

How to Get Your SIGGRAPH Paper Rejected

Jim Kajiya, SIGGRAPH 93 Papers Chair

Introduction

Everyone knows what acceptable SIGGRAPH papers look like: just look in the proceedings. When one only sees the accepted papers and not the rejected ones, it is easy to get the wrong impression of what it is that SIGGRAPH likes and doesn't like.

I've submitted a lot of papers that SIGGRAPH didn't like, as well as a few they did. Also, I've been on the papers committee a few times and know what it is they look for. This note tells you something about what happens to your paper as it goes through the reviewing process as well as what people discuss when they're trying to decide whether to accept or reject your paper. I'll try to tell you everything I've learned about the SIGGRAPH secret: What SIGGRAPH wants, and how you can give it to them so they'll accept your paper. I'll also talk about some of the flaws in the reviewing process and how you can protect yourself against them. Finally, I want to share some thoughts on the present course and the future of technical papers for SIGGRAPH.

Before we do this, I would like to say why SIGGRAPH reviews are done the way they are. There are two reasons.

The first reason is the principal feature of the SIGGRAPH conference publication that makes it very attractive: speed. SIGGRAPH is one of the few high-quality publications that can publish a paper in less than a year. In 10 weeks, SIGGRAPH can do what other major publications take 10 months to do. In a fast-moving field like computer graphics, this is crucial.

The second reason is that SIGGRAPH has chosen a very different quality strategy than most other conferences. While other conferences will accept papers of incomplete work in progress, SIGGRAPH has chosen to shoot for the highest quality papers of complete results. Because of this, 80% of submitted papers are rejected. The MacArthur Foundation is more generous with its "genius" awards than SIGGRAPH is with its papers. There are more MacArthur awards each year than SIGGRAPH technical papers.

The emphasis on both speed and quality makes the reviewing process for SIGGRAPH very different from of a journal or another conference. The speed and quality emphasis also puts severe strains on the reviewing process. In a journal, the reviewer and authors can have a dialog where shortcomings and misunderstandings can be resolved over a leisurely pace. Also, even if there are significant flaws in a paper for another conference, the chances are that strengths will overcome the weaknesses in the judging. In SIGGRAPH, if the reviewers misunderstand your paper, or if some flaw in your paper is found, you're dead.

The reviewing process for SIGGRAPH is far from perfect, although most everyone is giving it their best effort. The very nature of the process is such that many reviewers will not be able to spend nearly enough time weighing the nuances of your paper. This is something for which you must compensate in order to be successful. But I'll get to that later. First, let's talk about what happens to your paper.

The Reviewing Process

How does your paper get accepted or rejected by SIGGRAPH? Let's follow it through the entire process.

First, you work for months, slaving away at equations, hacking code, and feeding slides to your local photofinisher. SIGGRAPH fever rises to absurd heights at the last week: "Let's see I only have 105 hours until the deadline...". You put everything together, accomplishing superhuman tasks to make the Federal Express deadline at the very last minute. Your six copies are taken by the courier safely to the papers chair, then you-and everyone around you-collapse.

The next day, as you and hundreds of other morlocks around the world come out of sub-basements to blink at the first natural sunlight you've seen in weeks, is deadline day. Fully 85% of the 200 or so SIGGRAPH submissions arrive at the papers chair doorstep. Everyone else has worked until the last possible minute, too. The papers chair and several dedicated assistants then spend their long all-nighter giving your paper a number, entering it into the database, and typing and mailing a letter acknowledging receipt of your paper.

Immediately after this, the papers chair, along with two or three others on the papers committee, sorts through all the papers and assigns your paper to the pile for a particular senior reviewer. The papers program committee is made up of 25 or so of these senior reviewers. With the large number of papers, this partitioning process takes a full day.

One copy of your paper is retained by the papers chair. One copy is mailed to the secondary reviewer, and four copies are mailed to the senior reviewer. Thus each reviewer receives a large Federal Express box of your papers and video tapes. This usually happens a week after the deadline.

The senior reviewers receive a set of 14-18 papers. For half of these they act as secondary reviewer and for half as senior reviewer. As senior reviewers, they look at your paper and choose three additional reviewers-at least two of whom are external to his or her institution. The senior reviewer sends a list of these reviewers to the chair within two weeks.

