crystals

MDPI

Communication

Preventing Bio-Bloopers and XFEL Follies: Best Practices from your Friendly Instrument Staff

Christopher Kupitz * and Raymond G. Sierra *

Sample, Environment, and Delivery Group, Linac Coherent Light Source,
SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA
* Correspondence: ckupitz@slac.stanford.edu (C.K.); rsierra@slac.stanford.edu (R.G.S.)

Received: 5 March 2020; Accepted: 24 March 2020; Published: 27 March 2020

Abstract: Serial Femtosecond Crystallography (SFX) at X-ray Free electron Lasers (XFELs) is a relatively new field promising to deliver unparalleled spatial and temporal resolution on biological systems and there dynamics. Over the past decade, though, there have been a handful of results that have truly delivered on these promises. Why? SFX has many paradigm shifting techniques not seen in typical structural biology arenas, such as creating a concentrated slurry of microcrystals rather than a handful of *loopable* prisms worthy of a catalog photo. Then taking that slurry and high speed jetting them towards the vacuum or helium interation region to destroy less than 1% of your sample and waste the other 99. The literature is full of techniques and methods promising to cure what ails your experiment, yet as an instrument scientist will tell you –and a first author might admit after a few drinks at the conference happy hour—is that there are a lot more failures than the success we published, *results may vary*. We will walk through a best practices on how to prepare your sample and chose a sample delivery technique that will amerliorate your studies rather than undermine your hardwork and hopefully lead to better experimental planning and execution, inching you closer to that scientific goal and that call from Stockholm. This will be written in a more editorialized fashion and is meant to give the reader an idea of what to try or how they should be thinking. Welcome to SFX, now what?

Keywords: Serial Femtosecond Crystallography; X-ray Free Electron Laser; Sample Delivery; Sample Preparation; New Users

1. Introduction

This is not a literature review of sample delivery techniques or on how to crystallize your sample. You are reading this because you have been approved for a serial crystallography beamtime and it is your first time preparing for the experiment. We are experienced instrument staff that have performed our doctoral, postdoctoral and now our current work related to serial crystallography with X-rays. We by no means claim to be experts, but have seen many successes and many more failures. This is the advice we give to our users and our closest friends and collaborators. We hope that this fairly informal, editorialized piece helps put things together you have been reading in the literature or discussing with your collaborators. We have these discussions fairly regularly with users new and old, and figured it might help to jot some suggestions down so that you may come into a planning meeting or into your lab space more prepared and, at the very least, know which questions to ask and maybe how to troubleshoot. We will give some tips on how to plan and prepare, and even give some tips during your experiment. This is not a one-size-fits-all, but these suggestions are things we tend to repeat a lot. No, anecdotes are not data points, and correlations are not causations, but we figured this might be useful. Please do not take the tone as insulting or trivializing; rather it is simply trying to make the information more palatable and memorable in the tone the authors typically use with their users.

You might be a synchrotron crystallographer or you might have never touched a pipette before. Perhaps it is not your first time, but you're looking to get better results. We hope to show you some "best practices" on how to prepare for your first, or even second experiment. It wouldn't be a good introduction without some reference, so please check out the comprehensive textbook on SFX [1], written by many key figures in the field, instrument scientists (such as the authors), and frequent users. Of course, within that text, there are the requisite myriad citations to help you get started or to pore over in more detail. Again, we assume you have either read through the literature at this point, know some references or key players in the field, and are simply looking for that little bit of extra information that is not always found in a citation, but is carried around in the backs of minds from years of performing these experiments.

We really tried to address a concern in our field that there is a lot of trial and error that the instrument staff sees in experiments, but then does not always get reported out to the community in publications. This is not malicious, but likely from the strict publication limits on word counts and lengths from publishers, or the information is there but perhaps hidden or unclear to a first-time reader. Also, many times in science, we are quick to over report the successes and underreport the failures. In our community, where beamtime access is still precious, it does not make sense for all parties involved to waste time reinventing the same circular wheel, or worse, trying to see if that square-shaped wheel is any better. Instrument scientists have seen the different failure points before and try to guide the experiment in the right direction, but many times this falls on stressed ears or is not properly communicated. The instrument scientist(s) is(are) quite invested in the success of the experiment, despite them not being owners of the science. It is rare they would recommend something without the best intentions, but it is difficult to betray the trust of a previous group while explaining to the current group why their approach might not be the best. This is why we hope some of these best practices will prime the teams to think about some common pitfalls.

In this paper we will discuss proper protein crystal concentration for liquid jet experiments, aiming for monodispersity and filtering, crushing your big synchrotron crystals, selecting your sample delivery method, and when to call it quits and move on to the next sample. Some basic knowledge of the SFX technique is assumed in this paper, as it is impossible to cover everything involved in the technique in one paper. Again, this is not a literature review, remember. It is light on citations and scientific rigor. It's a simple *Letter to Our Newest User* or a *Frequently Asked Questions,* if you will. If you take anything away from this, it would be to listen to the beamline staff; we're not always right but we are highly invested in your success, even though it is not our science. Communicate with us early and often. Keep being a good skeptical scientist, ask your collaborator how a data algorithm works, or how a sample injector works; ask the beamline staff why they are suggesting one approach versus yours, and think how it compares to your sample. Fight the urge of *beamtime superstitions*. Listen to the experts, they have experience, but feel free to question, and don't be afraid to speak up. Sample delivery and preparation fall to the bottom of the hierarchy and many times a grad student or postdoc might be afraid to ask a question or challenge the more senior people on the experimental team. Most importantly, be candid with failures and caveats, in your papers, your presentations, postbeamtime meeting debriefs, and your conference coffee breaks. Just like Buzz Lightyear would not fly, but would *fall with style*, great scientific successes are sometimes simply *failing with style.*

2. Sample Injection

Currently, there is an illusion of choice when selecting the sample injection method best suited for the SFX experiment. One might like to choose his/her favorite injector, but in reality, the fluid properties of the slurry determine the best—nay, the most compatible—injection method. Sometimes injection is not possible with certain fluid properties, and therefore, the fluid might not be usable, a fixed target approach might be employed, or a brand-new delivery technique must be designed and prototyped. This high dependence on the properties of the fluids means that there is currently no one-size-fits-all injection technique. One of the most important parts of a successful experiment is the

collaboration between the sample preparation team and the sample delivery team. The experiment is largely dependent on the quality of the sample and the ability for that sample to be properly delivered to the beam. In an electrokinetic injection into a vacuum, for example, the properties, such as viscosity and even ionic strength, can make a difference.

One of the biggest differentiators between the delivery methods is how they handle the viscosity of the protein crystal slurries. High viscosity media should be extruded through an injector that extrudes the slurry at high pressures. Low viscosity media should go through a gas-focused injector. Variable viscosity media is handled by an electrokinetic, dual flow, or acoustic injection. The low-, mid-, to high-viscosity injection methods typically have high (over ~10 μL/min), mid (1−10 μL/min), to low (sub microliter/min) sample consumption rates respectively, which is of great importance to the sample preparation team. A high crystal concentration at dozens of microliters per minute means grams of protein with months of preparation time are necessary to run a few 12-h shifts, for example.

2.1. Estimating Sample Volumes

Although it is the properties of the crystal slurry that suggest which is the best sample injection method, many times. the quantity of sample might further restrain your choices; having only 50 μL of a water-like solution might not last long enough in a gas dynamic virtual nozzle (GDVN) to screen the quality of your diffraction data, much less solve a structure. In this case, you might *need* to mix your sample into more viscous media and use a high viscosity extruder or a fixed-target approach to effectively use such a small quantity of sample. At this point you should be reaching out to your sample delivery team and sample preparation team, and both should be working through options together; the preparation team should be thinking of how to produce more sample if possible, and at the same time, of potential viscous media compatible with the crystals, while the injection team is thinking of simple modifications to the injection method that might be feasible. For example, if it is 50 μL of uniform 200 nm crystals, then it is possible to consider making the inner diameter of the GDVN smaller than the standard 50 μm inner diameter, e.g. 30 μm. The running pressures will be higher and the chance of clogging will increase, but the fourth order dependence on this diameter change will give $(50/30)^4 \approx 8$ times less flow than normal. The team might determine that this is a worthwhile risk, or at least prepare some additional small diameter nozzles for the beamtime as a backup.

Both teams continue to test and iterate and go with trying to make it viscous, and have a backup injection method if the buffer exchange fails. The crystals might dissolve or become physically damaged when added into more viscous media if they were not natively grown in the media. The tolerated viscous media might not be homogenous or stable enough for an extrusion, the media might not be vacuum-compatible with an electrokinetic injection, the fixed target approach might not be available for your scheduled beamtime. While working with delivery and preparation teams, contacting the beamline staff for advice would be wise. Even wiser would've been to test these problems before the proposal was submitted (see Preproposal Characterization); this should not be happening at the beamtime or the two days before when the sample preparation and delivery teams first met. It is OK if there is no obvious answer from these tests and preparations, but discussing this early with the team and the instrument staff will give everyone the best chance for success. Figuring these things out during the stress of the experiment will lead to erratic decision making, multiple variables being changed at once, and increasing the likelihood of failure. Having tested and prepared before, even if it's for likely failure, will surprisingly result in calmer and more deliberate decision making, and can lead to more realistic experimental goals and create more beamtime success. The success of this example could be simply collecting a single diffraction pattern, rather than a full structure, to know whether the crystal condition is viable to continue refining the crystallization condition or whether a new injection approach is needed.

