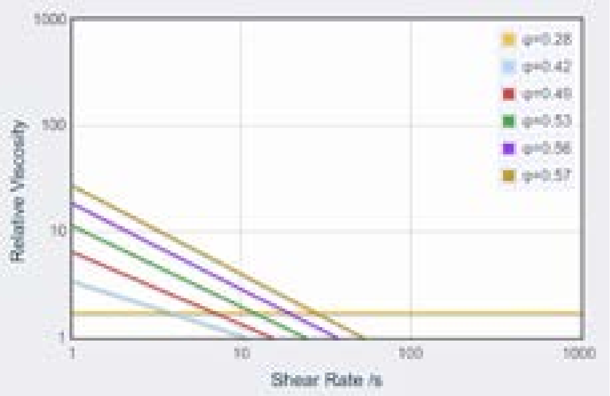

For a simple binary distribution it is well-known that there is an optimum size ratio and an optimum ratio of small to large particles. For a 10:1 size ratio with 33% of the

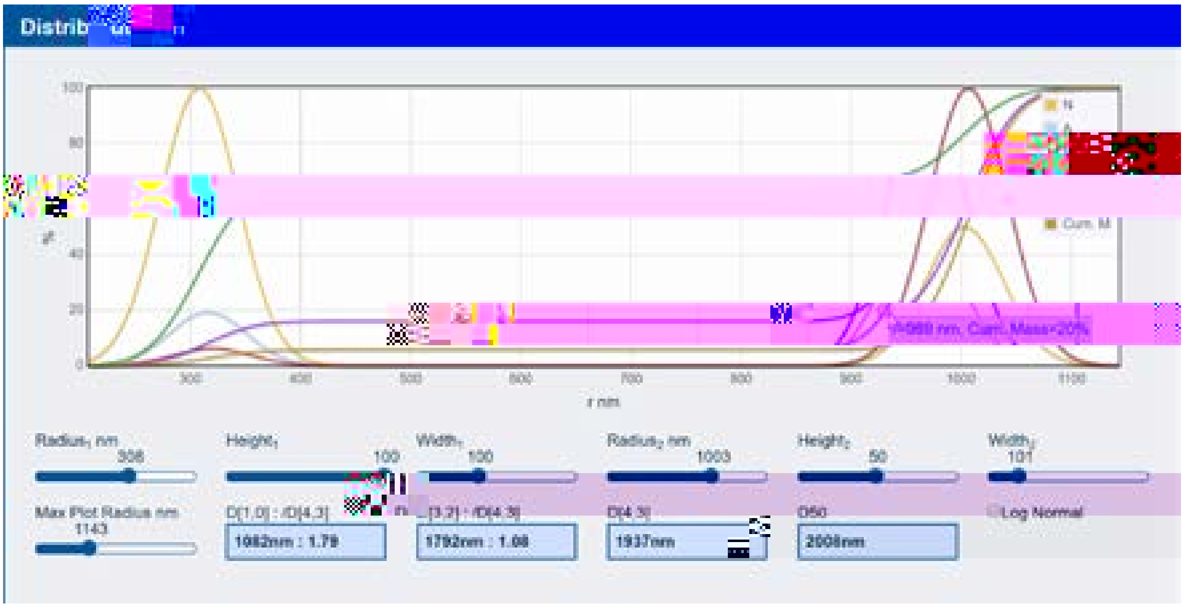

Radius₁ nm 308

Height₁ 100

Width₁ 100

Radius₂ nm 1003

Height₂ 50

Width₂ 101

Max Plot Radius nm 1143

D[1,0] : /D[4,3] 1082nm : 1.79

D[3,2] : /D[4,3] 1792nm : 1.08

D[4,3] 1937nm

D50 2008nm

☐ Log Normal

App 7-3 https://www.stevenabbott.co.uk/practical-rheology/distribution.php

Here we have a rather unlikely distribution, created only for visual convenience. We have a large number of 300nm particles and a set at 1000nm of half the number (the relative "Height" values are 100 and 50). They are the same 100nm nominal width. The yellow "N" for Number curve shows exactly what I have just described. The "Cum. N" curve shows the cumulative number and we find that 66% are below 400nm and 34% are above 900nm.

