

First-hand Account of “The Undefined Behavior Question” Incident

By Andrew Tomazos <andrewtomazos@gmail.com> www.tomazos.com - 2024-11-26

TO WHOM IT MAY CONCERN

I wanted to write this statement to set the record straight over recent events regarding my expulsion from the C++ Standards Committee over the paper title ‘The Undefined Behavior Question’. I will try to give an honest unbiased accounting of events as best I am able.

I have been a contributor to the committee off-and-on since about 2012. Around that time I created a well-known advanced C++ programming course and, as part of preparing the course materials, my C++ questions got hard enough that they ended up getting routed to committee members, and so I discovered and got involved with the standardization process through those connections. Over the years there are a few features of C++ I have had the honor of designing or contributing to in various ways, and while the process is sometimes “frustrating” (to put it mildly), it is quite cool and rewarding to be able to say “you’re welcome” when a co-worker uses a language feature that you worked on.

At some point I was made officially part of the Standard C++ Foundation delegation during some kind of reorganization. I’m not super-familiar with the procedural inner workings of the committee’s bureaucracy, and to be honest, I’ve never really cared. I’ve always been a volunteer (as are most members), and was never paid to contribute before or after this. I donated my spare time for the purposes of giving something back to the programming language I primarily used professionally during my career. I’ve reviewed and written many C++ proposals over the years.

In May 2023 I reviewed an interesting C standard committee paper that came out of the Undefined Behaviour study group, called **N3128 Taming the Daemons - Undefined Behavior and Partial Program Correctness**. (That paper would go on to be adopted into C23.) In 2024 I reviewed another paper called **P1494 Partial Program Correctness** that seemed highly related to that paper. I noticed that if C++ adopted the C paper, then it would render P1494 moot, and it seemed “safer and simpler” to me to do so. I brought up this point at a committee meeting, and as a result of that comment I was challenged by the chairperson to write a C++ version of the C paper. I contacted the author of the C paper and we worked together to co-author a C++ version of that C paper called **P3352 Taming the Daemons (C++ version) - Undefined Behavior and Partial Program Correctness**. All the paper authors of these three papers were in private technical discussion about them leading up to the recent Nov 2024 meeting in Poland.

In the lead up to the Poland meeting I also wrote a summary paper of the three papers mentioned in the previous paragraph. I felt the three papers were too complicated to present from a cold start, so I thought a primer paper could provide a good framework to get the main ideas loaded into peoples minds. It posed the key language design question the three papers are tackling which is "Should undefined operations be allowed to affect observable operations that happen before them?". I dubbed this question "The Undefined Behavior Question" and that was also the title of the paper: **P3403 The Undefined Behaviour Question**. I posted this paper in the usual fashion to the C++ standards committee.

At the meeting all four papers were routed to the SG23 Safety group for a preliminary discussion. SG23 discussed all four papers for several hours. As planned, I presented The Undefined Behaviour Question first as an introduction and then we discussed the other three papers using the language from the primer paper as a context and a framework. I would say this went to plan. SG23 effectively voted to change C++'s answer to The Undefined Behavior Question (technically this means we wanted to protect volatile accesses from subsequent UB, the same as I/O functions), which is what P3352 proposes: So SG23 also voted to forward my main paper P3352 to the next higher group which is called Evolution. No vote was taken on P3403 because it doesn't propose anything. That is, it poses and explains a question, and so is purely informational (it doesn't propose an answer to The Question, the other papers do) - but certainly no issues were raised with either the papers title or its content.

The next day we took the papers to the Evolution group. During the presentation of P3403 The Undefined Behavior Question there was what seemed to me to be an odd comment. I think I heard someone say "you should change the paper title". Upon hearing this I didn't really understand, I thought it was a joke about "undefined behavior" being a dirty word or something. So I giggled a little, playing along with what I perceived as a joke, and continued on with the presentation. The other papers were presented as normal.

The next day I received an email from the head of the Standard C++ Foundation delegation informing me there had been formal complaints lodged with the International Committee for Information Technology Standards (INCITS) about the interaction mentioned in the previous paragraph. After some back and forth I was informed by the delegation head that the reason for the complaint was "The reason people are upset is because they feel the paper title refers to Jewish Question", with a link to the wikipedia page of the [Jewish Question].

I was mortified, and felt physically sick. My first reaction was to immediately apologize profusely and to offer to change the paper title. I had no idea anyone would form such an association. I was informed of the technical procedure to change the paper title, and started to work on that.

But something happened as I was going to do that. I thought through options for drafting changes to the paper, and realized what I was thinking about was altering the semantic structure of the paper itself. The paper poses an important question about undefined behavior and so the title is correct, so to speak. I could change it to "An Undefined Behaviour Question" or "The Undefined Behavior Query" or a "A Question on Undefined Behavior", but it was

changing the way I was thinking about the paper, and it's not clear if any of those changes would alter the association and be acceptable. Is it the concept of asking an important question that is creating the association? Throughout the paper it refers to "The Question" as an abbreviation for "The Undefined Behavior Question". So even if I changed the title the paper would still refer to "The Question", and I thought to myself, would this satisfy those complaining or do they want that changed too?

