

On Contest Creation

Remarks on the Inaugural MAT

Dennis Chen

September 26, 2021

This is a reflection of the contests I've created under Math Advance for the last two years, primarily the Summer 2021 MAT. At the same time, it is meant to serve as a loose guide for anyone looking to create their own contest or refine their craft.

This report assumes you have some idea of what makes a problem worth doing and what makes a contest worth taking. In short, I assume you have **taste**. It also assumes basic common sense, because my goal is not to make a document for the 95% who can't figure out that they should announce their contest after they write it. It is for the contest organizers with real potential.

1 How to start

1.1 What's your purpose?

The mission of your contest isn't something you actively need to think about, but I think it's important to keep your contest purpose somewhere in your head because it informs your **test compilation values**.

Take the MAT, for instance. Most contests have primary goals of being a) educational and b) fun, and sometimes this comes at the expense of problem quality. MAT instead focuses on **egoism**; we strive first and foremost to make a beautiful test with only standout problems. I like to believe that our contest is also educational and fun, but it is undeniable our focus on **problem quality** supercedes most groups.¹

Even though none of these goals are at odds with each other, the difficulty of putting together a test compounds exponentially with each goal you have. So if you're not meeting your goals, you either have to tighten your focus, spend more time, or put less problems on the test.

¹Now that I think about it, MAT really is like Blue Lock. A contest full of standout problems and a team full of standout strikers. Perhaps I should use this analogy more often, at least in MAC.

Two of your contest series do not have to have the same philosophy. The MAT is produced like a work of art, meaning it is smaller and takes a lot longer to finish. The JMCs are meant for practice, so less emphasis is placed on egoism (making sure each individual problem is good) and moreso on cohesion (the test is good practice as a whole).

Your philosophy for a line of contests can also change between contests. For instance, the first MAT had a large emphasis on egoism. We wanted to put together the best test we could, and we made sure that the difficulty scale made sense on a **relative** scale. But next time we're going to put together a contest whose difficulty makes more sense on an **absolute** scale. This is especially because we're going to be moving locally, where the average contestant will not be as skilled.

1.2 Getting people

You need other people you can trust

This has become much less common recently, but there are still people who try to make contests alone. I think that is okay when you're starting out from a learning perspective.² But don't expect to be able to make a good contest alone.

Trust is a tricky word. Trust means you can confide in people without worrying about the fallout. But trust also means you put stock into their judgment. If someone you trust disagrees with you on something, it means you seriously consider the possibility they are right.

At a minimum, if you're working with people you trust, there's a higher chance that any mistakes you make get caught. Let me give a mathematical demonstration.

Example 1.2. Say that each member of your group of 3 people makes the correct judgment call 70% of the time, and if two or more people make the correct judgment call, then the group as a whole will act on it. Then the probability your group makes the right judgment call is $\binom{3}{2}(0.7)^2(0.3) + (0.7)^3 = 78.4\%$. Sure, it doesn't seem much more than 70%, but if you consider the probability you don't make the right judgment call, it's 21.6% versus 30%. That's a decrease in risk of almost 30%.

But there are plenty of other reasons you have to work with other people. Presentation is one. Unless you can consistently write all four subjects of problems, typeset documents, and do web development, you need other people to help, particularly with what's unrelated to the contest.

²But this is not necessary. Actually, I recommend against it, because you're going to waste a number of good problems on mediocre contests that you release. And unless you're really lucky and just happen to be a decent setter from the get-go, you also run the risk of tarnishing your reputation. And kids are quite harsh about contests these days.

And you get a more interesting problem bank to work with if you have more people. I think problem-writing and test-setting are somewhat separate skills. If you have more **good** problems to work with, it's self-evident that you'll compile a better test. Even if it's not significantly better, it's a lot more fun when you can consider another person's problems.

Speaking of fun: that's arguably the most important reason to work with others. Math contests at their heart are about the people and the interactions. If you're not even going to bother interacting with other people to make the test, how likely is it that your test will interact well with the people taking it?