The reviewers then each receive a copy of your paper, slides, and video. The reviewer reads your paper, evaluates it, and fills out the review form that eventually makes its way back to you. He or she may fill out the hard copy or may email the review back to the senior reviewer. The reviewer has four weeks to do this.

After the senior reviewer gets each review of your paper, a review summary is made and a score is computed. Copies are made of the summaries and reviews. The originals are then Federal Expressed to the chair.

The chair tabulates all the scores, sorts your paper according to score, records it in a database, and prints out a set of custom lists for each senior reviewer summarizing all the papers.

The following week, the paper selection meeting occurs. This meeting, where the fate of your paper is determined, lasts for two full days. If your paper is on the very bottom or very top of the list, very little discussion is given to your paper (unless the senior reviewer wants a short discussion by full committee). This no-discussion acceptance/rejection eats away at the top and bottom of the list until the density of discussion slows the process.

Then a "triage" session occurs. During this time, the senior and secondary reviewers, as well as others who might share expertise in the subject area, discuss your paper. They then decide to accept, reject, or discuss your paper. If they decide to accept or reject, your paper will receive a short summary in full committee session. But let's say they opt for discussion.

Toward the latter part of the first day, the triage session is over and the real work begins. About 60% of the papers could not be judged easily one way or the other. Yours is among them. So the entire committee discusses each paper and decides its fate. Often the discussion is postponed while more people read your paper and discuss it with the other senior reviewers. These papers are then discussed, often over dinner.

The second day is taken up with full committee discussions of your paper. I've been in sessions when some papers have been discussed and then postponed and then discussed again for five or six times. There's a lot of argument, some shouting, photos are passed around, and the slides are peered at. Usually the videotapes are viewed during the breaks. For difficult cases, summary letters are written to you that described the final opinion of the committee. At the end of the day, consensus has been achieved on all submissions, and your paper is either accepted or rejected.

After that, the disposition of each paper is double checked by the entire committee. All materials that go back to you are collected, and all copies to be destroyed are collected. People say good bye and rush off to the airport. Some stay to help the chair to group accepted papers into sessions for the conference, try to make up some sort of silly theme for each session, and to assign session chairs among the senior reviewers.

The chair then takes the database and generates acceptance or rejection letters and packages it up with any additional material to be sent back to you. You find out whether your paper was accepted or not in about 10 weeks' time.

If all this sounds like a scheme to exercise Federal Express, you're right. SIGGRAPH's Federal Express bills for this process run over \$3,000. That doesn't count your Federal Express bill, which in toto probably matches this.

The SIGGRAPH Secret

Just what is it they are discussing about your paper? Why are they shouting?

The SIGGRAPH paper selection meeting is an intense experience that only a few dozen people have ever encountered. It is not coincidental that the same people who sit on the selection committee will author many papers that appear in SIGGRAPH year after year. This is not because they're part of the "in" crowd whose papers are given favorable treatment-I haven't seen anything like that the times I've been on the committee. There are two real reasons. The first is that the program committee members are all accomplished authorities in their respective fields-they tend to do good stuff. The second reason, though, is due to their experience as a papers committee member. In this, they do have an advantage over you, an ordinary author, who hasn't been among the chosen few.

The advantage these people have is that they know what it takes to get a SIGGRAPH paper accepted. They know what the reviewers like and don't like. They know what kinds of things get discussed in the selection meeting. In short, they know the secret of what SIGGRAPH is looking for.

Review Criteria

What the technical program committee talks about when they consider your paper in their secret discussion is not really complicated. They discuss the questions in the review form you receive back with your paper.

They discuss what the reviewers said in their answers and whether they believed the reviewers. They talk about their personal answers to the review form questions concerning your paper. They sometimes are absolutely positive, or other times may admit they're unsure. Often times they want other committee members to read your paper and form an opinion. Several people who are intrigued may volunteer and enter into a small separate discussion on the various points in your paper.

The questions on the review form change slightly from year to year, but the basic thrust remains the same. If you know the questions asked on this form, you'll be able to predict what the discussion topics will be in the committee meeting. Let's look at the questions and see what kind of discussion goes with each.

