

Feedback

Qualcomm's GPU History

Though Qualcomm sold its own handset business to Kyocera in 1999, many of today's most popular mobile devices are powered by Qualcomm's Snapdragon mobile chipsets with integrated CPU, GPU, DSP, multimedia CODECs, power management, baseband logic and more. In fact the typical "chipset" from Qualcomm encompasses up to 20 different chips of different functions besides just the main application processor. If you are the owner of a Galaxy Note 4, Motorola Droid Turbo, Nexus 6, or Samsung Galaxy S5, then you are most likely a user of one of Qualcomm's Snapdragon chipsets.



RYZEN | GIGABYTE

Shop Now ►



EVGA
GeForce® GTX™ 1080 Ti
 w/ **iCX Technology***
 *FTW3 and SF2 models

- [PC Perspective Live! Google Calendar RSS](#)

Support on patreon

amazon
SHOP.
CONNECT.
ENJOY.

You May Like These




Anker PowerCore Speed 10000 QC, Qualcomm... \$35.99 ~~\$40.00~~ (538)	Nokia Lumia 920 32GB Unlocked GSM Wi... \$98.00 (61)
---	---

Ads by Amazon

Latest News Latest Videos

Build your own watercooling loop with the EKWB L360 2.0 Kit



Qualcomm's GPU History

Before 2006, the mobile GPU as we know it today was largely unnecessary. Feature phones and "dumb" phones were still the large majority of the market with smartphones and mobile tablets still in the early stages of development. At this point all the visual data being presented on the screen, whether on a small monochrome screen or with the color of a PDA, was being drawn through a software renderer running on traditional CPU cores.


But by 2007, the first fixed-function, OpenGL ES 1.0 class of GPUs started shipping in mobile devices. These dedicated graphics processors were originally focused on drawing and updating the user interface on smartphones and personal data devices. Eventually these graphics units were used for what would be considered the most basic gaming tasks.

[Continue reading Qualcomm History and its GPU \(R\)evolution.](#)

The mobile GPU market grew rapidly from 2006 into 2009. Qualcomm admits that these were the most difficult years of development, caused by the massive growth of the market. Though the company had talented engineers that were working on GPU technology, the speed of the market shift forced Qualcomm to look for help outside the company, to ATI (now AMD). Together they developed the Adreno 130 – an upgrade from Qualcomm's own in-house designed Adreno 120 GPU. The partnership expanded when Qualcomm licensed a GPU from ATI, called it Adreno 200, and later made some upgrades to create the Adreno 205. Eventually Qualcomm would purchase the handheld graphics division and its "Imageon" graphics technology. \$65 million dollars later, the deal was complete in early 2009.

Mark Papermaster on AMD's tiny things
 AMD Teases Ryzen Threadripper Packaging, Lisa Su for Scale
 CRYORIG's A40; two heads are better than one?
 Won't someone think of the children? Corsair will!
 PCPer Mailbag #1 - 7/20/2017
 Devil's Ivy, a voyeurs dream come true
 Podcast #459 - Threadripper Pricing, Liquid Cooled VEGA, Intel Rumors, and more!






During this complex but vitally important business division transition, the world of GPU technology was not standing still. In 2006 fixed function hardware and a new OpenGL ES mobile-specific API was being built. By 2008, fixed hardware had moved to the side in favor of programmable shaders, allowing for a more flexible environment. User interface was still the primary usage model for GPUs including the first GPU-accelerated composition in the form of the "Android Surface Flinger" but simple 3D games were starting to pop up too.

NVIDIA, a company built on GPU technology, has marketed and promoted its own Tegra processors on the stance that mobile GPU horsepower is critical and the company's expertise from the desktop markets will trickle down into the ultra-low-power fields. Only recently though has the GPU really been able to take advantage of the compute tasks that are executed on smartphones. The world of the mobile GPU is now starting to come into its own, proving and showcasing the importance of this particular portion of a typical SoC.

Bringing Modern Designs to Mobile

The next era of GPUs started in roughly 2010 and ran through 2012 and was even more disruptive than the previous. The use cases for GPUs on mobile devices was snowballing, starting with some major game engine developers outwardly discussing bringing console level gaming to mobile platforms and devices. GPU acceleration in the world of HTML5 and more advanced multimedia compositions required more processing power to support multi-camera configurations, overlays, windows, and visual effects. Everything in smartphones was happening *concurrently*, invoking engineering challenges that were greater in some ways than those that were faced in desktop PCs, particularly when considering the far more stringent battery and thermal constraints of mobile devices. Meanwhile, GPGPU (general purpose GPU) workloads were realizing their potential in the consumer desktop computing space; users, software developers and OEMs saw the potential benefits of highly parallel computing in a low power form factor like phones and tablets.