2.2. Testing Buffers

Testing a buffer, or mother liquor, without suspended crystals in a liquid jet is a good place to start. Ideally, you want to work as close to the experimental conditions as possible, but this is not always the case. Ultimately, nothing is as predictive as flowing the crystal slurry in the exact environment with X-rays—otherwise known as beamtime. If you've had a prior beamtime or screening time, then you are ahead of the game in terms of troubleshooting. If you are concerned with your buffer, test it in atmosphere or vacuum, and then with crystals. If it will not inject stably out of a vacuum without the crystals, the addition of a vacuum and/or crystals will surely make matters worse.

Reach out to your point of contact at the facility to inquire about any testing facilities available and when they might be available for pretesting. Testing at the actual instrument is very useful, but this is not always possible, since other experiments are going on. You might be afforded a few days prior to setup, which should be taken advantage of; but this might be too late to switch injection methods. Test early and test often when the preparation team makes changes. You can also try to test at your local facility, but remember: not all vacuums are created equal (See vacuum section).

Once, we had a beamtime preparation with a collaborator with a very precious sample. Precious in that there was not much to spare, even for the screening time, so asking for "spare" crystals was difficult. The injection method was a classic MESH injection, which typically has no sheath gas to protect the slurry from the vacuum but requires some modification of the slurry to handle not freezing once the slow-moving fluid enters the vacuum. The sample preparation team and the injection team were meeting ahead of time to figure out the best way to inject, which turned out to be: dope the mother liquor with up to 40% v/v of a *vacuo-protectant* while not perturbing the stability of the crystals. This worked by having the sample preparation team suggest likely modifications, probably from known *cryo-conditions* from the standard cryogenic synchrotron experiments, as well as a few other possible additives. The preparation team would provide the injector team with all of the guesses and the injector team would test the buffers alone in a vacuum. Concurrently, the preparation team would test the modifications on small batches of crystals and optically inspect whether there were any adverse effects, as best as possible. With infinite time and resources, one would screen conditions with X-rays (screening times, home source powder patterns, MX beamlines), but, at the very least, one can check optically, using birefringence or Second Order Nonlinear Imaging of Chiral Crystals (SONICC).

The teams iterated and found a great solution, a favorite of the MESH injection, the *sweet and sour.* A high-volume concentration of something *sticky,* such as 30−40% ethylene glycol, and then some salt and pH buffer. The problem arose that the salts were in high concentration and acted as the precipitant. The sample preparation team noted that this condition can sometimes cause phase separations, but it seemed to behave so well in the small in-vacuum test chamber that the team eagerly awaited the beamtime.

Come beamtime, there was only one sample, *the* sample. The team decided to take their time and troubleshoot with the buffer to not waste sample and ensure the kinks were worked out on the actual experimental setup—which was much larger and more complex than the small test setup at the injection team's home lab. All was good. Then, for the first time, the sample in the new slurry conditions were loaded. The second the sample exited the nozzle tip and was exposed to the vacuum it would fail from something precipitating out of the fluid, with it being so close to a phase transition.

"So why bother preparing then?!" you might have just yelled. In this case, the beamtime went smoothly. The screening time was short and sweet. They knew they tried and did not try to frantically come up with a new solution condition in the amount of time remaining. They had set themselves proper expectations from the preparation steps. In hindsight, the teams could have taken the new condition and doped it with another crystal (yes, lysozyme and other robust standards), or perhaps even some polystyrene spheres of similar size and concentration and tested one last time in a vacuum. This might have clearly shown that the buffer was too close to a phase transition to be viable when finally doped with the sample. This evidence might have persuaded the sample preparation team to try to dilute out some more salt, while keeping the crystals happy.

2.3. Caveat Injector

Look out for some of these common problems with injectors and crystalline mother liquors. This is not an exhaustive list of course:

PEG/PEOs. The most famous PEG is PEG3350, which is sold on your pharmacy shelves as stool softeners. What makes PEGs useful for softening your stool or hardening your protein into crystals is the way the long floppy chains shepherd the water molecules away from the polymers and towards the biological sample. This same principle is at odds with vacuum injection, where the vacuum is actively removing hydration from the meniscus. A meniscus with lots of PEG molecules driving water molecules away, while the vacuum ambient is gladly receiving them, will quickly cause problems. Problems that do not get better with the addition of crystals.

Also, these long chain polymers can easily increase the viscosity of your solution, so even changing values by a few percent without telling your injection team might be problematic and cause confusion and unforeseen failures. In certain instances, preparation teams have grown their crystals in high molecular weight PEG conditions and then done a buffer exchange in order to inject through a GDVN. Here, the robust crystal could handle the buffer exchange and injected well, whereas the injection could not handle the previous high PEG condition. Yes, it was lysozyme, but this technique might work for your system.

Solvents like ethanol and PEGs do not play nice. A capillary filled with PEG buffer should not be followed by ethanol, or vice versa. The two fluids interact and cause an instant crashing out of solution for the polymer, essentially making microfluidic concrete inside the capillary channel, rendering the tube and the nozzle useless. Concentric injectors, like the DFFN use ethanol as a sheath fluid to great success. They can run PEG fluids on the inside, but not always readily, so again, test ahead of time. Solutions with 40% PEG400 and 40% PEG8000 are quite different. Even test common concentration variations of PEGs to try and see if ethanol will cause a problem for your injection.

Non-Newtonian fluids. A Newtonian fluid uniformly deforms when a shear stress is applied, like water and ethanol. Non-Newtonian fluids do not, such as ketchup, toothpaste, and paint, which can deform faster or slower depending on how hard or fast you shear, and can be unintuitive. Many viscous media try to replace the monoolein in LCP injections (also non-Newtonian), but the new media can have new properties, such as being viscoelastic, or slow to start flowing, or not responsive to stopping the flow; do not assume they will behave the way your previous LCP injections have worked, test them ahead of time.

Long chain polymers can be viscoelastic, which might not be problematic for the injection testing or the crystals, but might be very problematic when hit by the X-rays. Here, the long viscous extrusion is fine until the stream elastically springs back like a broken guitar string *plucked* once by an X-ray pulse. This is tough to anticipate without having X-rays; however, when trying new injection conditions, it is advisable to test your crystals in other candidate media to have as backups. Check the literature; there are many, but remember to be discerning and see how the sample they used compares to yours. Test it. Reach out to the authors and ask for failures or tips. Beamline staff can be quite candid in this regard, as injection failures lead to long, dragged out, unsuccessful experiments.

Detergents. Detergents are sometimes necessary in the carrier media. The problem with detergents is they can trap bubbles and make things *foamy*. Ideally, if the detergent is not necessary, try to remove it for a vacuum sample injection. This is not always possible. Try a *less foamy* detergent, if possible. If the detergent cannot be removed, then try to outgas the solution(s) that make up the crystal slurry as much as possible, and be careful not to vortex or harshly agitate the solution or slurry to prevent the creation of newly trapped bubbles. Seeing foam at the top of a vial is definitely a sign of a *foamy* solution and trapped bubbles, but even worse are solutions where the bubbles are small enough to not be seen by the naked eye, but which begin to show once the sample is being injected, by having random sputters and injection instabilities, for example.

When switching between samples, there might be some air introduced or trapped in the front of a new tubing line. Sandwiching this bubble between detergent-containing slurries can cause more

bubbles that lead to injection instabilities and reduced hit rates. Bubbles at T-junctions can be picked up by the detergent and ruin the stability. Also, flowing too slow can cause the fluid to begin to outgas and create bubbles that get trapped as well, and thus ruin the stability. It might be best to reprime lines (good microfluidic hygiene), and to try some of the techniques here to minimize the foamy detergent effect.

High Salts. Salts are ever-present in all injection methods and myriad crystal conditions. Very few proteins form well-ordered crystals simply in just water, except unicorns like photosystem I nanocrystals. Salts can sometimes trick you into thinking they are protein crystals—at this point you should have hopefully sorted this out and are not writing beamtime proposals about salt crystals masquerading as protein crystals. The problem is still that there can be salt nanocrystals that naturally form when dehydrated further in vacuum. This is detrimental to the experiment because the salt diffracts even better than lysozyme and more intensely. If you cannot remove the salt, the instrument team will likely attenuate the beam and the data team might not be able to handle the intrusion of the salt data on the other data. This is a tough one to deal with, and might require some troubleshooting at the beamline. Small salt crystals might not be evident in previous vacuum injection tests, and are not evident until data is collected. Once you find a condition, try to iterate on the salt and buffer conditions and try to minimize them, especially if they are high concentrations, such as hundreds of millimolar or more and/or contain multiple salts. Also, if you know you are near a precipitation point of the salt or near a phase transition—this only gets worse when crystals are added and introduced into vacuum—consider diminishing the salt concentration. A prime example is ammonium sulfate, which constantly appears as a precipitant in the high 2 M concentrations. After adding the pH buffer, another salt, the proteins, maybe a detergent, there isn't much water left to keep things hydrated, which might be why the protein falls out of solution and into crystals, but maybe is too close to be stable for injection. The same as before with PEGs, sheathing a fluid like this with ethanol in a dual flow injector can still cause salty precipitation; however, one can be clever and try to see if the ethanol sheath is replaced with perhaps a water sheath to aid in keeping the exposed fluid surfaces sufficiently hydrated.

2.4. Testing Slurries

Ideally, every experiment would have a thorough testing of their sample in a liquid jet so that the jetting parameters are known well in advance of the actual beamtime. Unfortunately, we also know that scientists are extremely protective of their precious sample, and the very idea of using it for a liquid jet test horrifies most of them. Don't think of this as a wasted sample but as an investment in your future beamtime success. Something to remember though, as the existential dread surrounds you at the thought of wasting sample. A simple test now will not only save you time during the beamtime, but will also save you sample in the long run. When you are attempting to run the injection for the first time for a sample at the beamtime, inevitably things go wrong, and inevitably in the attempt to correct this, more sample is lost than what would have been used to run a quick injection test in the first place. It should be considered good practice to simply test initially and consider it an investment in your future beamtime, and happiness when you aren't dealing with as many problems during your beamtime.