1. Briefly summarize the paper.

This question really is a sanity check to make sure the reviewer understood the paper. The most dangerous mistake you can make when writing your paper is assuming that the reviewer will understand the point of your paper. The complaint is often heard that the reviewer did not understand what an author was trying to say. Remember, SIGGRAPH operates under the twin constraints of speed and quality. If you have quality, but it can't be recognized by reviewers who are in a hurry, you'll get rejected.

2. What does this paper contribute to computer graphics?

This question often generates the most discussion. Is your paper a pioneering new direction? Or is it just a small delta over previous work? The collective memory and knowledge of the papers committee is truly awesome. Obscure work that has appeared in a seemingly unrelated journal, or work embodied in some commercial product is at the collective fingertips of the committee. Nearly any facet of computer graphics, no matter how small, seems to be known by someone on the committee. Thus, your work is judged against a very rich context and history.

Your paper will get rejected unless you make it very clear, up front, what you think your paper has contributed. If you don't explicitly state the problem you're solving, the context of your problem and solution, and how your paper differs (and improves upon) previous work, you're trusting that the reviewers will figure it out. Don't try to make the reviewers dig it out from inside your paper. Maybe they will, or maybe they won't.

3. Is the paper stimulating?

Is your paper likely to create a new direction for research in computer graphics? Are people going to read your paper and want to extend your ideas? Are they going to read your system paper and say "Yes! I've been wanting to implement something like this, and now I know how." Is your application paper going to make people talk about your great new way to use computer graphics? Will your algorithm be implemented by dozens of people to become a standard widget in the graphics toolkit? Or is your paper a dead end? Is it just going to take up pages in SIGGRAPH, not be read or referenced, just drop out of sight?

Again, stating the problem and its context is important. But what you want to do here is to state the "implications" of your solution. Sure it's obvious....to you. But you run the risk of misunderstanding and rejection if you don't spell it out explicitly in your introduction.

4. Is the paper of interest to the SIGGRAPH audience?

Does your paper solve a long-standing problem that people want to know how to solve? Is your system or application interesting to a wide swath of the audience? Or is your paper so narrow that only ten people at the conference will care about it? When you speak will the auditorium be packed, or will everyone leave?

Well, to get rejected, pick a subject no one cares about. But, if your subject has less than obvious application to a wide range of graphics problems, you'd better figure out how to say it convincingly in your introduction.

5. Is the paper well written?

Your ideas may be great, the problem of burning interest to a lot of people, but your paper might be so poorly written that no one could figure out what you were saying. If English isn't your native tongue,

you should be especially sensitive to this issue. Many otherwise good papers have floundered on an atrocious text. If you have a planned organization for your discussion and you not only stick to it, but tell your readers over and over where you are in that organization, you'll have a well written paper. Really, you don't have to have a literary masterpiece with sparkling prose.

6. Can an experienced practitioner in the field duplicate the results from the paper and the references?

This question often gets people shouting in the committee meeting. Basically the question is about completeness. Your paper may be doing something very interesting, of obvious importance to graphics. But your paper leaves something out. Your description of what you're doing is so sketchy and abbreviated that no one will be able to do the same thing. The key purpose of a technical or scientific paper is that it contains enough information so that an experienced practitioner, say, a graduate student in graphics, can reproduce the experiment. If you've not explained enough about how you do things—even if you think it's just obvious—then it's quite likely your paper will be rejected.

7. Should we accept this paper for SIGGRAPH 93? Why?

This last question is the final recommendation about acceptance. This recommendation is tabulated to make a score for your paper that determines where in the sorted list your paper will find itself. I used to think that if just one reviewer didn't like the paper, you'd be dead. But since I've been on the committee I've found that that's not true at all. I've seen some rejected papers that have had four "accept" recommendations and one "maybe." This is because the committee doesn't blindly follow the scores at all. They really discuss the merits of each paper. A paper might be a solid technical paper, written by well-known names, but it might be boring. It might be just so small an advance over existing techniques that it's not very exciting. The committee has a detailed discussion trying to isolate a new twist in the paper. The discussion goes back and forth about whether the new twist is obvious or not. Even though it gets favorable reviews, the committee decides to reject.