Now look at the M for Mass (or it could be V for Volume) curve. There is a tiny peak around 300nm and a huge one at 1000nm. The cumulative curve tells us that only 6% of the mass is below 400nm and the 94% is above 900nm. The A for Area curves are in between.

If your focus is on volume fractions then the M curve is what you need to know. If you are concerned with yield stress which, according to the YODEL model is proportional to 1/r then because 66% of the radius is in the 300nm range the yield stress will be nearly 3x larger than if you thought that the particles were 1000nm. If you are concerned with the amount of dispersant needed, or are worried about flocculation which depends on radius and surface area, then the A data are important.

So what is "the" size of these particles. If you take the volume average, ("mean diameter over volume"), this is ... And here we hit a typical problem. We have specified inputs in terms of radius and particle sizes are traditionally measured in terms of diameter. So the volume average, shown as D[4,3] is 1937nm, just below what you would get if you had a pure 1000nm monomodal powder. The number average, D[1,0], is 1082nm, a factor of 1.79 smaller than the volume average. Many users like D50 which is a median based on the cumulative volume distribution - half the particles have a diameter less than this value. If you prefer a different different

7.4.1 The Péclet number

You often come across the Péclet number, Pe, in discussions of particles and shear rates. This dimensionless number describes the ratio of the movement induced by the flow to the random movement caused by diffusion. It is often considered as a sort of Deborah number. At high numbers >1, the particle motions are dominated by the flow, at low numbers <1 then diffusion dominates. We can express it in terms of $\dot{\gamma}$, radius r and diffusion coefficient D or viscosity (of the background fluid) η:

Equu. 7-4

$$Pe = \frac{r^2 \gamma}{D} = \frac{6\pi\eta r^3 \gamma}{kT}$$

η cP 10

$\dot{\gamma}_{max}$ /s 18

r nm 253

App 7-5 https://www.stevenabbott.co.uk/practical-rheology/Peclet.php

7.5 Shear thickening

When we pour a somewhat viscous cornstarch dispersion into a bath it is a reasonably normal liquid, yet if we try to shear it suddenly by jumping on to it and swiftly walking, it is highly viscous, i.e. it has shear thickened. If we pause to reflect on the mechanism while in mid stride, we sink to the bottom.

The main problem with the classic explanation for shear thickening is that it can be argued both ways (and often is).

The conventional starting point for the explanation is flocculation - the fairly delicate self-association of particles that is explained as being due to a local minimum in the DLVO balance of van der Waals attractions and steric or charge repulsions. Those who have forgotten their DLVO theory can find a handy app at https://www.stevenabbott.co.uk/practical-solubility/dlvo.php. The good thing about flocculating systems is that because the particle-particle attraction is only modest, it is readily broken apart by shear. Shake or stir a flocculated system such as a gel and it reverts to being a normal fluid. So this argues for shear thinning.

To argue the opposite, take a fairly concentrated dispersion and add extra energy via shear. This forces the particles into more intimate contact than they might like, creating a flocculated system with a much higher viscosity. This is the most popular explanation for shear thickening, even though it is at best limited and at worst wrong.

Both effects are entirely reasonable, so why does shear give us either shear thinning or shear thickening?

It turns out that most systems show both types of behaviour, with the shear thinning effect winning in terms of flocculation. The hydroclustering model mentioned earlier, not flocculation, is the most popular explanation for the sudden increase in viscosity in many systems. The model reflects many of the confusing aspects such as the strong effect of particle shape, e.g. where silica spheres are rather immune to shear thickening while fumed silica (irregularly shaped) shear thickens easily.

The basis of hydroclustering is that when particles are close, the fluid flow around them can either send them flying apart (i.e. break up flocs) or start to form "closed loop" flows that correlate the motions of the particles. As this correlation extends to other pairs, long-range order is built up which seriously gets in the way of the fluid flow, giving the high viscosity.

In fact there are multiple mechanisms for shear thickening. The "dilational" mechanism is entirely opaque to me but seems to apply to some systems. It provides a very sudden onset of super-high viscosity and is the sort of thing that might be useful for new types of, say, body armour. And what is supposed to be a classic example, the corn starch, works via a log-jam effect not at all connected to the hydrocluster and dilational mechanisms.