Interestingly, Qualcomm and others stated that during this time that the importance of graphics heavy benchmarks should not be overlooked. While very often OEMs and SoC designers will lament the unfair or unreasonable impact that benchmarks can have on sales, mobile processor companies generally believed that the improved benchmarks from professional graphics benchmark publishers like Kishonti, Rightware and others forced mobile OEMs to focus on application performance when selecting processors for mobile devices, rather than just making decisions based on theoretical hardware specifications. All benchmarks should be used and interpreted with care, since no single

benchmark can approximate the relative device performance for every conceivable workload. But it is likely that without these applications pushing hardware vendors to improve their frame rates for heavier game workloads, mobile devices would be much less capable today.

Compute for mobile devices first emerged in 2013



■ **Moto X "Windy Day" Application**

- Adreno 320 subdivides surfaces with OpenCL accelerated Pixar "OpenSubDiv" APIs
- Concurrently renders the resulting 3D graphics (first mobile CL/GL Interoperability)



■ **Android 4.4 Gallery Application**

- Utilizes Qualcomm's new GPGPU accelerated RenderScript implementation
- Qualcomm was first to market with GPU accelerated Renderscript, with Adreno 330 in Snapdragon 800 for Nexus 5


Other outside factors continued to push the importance on the GPU forward. Screen resolutions were increasing to HD and beyond, increasing the pixel processing power necessary for smooth and fluid motion dramatically. Feature phones were slowly fading away during the 2012 timeframe, forcing SoC developers like Qualcomm to integrate GPUs not just in high end processors, but into the lower-end markets as well. For this to be successful, improved power efficiency was not just desired, it was necessary for basic functionality.

Qualcomm started developing new programmable OpenGL ES 2.0 capable GPUs and its development team tripled in size in the span of only a couple of years. The Adreno A225 was perhaps the best example of a GPU built to address this rapidly changing market. It added support for the latest API specifications including DirectX 9 and OpenGL ES 2.0 and was one of the most power efficient GPUs in the mobile space. Built into the Snapdragon S4 and S4 Plus SoCs, A225 powered some of the world's most popular devices including the Nokia Lumia 1020, the HTC One X, Droid Razer M and the Galaxy S3 to name just a few. Qualcomm's dominance in a market that only six years ago hadn't existed was taking shape.

The next-generation of Adreno architecture, the Adreno 3x series (abbreviated A3x), actually has its beginnings in development prior to the acquisition of AMD/ATI Imageon, but was also heavily influenced by a new GPU architecture codenamed "QShader". The result was an OpenGL ES 3.0 capable GPU that transitioned away from A2x's VLIW shader architecture, to a much more flexible scalar-based one. It was also designed with GPGPU computing purposes in mind, and was another success for Qualcomm, since it was a GPU architecture that scaled particularly well from low to high tier, finding a home in various Snapdragon 200, 400 and 800 parts.

Despite the dominance that Qualcomm held on the mobile processor and mobile GPU market as late as 2012, there were outside forces that began to put pressure on the company to increase its development resources once again. Newcomer to the mobile processor space, but marketing master NVIDIA, announced the Tegra processor - a mobile SoC with a focus on the GPU. Though the company had very little market share and experience with anything other than large, power hungry graphics chips used in laptops, desktops and workstations; the promise of a mobile chip and GPU built by a company with such a pedigree was exciting with the media and some OEMs taking note.

Another company would also throw its hat into the ring, one with significantly more potential impact, though equally little mobile experience. By 2010, Intel was beginning to see the writing on the wall, which was that the mobile space, including smartphones and low power tablets, were the next frontier of computing. Taking a completely different approach than every other company competing for these segments, Intel bought Infineon's wireless unit and subsequently would attempt to bring x86 (rather than ARM) into the fold as well as a unique GPU implementation. Intel still has yet to truly deliver on its claims to enter and make an impact on the mobile market, but a computing giant this size should not be overlooked. Qualcomm needs to continue to push innovation forward if they are to maintain the lead from this new player.



A collection of GPU technology demos running

Work on GPU hardware continued, but Qualcomm also took note of some of the advantages that this new competition had over its current support system. Software development teams were ramped up and an outreach program for gaming developers implemented, including specifically the developers working on game engines like Unity and Unreal Engine. The sales pitch from this new group was easy – game developers want their games to work on the widest array of devices to increase sales potential and Qualcomm's Adreno graphics was the most popular GPU in the mobile space by a wide margin. Work with them, optimize their engine for Adreno and Snapdragon, and enjoy the instant benefits where it matters most – the wallet. So far it has worked and partnerships with UE and Unity were forged, among others. Qualcomm also brought in its own team of game developers to effectively start an in-house engine aimed at offering support to other teams, and to work on new visual effects libraries to share with the community.