On the other hand, a paper might have a really neat new idea. That idea may open up a whole new line of work. But the paper is badly written, and it doesn't really explain things enough so that someone without a Ph.D. in mathematical physics would be able to do anything with it. Because of this, all the reviews are bad. Someone says that one of the authors is a responsible person and will probably rewrite the paper into something decent. Someone else says that there's no guarantee that anything at all will be changed, then the proceedings will have this horrible paper in it: why not reject and wait till next year? Finally the committee votes, it passes by a narrow margin. Thus the committee has decided to gamble on the authors to fix the problems once they're pointed out. Sometimes the gamble pays off; sometimes it doesn't.

All this brings up a phenomenon that happens inside the paper selection meeting. Often a committee member may take up the cause of getting your paper "in" and argue for acceptance of your paper. Tom Sederberg, the SIGGRAPH 91 chair, has called the people who can ferret out the good features of your paper "paper champions." On the other hand, there may be a committee member who is very articulate, forceful, and negative, who argues against your paper. They look for and find flaws in your paper, they sway the committee to reject your paper. Ed Catmull, the SIGGRAPH 92 chair, has called these people "paper killers."

One job of the papers chair is to see that the committee is staffed with people who are paper champions. We want to avoid paper killers.

So that's it. That's what goes on in the discussion. I must admit that as a paper author I've been guilty of screwing up on almost all the points mentioned in the review criteria. My long string of paper rejects have been due to repeated deficiencies in not stating the problem or its context, not explaining why the subject is

interesting, writing disorganized papers, and leaving out key points that I thought were obvious. And just writing stuff that was plain hard to read, so that some of the reviewers just missed my point.

Mistakes

The characteristics that make SIGGRAPH so attractive - speed and high quality - also make SIGGRAPH an imperfect vehicle for technical dissemination of graphics ideas. The review process is far from perfect. The chair needs to get your paper quickly distributed. The first mistakes are made right there: among the 200 or so papers, some are just sent to the wrong senior reviewer. The senior reviewer may not carefully read your paper and ask the wrong people to review it. Those people may not read your paper carefully, they misunderstand it. Finally, you may have your paper attacked by a paper killer that the chair mistakenly appointed.

How can you protect yourself against these mistakes? You must make your paper easy to read. You've got to make it easy for anyone to tell what your paper is about, what problem it solves, why the problem is interesting, what is really new in your paper (and what isn't), why it's so neat. And you must do it up front. In other words, you must write a dynamite introduction. In your introduction you can address most of the points we talked about in the last section. If you do it clearly and succinctly, you set the proper context for understanding the rest of your paper. Only then should you go about describing what you've done.

Another point is why rendering papers have an advantage in SIGGRAPH. If you have good-looking pictures, you've got your foot in the door. SIGGRAPH reviewers are like everyone else. They first look at the pictures in your paper. If your pictures are really good looking, they're going to go to some effort to find out how you did them.

You can use those pictures in another way. Ivan Sutherland once told me that Scientific American articles are constructed so that you can get the point of the article just by reading the captions to the illustrations. Now, I'm not suggesting that you write a technical comic book; but you should take a look at those SIGGRAPH papers you were initially attracted to and see how they went about getting their point across.

Unless you write about a very limited subject, or unless your results are technically incorrect, rejection has very little to do with the subject of your paper. It has a great deal to do with how you wrote your paper. After all, if everyone misunderstood your paper, you might consider that it might not be quite as clear as you thought. Reviewers are in a hurry: you have to get your paper just right or you will suffer rejection. Rejection doesn't come from the subject area, it really just comes from an imperfect understanding on both sides.

But on the whole, it's a very noisy process. The SIGGRAPH review is done quickly, by the best people the chair knows, and by the best people they know, with everyone earnestly committed to put out the highest quality proceedings possible. Mistakes are sometimes made.

What SIGGRAPH Wants...

There seems to be a number of prevalent myths and misunderstandings about what it is that SIGGRAPH wants and doesn't want for its papers. Each year, the papers that appear in the proceedings appear to be more and more technical, about narrower and narrower areas. I've spoken with many people who've been concerned about the path that the papers sessions for SIGGRAPH have taken.

I fear that this trend is all too real. I'm very worried about it. I believe that the papers sessions at SIGGRAPH are in trouble. Only about 10% of the technical program registrants go to the papers sessions. Sometimes fewer than 200 people are in attendance at a paper session. This tells me that very few people find the SIGGRAPH papers interesting anymore.