As stated previously, testing stations may already be available at the facility you are heading to, and while testing two days before the experiment isn't ideal, it's certainly better than testing during the experiment. At the very least, test your buffers in the injector, and attempt to test your buffers with some surrogate, such as a lysozyme, thermolysin, or granulovirus of a similar shape, concentration, and morphology to see how things might change.

Please PIs, have pity on the sample team: preparers and injectors. Despite all of this pretesting, there can still be unforeseen failures. Nothing is going to fully simulate the experimental chamber geometry, with your current crystal batch, with a violent dose of X-rays spreading debris hundreds of micrometers away from the delicate injection aperture. However, the testing is still worth it as the team learns injection parameters that are common and tolerable, signs that things are about to become

unstable and how to mitigate. Hours of beamtimes have been wasted on hoping the injector will "unclog", "clear up", "make less ice", "get more stable" when a practiced team would have recognized this was an impasse and a new nozzle or crystal batch should have been loaded instead (15–30 min proper intervention versus hours of subpar, useless data).

3. Sample Loading

3.1. Sample Reservoirs

A deceptively time-consuming and important part of the entire process is loading your precious sample into reservoirs. There are many types of reservoirs out there. Which one is the best for your sample? Which reservoir is least likely to fail and spill the sample you've been working for months on all over the floor? As mentioned above in the Sample Injection section, it is critical to have solid communication with your injection team to have all of this worked out ahead of time. Some teams may have preferred methods that they've tested their injectors with, others will direct you to ask the beamline staff what the options are and what the general preference is. Again, we cannot stress enough that solid communication IN ADVANCE with all parts of your team is critical.

Some general suggestions when loading reservoirs seem obvious but are a constant source of frustration at almost every beamtime. When loading the reservoir, make sure there are no air gaps, bubbles, dead volumes, etc. These take time to push out of the system once everything is loaded, and very little is more frustrating than waiting to see crystals, seeing nothing, assuming something clogged, opening everything back up, redoing the whole process, only to open up your reservoir and find that you didn't push out the air, or that it had too many bubbles in it. Don't forget to properly prime and remove all bubbles from the nonsample line. Having to compress a bubble first is not efficient and causes instabilities as the bubble relaxes and recompresses as the plunger moves. Also, make sure all the connections and pieces you are using to a reservoir are placed on properly and sealed tight, while not overtightening the fittings. Wondering why your liquid jet died only to walk in and find your sample on the floor because someone has forgotten to tighten (or overtightened and broken) the sample side fittings has caused more than one headache during a beamtime. Talk to your injection team and set up some practice runs so that everyone knows what they are doing, and it becomes automatic, because at hour 9 of a 12-h shift on day 3, you'll want things to be automatic.

3.2. Antisettling Devices

Crystals are semi-large objects that are suspended in solution, right? Not really, they aren't suspended, they are floating, which means eventually they will sink, unless they are the same density as the media, the odds of which are not in your favor. A lot of factors play into how fast they will sink, but at the end, unless you are one of the lucky few who's buffer is right at the neutral buoyancy point of their crystals (and they have their own problems: how do you spin down and concentrate neutrally buoyant particles?), your crystals will eventually settle down.

Why is settling a problem if the crystals are technically already settled out of solution by definition. The problem is geometric. Imagine your bathroom sink full of water as a sample reservoir. The drain at the bottom represents the fluid connection taking your reservoir to the sample injector. Now imagine a single piece of hair floating in the sink water representing one protein crystal. Drain the sink water. It is not guaranteed that the hair will make it to the exit. The fluid mechanics typical of the reservoir flow are different than the sink analogy, but it is illustrative nonetheless. If the sample, or hair, is not drawn into the exit of the reservoir, then it will not get to where it needs to go in your plumbing. If you place your sample reservoir horizontally without any trapped air bubble, the fluid might be able to exit, but once the sample settles too low in the reservoir, it might not make it to the exit. Please don't just point the reservoir so the sample flow is down, it is one of the fastest ways to clog your injection set-up you will find. Imagine a sink full of hairs all rushing down towards the exit—that's a clog waiting to happen.

To prevent this sedimentation, you can rotate the reservoir so that the settling is always occurring in a new direction. The best methods for antisettling have been discussed, invented, examined, and then reinvented again so many times since the start of SFX that we can't even keep track of what version of which antisettler we are using anymore. The big thing to take away is that antisettlers exist, and you should be using one. None are perfect and everyone has favorites, so work with your teams to find what works best. Whichever beam line you are going to for your experiment, almost assuredly can provide you one, and if your injection team has a strong preference, then they probably have a version of it available too. Always make sure your sample reservoir is attached firmly, the temperature is set to what you want it to be, and the antisettler is turned on. Perform this check EVERY time you change reservoirs, or even are in the hutches in general. It is easy to forget to turn things back in when are trying to rush in and swap a sample reservoir and race against the clock. Make sure reservoirs are labeled properly even before loading; your sample injection team might do the mounting, and they may accidentally put it down while troubleshooting something and incorrectly load something. If it was accidentally turned, the crystals might not be suspended, and now everything is clogged, or no data was collected.

This is also something that is very easy to test ahead of time, yet no one spends the time doing it. Load up your sample on the antisettler you will be using and run your sample. Here, you can run a tube from the reservoir into a covered tube so that you can reclaim your sample—i.e., no sample was wasted in this pretest. Monitor it over time and see if you notice obvious changes in the turbidity, or take drops under a scope and compare over time. Even if you do not have access to a syringe pump or are sample limited, it would be useful to leave an aliquot on a benchtop and monitor it over time. Note how long it takes things to aggregate, flocculate, and/or fully sediment out. Once something has settled more than half the height of its highest level, it has sufficiently settled to not properly exit its reservoirs, in most situations. This is useful information when troubleshooting during a beamtime. If your sample takes minutes to settle on the benchtop, then the antisettler would have to be on; maybe you're not seeing hits on this new sample because the crystals have already settled, and not because they are poorly diffracting crystals. Check if the antisettler is on. Reload a fresh batch. If you know that they take hours to settle on the benchtop, then maybe your poor hit rate is from bad diffracting crystals, or there is a leak or clog. Use other diagnostics or pull the sample and move on, as long as you've patiently waited for dead volumes to clear etc. More than 15 min after a new stable liquid jet has formed, it is wishful thinking that good crystals are still on their way. The injection team can be monitoring for pressure spikes or drops to indicate clogs or leaks, for example. In a pretest with crystals or a postbeamtime troubleshooting test, it could be useful to collect some of the ejected fluid and make sure that crystals are passing through. Colorful crystal suspensions might help you to diagnose too. If the crystals are red but the fluid you collect is clear, then they are getting stuck somewhere prior to the final nozzle exit. If you are collecting red fluid then they are making it through, but check them under a scope for quality assurance.

4. Preproposal Characterization

It is common to come across a paper or meet someone at a conference and come up with a great SFX proposal. Sometimes, you might even get a *cold call* request to join a collaboration because your sample or technique seems interesting. This is great and can lead to great science! However, it can be problematic if this was the first interaction and the second interaction is at the experimental beamline. One of the most important team ups to get to meet in advance is the sample preparation and delivery teams, as they will likely need to interact and iterate before shipping samples or even arriving at the facility for the experiment.

If the first time your teams are talking to each other, as both a large group as well as individual subgroups, is at the beamtime, you have already put yourself in a position for a very trying beamtime. You should have meetings well in advance so that your entire group knows each other, and the group also knows who is responsible for what. This is useful because all the individual groups, sample

preparation, sample injection, data analysis, laser, etc. should also be cooperating with each other outside of these large group meetings. The preparation team needs to be collaborating with the injection team well in advance, even if the injection team is the team from the facility, to find a way to test your samples, or at least your buffers ahead of time so you are aware of the jetting conditions. The analysis team should be going in with a rough idea of what the protein and crystal they are analyzing will be. The laser team needs to integrate their systems around both the preparation team and the injection team so that everything flows, and you limit the number of issues presented at the beamtime. Having all of this worked out in advance will not create a stress-free beamtime, but a greatly stress-reduced one. There should be larger meetings with all parties involved, of course, but many times this important pairing of preparation and delivery is forgotten, and much time is spent on planning laser pulse powers, data analysis, detector calibrations, etc. This is a lot of placing carts before their horses. Lasers don't get shot until the liquid jet is stable. Data is not collected without a stable liquid jet, or at least a quasistable jet.

For example, a complicated pump-probe experiment might focus a lot of discussion on how to interface all the lasers into the experimental geometry. When beamtime arrives, a complicated custom, laser-coupled injection system arrives; so complicated that troubleshooting and fixing common injection issues are compounded by the system's complexity. In reality, two setups should've been employed: (i) a simple laser-less injection method to troubleshoot the injection problems and collect a high-quality *dark* (or unmixed, or reference) structure, which is necessary for all time-resolved experiments. And (ii) the more complicated pump-probe setup, now with diminished complexity, since the sample injection problems have been sorted out. Yes, the goal of your experiment was to observe where that elusive electron goes 10 fs after the laser light comes in, but it is highly unlikely you will get that answer without a properly vetted concentration, delivery system, and high-quality reference structure.