[The Present and Future of Qualcomm GPUs >](#)

Review Index:



Comments

[Post a comment](#)

June 22, 2015 | 11:45 AM - Posted by Heavy (not verified)

great article

[reply](#)

June 22, 2015 | 12:02 PM - Posted by Snake Pliskin (not verified)

Can they run Crysis?

[reply](#)

June 22, 2015 | 02:44 PM - Posted by BillDStrong

If you have a source license for the CryTek Engine, and all the assets for Crysis, then yes, yes it can, now that CryTek Engine supports Linux. Not to mention Windows 10 support for different platforms.

[reply](#)

June 22, 2015 | 12:08 PM - Posted by Anonymous (not verified)

I'd like the manufactures of mobile GPUs to provide block diagrams of their GPU products including any optional components that OEMs may license and use, like dedicated Ray tracing, decoder logic units, etc. And for once I'd like reviews of technologies to be less testimonials of the product and more of a comparison and contrast with the competing products, including the complete block diagrams of the competing producers' GPU products. I'd also like more generic naming used for all these specialized GPU units that are given Trade Names/Brands to see if the functional units have comparable functionality across all brands of mobile GPUs.

If at all possible utilize the generic computing sciences name for the hardware functionality instead of the manufacture's trade/marketing terminology or put an annotated parenthetical reference generic name/term next to the marketing brand/term, for example SMT for Hyper-Treading, SMT(Simultaneous multithreading) is the proper generic terminology with which to compare CPUs with SMT ability when discussing/comparing CPUs, and GPUs are even more full of these trade names in an attempt by the devices manufacturer to differentiate, and sometime obfuscate, their products from the competition and confuse the consumer.

There is too much marketing speak and to little educating in most online technology websites, except for the professional trade journals which are behind pay walls, and most professional trade journals maintain extensive dictionaries of computing terminology as well as disambiguation entries of trade/marketing terminology translated to generic computing science terminology, to allow readers to properly compare different hardware systems between different manufactures.

The Mobile GPU makers/licensors are not providing sufficient data sheets, and diagrams for consumers to make an educated decision on just which of the mobile GPU, usually integrated in an SOC, has the largest feature set, that and there is little technical information to be had online, except the behind the pay-wall variety, or the occasional Hot Chips Symposium, where the professional engineers utilize the proper computing science terminology, although the marketing monkeys are trying to ruin hot chips by forcing their engineers to use the marketing terms.

There really needs to be a good online reference for computing science terminology, as well as proper technical documentation, Wikipedia is piss poor with their "technical" information on the various GPU/CPU/SOC processors, and processor pin-outs, as well as block diagrams, including a dedicated disambiguating section that translates marketing/trade terminology into proper generic computing science terminology.

[reply](#)

June 22, 2015 | 12:18 PM - Posted by craig lancaster (not verified)

No pleasing some people, maybe you will be better served if you use other sites that to us mere mortals are difficult to comprehend. NO offense to the article writer but at least i can understand most of what he says.

[reply](#)

June 22, 2015 | 04:37 PM - Posted by Anonymous (not verified)

The post is not specifically directed at PcPer, which is one of the better benchmarking/computer news sites outside of a pay-wall, but it would not hurt for these technical websites to pool their resources and get a service to properly document the complete technical details on the products that they review, including some educational articles once in a while to help the non technical readers better understand the technology nomenclature and definitions, and educate their readers to able to differentiate the marketing brand obfuscation and the actual computing sciences terminology for the various CPU/GPU/other hardware than needs to be compared and contrasted. This may be a consumer products oriented sight, but the technology is very complex and with the marketing and MBA types in charge for the most part of some very large technology companies, the tendency is more to confuse rather than inform on the part of said companies.

A lot of the high tech electronic devices have been commoditized and are marketed the same way bulk laundry soap is marketed, but in order for the consumers/readers to have any possibilities of making an informed decision proper education and review methods need to be used. This includes defining Acronyms, and disambiguating marketing "technical" terms/branding with the proper computing sciences terminology, so readers can properly research the specific technologies on CPU/GPU/SOC/other computing systems.

Too many technology websites are becoming little more than extensions of marketing departments, and click bait journalism is rampant on more than just a small number of websites.