For some years, people thought of the papers sessions as almost exclusively about rendering - SIGGRAPH as "SIGRay" or "SIGRadiosity." Or people have viewed the papers sessions as valid only for those papers that have been about "pure" graphics. Almost everyone agrees that the papers are the exclusive domain of the academics, exploring esoteric and obscure corners of graphics.

I believe that the reason for this alarming narrowing of SIGGRAPH papers is a dangerous positive feedback loop. You see, people can't see what papers are rejected. They can only see the papers that are accepted. Thus when you look at a proceedings you see a certain set of papers and you say, "Ahh,...THAT'S the kind of thing that SIGGRAPH wants." So, if you have an idea for a paper that isn't like the kind that have been appearing in SIGGRAPH for the last ten years or so, you wouldn't send it in to SIGGRAPH. You say, THIS is not really what they want at SIGGRAPH anymore, they want THAT. If you are brave, do submit to SIGGRAPH, and your paper becomes a casualty of the 80% rejection rate, you feel that SIGGRAPH really doesn't want your type of paper anymore. Thus you don't send anything in to SIGGRAPH about that subject again.

Well, the papers committee and the papers chair don't really determine what SIGGRAPH publishes. The authors who brave the SIGGRAPH review process are the real controllers of what appears in SIGGRAPH. The committee can only select among the papers that are submitted. Consider this: if there are 150 rendering papers submitted, only two systems papers, one interactive techniques paper, and no applications papers submitted, what will the proceedings look like? Then everyone will say, "See, SIGGRAPH only wants rendering." But what really happened is that SIGGRAPH "rejected" 127 rendering papers, and rejected only one systems paper, and didn't reject a single application paper!

How Can Papers Sessions Be Fixed?

Is there a way to make a kinder, gentler SIGGRAPH? Can something be done about the 80% rejection rate? Actually, something has been done about it. Several years ago, there was an institutional constraint on the papers session and proceedings to fit in a single track. Because of this, there was a limit on the maximum number of SIGGRAPH papers that would be accepted, no matter how many good papers there were. During those years, one thing that was watched very closely was the number of papers that were accepted as the paper selection meeting progressed. As the limit was approached, people tended to get a bit more critical of flaws in the paper under discussion. Almost as a confirmation of the policy, the limit was never reached. Meanwhile, the number of SIGGRAPH submissions (and rejections) steadily increased. Today, that constraint has been lifted. There is no longer any limit on the number of papers allowed. And pleasantly, I found that during the last meeting, concern about the number of accepted papers was not a big issue. In SIGGRAPH '92, no parallel sessions happened to be required. We're still under the old limit. But now, the number of papers accepted is solely determined by the big issue. The big issue, of course, is "Is it above the [quality] threshold?"

It is foolish to capriciously tinker with the speed and the quality of SIGGRAPH in the hopes one might fix the serious positive feedback focus problem. Frankly, I'm afraid to make big, sweeping changes in a process that works well a lot of the time. However, you'll note that this year there is a new class of papers: long papers. Only systems and applications category papers will be admitted in the long class. Andy van Dam has pointed out that the page limit for regular papers favors research papers. A research paper can usually state the problem, its context, and the solution in a short space. A system paper needs more pages to do this and it must also describe the experience that the builders have had with the system. Eight pages was just not enough room to write a decent system or application paper.

The root cause of the positive feedback loop, however, remains. It is self-censure. People just won't send in papers on subjects they think SIGGRAPH doesn't want. I can understand this: even after all my SIGGRAPH rejections, it still hurts.

The entire reason I've written this document is to try to break the loop. I want to communicate to you, and I want you to communicate to your colleagues, that SIGGRAPH is in the business of publishing good technical work in graphics of "all" flavors.

SIGGRAPH really does want papers about user interfaces, visualization, graphics hardware, graphics software systems, interactive techniques, displays, innovative applications, video games, combined graphics and sound, hypermedia, virtual reality, typesetting, color, paint systems, image and video compression, image and video processing, and how to make pictures that aren't just pretty but say something too.

Sure it's true that they've rejected papers in all these areas over and over again. But, it's also true that they've rejected 10 times as many papers on ray tracing. The narrowness of the technical focus of the papers can be fixed only if you and your colleagues send in quality papers about a wider range of subjects. My earnest hope is that the SIGGRAPH 93 technical papers program will not just be about modeling, rendering, and animation.