Typically, the more senior members of the collaboration are the professors and PIs driving the science, the laser scientists, the data analysts. The tasks of sample preparation and sample delivery typically fall to graduate students or postdocs. Many times, those lower in the pecking order of a collaboration are not quick to speak up during collaboration meetings, or might not know what the right questions are to ask. If you are one of those, speak up! The sample preparation and delivery team should not be meeting for the first time at the facility. There should be shipping of samples, testing of samples, discussions, and iterations in order to have a successful beamtime. Ask for that extra piece of equipment to make harvesting easier; a doubling of your efficiency can payoff greatly come beamtime. More sample stock or better injection diagnostics before the beamtime can pay huge dividends. Talk to the facility and see if such resources are available, or if they can be made available. So, reach out to your counterparts, ask your PI for their contact info, they will appreciate your can-do attitude. Reach out to your PI and your facility point of contact for more help figuring things out. Troubleshooting during beamtime is not the most efficient use of the time. Many of the facilities welcome you coming early, sometimes even months ahead of time, in order to test and troubleshoot without X-rays in order to be as prepared as possible. One of the most neglected aspects of beamtimes currently is preproposal and prebeamtime characterization, and sadly, it is also one of the easiest to fix. Considering how hard it can be to get beamtime, and how much time and effort you are already putting into the experiment, it should come naturally to test everything in advance to guarantee as much success as you can. Test your crystals in advance, even just powder diffraction will give you a basis for determining lattices and other important information. Test your buffer and crystals in injection set-ups, if your injection team doesn't have a set-up. Or if you are using the facility's injection team, then see if they will allow you to send them samples to test, or travel there days or even months early to test them yourself to ensure you understand your injection parameters to limit downtime. Any experiment that is done before the beamtime is one more that you don't have to do during the precious beamtime that you have.

All the points listed above are also important for another reason. Beamtime is still very limited, even with more beamlines opening around the world. Everyone is still competing for a very limited

selection of beamtime. Talk to the facilities before you even put in a proposal; information like that stated above about the diffraction and feasibility of injection may be just what is needed for your proposal to stand out from the pack when it comes to deciding who gets beamtime.

5. Preparing for Your Beamtime

5.1. Crystallite Concentrations

The nature of the serial crystallography experiment requires thousands of indexible diffraction patterns, each from a new crystal. This means that for macro crystallographers, the dozens of millimeters to submillimeter-sized crystals are no longer sufficient; a slurry of well over a million crystals on the order of tens of micrometers to hundreds of nanometers are what is required. Years of harvesting crystals from trays has not prepared you for this.

First let's establish a realistic baseline of what the goal is in terms of concentrations of crystals. You will not collect more data than the optimum of the X-ray source's repetition rate, the detector's rate, and the speed of your injection method. Currently, most serial data is limited by the detector rate. At the LCLS, the CXI instrument has detectors that can read up to the source's rate of 120 Hz. The MFX instrument can use a similar 120 Hz rate detector, or a detector with better dynamic range, balanced with appropriate pixel binning, can be set to 30 Hz. Depending on your experimental needs, $\frac{1}{4}$ of the data rate might not *cost* your experiment in data collection speed, as the data quality might be better. Talk to your facility point of contact to know which is best for your scientific needs. At the high repetition rate European XFEL, the detector technology is continuously improving; however, it is also a limiting factor. At these rates, liquid injectors are still able to replenish fast enough to successfully utilize the pulse trains and detector rates; however, using slower fixed target systems might be what limits your maximum data rate. *The slowest rate of your detector, injector, or X-ray source will be your idealized highest hit rate possible.* Essentially, this number is the number of crystals you will need for your entire experiment—assuming you hit every crystal, which you will not, unfortunately.

We will discuss sample delivery methods later, but let's assume you have an ever-present sample stream. That means that whatever sample delivery method you choose is infallible and is always in the path of the incident X-rays (we will get back to the likelihood of this later). This means that it is up to your particle concentration to have one particle in the beam at a time. *A good heuristic is a 10% hit rate*. This means that 1/10 of the detector images have a diffraction pattern on them and not just a blank (missed the sample delivery) or a solvent ring (sample delivery made it but there was no crystal in the probed volume).

A hit rate between 1–10% is passable and can get you a structure. If you are screening for crystal quality or the goal of your experiment is to solve one structure, then this can be sufficient. If the goal of your experiment is to solve multiple structural intermediates, each intermediate is its own separate structure, so you might run out of time. If the hit rate is below 1%, you should consider intervening. The sample might be settling, might be too dilute, or might just have poor diffraction quality. Remember that optically looking like a crystal does not mean that it diffracts well. Read below to know what to do *during* the beamtime if the hit rate is low.

A hit rate higher than 10% is great, although too much can be problematic. If the hit rate is too high, it implies that you will have multiple samples in the beam at the same time and it will become harder for the data analysis software to index the image. Although data analysis packages are continuously being improved, and many say they can deal with multiple hits on one image, it's more of an exception than the norm to want multiple lattices in one frame. Rarely does the data analysis team ask for more multi-hit images so they can solve the structure faster. So, 100% hit rate might feel great, but you can still leave with no structure. Go dilute your sample, aim for 2/3 hit rate (based on Poissonian statistics).

A typical rate of indexing is approximately 50%. This number can be higher or lower depending on crystal quality and multiple hits. Unlike the hit rate, this number should aim to be 100%, but this isn't necessarily expected; however, low indexing rates can be problematic and can indicate bad detector

geometry, bad crystals, or perhaps incorrect or new unit cell parameters. Maybe the synchrotron unit cell is not the same as this new microbatch approach, for example. A powder pattern from a home source might help you know the new batch's unit cell prior to the beamtime so the data analysis team can troubleshoot better.

Indexing depends on your crystals and knowing the detector distances and the detector metrology. Nowadays, most beamlines will have properly calibrated detector geometry files before your experiment. Unless something has happened physically to the detector since the calibrated file was made, you will likely not need to run a calibration standard. However, many times an analyst team might require this calibration anyways. The distance from the sample is the only parameter that might need calibration, and can be done in the data processing phase. If in doubt, leave some time aside in your experiment to run a known standard through a liquid jet, such as lysozyme, or a fixed powder like silver behenate. The powder loaded into a capillary is typically faster and gives rings quickly for calibrating, but might not give high resolution data to the edges of the detector to calibrate those far panels. Consult with your teams to see what makes the most sense. If you're looking for 3 Å structures, then you don't need the corner pixels calibrated for 1.5 Å data you may not get. As a side note, our X-rays do not perform miracles. If your crystals diffract to 7 Å at the synchrotron, it is unlikely they will jump to 2 Å in our beam. Unless the poor resolution is from the cryogenic condition or vitrification process at a synchrotron, it is unlikely that our photons will give you the resolution bump. A poorly diffracting crystal is a poorly diffracting crystal by any other source, so it is wise to set appropriate expectations.

Running a standard for detector calibration can still be full of the same troubleshooting problems as your main sample of interest, so beware. If you have your heart set on running a standard for calibration, it is unpopular but advised to run this first. Even tried and true lysozyme can have issues and might not give sufficient high-quality diffraction to create a virtual powder pattern to calibrate the full detector. Also, once this phase is done, the data analysis team has a more reliable calibration and can more quickly process the data and give feedback on how much *good data* has been collected, and how much more is needed. In most beamtimes, they jump into main samples first and put off the calibration until the end of the first shift or at some later point. What ends up happening is some data is collected on interesting samples, and quickly the PI wants to know about the data quality. The analysis team cannot give faster feedback because they are trying to play with nonstandard data, and are unsure if the unit cell parameters, or detector distances, or bad geometries are to blame. If they had a quick test structure from the beginning, they could optimize their scripts and start to understand the detector distance better, and now be ready to answer the PI when new interesting data starts to flow in.

So, bringing this all back together to answer the original question of crystal concentration: take the max possible hit rate of your experiment and multiply that by a 10% hit rate; of those 10%, assume that 50 % will index. This is the number of indexed patterns you can get, which should be in the thousands; *10,000 indexed patterns per structure* is a good estimate when planning. If you have studied this sample before, during a screening time or at another FEL, then see how many indexed patterns were used last time and work backwards from there. For a 120 Hz instrument, using a 120 Hz detector, the max data rate is over 5 million shots for a standard 12-h shift of constant data collection, which is rare with sample swaps and troubleshooting. Let's assume 50% of those shots are taken so that our estimate is conservative, with a 10% hit rate and a 50 % indexing rate on those shots. This means that during a 12-h shift at 120 Hz, over 100,000 indexed patterns can be realistically collected on a fairly successful experiment. This is a thorough collection and will likely have higher confidence in the collection statistics to the highest resolutions observed (see diffraction quality section below). This high quantity might be necessary for difference maps in mixing experiments, but might not be necessary to solve a novel structure quickly. In other words, these numbers imply you can solve one structure really well, or perhaps 10 quality structures with 10,000 indexed patterns each. Changing to 30 Hz or to 3000 Hz can quickly slow down or speed up this collection. Lower repetition rates are fine for solving single novel static structures, but can be taxing when aiming for intermediates and high-quality difference

maps are sought; consider this when planning or applying for beamtime, and know what's realistic and what is not.

Implied in our previous calculation was that we needed over 2 million crystals in order to get to 1 high quality dataset or perhaps multiple good quality datasets in one 12-h shift at a 120 Hz experiment. In order to know the particle concentration, we need to know a typical probed volume. The probed volume is the X-ray focus and the thickness of the liquid jet, for example. In most cases, these dimensions are on the order of 1 μm, i.e., a probed sample volume of 1 μm^3 or 1 x 10^{-12} milliliters. This does not mean that you want a particle concentration of 10^{12} particles per milliliter. Imagining a 1 μm particle in a 1 μm^3 volume would imply a perfect chain of samples, which is not possible nor desirable. This literally implies an extrusion of sample with no liquid in between particles for a continuous X-ray source; this extrusion is otherwise known as a clog. The X-rays are destructive and can impact nearest neighbors, but are luckily out of the way by the time the next X-ray pulse comes in to probe the next sample at the common repetition rates available today. The fact that the sample is flowing, and the X-rays take time to *reappear* at the interaction region, implies that more dilute samples are acceptable. We typically recommend 10^8 particles per milliliter, but who's counting?