The total number of sponsored article content has sky rocketed, along with reviews that only talk about a single manufacture's products without any direct feature for feature comparisons of any competitor's equivalent product technology. It's become more difficult to obtain the proper data sheets and specifications on especially the mobile SOC GPUs and their specific technologies compared to the more thorough analysis on the top end gaming GPU SKUs.

[reply](#)

June 22, 2015 | 12:23 PM - Posted by willmore (not verified)

Well said. PCPer is for users, not for marketing.

Too many marketing terms and I start looking for the "paid promotional article" tag hidden somewhere.

[reply](#)

June 22, 2015 | 02:56 PM - Posted by Anonymous (not verified)

This is why benchmarks exist. It would be nearly impossible, even for someone knowledgeable in the field, to predict how all of these features will work together to enhance the user experience. Going into low-level architectural detail isn't going to be useful to most users.

[reply](#)

June 22, 2015 | 05:15 PM - Posted by Anonymous (not verified)

Software benchmarks are the most gamed to deceive statistic where computer hardware is concerned, and nothing replaces a good thoroughly documented hardware data sheet, that gives definite SP/ROP/etc. counts on GPU hardware, as well as the proper block diagrams describing the complete workings of a GPU/CPU/Other processor, or at least links to data sheets/technical manuals that do have the most complete information, without revealing any trade secrets.

The low level hardware is the most important to see in order to at least have an estimation of what the device has compared to its competition, everybody knows that whatever SOC/GPU/CPU may have differing characteristics when placed into an OEM's final product, especially in mobile/laptops products where the device's thermal settings may be lowered to run in the mobile form factor.

I'm a big fan of requiring mobile CPU/SOC OEMs and manufactures to provide some testing rigs/mules for reviews of their SOC's, custom and otherwise, so that the CPU/GPUs and SOC's themselves can be properly put through their paces, outside of any eventual devices the CPUs/GPUs/SOCs may be placed in, if it works for the big gaming rigs, it should work for the smartphone/tablet SOC's, and believe me there are testing rigs/testing mules that can do the job. The industry uses them for internal device development before the phone/tablet product designs are finalized.

Even among phone/tablet systems the mainboards are fairly standardized, maybe not the shapes and dimensions of the mainboard PCBs, but the platform controller chips, and chipsets are fairly standardized, and the testing rigs/testing mules are used to put the mobile SOC's through their paces, every bit as thorough as the gaming rigs are tested and even more thoroughly with electrical usage and such.

So benchmarks: AnTuTu, and such gamed by device manufactures makes me mistrust single benchmarks, it's already to a point where device manufactures, and SOC manufactures should be required to send their devices to independent third party testing labs to have the SOC's tested individually and also in the respective devices and the information made public, in or for the devices to be approved for sale. The FCC does this to a degree, but the information is difficult to find, and the Department of Energy as well as the EPA does testing, but some form of impartial standardized testing by an outside lab is in order for the entire mobile devices industry, as well as the PC/Laptop industry.

[reply](#)

June 22, 2015 | 06:47 PM - Posted by Anonymous (not verified)

I wouldn't trust synthetic benchmarks, but what better gauge of performance can you get than actually running the applications people are interested in?

[reply](#)

June 22, 2015 | 06:59 PM - Posted by Anonymous (not verified)

"Software benchmarks are the most gamed to deceive statistic where computer hardware is concerned, and nothing replaces a good thoroughly documented hardware data sheet, that gives definite SP/ROP/etc. counts on GPU hardware, as well as the proper block diagrams describing the complete workings of a GPU/CPU/Other processor, or at least links to data sheets/technical manuals that do have the most complete information, without revealing any trade secrets."

The tech specs of these devices obviously make a difference, but they are actually often not useful for comparisons. There can be other bottlenecks in the system. Testing the SOC independently of the device it goes into is also not that useful since the final device will have its own specific thermal characteristics and possibly other bottlenecks. A lot depends on what screen the SOC is paired with. Apple does not disclose much info on their SOC's, but it

doesn't really matter since we can run test and see how it really performs in the applications we are interested in.

[reply](#)

June 22, 2015 | 12:12 PM - Posted by Anonymous (not verified)

good read :)

[reply](#)

June 22, 2015 | 12:21 PM - Posted by willmore (not verified)

Don't forget that they're also the company that pushed CDMA in North America because they had it locked up with patents--even though none of the rest of the world used it because it was rubbish. This left NA and good chunks of Central and Southern America behind the world in cellular development.

[reply](#)

June 22, 2015 | 02:13 PM - Posted by Rickard Eneqvist (not verified)

CDMA isnt rubbish, it has a lot of beneficial features that the GSM standard did now have. Dont blame the whole network intercommunication problems on the technology, that's a basic issue with creating standards in the USA/Rest of the world scenario.

