


## Exposing OpenCL on Android: Q&A with Tim Lewis of ZiiLabs



ZiiLabs has been offering an early access program for [OpenCL SDK](#) since last year. This program was very selective in choosing developers and little news has been put on their webpage. Now they are planning to make their Android NDK a standard component, it's a good time to ask them some questions. GPGPU-consultant Liad Weinberger of [Appilo](#) also added a few questions.

The Q&A has been with Tim Lewis, director Marketing and Partner Relations of ZiiLabs, who has taken the time to give some insights in what we can expect around accelerated computations on Android. ZiiLabs has been better known as 3DLabs and has reinvented itself in 2009 (you can read the full history [here](#)). Like other companies in the ARM-industry they mostly design chips and let other parties manufacture devices using their schematics, drivers and software. Now to the questions.

**Vincent Hindriksen:** You have had the early access program for a year now. Why did we hear so little?

**Tim Lewis:** Since the initial launch we have been working with a small group of partners via the Early Access Program to help us ensure that the technology we release to mainstream developers meets are quality standards. The partners have provided invaluable feedback and we are now looking forward to releasing it to a wider audience later this year.

**VH:** The last few months the number of searches on my site for ARM or Android and OpenCL has increased a lot. Do you see an increased interest too?

**TL:** Yes, we see a growing interest in finding ways to enhance the CPU performance of the ARM cores by taking advantage of the programmability and floating point performance of the media processor.

**VH:** Are drivers available for Android only, or also for Windows 8, Windows CE and other Linux flavours?

**TL:** We are focused on Android.

**VH:** Is the emulator using the PC's OpenCL-capabilities? And can you tell something about developing OpenCL-software for a ZiiLabs device?

**TL:** In terms of programming OpenCL for our processors we map physical processing elements to work items and allow the programmer to specify a work group size of up to 8 processing elements (the number of processing elements in a cluster). The runtime then maps the whole problem onto the array by iterating over the problem as many times as required to run the kernel on the specified problem size (total number of work-items). If the programmer specifies a work group size of 1 then the runtime maps 8 work-groups onto a physical cluster. If the programmer specifies a work group size of 8 then the runtime maps 1 work-group onto a physical cluster. And we can of course run multiple work-groups concurrently as we have more than one cluster, in the case of ZMS-20 that's 6 clusters and ZMS-40, 12 clusters.

The OpenCL cross compiler is supported on Linux based PC hosts, but does not currently support any native PC OpenCL implementation.

**VH:** As OpenCL is very low-level, how does it handle crashes? Do you have tips for developers?

**TL:** We have extended debug tools to help developers write and debug programs.

**VH:** Many ARM-chips have specialised silicon to do multimedia-computations like encoding/decoding video. Does OpenCL make use of this or only the GPU?

**TL:** The ZiiLABS ZMS processors use our general purpose StemCell Media Processing SIMD array to offload the ARM from all media intensive tasks such as video encode/decode and OpenGL ES. Traditionally we have used hand-written Microcode which has been optimised over the last 10 years to create these key software components. However for new components and to enable 3rd party developers we are increasingly turning to OpenCL to fully leverage the power of the StemCell array.

**VH:** Do you think OpenCL (or any comparable technology) will make specialised silicon replace with more processing-cores?

**TL:** As our architecture is based around a fully programmable array of floating point processing units, it should come as no surprise that I strongly agree with the statement that specialised silicon will be replaced by programmable cores. There will however always be a place for specialised, fixed function silicon or specific components that do not lend themselves to a SIMD approach. However, as mobile SOC's become more complex and are required to perform more PC-like tasks so the need for flexibility in terms of features and performance increase. The emergence of the VP8 codec is a good example. Fixed function devices have to go through a rev of silicon before support for VP8 can be added, whereas we could add fully optimised 1080p support very quickly by simply updating our 'codec' program. And when running any single function we can dedicate more of the available silicon to that task, which helps us achieve better peak performance. As we move forward we expect customers to be able to implement their own proprietary software components that leverage the performance of the media processor via OpenCL.

**VH:** Android has Renderscript compute as an alternative to OpenCL. How is your support for RenderScript and if so, does it work together with OpenCL?

**TL:** We are committed to 100% Android compatibility, so we support Renderscript as well as offering OpenCL.

**VH:** What do you think of the upcoming "battle" between RenderScript, CUDA and OpenCL?

**TL:** Developers will drive this and our goal is to put the core technologies into their hands so that they can make the right decision for them, given that we will be focusing on Renderscript and OpenCL.