The easiest way to estimate the proper concentration is to set up your crystal slurry, load it into a vial, and allow it to sediment over time. The precipitate should take up approximately 10–50% of the volume of the slurry. More is typically better, but can lead to injector clogs. It is best to titrate up in concentration, rather than overconcentrate and clog sample injection equipment, which can eat up more of your experiment time. It is better to estimate your concentration, observe your hit rate, then go back and increase your concentration. At first, it is hard to decouple poor crystal quality from a low hit rate. Use as many diagnostics as possible when troubleshooting a no/low hit rate. Inspect the reservoir afterwards and see if there are any obvious leaks. If you injected 500 μL of sample, are there 500 μL missing from the reservoir? If you have more sample left than expected, there might've been a clog. Remove inline filters and inspect lines for clogs. Make sure the sample flows out of your reservoir's final tubing connection right before it enters the injector's capillary. Make sure the sample flows through the reservoir and out of any valve or equipment before it enters the injector. Use any inline diagnostics, such as viewing inside the capillary, to see if the sample is flowing. Make sure the antisettler is operating properly. If your slurry has a color or is "milky", make sure that the same turbidity is seen dripping out of the reservoir as what you put into it; if not, things might be settled. All of these types of diagnostics should be tested before arbitrarily increasing the concentration for the crystal slurry, as this quickly leads to clogs and irreparable damage to the sample injector, and can cost more time than actual successful data collection.

It is likely that you already have a mother liquor for your crystal, whether you're trying some new recipe from the literature or you have years of experience growing this crystal in the exact same recipe. It is also useful to know any other potential crystallization conditions for this crystal, such as the cryogenic conditions or any others, you might need to change the recipe in preparation for the experiment, or in the middle of your beamtime. Most of the beamtime planning will unfortunately depend on this recipe. It is the mother liquor which causes the protein to precipitate and form into the crystals. It will also give the flow properties to the slurry, which will dictate the best injection method.

Look at your crystals under a microscope and take photos of your dilutions and compare when possible. It is also useful when troubleshooting hit rates and dilutions to inspect the sample under a scope before and after loading in the reservoirs. Use a hemocytometer to get a count and aim for at least 10^8–10^{10} particles per milliliter. Smaller crystals will need higher concentrations, while larger crystals will not be able to concentrate as high.

5.2. Harvesting

Of course, the most brute-force method is setting up hundreds of wells on plates, letting your crystals grow in 5–10 μL droplets, aspirating the droplet gently into a pipette tip, and transferring it over to a larger container which will now become your stock solution of crystals. This can be quite

tedious and time consuming, but might be the only method available. Of course, this assumes that your crystals grow in small enough sizes (<20 μm) and sufficiently high quantities per droplet to be viable. This can work from soluble proteins to membrane proteins grown in viscous media. It is not ideal, but it works. Be gentle when harvesting; it is best to use a large pipette tip with a larger opening, even if you are only aspirating a small volume, so the crystals are minimally disturbed.

5.3. Batching

Some groups have discovered batch protocols for their proteins. Lysozyme crystal slurries can be made by mixing the protein and precipitant into a container and stirring at a certain temperature in order to create a uniform size, depending on the chosen temperature. This is not always possible, unfortunately, for all proteins. A way to test for batching recipes is to screen in oil-covered droplets. If any droplets give a potential cascade of small crystals, this might scale up. Take the same recipe and fill a large vial or tube with the ingredients and gently agitate, like on the orbital shakers found in biological labs. The larger volume can lead to gradients in concentrations, and thus, the gentle agitation keeps the solution as homogenous as possible, replicating the small volume of the oil-covered droplet. Ideally, the slurry will reach a steady state of uniform crystals suspended in a solution.

5.4. Crushing

Sometimes you do not have the time or resources to rediscover a whole new recipe to make perfect micron-sized crystals, but you're a pro at making large crystals or imperfect polydisperse slurries. Here, crushing your large crystals into micron sized ones may be the perfect answer. The process can be similar to *seeding* in conventional crystallography, where you crush a larger crystal with a vortexer, or a bead beater, to create small fragments that in traditional crystallography serve as nucleation points for new crystals. Here, they should be the perfect size for injection. Crushing with commercial beads or disruption beads has shown great success.

If the crystals are in a slurry but are too large or are too polydisperse, then passing them through a stainless steel or PEEK, frit filter, commonly used in liquid chromatography setups, can be useful. A 20 μm pore size on the filter is likely sufficient, even if you are trying to make more monodisperse crystals that are 10 μm in size. The torturous paths of frit filters will likely crush crystals down to size. Be careful that your particle concentration is not too diminished. It is a filter, and not everything will pass through, so you might have to concentrate more after the filtration process. This can be a rough process and might damage your crystals too, so beware.

5.5. Filtering and Concentrating

Diluting your sample slurry is straightforward as long as you make sure to maintain the mother liquor, so the crystals remain pristine. It is important to note any changes in the dilution or solution makeup to the sample injection team, as even concentration or other fluid properties such as viscosity or conductivity might affect the stability of the injection. On the other hand, concentrating or filtering are not trivial feats with the crystal slurries. If the crystals are robust, then spinning the slurry down in a centrifuge and removing fluid is a common way to concentrate. However, not all crystals can survive such harsh treatment, especially if they have large unit cells and/or solvent channels making them more prone to damage. If concentrating is still necessary in these cases, then perhaps timed sedimentations can replace the centrifugation.

Filtration is also a necessary step in sample preparation. Having mostly 5 μm crystals with an occasional 25 μm crystal population can lead to injector clogs or sufficient high intensity diffraction events that the data quality suffers from a forced attenuation of the beam. It is always best to have as monodisperse a sample population as possible. Using the timed sedimentations can be used as a means of filtration, for example. Quick sedimentation times might have dense, larger crystals quickly fall to the bottom of the vial allowing for a more uniform slurry of smaller crystals to be left behind in the top fluid solution. You might have to concentrate this solution with longer sedimentation times.

If the crystals are small, then frit and mesh filters can be used, and they potentially minimally damage the crystal quality. For example, submicron crystals will likely pass through a 20 μm frit filter without a problem. If they are delicate and/or not small, then patterned filters with uniform laser cut holes or weaved materials can allow a gentler transmission of crystals to occur through the filter. Remember that filters eventually need changing and can reduce your particle concentration. If the sample concentration is high but the hit rate is low, check the upstream side of an inline injection filter for visible buildup and replace the filter. Removing the filter entirely risks damaging the injector, but is worth it as a last resort in troubleshooting. Prior to this, test whether the turbid slurry is coming out downstream of the filter in order to see if the transmission through the filter is sufficiently high. Consider inspecting the concentration upstream and downstream of a filter under a microscope if possible.

When filtering, it's best to remember that hydraulics are your friend. Even if you must filter 5 mL, it's better to use a 1 mL syringe. The smaller diameter of the 1 mL applies more pressure for the same applied force. In fact, most operations involving a syringe at the beamline are likely better with a 1 mL syringe, despite the larger volume of other syringes.

5.6. Vacuum Sucks

Nothing in the experiment, except for the photons, really wants to be in a vacuum. If you are not in dire need of this vacuum, then do your experiment in air or in helium. Now, if, for whatever reason, you decided to still go the path of vacuum, through experimental necessity, beamtime availability, or pure masochism, we will talk about some problems to expect with a vacuum, and hopefully how to prepare and avoid them as best as possible if this is your first time.

Read the above Testing Buffers section with some vacuum-adjacent warnings about testing different flavors of crystal slurries. Unfortunately, many times, the thing that makes the mother liquor ideal for making a protein precipitate out of solution and form a crystal also falls out of solution in vacuum and fails the liquid jet testing. Look out for things that might dehydrate readily, as this happens more quickly in a vacuum. Remember that the vacuum conditions typical for SFX are about 10^{-4} Torr to 10^{-7} Torr, and so, many things will evaporate in this pressure range, but not all! Actually, most of the things you are planning on flinging into vacuum will likely not evaporate. Sure, the water or ethanol will evaporate or sublimate away, but the unused protein crystals, the PEG, the salt, and the detergent will form a nice gunky chunk on whatever it lands on. The X-rays will pull double duty and destroy some things for you, but will send some of that debris back onto your nozzle tip and onto nearby surfaces, such as optics, in order to further conspire against the success of your experiment.

Some teams have come up with ways to convince themselves they are optimally doing their best to fight the onslaught of vacuum problems. Some teams methodically switch the sample on for data collection for a few minutes and then, no matter what, switch to water for a short interval, the idea being that all the sticky things might get potentially rehydrated from the water and move further away. Some try to clean debris on a nozzle by forcing a controlled icing event, which works well for old-style glass GDVN nozzles, but is not so successful with glued polymer 3D-printed nozzles. The sudden flash freeze of the fluid in vacuum expands violently, which is useful to clean debris off the nozzle, but can damage soft squishy materials. Others ride the debris buildup until the instrument scientist finally throws in the towel and deems a nozzle swap is necessary (best case) or a cleaning of the chamber (worst case), since this requires about an hour of downtime to properly vent, clean, and repump, assuming nothing goes wrong in between. Others introduce additional solvents, like ethanol or water sprays, to try to clean the buildup with minimal intervention. Remember that the more mass you add to the vacuum chamber, whether it be gas or liquid, the more the vacuum turbopumps have to work to expel that mass. Therefore, the chamber pressure creeps up, and ultimately, the reason for needing the vacuum, like the detector or the background signal, become unacceptable and the experiment has to stop. Unfortunately, many of these things are triages to a larger problem that is hard to beat, and thus, a constant scourge on the vacuum experiments with biology at XFELs.