I think a lot of math contests these days are not fun. When I look at the average contest, I'm usually not trying to decide which problems I want to do first. I'm just trying to find something I don't want to put off until later. You can tell when the contest-setting process is cold, because the problems will feel the same way: just like a list of things to get through. So make sure you enjoy yourself throughout the process, and don't be overly bureaucratic, because where's the fun in that?

Finding people

This part is the hardest, because you can't go out and send advertisements to everyone. You need to be deliberative about who you pick, but at the same time, your group has to form spontaneously. This is because, at any point in time, people have different ideas on what they want to do, and your job as a leader is to find people who want to write contests, take that feeling, capture it, and magnify it.

In short, you have to find people without trying too hard to look for them. In fact, the rule of thumb I use is that if I have to ask more than once, they're not interested enough to be on the team anyways. ³

If you can't find enough people, then don't run your contest yet. This doesn't mean you can't make preparations or write problems on your own. It just means that you should be patient.

I've seen a lot of people run crappy contests with a couple of standout problems. This is a real shame, because if only they'd waited longer, they might've been able to put together something good, rather than something mediocre.

The way you find people to work with is pretty simple. Get to know more people. That's how you get a wider base of people to draw from that trust your judgment.

When you're working with people, it's important to have a **center of gravity**. The chain of leadership needs to be clear, and it also needs to be intuitive. Arbitrarily selected leaders generally don't tend to work well.

It should be clear who the primary leader and representative of your group are. They don't have to be the same person – for instance, I'm not the one sending most of the official emails in Math Advance – but roles should naturally fit the people they are

³I'm in a position to expect that enough people will want to work with me to ask themselves. So that's the rule I tend to go by, especially since we already have plenty of people working on Math Advance projects.

assigned to.

There's a reason I say roles fit people, instead of people fit roles.

It's a typical misconception that you should find the best person for a certain role, and repeat for however many roles you ave. It gets further magnified by the official structures of clubs and non-profits, where they treat roles as permanent and people as interchangeable parts.

Its roles, rather, that should be frequently rotated out as everyone's interests and circumstances change.

So don't try to look for people because you think they'd make a good secretary, or a good treasurer, or whatever. Look for people who seem interesting and are interested in contributing to your group. Even if only half of them do any serious work, it'll far outshine any robotic organization put together by placing people in roles, rather than giving roles to people.

1.3 Accidents

When it comes to putting people and a purpose together, it has to be natural.

The best way to have something be natural rather than forced is to have it happen by accident. By definition, if something happened on accident, it couldn't have been forced.

When I say "by accident," I don't mean it in the same way that you'd trip over a branch by accident. But nobody who has built a large project has done it on purpose. It's usually something akin to a snowball effect: you have an idea, act on it, take the natural next step, and repeat, collecting experience and growing larger in the process.

It's not that it's a bad thing to have an idea of what you want to do. But if you start with a grandiose idea early on, your perspective on what's possible is going to be quite limited.

So don't steer too hard in any direction at first, don't worry if you've got no plans at all, and just wing it. You'll figure it out eventually.

2 Test-setting

This advice is going to be very abstract.

I assume that you've got most of the concrete stuff down already. If you don't know how to write a good problem or assemble a balanced test, you should learn those first. As the saying goes, "you have to learn to walk before you learn to run."

Once you've gotten a committed group of people that all want to work with each other, you're going to want to consider how you want to assemble your tests. As you're forming your group, you should already be creating a solid base of problems. ⁴

⁴I think your prospects are a bit sketchy if people are not actively doing this at the start, at least. How enthusiastic people are in the first month or two generally tends to be an upper bound.

There are two broad approaches you can take to determining your test philosophy. The first is letting your existing problems dictate how you build your test around them, and the second is finding a vision of a test and choosing your problems around it.

These two approaches aren't something to pick and choose between. They work in tandem. Of course, you can be more influenced by one approach than another. But neither of them is better.

For instance, the MAT was a case where the test philosophy was determined early on, and problems were mostly evaluated based on how suitable they were for the contest. But, at the same time, the **style** of the problems people created also affected our test philosophy, most notably the tiebreakers.

The problems and the test philosophy mutually affect each other. So **both have to be solid** when you're starting out.