**VH:** From your page the typical usage seems to be media-processing, but what typical applications did you have in mind? (Note: I indirectly ask for the strengths and weaknesses of your processor)

**TL:** If I had to pick one areas where we are seeing most interest it is in implementing proprietary image processing algorithms, including enhanced face-tracking and object recognition but the interest is much broader including enhanced audio processing and general floating point compute.

**VH:** You chose for the full profile and not the embedded profile. Is there a reason behind that?

**TL:** We come from a PC graphics background, so because are chips can support it, we chose to support full.

**VH:** You are claiming 26GFlops of compute power (A dual ARM A9 has 6-10 GFLOPS). Did you use LinPack or made your own test?

**TL:** We have our own internal tests and this figure of 26 for ZMS-20 is almost doubled with the ZMS-40.

**VH:** X86-GPUs work with a wave fronts of 64 workers (AMD) or warps of 32 workers (Nvidia). Does the ZMS-20 (48 cores) work with a comparable concept?

**TL:** I refer you back to the answer to the question on developing OpenCL programs.

**VH:** Extensions like OpenGL memory-sharing seems obvious, but which extensions are actually available?

**TL:** We are still finalising the extensions we will support but OpenGL ES and direct access to Camera sensor data seem of particular interest to our partners.

**Liad Weinberger:** What are the typical (idle/norm/stress) power requirements of each of your OpenCL supporting chips.

**TL:** We target the mobile space, so typically we are under 1W of total system power and looking to 10+ hours of HD video playback from a tablet sized battery.

**VH:** As the Watts per GFLOP is lower when using GPUs, the battery can be spared. Do you have done any tests on battery-endurance when using OpenCL and when using the processor only?

**TL:** We don't have any specific results we can share, but certainly media intensive tasks on using the array extend the battery life compared to running the ARM.

**VH:** Being on batteries, is it better to use all cores to the fullest or have some focus on avoiding peak-consumption? In other words: how could the batteries be spared while getting all the work done?

**TL:** We support a variety of power saving depending on the task and workload, including reducing the number of clusters being used or simply reducing the overall array speed.

**VH:** What can you say more about the market position of your products?

**TL:** We are a strong player across a wide range of markets from surveillance cameras and portable media players to tablets and embedded platforms. Over the coming months I expect to see our customers making some interesting announcements in both the embedded and tablet space.

**VH:** How long do you expect to have the advantage of being first in offering OpenCL for embedded devices?

**TL:** We expect to keep pushing the limits of what can be achieved on a mobile processor. Our recently announced ZMS-40 doubles the array size to 96 floating point cores, which, through OpenCL, puts an enormous amount of compute power into the hands of developers targeting the tablet and embedded markets.

**VH:** Which device do you suggest for developers who want to start developing for OpenCL on Android?

**TL:** As the core technology provider we have to synchronise our software releases with those of our platform customers. I would expect our customers to be announcing suitable platforms toward the 4th quarter of the year.

**LW:** Could you give information about the amount of OpenCL-capable ZMS chips on the market? Or next year?

**TL:** We do not disclose market sensitive data about volumes. What I can say is that we continue to see an amazing increase in the take-up of our processors as user requirements for low-energy, high-performance processors across a broad range of markets including high-volume consumer segments such as tablets, connected TV and the remote medical sectors.

**LW:** Is such board available if one cannot commit to an amount of units end up purchasing?

**TL:** We expect to support platforms that are available to end-users and so have no volume requirement.

**VH:** Is there something you want to share with us, which not got mentioned?

**TL:** Not at the moment, other than to thank you for the questions and for developers to continue to ?watch this space? as we prepare to roll out the next phase of our OpenCL program.

**VH:** Thank you very much for your time. Can readers ask questions in the comments, so you can answer them?

**TL:** Yes. Please email me directly at [tim\[dot\]lewis\[at\]ziiilabs\[dot\]com](mailto:tim[dot]lewis[at]ziiilabs[dot]com).

We left out questions how the product was compared to Nvidia Tegra, ImTech PowerVR and AMD G-T40N, as ZiiLabs does not comment on competitors' products. When the products of ZiiLab's partners hit the market end of the year, we will be able to tell more. Be the first to know when that happens: [ZiiLabs@twitter](#) and [StreamHPC@twitter](#).

If you want to see more, watch the below Youtube where Tim Lewis demoes what is possible with current generations of ZiiLabs-hardare.

YouTube Video: [YouTube.com/watch?v=\\_ljSkCgEV-U](https://www.youtube.com/watch?v=_ljSkCgEV-U)