5.7. Ways to Fix Vacuum Issues

Are you sure you can't do this experiment in atmospheric conditions? Fine. Here are some things you can do to practice and prepare for the vacuum experiment. Again, reach out to your facility contacts and see if there are test chambers or opportunities to come out and run some tests in the actual vacuum chambers well before your beamtime, especially if you are trying new injection technology.

Not all vacuums are created equal. Small test chambers might give you an idea, but the small roughing pump or the tiny turbo attached to it might not really get you to the pumping speeds felt at the beamline. Consider where your pressure gauge is and how small your vacuum chamber might be. It is possible you're reading the great vacuum nearest your vacuum hose but the tight geometry in your small vacuum setup might really be closer to the vapor pressures of your fluids, and thus, not show failures seen in the actual beamline vacuum.

Find the leakers. So, you are still committed to a vacuum? Well, that comes with great responsibility. Something in your experiment needs the vacuum, and sending in extra fluid from leaky tubing is not helpful. In fact, new injection methods need to be vacuum tested, and even old injection methods mounted onto a nozzle rod for the first time should be vacuum tested. If the injection rod cannot hold at least 200 mTorr of vacuum or better, then there is a problem and it will leak further and raise the chamber pressure high enough to make the operation of the detector nebulous. This is at least true for injection into the CXI endstation; this might be at different points for other instruments. All together now: check with your facility point of contact. The first troubleshooting everyone thinks to check is the exposed unconnected capillary tubing from the nozzle. These are long lines with small inner diameters; a 0.5 mm inner diameter, over a meter long, has held up vacuum on a well-sealed system. Ideally this will seal up and be fine, but the exposed 50 µm or 150 µm inner diameter capillaries are not what is causing the load-lock antechamber to not hold vacuum. Check for loose microfluidic or swagelok fittings, without over tightening! The mounting to the nozzle rod might need further tightening or teflon tape to create a better seal. At this point, you have wasted about 15 min checking and wondering and discussing. Pull the stick and mount your second stick. Troubleshoot while data is being collected. Always have a second injector mounted on a nozzle rod and ready for situations like this.

Tighten, apply teflon tape, and worse case, apply vacuum friendly glue to seal any holes. Sometimes an injector requires glue and it is not vacuum tested, or over time, the glue cracks and creates a leak. Seal and try again. There should be a test stand near the beamline to test the vacuum readiness of the rods, to not be surprised once you rush into the hutch to swap nozzles.

Vacuum cleanliness. When initially preparing the nozzles, many times, fluids can be introduced into the capillaries, such as fluid from polishing the capillary to a tapered point at the tip. Ideally, the sample injection team has tested all the nozzles ahead of time, at least with a simple solution. Even building the nozzles in your home lab during a humid summer can lead to moisture being trapped in the center of long capillary lines, which is evident only in vacuum. When you mount your injector onto the rod, rinse out any fluid inside; hopefully you didn't leave salty, sticky protein slurries in capillaries for months after your last test. If you did, go ahead and throw your nozzle out now and save everyone some time. Clean out with an appropriate solvent, like water or ethanol. Remember to think of the compatibility of the solvent and the last thing in the capillary line (think ethanol precipitating PEGs and permanently clogging a capillary line). Rinse with water; at least 10–20 drops should fall from the nozzle tip. Then, push the water out with ethanol, with a similar number of droplets. More drops are always better, if you have the time and patience. Yes, use a 1 mL syringe for these to push with enough hydraulic pressure. After the water and ethanol, grab a filtered air line (not canned air) to push out the remainder of the ethanol and its vapors, and flow clean gas for as long as possible, at the very least until you are no longer wicking away fluid from the nozzle orifice. When you're dabbing the tip of the injector, try to make sure to use a gentle dabbing action with a lint-free wipe. No, lab wipes are not lint free! You can use them as they are easy and readily available, but just be careful to wipe away from the opening of the nozzle, and maybe you should inspect for lint under a nearby scope afterwards.

Gas is your friend. The point of the gas sheaths on many injection methods is to keep the meniscus from freezing, and so that errant fluid is properly pushed out of the nozzle or jetted. Before turning on any vacuum, connect a low flow of helium on the gas line, such as 50 psig, i.e., just enough to keep the meniscus protected from flash freezing any fluid, but not overload the antechamber vacuum. The load-lock antechamber of CXI can typically handle this low flux of gas. If you cannot achieve a vacuum, it is unlikely this gas load is what's stopping you. Check with the staff. The order of operations for vacuum injection is to connect gas, turn on gas, turn on rough vacuum, THEN connect the liquid line last. Follow this in reverse when disconnecting. Disconnect the liquid line before modifying anything about the gas flows, connect the injector's liquid line to another gas line—without disconnecting or turning off the injector's gas flow line still—and purge it clean for a few minutes. Then you can turn off the gas, disconnect fittings, and remove the nozzle safely from vacuum. This nozzle will likely be clean and ready to start operations on the following shift if shut down properly. Make sure to follow any protocols and ask the staff for help. Don't be shy, they are very willing to help and would rather be asked then deal with the fallout.

One group was loading their injector nozzle when suddenly the PI came in and requested a prebeamtime meeting to discuss the plan. Accidentally, the helium was left on with no active vacuum removing the gas as everyone ran to the prebeamtime debrief. As a result, a piece of steel vacuum hose was ruptured causing quite the delay and the literal heavy lifting of new equipment for the experiment to continue. Protocols are your friends when you're literally jetlagged from driving a liquid jet and troubleshooting during a sleep deprived beamtime.

6. Things You Need to Know During Your Beamtime

6.1. Sample Screening

The first steps in every experiment should be to *SCREEN YOUR SAMPLES*. Even if it's the same aliquot from last month, or the same exact recipe from the last successful beamtime, there is no telling what could have gone wrong or what might have been different. We've seen the same person measure their precipitant in percent weight per volume one year and the following year do the same percentage in volume per volume. Yes, it matters. Your crystals might not diffract with the same quality, or the injection method might not be as stable. We have seen users create a fresh batch of slurry using the nearest stock solution of PEG. It turned out, after some investigation, that the stock was old and hydrolyzed and was at a much lower pH than expected. Even minor perturbations such as contamination, or differences in pipetting or stirring, or vortexing from preparer, can make a difference. You might not need to screen at the start of each shift of a five-day experiment; shift-to-shift variations can exist, but are hopefully rare—although please be scientifically skeptical when things start to go wrong.

Crystals can keep growing, change their size distribution, or get stickier from noncrystallized proteins in solution. Unless you've performed a study on how these crystals age, do not assume they are the same crystals from when you found them on your screening tray months ago. Check a drop quickly before and after loading into and out of a reservoir, or, as the very least, at the beginning and end of each shift. Take a photo; this is great for records to compare concentrations and quality when tomorrow's batch is mysteriously underperforming. These photos go great in talks or supplemental figures, or even insets of figures.

6.2. Beamtime Interventions

While studying the hit rate, you will be diligently staring at a stream of near-live images collected from the detector. These are typically displayed at 1 Hz; however, this is likely representative of your data, i.e., if you are seeing a poor hit rate at 1 Hz, it is unlikely that the hits are lurking in the other 119 shots you are not seeing that second, and you are consistently seeing the missed shots. If you're not seeing hits within 5–10 min after loading a fresh reservoir and getting it to inject stably, it is time to

troubleshoot and intervene. Check that the antisettler is turned on, make sure you see crystals flowing and the delivery is working properly, as there might be a clog. Check for high or low pressures on the injection system, fresh reservoir, etc.

While determining the hit rate, make sure you are hitting the injector stream. Depending on the wavelength, detector distance, and experimental geometry, you may or may not see a solvent ring. Typical detector distances (50−100 mm) at ~10 keV wavelengths will have a solvent ring somewhere in the middle of the detector area. This ring should show up in every shot when you are stably hitting the liquid jet. Depending on the injection method, other artifacts, such as a strong horizontal streak from the liquid jet, might show up, or backgrounds from a solid target mounting surface. A monoolein ring will be smaller than the ~3.5 Å water ring, and a dual fluid injector might show two rings for each solvent. Unless you are studying the solvent, the rings are merely references to show that you are maximally hitting the particle stream with the X-rays, and it is likely the beamline staff will focus on maximizing the presence of the solvent ring. Small micron-sized jet streams and micron-sized beams have an obvious optimal overlap. When shooting thicker extrusions that can be 50 µm or wider, then optimizing the position of the liquid with respect to the X-rays is more important; the brighter the solvent ring, the more centered the shot.

It is possible that 1 µm crystals are preferentially hiding in a 50 µm thick gel-like extrusion, but it is very unlikely that they will consistently form on the periphery of the extrusion and stay there and never be seen by the X-rays, as they are randomly distributed in the slurry. A 1 µm X-ray focus hitting a ~5 µm stream of similarly sized particles means they are not likely hiding from the X-rays. You are more likely need to concentrate, without clogging, or you have poorly diffracting crystals.

6.3. Diffraction Quality

If you need 1.9 Å data in order to answer your research question but you are not seeing diffraction data beyond the solvent ring, then increasing the transmission is not going to help; you need better crystals. The increased transmission also increases the background noise, so while it seems like it might help, it probably will not payoff. It is possible that there is too much background from the sample injection or the ambient environment, but asking for more beam intensity will not fix the problem. Consider switching to a new sample. Even if it's a fresh batch of the same sample, it might be markedly better.