How do you do this? By finding people who have good taste, or, at least, finding people who'll develop good taste. Your initial philosophy will look like a weighted average of everyone's individual philosophies. And your initial problem base will be the same. At some point you will have to figure out how you synthesize these perspectives, but before you get there, make sure they are worth synthesizing to begin with.

Figuring out how to synthesize everyone's perspectives, at least initially, is really hard without a center of gravity. You have to have someone calling the hard shots as well as the easy shots that could go either way. Otherwise you won't really get anywhere; everyone'll be stuck in limbo.⁵

Once you've got everyone going, inertia will do the rest for you. I say inertia in the sense of physics. Once you apply enough force, you'll have a fast velocity, even if you stop applying the force for a while. Of course, friction will slow you down. So you have to keep pushing, even if it's just a little.

After you run your first contest, your test philosophy may have to shift depending on how well people did, the demographics of your participants, et cetera. I call this factor **reality**.

Much like how you can't change the fact that your favorite baseball team lost, you can't change the fact that your target audience is mostly made of younger students. So you have to account for that.

I think that people could benefit from make their tests easier. On average, test-setters make their tests too hard. But that's not the main reason I recommend leaning on the side of too easy, rather than too hard.

I think it's better to make your contest too easy because if your problems are interesting, strong contestants will still enjoy the competition. Whereas if your problems are too hard, it doesn't matter how interesting your problems are: **the bottom two thirds of the contestants won't understand them.**

⁵People will protest, "But what about my democracy?" as if they fear that having someone clearly calling the shots will make their opinions go unheard.

If this is the situation you're in, give up while you're still ahead: if you can't trust your center of gravity, you don't have one.

I get the feeling: it's really hard to resist putting on that really good problem that also happens to be really hard. It's especially hard when a large part of your philosophy is based on egoism.

But any contest held on a sufficiently large scale needs to at least have its first third be approachable. I'd even recommend that the first two thirds be approachable, because you want everyone discussing all of the problems. Otherwise, why even put the problems that nobody discusses on the contest?

3 Logistics

Here is the long and short of it. You need

1. A website,
2. sponsors,
3. and outreach,

in this order. Most sponsors will ask for a website when evaluating your email, so you absolutely have to finish that first, though you can work on email drafts before then. You could start outreach before you get sponsors, but I strongly advise against this, particularly because a signup period of two or three weeks is ideal for online contests, and that is not enough time for sponsorship queries to be finalized.⁶

What follows are specifics on sponsors and outreach, as well as some general tips on time management. If you want to learn how to make a website, read the docs, they'll teach you much better than I can.⁷

3.1 Sponsors

If you're not a 501(c)(3) non-profit, your opportunities are a bit limited. In particular, most companies, particularly large investment banks, cannot send you money. But you can still receive merchandise, and even without the 501(c)(3) status, you can still get an impressive array of prizes.

For obvious reasons, we won't be including the list of sponsors we contacted, or any of our communications with them.

Here's how the process roughly goes for each sponsor. Repeat this for whichever potential sponsors you plan to contact.

1. Find the email addresses for whoever is in charge of publicity at the company.

This tends to be the hardest step, because sometimes there are no clear steps you can take to get them. There are a couple of companies whose public relations representatives are more difficult to find; Jane Street and Citadel come to mind.

⁶I would imagine that the ideal signup period for in-person contests is about a month more, i.e. three or four weeks.

⁷Feel free to send me questions though, if you've done your own work.

2. Send an initial proposal to them, with a bit of information about your contest.

Note that you might be asked to repeat information later on. The point of the initial email is just to provide enough information to signal your legitimacy to them.

This is the hardest filter to get through. Most of the time, if you don't get a potential sponsor, it'll be because they didn't respond to your initial email. It's very rare for someone to explicitly say "no," so any news is good news.

3. Respond to any further queries and make sure to prominently display your sponsor on you website and any advertisements.

This is by far the easiest part of the process, so I won't elaborate too much here.

The process outlined here seems quite simple, and that's because it is. But because the first step can be nebulous and intimidating, people start this process a lot later than they should. So know that the process is really not that hard.

