AMD to launch new 7nm Navi GPU, Rome CPU in 3rd quarter

Enlarge / AMD CEO Lisa Su, holding a Rome processor. The large chip in the middle is the 14nm I/O chip; around it are pairs of 7nm chiplets containing the CPU cores.

reader comments

94 with 55 posters participating

Share this story

In its earnings call, AMD offered a little more detail about the launch of its next-generation processors, built using the Zen 2 architecture and TSMC's 7nm manufacturing process, and new GPU architecture, Navi, again built on 7nm. Server-oriented EPYC-branded chips (codenamed Rome) should be shipping to customers in the third quarter of this year, and so too will Navi-based video cards.

In November last year, AMD outlined the details of the Zen 2 design. It makes a number of architectural improvements to shore up some of Zen's weaker areas (for example, it now has native 256-bit floating point units to handle AVX2 instructions; the original Zen only had 128-bit units, so it had to split AVX2 workloads up into pieces). But perhaps more significant is the new approach to building the processors. Zen used modules of four cores (handling eight threads), with two such modules per chip. Mainstream Ryzen processors used one chip; the enthusiast Threadripper range used two chips (first generation) or four chips (second generation), and the server-oriented Epyc range used four chips. Each die is a full processor, containing the cores, cache, memory controllers,

PCIe and Infinity Fabric connections for I/O, integrated SATA and USB controllers, and so on and so forth.

Zen 2 will continue to use multiple chips, but this time the chips will be more specialized. There will be 7nm chiplets, each containing CPU cores, cache, and Infinity Fabric links, and a 14nm I/O die, containing memory controllers, Infinity Fabric connections, and SATA and USB controllers. The 7nm parts should be able to achieve higher clock speeds and lower power consumption than their 14nm predecessors. The parts on the I/O die, however, generally don't benefit from higher clock speeds. In fact, they can't—PCIe, USB, SATA, and even memory, all need to run at predetermined speeds, because their performance is governed by the bus specification. The extra performance headroom that 7nm would offer is wasted. By keeping these parts on 14nm, AMD is likely able to cut costs (because well-established 14nm manufacturing should be cheaper than the newer, more advanced 7nm).

The Rome processor will have up to eight core chiplets, for a total of 64 cores and 128 threads, and it will support up to two sockets. These parts are already sampling and will be ready for a third-quarter launch. AMD didn't, however, make any mention of the mainstream 3000-series Ryzen chips. There is speculation that these will be announced at Computex later this month, but for now it seems that AMD's focus is on the more lucrative server market.

Also using 7nm is AMD's new Navi GPU architecture. AMD is already building 7nm GPUs, such as its Radeon VII using the Vega architecture. Unusually, the company said that the first Navi parts will be priced below the $699 Radeon VII. This suggests that rather than launching with a top-end GPU, AMD will instead start with a mid-range, mainstream part. AMD has disclosed very little about Navi, refusing to answer whether the chip will offer any kind of ray tracing acceleration comparable to that found in Nvidia's RTX parts. But with a launch next quarter, the company is going to have to spill the

beans sooner rather than later.

https://arstechnica.com/?p=1498591