

The 12 latest Twitter Poll Results of 2018



Via [our Twitter channel](#) we have various polls. Not always have we shared the full background of these polls, so we've taken the polls of the past half year and put them here. The first half of the year there were no polls, in case you wanted to know. As inclusive polls are not focused (and thus difficult to answer), most polls are incomplete by design. Still insights can be given. Or comments given.

Below's polls have given us insight and we hope they give you insights too how our industry is developing. It's sorted on date from oldest first.

It was very interesting that the percentage of votes per choice did not change much after 30 votes. Even when it was retweeted by a large account, opinions had the same distribution.

Is HIP (a clone of CUDA) an option?

AMD has worked on their implementation of CUDA for quite some time. It's rather simple to do 80% of the compiler part, but then come the weird functions that might only be there for backwards compatibility. Add to that the libraries which needed to be optimised for AMD GPUs.

It's July 2018 and the time to port software using a Python-based tool called "hipify" is a lot less then when the tool was first created. But how is the current status observed? AMD might see it as an option, but do others share that idea?

<https://twitter.com/StreamHPC/status/1014959061807267841>

50% Always Nvidia with CUDA**11% Always AMD with HIP****39% The fastest GPU with HIP****28 votes**

The fastest GPU is obviously the best solution, and would have be much higher if it said "The fastest GPU with CUDA". I wanted to know how HIP was observed, so framed it into CUDA and HIP. And you see HIP is not seen as a good solution, increasing the votes for Nvidia.

Out of scope: SYCL is the answer from OpenCL to make a real alternative to CUDA.

Intel discontinues Xeon Phi

We've been a (sometimes loud) non-fan of the Xeon Phi, as sarcasm was used when benchmark results were discussed. Meanwhile various scientific papers were written with great info on the accelerator, and numerous HPC-centers were happy to build these in. We remained skeptic and assumed Intel gave those cards away for free (or even paid for the electricity costs). We were hopeful to see the results of the socketed XeonPhi and were ready to invest in it (a complete system for developers was ?4000), but Intel gave up and pulled the plug.

There is no better moment to ask if others when the whole line is discontinued. To keep the focus on Intel, the question was not focused on competition.

<https://twitter.com/StreamHPC/status/1015707476879167489>

44% It simply did not perform11% Replaced by FPGAs29% Replaced by 20+ core CPUs16% No, it worked perfectly..55 votes

The results are clear, that we were not the only ones who disliked the architecture.

FPGAs are unexpected, as we've seen even worse results on FPGAs when porting the XeonPhi-code, even when taking a lot of time to fully optimise the code for the architectures. But indeed there are a (very) few examples where FPGAs are a better option. But so there are people who said the XeonPhi performed perfectly - simply cases where we did not work on ourselves.

Multi-core CPUs are actually quite close to the Xeon Phi, as the architecture was sometimes described as nothing more than 72 Pentium cores. As the latest CPU-architectures

So please, **dear researchers**, stop writing great articles about new accelerator-cards, but always (read: **ALWAYS**) benchmark your algorithms on various other (recent) GPUs/accelerators. No access to such machines? Ask [Professor Simon McIntosh-Smith's group](#) and/or us, to benchmark your code.

Shitty code is impossible to fix?

When searching articles around code quality, I came to this statement. It's not an exact quote, but I wanted to make it a statement.

<https://twitter.com/StreamHPC/status/1019502311100616704>

91% Agree9% Disagree**47** votes

[Reddit source.](#)

Most coders who worked with the so called "shitty code" think that in most cases only a rewrite will work. So when the request is to make the code faster with something like magic, then we only have the overclocking opportunities left. If that does not work, then nothing else than rewriting works.

Shitty code is not bad on itself. Delaying rewriting is.

Intel on 10nm - Is there an escape route?

Lots of rumours that Intel could not get their 10nm, while the others could, even delaying their 10nm products to 2020. Serious news, if true. Given that Intel's leadership was not really good for quite some years, the poll was also a bit about how one thought this would change on the short term.

Using another foundry would solve the problem that was predicted (and as time of writing coming out): IBM, AMD, ARM, everybody getting a chance to get a larger part of the cake.

<https://twitter.com/StreamHPC/status/1027458173874122757>

53% Patience till 2020 **15%** Surprise breakthrough **6%** Use TSMC, GF or Samsung **26%** Something else **34** votes

Patience is not the Intel I used to know, but neither such delays in technical advancements.

Metal only on MacOS

So Apple kicked out OpenCL, and was pretty clear about what developers should do: rewrite everything in Metal. As there is not even porting-tool made available this was quite unfriendly. So what did the people think, who should do the work?

<https://twitter.com/StreamHPC/status/1035399749711605761> **13%** Yes **67%** No **13%** Yes, via Vulkan+MoltenVK **7%** OSS, not by me **45** votes

Those who only had a few kernels, could quickly port it and would not have much work in maintenance. Those who found MacOS important for their software, would do the porting. Several put their hope on Vulkan-on-top-of Metal. A large part simply objected.

Intel stuck on 10nm - should they skip 10nm?

Intel always had the technical advantage, but that era seems to be over. The rumour is this time they realised themselves that they're lagging and 10nm efforts have stopped, but soon Intel officially reacted that this was false information.

A rumour starts often by unsatisfied (higher up) employees, so we do can conclude that this discussion has actually taken place.

While they're being attacked from all sides, I'm a bit surprised they don't seem to show strong leadership. All I see is this "AI software also runs on our processors", which doesn't sound great to me. I love it when hardware companies understand that software is as important, but when hardware gets secondary then I think something is very wrong there.