6.4. Inequalities in Crystal Preparations

When troubleshooting, be open to the fact that the same recipe does not guarantee the same sample quality. Different ingredients, different equipment, and different personnel can have enough variations in the techniques and preparation that the quality can vary. Old PEG from the storage locker can be hydrolyzed and have a much different pH, thus affecting your *fresh* batch of crystals with the same recipe as last year's experiment. Even growing a large batch of crystals and loading it into different reservoirs can lead to some variability. Pipetting on the fourth night of experiments can be sloppier than when well-rested. Crystals can continue growing. One reservoir can be at a different temperature or held in a hand longer than others, causing a temperature rise that can make things better or worse. The first reservoir loaded can come from a homogenous suspension of crystals while the second reservoir has already begun to separate by size and have a subtle size variation. Many groups have fought the urge to load a fresh sample reservoir hoping that the current ones would get better, only to look back in hindsight with regret when the new reservoir is inexplicably better diffracting than the seemingly identical first batch that was failing. Think about an explanation later while the data is pouring in, not while no data is pouring in and you're hoping for a miracle.

Typically, it's the last thing you want to hear, but sometimes your crystals are just bad. The sooner you realize this, the sooner you and your team can readjust. Many hours of beamtime have been wasted chasing better data on the same reservoir/batch/recipe/aliquot that "...worked perfectly last

time?!", "...is the goal of the experiment . . . ", "...should be our best batch . . . ". Like the song says: "know when to fold 'em."

6.5. Practice Good Microfluidic Hygiene

Always be on the lookout for bubbles. Prime things with the proper fluid to get rid of any air bubbles, and prime with gas when trying to get rid of that last bit of liquid in a line before introducing it into a vacuum system. Bubbles introduce pulsatility in a pressurized fluid system. When the bubble gets into vacuum, it can feel larger as the gas expands. When doing a mixing experiment, prime the junction with the proper fluids so there is not a bubble at the junction. When loading your reservoirs, treat them like syringes and make sure there is no bubble trapped between the plunger and the fluid. Load things vertically and think of narrow channels and how a sticky fluid might trap a bubble you might not be seeing. Yes, your sample is precious, but push the plunger until you see a drop fall out of the reservoir to ensure everything is primed and all bubbles have exited from the top—the sub-microliter droplet lost will be worth it when you don't waste over a milliliter of sample on clogs and other preventable issues. Even if your experiment is in atmosphere, the pulsatility from bubbles can be damning.

You're gross, you should always consider using gloves to not get things on the nozzle tips or capillary entrances. Don't use powdered gloves, the powder on them gets everywhere and will eventually lead to a clog in something that you are working on, usually at the most inopportune time. The labs and beamlines where the nozzles are made are much dirtier than you'd think. Walking from the preparation area to the beamline can cause dust or dirt to get captured on the nozzle or capillary entrance. Takes care when cleaning nozzles, and a special note: Kimwipes are not lint-free wipes. They can be used in a pinch if lint-free or lens wipes are not around. Consider using the corner of the wipe and wetting it if you need to wipe something off a delicate object, like the nozzle tip, and leave minimal debris behind. Be careful with compressed air as well, as it can be unfiltered and dirty. Even when filtered, it can also be a blast of high pressure strong enough to break a delicate glass capillary tip. Polyamide-coated capillary is surprisingly tough, but if you have uncoated or stripped the capillary, then it becomes extremely delicate.

Finally, please stop overtightening the various parts of your microfluidic devices. These fittings are designed to hold hundreds to thousands of psi and can be reusable, if not overtightened. Many users fear leaks, are inexperienced, or like to show off to colleagues how strong they are, but the rule of thumb is to get the fitting *finger tight*, which means with minimal effort, you can no longer turn the fitting. Then, with maximum effort, give it a quarter of a turn more. Do not use your mouth or ask the burliest member of your team to help you crank on this. Most importantly, don't go grab a tool to help you crank down on any fitting. If a tool is needed, it would have been given to you. In the extreme, you can break things by overtightening, costing time and effort replacing and repairing them during the experiment. At the very least, this wastes another polymer fitting that now has to be thrown out instead of being reused. Overtightening progressively gets worse over time on the same fitting, and each time you retighten, you have to tighten further just to create the same seal. Talk to your injection team, they built and run the set-up. They will know how tight everything needs to be and can teach you what the best practices are for using the system for your beamtime. Overtightening is a very easy problem to have, and an easy problem to fix.

6.6. Be Wary of Your Tubing

Any beamtime uses multiple sizes, types, and lengths of tubing. You wouldn't think that this was something that you need to be concerned about, but being unaware of it can cost you beamtime. The standard tubing cutters use razor blades which can go dull rather quickly. It can take as little as four cuts before the razor blade can be dull enough to go from cutting the tubing with a nice flush cut to simply pinching the tubing off and completely erasing the hole, thus stopping your entire experiment. Don't cut tubing with scissors or pliers, as they can crush the tubing wall. Your blade should not have

visible wear. The common practice with both polymer tubing and glass capillary tubing is to cut a length longer than necessary with a cruder cut, and then to make a more precise cut at each of the ends of the tube. Talk to your beamline staff about the proper tools.

One simply does not cut a glass capillary. You score the capillary. The staff will have something with an edge to score the tip. Do not use scissors or blades used for polymer tubing. Ask for help or look up how to score capillary tubing. The idea is to gently rub the edge of the tool on the surface and score it, without crushing it. Then, pull apart the pieces of capillary along the length of the tubing. The brittle glass or silica will shear perfectly and have a smooth surface instead of a jagged or angled surface. Score the capillary on a soft surface to not crush it; the edge should not cut you, and thus, your finger can be the soft surface. Ask for help, practice, and check your work under a microscope. During the experiment, scoring a fresh capillary tip might be necessary at the top of an instrument or in less-than-ergonomic conditions, and require practice so you're not crushing capillaries but making smooth cuts. Bad scores can cause leaks, dead volumes, or worse, cause glass shards to get in the capillary diameter and irreparably clog the injector.

6.7. Diagnosing a Clog

Clogs will happen, it is a fact of the experiment that can't be completely avoided. What's important is to understand where, why, how to minimize, and how to fix them so you have the least amount of downtime. Clogs tend to happen at tubing transitions such as unions, adapters, and other fluidic connections. You can try to minimize fluidic connections when possible to decrease the chances of clogging. You can slowly check for clogs by working backwards from the nozzle tip towards the reservoir. The tubing and connections upstream of the nozzle capillary entrance are typically easily replaceable if a clog gets stuck in the line. However, if the clog makes it into the injector capillary, it will have to travel over a meter through a ~50 μm opening, thus ruining the handcrafted injector and costing lost data collection time. To check where a clog is, simply undo the connection and see if your turbid sample is flowing through it. If it is, then the clog is after that connection; if it's not, then it is before.

Use a pressure monitoring chart to monitor the applied pressure on the injection system. Ask your friendly beamline staff how to do this. Most injection methods should reach a stable applied pressure when operating. Sudden drops and spikes in pressure might indicate problems. A sudden drop to zero indicates a line has likely popped from its fitting and is now leaking. If this was preceded by a large pressure spike, it could have been from a clog downstream of the burst. A properly plumbed system might be able to withhold over 2000 psi, so a pressure spike without this pressure relief is still indicative of a clog. Be aware that a steady climb in pressure can also mean that the reservoir's plunger has reached its end of travel. If you are not expecting to be out of sample, then it is likely a clog has formed or is forming.

Sometimes a clog will naturally pass through despite the increase in pressure. If the liquid jet performance (i.e., the hit rate) is not diminishing, then perhaps let it pass. However, if the performance is being greatly affected, then you should try the Bothe method of clog prevention, named after the user that showed us the technique. Once you believe a clog is forming, quickly release the vent valve on the front of the HPLC (hopefully this is outside the hutch near you and does not lead to delays, i.e., by breaking into the experimental hutch). Enter the hutch and disconnect the injector's capillary from the rest of the system. Score off a few centimeters of length from this capillary entrance, hopefully excising the newly-formed clog entering the capillary. Troubleshoot the upstream connections and make sure everything flows fine to the final union. Then, reconnect, and hopefully the clog was captured and removed. If the injection does not resume normally, then go ahead and stop wasting more time and swap nozzles. If this is your last nozzle, then continue troubleshooting and hope for the best.

6.8. Time-Scale for Events

There are certain procedures that you just know are going to have to occur during your beamtime. Some of these will be fast, some of them not so much so. It is important to have a basic understanding of how much time certain procedures will take so that you can plan for them when they do eventually occur during a beamtime. A perfect example of this is that changing the flow rate to zero, or even down to half the current flow rate on an HPLC, does not stop the flow or slow it down instantaneously! Be patient, use a pressure plot to see when the flow stabilizes, or a flow sensor. Nothing happens instantaneously. Other procedures that will almost assuredly happen are things like nozzle changes (15 min if you're fast, better plan for 25), reservoir loading (15 min), venting the chamber (30 min), cleaning the catcher (15 min), and pumping back into vacuum and resuming (30 min). These are approximations and can take more or less time depending on custom configurations or unknown problems.

There is a very notable procedure that you should be aware of: the first day start-up. Unfortunately, you don't start the first day and immediately start collecting data; the beam needs to be brought to the parameters you requested, then tuned for the best energy, then it needs to be aligned inside the hutch, then with your liquid jet, and a myriad of other small actions that all take time. Assume that this procedure will take between 2–4 h. This same procedure also occurs at the start of every one of your shifts. Thankfully, the first day is usually the worst, and on subsequent days (assuming you don't request any major changes), usually only takes an hour or less, barring unforeseen circumstances.