You should be sending multiple sponsor emails in parallel. Because the publicity departments are quite busy, they might take more than a month to get back to you. I would say that on the tail end, some of our sponsors took about two months to confirm they were sending prizes.

3.2 Outreach

Once you're done getting sponsors, it's time to announce your contest. I recommend you make sure that you don't have any scheduling conflicts with other contests planning to host on the same day, particularly if they also seem legitimate and well-thought out.⁸

When it comes to online outreach, there are a couple of things you can do.

1. Make a post on the Art of Problem Solving Fora.

This option has become much worse over the past year or two. The quality of the fora have degraded due to internet trolls and general immaturity, and the establishment of the Peer-to-Peer subforum has greatly reduced visibility of student-run contests and programs. Regardless, it will still get your contest a decent amount of traction, and is not too hard to do anyways.⁹

2. Math clubs.

⁸If you are a prospective contest organizer, I strongly recommend you contact me for access to our contest discussion group. One of the things we do is set dates for contests and make sure there are no scheduling conflicts, which traditionally has been very hard because the only time contest organizers could learn the dates of other contests was when they publicly announced it.

Thus, one of the main benefits is that we set tentative dates internally rather than externally.

⁹If the AoPS post is getting hard to write, you're probably taking too much time. Publicity is the sort of thing where you get diminishing returns after a certain point, and for the fora, this point is not very far.

Send out advertisements to your local math club! While each math club may only have 5 to 10 people interested, if you have people from 6 different schools, you can get 30 to 60 people just by copy-pasting your initial advertisement.

This is limited by how many people are in leadership positions in their math clubs. I do suspect that out of people interested in writing math contests, a large number of them are also in charge of their school math clubs, or at least are on good terms with the people who are.

3. WeChat.

This is a pretty drastic choice to take, and unlike all the other items on this list, you could justify not sending advertisements on WeChat.¹⁰

This is because your main audience is explicitly parents, not students.

Sending advertisements on WeChat will drag your average grade and score down, and many of your participants may not enjoy the experience. On the other hand, more casual math students or competitors outside your reach might learn of the contest this way. I know there were a couple of contests I signed up for because my parents found out about and enjoyed, but I suspect that this may not make up the majority of the competitors you get.

These are not steps you have to complete for outreach or a “pick one of the above” situation. You should do all of these things, and you can complete them in any step, though it usually makes more sense to write the AoPS post first so word of mouth can spread on its own for a little.

3.3 Time management

By the time you publicly announce your contest, you should have everything finalized. Now, this is not always possible and sometimes you want to strategically push out a public announcement just to force some final checks to happen. But be aware this is a dangerous thing to do.

At no point should you announce a contest if there is no way you can finish on time. Nor should decide early on that you can finish on time, especially not in the face of contradictory evidence. Your reputation as a contest host is important, and you should not throw it away by haphazardly announcing contests that you know you cannot hold.

Just as a sufficiently complex computer system should not be single-threaded, nor should your contest preparation be. I understand that it is hard to be confident your contest will even exist when you don't have a test composed or many problems written. But it is best to start handling logistics as soon as possible. That does not mean you should immediately begin contacting sponsors before you've even assembled a draft,

¹⁰To be clear, I personally think it's a good idea to send advertisements on WeChat. We did this for the first MAT and plan to continue doing this for our next contests.

but nor should you wait until you've checked the contest 50 times to get your website started.¹¹

When it comes to scheduling, you want to host your contest in a period when interest is at its peak, especially if it's the inaugural contest. A couple of good times are early June, late July (before school is going to start but close enough to the start of school), and January/early February. Early and mid October may become good times as well, because the AMCs are moving to November.

But besides that, don't try to force your contest date. I don't recommend shooting for symbolic dates unless you've left yourself much more time than you need.

Something really common I've seen with respect to timing is **announcing the contest before the test itself is ready**. It shocks me that this even has to be said. If you aren't ready to run the contest in some form, don't run it to begin with! In fact, I would recommend that you refrain from setting any dates until you've finished compiling the test.

Generally you'll want three to six months to assemble the problems, depending on how much attention to detail you want to put inside your test. It may take longer if you're working out the kinks of your contest format. It took us about eight months to finalize the MAT, though a large reason for that is because we had a pretty far release date and kept revisiting the problems before then.