<https://twitter.com/StreamHPC/status/1054730988759191553>

48% Yes14% No38% It's complicated98 votes

Skipping 10nm is not as simple as "skipping 32nm" would be, but I see the high number as a cry for Intel getting their act together.

Who's the master of tools?

Nvidia is well-known for their tools. But Intel, AMD and the others have not sit still. Due the maximum of 4 options, I focused on server/desktop GPUs.

There are two ways of using tools. As a beginner, when you have not really an idea where to look. And as a senior, when working on a large project. When you are doing MPI+GPU, then good tools are crucial.

<https://twitter.com/StreamHPC/status/1055553754328645637>

67% Nvidia30% AMD0% Intel3% Other30 votes

As expected Nvidia came out highest. AMD has built up trust with ROCm. Intel has quite a good suite, but got an unexpectedly low 0 votes. It might be because its products are paid for, or because Intel GPUs (or Xeon Phis) are not used in the large-sized installations (as was the question)?

Anyway, Nvidia remains the example of how it should be done.

How should we share our knowledge?

Sharing our knowledge actually does not reduce the number of projects and has a good chance we're educating a future colleague. We're therefore confident we should keep sharing our knowledge. but what knowledge exactly? I think the most important are the techniques, then the language. Still the language is a good part of the technique.

I was curious if CUDA+HIP would get more votes than Nvidia alone, as there are good reasons to make code work on GPUs from both vendors. Second I was curious how OpenCL would do in comparison. "Generic GPGPU" was added as a "other" vote.

<https://twitter.com/StreamHPC/status/1056474875362246656>

21% Nvidia CUDA + AMD HIP48% OpenCL on GPUs+CPUs22% Generic GPGPU9% Nvidia CUDA only58 votes

A remark was "SYCL + OpenCL", which indeed is a good option.

As we've started with OpenCL, this could have influence on the number of votes. We do a good share of CUDA-projects, but that's not a well-known fact.

RedHat being bough by IBM

We're not really fond of RedHat/CentOS. The projects that required this Linux the time put in work-arounds and hacking was higher than Windows-projects, and far higher than projects that required a modern Linux. We like to put our time on the important things, not OS-related problems. So when it got known that IBM bought RedHat, the obvious question was "why?". It was said the Cloud was important, but for that you would not need to buy a complete distribution. In the end, the question would also be on the Cloud, if users are happy with their Linux distribution.

<https://twitter.com/StreamHPC/status/1056975447756496902>

15% Red Hat / CentOS **72%** Debian / Ubuntu / Mint **11%** Arch / Manjaro **2%** Suse / OpenSuse **61** votes

Intel/Altera support their FPGAs only on CentOS/CentOS, so the results could be skewed by those who got used to RedHat. It's higher than we expected, as we've not spoken any power-user who likes the distribution or used a more positive description than "it does the job".

How AMD's offering is observed

I got curious how AMD then was observed as a whole, and thus sent out a broad question.

<https://twitter.com/StreamHPC/status/1058714993640828938>

3% Already far better **10%** Truly competative **18%** Closing in fast **69%** Still year(s) behind **73** votes

One reaction was clear: two years in Gaming, a year in compute.

I had a follow-up question on AMD's software, but they released their new CPU "Rome" and 7nm GPU "MI60" this week, so a bit of bad timing.

It came to my ear a lot that most people don't know what AMD is actually doing. Adding the strong focus on CUDA by researchers (who got free GPUs from Nvidia), and the lack of strong messaging by AMD, and you get a lot of "I didn't know"s.

It seems to be a large part around AMD's software offerings, and since a lot has changed this seems to be because of inefficient marketing. According to a message on a forum HIP is a tool to converts CUDA to OpenCL, as AMD can only run OpenCL. As nobody corrected the claim, it seems to be the common knowledge.

The news of "new CUDA version released!"

With CUDA 9 and 10 there was no "I can't wait to get my hands on it", except by starters. This was because hardly any new features

were introduced. So the question came to mind how this was observed. Did the developer have piece with this?

<https://twitter.com/StreamHPC/status/1060468918500909056>

13% Yes, vastly improved**40% Yes, but only a little****30% No, it's all AI now****17% CUDA is sorta finished****30 votes**

When it was asked as the question for AMD ("Did you see a lot of improvements?"), you see there is one yes-answer and 3 no-answers. The yes-answer I don't really understand - I think it refers to the hardware and the tensor-cores being programmable. Most chose for "only a little progress", which is a neutral answer. Then there is the complaint that Nvidia is shifting its focus to AI, then there is the "I'm ok with it".

The question is how GPU-developers deal with it. Many are very smart people who like to learn new things, so what would be the next objective to learn for them? Another thing is that now OpenMP and several other languages have improved GPU-support, which will solve the needs for a large group - especially in 3 to 4 years when these languages have progressed even further.

Then there is AMD with HIP, their CUDA clone, who is helped a lot with a non-moving target.

For those who know AMD HIP, seen the improvements?

We've worked with AMD a lot to make high-performance libraries for their GPUs, and have seen the advancements made in their new compilers and drivers. But how was it observed by others?

<https://twitter.com/StreamHPC/status/1060468255016542208>

50% Yes, (very) noticeable**11% No, still too alpha-stage****0% Hipification fails still****39% What is this AMD HIP?****28 votes**

Upcoming polls

We've got various new questions we would still like to ask you. We hope that you join in, so we can all get to know the industry a bit better poll-by-poll.