Remember that GSM went to WCDMA which has its roots in the CDMA technology.

[reply](#)

June 22, 2015 | 02:02 PM - Posted by Anonymous (not verified)

Well, this is certainly very enthusiastic.

[reply](#)

June 22, 2015 | 03:00 PM - Posted by Anonymous (not verified)

Is running games at 4k on a mobile device really necessary? High pixel density is great for text, but for images or video on such a small screen, I doubt it would be noticeable. For a small screen I wonder if it would really be noticeable if it was rendering at 4k or just rendering at 1080 and scaling it to 4k.

[reply](#)

June 24, 2015 | 01:07 AM - Posted by liam (not verified)

Do you really have to game on your mobile?

[reply](#)

June 23, 2015 | 12:58 AM - Posted by nobodyspecial (not verified)

You might have wanted to mention they all may be banned soon, as they've been stealing NV tech for years ;) The markman hearing showed the judge favoring 6 out of 7 nvidia patents. That's a pretty clear sign he thinks Nvidia has a great case against samsung and Qcom and likely (no matter what ITC people think) that a 12 person jury would see the same, especially when they will be considering a 23-30Billion profit machine in samsung (6B for Qcom) stealing from an American company with ~600mil profits. It is clear mobile is now doing stuff that desktop has been doing for a decade, so someone is stealing without paying the patent owners (NV and likely even AMD in some ways) who blazed the trail 10-15 years ago. Note the patents NV sued over are from 1999-2001 when the tech that is being used on mobile was created. Patents for this stuff is created long before the products hit (IE we won't see what NV's working on in the last 5yrs until Volta and beyond probably), but unlike a patent troll NV/AMD are actually USING the tech in their products for a decade on desktops.

Good luck explaining to a jury how you're doing it different than the people who've been doing it for 20yrs. You will be trampling on NV/AMD (maybe some intel) patents for decades to come most likely also unless they come up with some magically radically different way to get pixels onto a monitor. There is a very good reason Anandtech called it the wild west of patent infringement in mobile. It's time to pay up for all these leechers. An NV win might actually lead to a case for AMD which they could really use to help fund R&D for the future (which has gone down for the last 4yrs, while losing \$6Billion in 12 years). They are now playing PC/console ports on mobile directly, so I can't wait to hear how they'll say they do it different while running the same exact games that came on consoles/PC's over the last decade+. It's also worthy noting NV has been trying to get them to pay for 2+ years (IE, willful infringement after being told to stop or pay up!). This is worse than the INTEL case which ended in 1.5B to Nvidia for the same stuff once the chipset agreement was broken (intel wasn't WILLFUL, it only happened due to breaking an agreement).

[reply](#)

June 24, 2015 | 12:59 AM - Posted by liam (not verified)

The worst that happens is that Qualcomm takes a bit from its warchest and buys nvidia.

[reply](#)

June 24, 2015 | 12:56 AM - Posted by liam (not verified)

Adreno is still vliw4.

<https://github.com/freedreno/envytools/blob/master/rnndb/adreno/a4xx.xml>[reply](#)

June 29, 2015 | 09:55 AM - Posted by chizow (not verified)

Interesting piece Ryan, I think it does show Qualcomm really "caught up" with their graphics performance and design right around the time they collaborated with and acquired ATI's mobile division.

It would've been a nice addendum to see the nature of Qualcomm's deal with ATI at the time. Clearly they were licensing AMD's Imageon IP in the earlier collaborations, but it doesn't look like any of that IP was transferred or continually licensed to Qualcomm when they acquired Imageon. That certainly makes sense, as 65M seemed like a song at the time for an entire mobile graphics division (and still does), but its more and more obvious no IP changed hands, nor did any perpetual IP license from AMD. It was really just a transfer of staff and working knowledge.

That really would be the only way Nvidia would have gone this far in their litigation against Qualcomm; if any of AMD's IP transferred or was still being licensed, that would have put an end to Nvidia's litigation full stop.

[reply](#)**Post new comment**

Your name: *

E-mail: *

Comment: *

- Lines and paragraphs break automatically.
 - Allowed HTML tags: <a> <code> <dl> <dt> <dd> <blockquote> <p>

- Web page addresses and e-mail addresses turn into links automatically.

[More information about formatting options](#)By submitting this form, you accept the [Mollom privacy policy](#).[Home](#) | [Reviews](#) | [News](#) | [Forums](#)PC Perspective [About](#) | [Twitter](#) | [RSS Feed](#) | [Reviews feed](#)

Copyright © PC Perspective 2000–2017