The format for the MAT took about two weeks to finalize after I came up with it, though we had a more liberal idea for tiebreakers for a while. Because our contest format was so conservative, it was finalized relatively quickly, but I recommend that organizers intending to write contests with more rounds (such as team rounds) give serious thought to time restrictions, format, and the difficulty spread.

When it comes to logistics, the primary concern is not how long it will take you to get things done. It's getting them done to begin with.

If you plan on making a website, make sure that at least one of your **core members** actually have experience working with a Javascript framework. If this is not the case, React is by far the most popular framework in the math contest community.¹² Learn it while you still have the time to.

If you don't plan on making a website properly, I don't think it makes a lot of sense to look for sponsors. Sponsors are meant to make your contest and your organization seem more legitimate and appealing, and it really makes no sense to put in all the effort to look for sponsors and no effort actually making a website. Nonetheless, if you really want to look for sponsors with a Google Site, I don't think a crappy website will hinder you too much.

You should bake in at least two months when it comes to sponsor emails; the turnaround can vary drastically, and that's about how long you'll need to finalize your

¹¹And if you **do** check your contest 50 times, you tend to grow complacent afterwards. So it doesn't do you a lot of good anyway.

¹²The stack I personally use is Node JS + React JS + Next JS + Tailwind CSS, and it's also what we used for MAT. We used Firebase for MAT, but I cannot personally recommend it one way or another because I haven't used it myself.

(Please don't ask me about backend, I don't know backend.)

sponsor list. Every student-run contest so far, including ours, has gotten more sponsors during the advertisement period. I think this is not ideal for the organizers, though it is workable. It creates a logistics time crunch, which is really no fun.

Empirically, people have given themselves less time than they actually needed for sponsors. So start about a month earlier than you think you'd need to at the latest.

📍 4 Test-solving and final checks

You should get a number of people with different skill levels **that have had minimal contact with the test** involved in the test-solving process. Test-solving is a process that means different things to different people, and no two groups use test-solving for the same purposes. Here are some common ones.

1. Validity checking – checking that all of the problems are correctly stated, the answers are what the contest-holders think they are, etc.
2. Quality checking – checking that the individual problems are good and the composition is good.
3. Difficulty checking – checking that the test is reasonably hard to be interesting to most of the target audience, and more commonly, that it is not too hard for the target audience. This is a lot harder, especially for less established contests, because they aren't exactly sure who their target audience is.

Whatever the case you should **never use testsolvers as a substitute for test-setters**. If you want help with the test or even individual problems, do not get testsolvers. Testsolvers are supposed to be the final check before you push a contest to production (i.e. print the papers).

4.1 Common mistakes

I recommend you read **Careful!** at the very least. If you work through the problems, you'll have a better idea of what I'm trying to communicate.

These mistakes are why you get people from outside to perform thorough **validity checks**. Some of these are so obvious that it almost seems impossible that an entire group of people missed it.

Here are some mistakes that MAC has made. The bolded parts are added fixes. All problems mentioned will be spoiled, so if you want to try them first, do not read the rest of this subsection.

Problems are from the 2020 JMC 10 and the Summer 2021 MAT.

Example 4.2 (JMC 10 2020/11). Jack has eight sticks of different lengths in the set $\{1, 2, 3, \dots, 7, 8\}$. How many nonempty subsets of these eight sticks can Jack choose so the range of the lengths is at most 4 meters?

There is nothing wrong with the formulation of the problem. But to get our answer, all of us did casework on the range, which we wrongly assumed had to be between 1 and 4 (inclusive). It turned out that we forgot about sets with size 1, which had range 0. Thus the correct answer was never on the test! Interestingly enough, none of the people who submitted pointed this error out and we only realized a year later.

There's a number of takeaways to be had here, but I think the most important is you shouldn't solve and write up at the same time, particularly for casework problems. Sometimes the elegance or conciseness of a solution can disguise the fact that it doesn't really cover every case properly.