For the sample preparation and injection teams, always have a backup. Have a backup person to relieve you from long shifts. Most other teams need to stay off-shift growing more sample or testing new injection conditions. Always have a second injector mounted. Always have a second reservoir loaded, or at least ready to be loaded; you never know when the data team says you have enough data on this sample, or the sample clogs, or the beam drops and it becomes a good time to change the nozzle rod. Bring yourself some snacks too. You never know if you'll be able to break away from the experiment, but a good PI keeps their team well fed.

7. Things you Need to Know After Your Beamtime/Shift

The Importance of Cleanliness

After a 12 h shift, the last thing anyone wants to do is more work. Unfortunately, for the sake of the beamline scientists, the injection team, and your future selves, it is worth it to spend an extra half an hour and ask what you can do to help. From the sample preparation side of the equation, there are reservoirs with samples to recover, reservoirs to clean, workplaces to keep clean and organized, and a log book to make sure everything is labelled appropriately, because we guarantee you that you won't remember it perfectly at the start of the next shift, let alone a month later when you look at it again. From the sample delivery aspect, there are injectors to clean and remove, switchbox ports to clean, all to make sure you are ready to do the same thing the following night. The data analysis also has to make sure their analyses are all cued up and will be running while they sleep. Therefore, if you don't have a specific thing that you know you are in charge of, ask. Everyone will appreciate an extra helping hand cleaning/drying reservoirs or controlling the HPLC's to clean everything in order to get to dinner/breakfast or bed faster.

8. Mock Beamtime Plan

It is important to go into a beamtime with a rough plan of how you will be spending your time. This won't and shouldn't be strictly adhered to, as flexibility is required, but see Table 1 for an idea of how you might want to consider your time being spent for a time-resolved experiment. This is based on our experience with a 120 Hz source. Lower repetition rates might take longer to achieve your goal, so plan accordingly.

Table 1. An example timeline of how you should plan beamtime allocation. This can obviously vary greatly depending on your particular experiment and sample.

Task	Hours Necessary
Beamline calibration	4–6 h on first shift, and 1–2 on subsequent shifts, unless you have something complicated
Test all samples and sample permutations that may be of interest	2 h per brand new sample; if it does not work within the 2 h, move to bottom of list and move on to next sample
Collect one really good dataset on reference structure	6–12 h depending on hit rate and quality
Align laser/mixing device for proper interaction region	6, laser alignment might require some time daily to ensure pointing and power, so account for this
Collect a data set at time-point X	8
Collect a data set at time-point Y	8
Total Time	36–60 h is typical

9. Avoid Beamtime Superstition

A wise scientist once said, "Don't stop collectin'!" Many times, the experiment finally starts working after hours of troubleshooting and dealing with unforeseen problems, but on a less than ideal sample or a reference structure. Suddenly, the control room is buzzing with energy again, PIs wake up and start asking questions, and suddenly someone is proposing to switch to the holy grail sample because "we've wasted/lost so much time already." More often than not, this is followed by numerous injector failures/clogs and ultimately loss of beamtime. Think carefully before you stop collecting data, even on a less-than-ideal or -favorite structure. Now you are without a complete structure on the working sample and have no data on the main sample. Worse, you cannot go back to the mediocre sample and finish that dataset because the running injector is irreparably clogged and the backup injector is having issues getting into vacuum, and the third injector has an unstable running mode and is giving a much lower hit rate. "Why did we stop collecting?" someone asks, rhetorically or not. We all know the answer was desperation. Hopefully, the following gives you some best practices and warnings on what to do during your beamtime to make it as successful as possible. You should be writing the manuscript and grants to fund more work on this project on the flight back, not thinking about what could've gone better.

9.1. Know When to Fold 'Em

A good card player knows when to fold and quit while they're ahead; however, we just told you to keep collecting. Like any good advice, it's vague, contradictory, and open to interpretation to fit any situation at the right time; thus, we want to remind you to stop collecting useless data. If you're seeing no hits, move on! No matter how perfect this batch was supposed to be. Properly log this. Loading a fresh reservoir of the same batch of crystals is a good place to start. Everyone rolls their eyes and then are pleasantly surprised when the *bad sample* or unknown sample is diffracting significantly better than the *wunderkind* sample of the last beamtime. Leave the egos at the Guesthouse and get ready to concede that your sample was not the best or your nozzle rod wasn't as reliable. It's OK; time to make the experiment work.

9.2. Stop Wasting Beamtime!

We literally dump the unused photons and you don't get them back. If your ideal sample is not ready or you're undecided on what the next steps is, leave a boring sample running. A few hits are infinitely better than zero hits. Zero data is collected when the beam is shuttered because the PIs are arguing and discussing with the beamline scientist whether they should be at 25% or 40% transmission, or whether the laser pointing should be checked or not. It is worth discussing things and making

informed decisions given the limited resources, but some data is better than no data. This data might be usable for the data team to test a new algorithm or check the detector geometry. It might be good practice for a new person to learn the data processing pipeline. It might be surprisingly useful later, e.g., the structure of a protein standard used for calibration can be thrown into a publication as yet another structure solved during your beamtime. It can even serve as a practice for the instrument staff to "clean up" the beamline to give diminished background scatter which is better for the rest of the experiment. This can be used to help figure out laser timing, or check if it is still correct. Got it? Always be collecting unless it's trash. You have other options and had high hopes for this sample!

9.3. Proper Logbook Keeping

This job always trickles down to the undergraduates or least occupied person at the moment. Let's be honest. If you're the best crystallographer, decorated liquid jet pilot, instrument scientist, data hacker, or PI, you are busy doing your job. However, the person in charge of the logbook is one of the most important, but least thought of, during the experiment. Afterwards, when you are processing the data or writing the paper, you will be grateful for a meticulous logbook. Experienced groups have figured this out the hard way, but new groups overlook it. This person should be actively asking the instrument staff about the details of the run, data analysts about the hit rates, sample preparers about the conditions and batch name, the injectors, what reservoir and loading time, and flow rates and nozzle flavor. They should make comments on what changed or is changing. They should also be asking the staff and collaborators to take screenshots of key figures and parameters. They should be taking pictures, they should be coordinating the shared spreadsheet the data analysts like to use, as well as the facility's experimental logbooks. Don't sleep on assigning this position. It should be a PI and maybe a senior, experienced person who knows what to log and what to ask, preferably someone that is not busy writing a funding proposal before the deadline or is not sleep deprived and nodding off. Have a backup, so when the logger is needed to make an important decision, or simply wants to use the restroom, someone can catch the nuance that might be needed later. Many times, later in the experiment, the question is asked: "which batch, nozzle rod, (insert parameter here) was giving us that great hit rate? How high was our highest hit rate? Let's go back to that condition." This should all be in the logbook.

9.4. Caveat Emptor

When preparing for an experiment, it can be overwhelming with all the techniques and methods. Keep your wits about you though. Just because a method works with lysozyme, granulovirus, or aquaporin, doesn't mean that it will work with your sample. Look in the methods section, supplemental info, or reach out to the authors and ask more questions. How do the crystal sizes compare to yours? What about their morphology, or unit cell size and solvent content? We all know lysozyme is indestructible, in many ways thanks to its tight unit cell; hence, buffer exchanges, physical agitation, temperature shifts, and changes in environment are no match for the HEWL. This does not necessarily translate to your crystal. I imagine if your crystal was this robust, you would not be reading this and would be busy writing the beamtime proposal or drafting the manuscript for your high-impact publication. More likely, the techniques might need testing and some slight customization and refinement to work well with your system. This is why it is important for you to reach out to the facility staff, come test at the facility or your collaborators' facility as early as possible, and be willing to adjust and test some more. Learn from your collaborators and ask questions, see why they think this approach is best and see if you agree. Tell them if you don't or have a different idea. Many times, there is not sufficient due diligence between the subgroups and things don't work until the postbeamtime meeting when everyone puts it together that the injection group is using a smaller diameter tube that might be clogging the new batch of crystals, for example.

10. Conclusions

Hopefully by now, we have inundated you with various ideas, anecdotes, and not-too-controversial hot takes, some of which you have probably thought of in the past in planning for your experiment, and some of which might be new to you as the idea had never crossed your mind. If you have taken nothing else away from this paper, however, we would like you to focus on two things: communication and characterization. Communication, not just between your whole team, but between you and the facility, will be essential in making your experiment a success. The facility people are there to make your experiment as successful as possible. Use the resources available to you, and ask every question you can think of. At the end of the day, every person who takes part in the experiment wants one thing: for the experiment to be a success. Contact early and often!

Characterization is a bit harder; some groups may not have access to the equipment necessary to perform all the characterizations we have mentioned in this paper. To this, we would like to redirect you to the first point: communication! The facility might have exactly what you need, and if they don't, they might know who you could ask for help, or how to get it for yourself. At the end of the day, the extra work put in by everyone before the experiment, or even in the proposal, will hopefully make a much more exciting success story that we can report, rather than another unreported failure.

Author Contributions: This manuscript was written in full collaboration between R.G.S. and C.K. All authors have read and agreed to the published version of the manuscript.

Funding: C.K. and R.G.S are supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract No. DE-AC02-76SF00515.

Acknowledgments: We would like to thank all the users whose stories we used in this publication, to help make our jobs easier. Thank you SED team and all of our brothers and sisters in arms at other facilities.

Conflicts of Interest: The authors are both employed by the Linac Coherent Light Source.

Reference

1. Sébastie, B.; Petra, F.; Mark, S.H. *X-ray Free Electron Lasers*; Springer International Publishing: Basel, Switzerland, 2018.