Then I became a little angry. I suddenly got this feeling as if something wasn't right here. After thinking things through, I made a moral judgement that I hadn't done anything wrong. The right thing to do was to respectfully refuse the request and to not explain myself - as I don't want to live in a world where someone is not permitted to form and ask important questions without explaining themselves. I have a moral duty to act locally in a way that matches with how I would want the world to be globally.

You see, I absolutely believe the question "What is the appropriate status of Jewish people?" is reprehensible and immoral. That question presupposes it is moral to discriminate based on a person's ethnicity or religion. I find that kind of discrimination disgusting, as individuals should be judged based on the content of their individual character and their individual actions, and not based on their "group identities". That kind of discrimination in the extreme leads to apartheid and genocide, as history repeatedly teaches us. I am not a religious man, but I believe in "love thy neighbour" and "do unto others as you would have them do unto you". I can think of no greater violation of those two rules than apartheid and genocide.

However, I think the question "What is the appropriate status of Undefined Behavior?" is a perfectly valid and appropriate technical question, as program behavior is not people, so discriminating program behavior based on whether it is defined or undefined is morally acceptable. This is easily shown: Removing undefined behavior from a programming language is not an act of genocide.

So whether a question about the status of something is moral or immoral, depends entirely on what the subject of the discrimination is. If the subject is an ethnic or religious group of people, that is immoral. If the subject is a class of operations in a programming language, then that is fine. We should not ban asking all important questions about the status of things, because such questions have been inappropriately posed in the past about ethnic or religious groups of people.

I wrote a short note entitled "Historically Insensitive Paper Title" to the Evolution group explaining briefly what had happened and my decision. I apologized for not understanding the request initially, but I informed them that I had decided that the morally-correct thing to do is to respectfully refuse the request.

I was then informed by the head of the delegation that they "fully understand that principles are important, but in this case the principles in question also affect the Foundation. If there is no way of mending the situation, we will have to politely excuse ourselves from being drawn into the

discussion by rescinding the invitation to be our alternate. At the moment, we are at a place where the title is causing offence and objections from multiple individuals, and that is not acceptable for us.“

I urged the head of the delegation to reconsider their ultimatum, but they took it as me refusing to change the paper title, and said “I'm afraid the Standard C++ Foundation is no longer able to sponsor your WG21 membership.” and I was expelled from the various WG21 services and uninvited from the committee meetings.

I went public with this story, because I feel it is a story of public interest. In the aftermath of that, there have been a lot of false rumors circulating about what happened. The “other side” of this story, the Standard C++ Foundation, have remained tight-lipped, simply claiming “the content of the complaints are confidential” and so they can't talk about what happened. This has created a vacuum where idiotic rumors have spread. People have suggested that the “real issue” was the quality of my work, or long-term behavioral conduct problems, or even more ridiculous that it has something to do with ChatGPT. I've been called a “grifter”, a “troll”, a “swine of a man”, that my work is a “denial of service attack” on the committee, that I've been trying to “use ChatGPT to gain power” - whatever that means - and that all this was just “the last straw” and a trumped up excuse to get rid of me. Frankly, that's a load of rubbish: This is the first time a formal complaint has been issued about my conduct to INCITS and forwarded to me. I strongly believe such a complaint would have been issued against anyone posting a paper with the title ‘The Undefined Behavior Question’, regardless of the author or the paper's content.

Some people have been claiming that if they were in my shoes that they would have just given into the request and changed the paper title - like what's the big deal. To those people I say: you weren't in my shoes - so you don't know what you would have done.

There have been a couple of comments from other parties that have first hand knowledge of some of these events that I want to highlight:

“As the person who wrote the original paper for C ended up co-author on the main paper of Andrew (but not of the paper in question). I do not think this paper was written by ChatGPT. Andrew also denies it (on Reddit). The technical content seems ok to me and when this paper was discussed in a subcommittee of C++ which I attended as a guest, nobody complained about the title or the content. I agree the title is unfortunate.”

Bjarne Stroustrup commented privately to me “Unfortunately, I have not gotten around to reading that paper. I don't see why anyone would be offended by that title nor do I see why you would not change it if someone did take offense. I think that we have better things to do with our time than to fight over something like this.”

Herb Sutter said privately to me “We've been glad to have you participating in WG21 as a free guest for many years. If your current sponsoring INCITS member isn't going to continue to be

responsible for your participation then you're welcome to rejoin WG21 via a different INCITS member."

Well, thank you for having me as a "free guest" for all these years, and not asking me to pay for the privilege of donating my spare time to improve the C++ programming language. I won't be applying to rejoin WG21 through a different INCITS member, as I've decided to wash my hands of the C++ standards committee. I feel unfairly treated over this incident, so I don't want to reward that - and besides - my motivation to contribute was always hanging by a thread anyway, given how frustrating the standardization process is to begin with.