How do you stop shear thickening from happening? The hydroclusters can only form when the particles are close, so their tendency to flocculate is an indication of problems ahead even though flocculation is not the genuine explanation. Typical stabilising dispersants such as polymers attached to the particles can enforce a large-enough distance to counter both flocculation and hydroclustering. Obviously, shear thickening is more likely as you approach ϕ_m so multimodal systems will be more resistant at a given ϕ.

If you really want to understand all this then you need a good Stokesian Dynamics program, Apparently it all makes sense if you do. Ever optimistic I wondered if I could create a good-enough Stokesian Dynamics app. It seems that it is way beyond my capabilities and the processing power available to a typical app.

One final word about shear thickening. It seems that a reliable method of implementing an academic flame war is to express a strong opinion on shear thickening. I therefore issue a disclaimer saying that what I have written is some sort of attempt to assimilate what Hoffman, Brown, Wagner and the other greats have said, and that you are entitled to your own opinion about what is really going on.

7.6 Can we apply particle rheology to the real world?

I once had to give a training course on topics that were to be applied to highly-filled particulate formulations. Although I could speak with confidence about many aspects of these systems, their rheology was beyond my then knowledge. So I had to give myself a crash course in particle rheology and out of that came the need to write many more apps and, from that, Practical Rheology. During the training, I was struck by how useful it was to be able to keep going back to the principles contained in the apps, along with the key principles of entanglement and relaxation times to cover the very different timescales in different parts of the processes.

As so often happens in science, once you learn something new, you find it useful in many other situations. Liquid chocolate and liquid cement are very different materials, yet the physics that controls the tricky problems of handling them through production are identical. I just had to learn that "super-plasticizers" in cement perform the same function as lecithin in chocolate (they each reduce the particle-particle E value used in the YODEL model) to be able to apply the physics of cement to the physics of chocolate.

Any models that are simple enough for me to understand and appify are unlikely to be a perfect representation of a complex system. What they do is allow "science-based formulation". Instead of saying "Increasing particle concentration increases viscosity and yield stress" which is both true and rather unhelpful, we can say "Increasing particle concentration over this range of concentrations will tend to have *these* effects and to control them we will need to focus on parameters A, B and C. Because parameters A and B are not, for various reasons, adjustable, that means we must focus on C. So what do we know about C in this case ..."

So the answer to the question in this section is a clear: "Yes".

8 Summary

This is a guide, not a book. It represents what I wish I had known over my formulation career. My hope is that it brings out a few messages about which I am passionate.

- Measuring a few viscosities or even a few G':G" values is not good enough if you want to optimize your formulation efficiently or create imaginative new types of formulations.
- Rheologists do us and themselves a great disservice by choosing to create a bewildering number of plots of the "same" basic measurement. They should be encouraged to plot fewer variants in order to reach a more appreciative audience. And we should use the apps to help translate between confusing plots
- Modern rheometers make it easy to carry out many different measurements - with the downside that we can get swamped by alternatives. So we need to think about the links between key formulation properties and the measurements that can provide insights.
- A lot of polymer-based formulation and associated rheology can best be understood via entanglement. This shows up in different forms in rotational and oscillatory rheology
- Interconversions between (at least) 6 different ways of looking at the same data should become standard for all of us. The interconversion app is one indication that this is both possible and desirable.
- We tend not to think in terms of timescales and the processes that control them. With a few ideas such as Deborah number and WLF and with measured timescales from oscillatory and thixotropic techniques we can better understand how to optimize the conflicting demands of different parts of our real-world processes, especially via relaxation spectra.
- Interconversions should allow us to mix-and-match types of measurements that have their own limitations, allowing us to fill out our understanding of our formulations over a wider range of timescales to better understand performance across a range of real-world processes.
- The trick of providing guesstimates of shear rates and timescales allows us to know from our measurements what our key properties will be under those conditions, or to extend our measurement range in order to determine them.
- Our inability to map between rotational, oscillatory and elongational measurements is deeply unfortunate and a grand challenge to rheologists
- Particle-based formulations can be made more comprehensible via a small set of key parameters.

Because this is an eBook, I will be delighted to update, revise, correct errors, or write new apps. I can only do this with your help. Feel free to email, phone, Tweet or LinkedIn - you can easily find my details on-line.