Example 4.4 (MAT 2021/5). Let a, b , and c be **distinct** positive integers with $a + b + c = 10$. Then there exists a quadratic polynomial p satisfying $p(a) = bc$, $p(b) = ac$, and $p(c) = ab$. Find the maximum possible value of $p(10)$.

In short, the problem is solved by doing some algebraic manipulation and ending up with $p(10) = ab + bc + ca$.

There is a procedure that contestants go through every time they solve a "find maximum" problem. First, they find the upper bound, and second, they verify it exists and makes sense. They also make sure to watch out for redundant information.

As a problem-writer, even if you are **confident** in your answer, you should still substitute your variables into the equation and make sure the result makes sense. For **integers** a, b, c , $ab + bc + ca$ is maximized when $a = 3, b = 3$, and $c = 4$ (or some permutation thereof).

But consider the common mistake made in AIME I 2020/11. If $a = b$, then there is a level of redundancy as the conditions reduce to $p(a) = ac$, $p(a) = ac$, and $p(c) = a^2$. Since two points do not uniquely determine a quadratic, $p(10)$ could have been anything. Some contestants did not notice our mistake and put 33, which they were awarded points for. This is why the "distinct" condition was necessary, and the only reason we didn't catch this mistake is because we never checked our construction **from the beginning**.

In both of these problems, our issues were not related to the main idea of the problem. Rather, it was because we did not check details closely enough. This is likely because many of us got conditioned into evaluating the merit of a problem, rather than the contents, when thinking about it as we grow more familiar with our problem base. Outside perspectives are not painted with such bias, which is why it is important to seek and consider them.

5 Holding the contest

This part is the hardest, because unless you are extraordinarily lucky, the cold, hard reality that you have made a mistake on one of the problem statements will slap you in the face.

There are two sorts of mistakes you can make on a problem. One is an **answer** mistake. Maybe you forgot to consider an extra case, or added two numbers wrongly, or whatever. It's inconvenient and a little embarrassing, but it's usually not a big deal: you can just change the answer key and regrade all the tests. ¹³

The other is a **question** mistake. For some reason or another, the question cannot be solved as stated. The most common reason is ambiguity, but sometimes the cause will be forgetting a special consideration. ¹⁴

Whatever the case, this is the worst kind of mistake, because it means that contestants cannot solve it and it will interfere with their contest experience.

If you can catch the error as people are taking the contest, just issue an errata. If it's an issue of ambiguity, you may have to accept multiple answers. In some cases you will have to throw out the entire question. I only recommend you do this if **no sensible interpretation exists**.

If you're hosting online, technical issues are a whole other beast. If this is the case, the best you can do is send out a Google Form with PDFs of all the problems. Be prepared for this beforehand, just in case, though I don't think it'll be likely.

When you run into issues, whether it's on the contest or the tech side, **you need a fast and convenient method to communicate with all contestants**. Make sure you set this up before the contest. And remember that errors aren't the end of the world. ¹⁵

6 Curtain calls

If you've managed to successfully hold your contest, congratulations. But you aren't done yet.

After you've held your awards ceremony, the first thing you should do is get results posted. And do it right away. I'm sure there were times when you competed in contests and had to get results through second-hand word of mouth. This is not a good way for results to disseminate, and they could be distorted in the long game of telephone that ensues. Post the results before people stop caring, not after.

Next, you want to put up the problems. Again, you want to get discussion in before people stop caring, not after. I would recommend posting official solutions once the initial discussion has died down. Posting solutions too early tends to kill discussion, even if the competitors would appreciate them.

Once you've taken care of everything urgent, make sure you do some cleanup publicity-wise. You don't want your website to say "the contest will be held in July 2021" when it's already August or September. ¹⁶

¹³JMC 10 2020/11 is an example of an answer mistake.

¹⁴MAT 2021/5 falls into this category.

¹⁵I cannot think of a single contest that has managed to entirely avoid mistakes. The AMCs come close enough, but as far as student-run contests go, know you're in good company if you end up having a couple of errata.

¹⁶Funnily enough, writing this sentence is what reminded me to change this on the MAT website.

Finally, thank your sponsors! You're going to be working with them in the future, so there's no harm in reaching out and letting them know how the contest went.