

Volume XII, Number 40 – October 4, 1999

GigaPixel's 3D tiling engine – GigaMan lives

- 5 Gpixels/sec. fully lit and filtered
- all APIs and AGP
- proof of concept viewable now

Our pursuit of pixels this week took me to Silicon Valley to see GigaPixel's gate array run some real world applications (well, not really real world; they were of *Quake III*). GigaPixel, for those of you who failed to find out anything at Siggraph, is an IP company with a technology based on a rather robust tiling engine.

Tiling of course is not a novel idea (being brought into prominence by Microsoft's Talisman initiative in '96). Problems with previous tiling architecture attempts, however, included incompatibility with standard 3D APIs; polygon sorting was required; a limited number of polygons could be handled per tile or frame; there were transparency problems and difficulty with Z-buffer reads; and those annoying texture coherence problems kept popping up.

GigaPixel claims to have overcome all those problems and they have 48 patents in process to prove it. The thing that GigaPixel brags about the most, and with good reason, is their full-scene anti-aliasing, which they can provide at a resolution with no performance penalty, and no increase in memory bandwidth or size. George Haber, GigaPixel's President and CEO, is quick to point out that competitors rendering at higher resolution do not solve aliasing problems because motion artifacts are still present — that approach requires more memory and bandwidth, and some aliasing still exists. "It's extremely expensive in memory size and bandwidth," he said. "You know," he added, "aliasing is like a scratched record, it bothers you."

Haber points out that the typical 4X rendering and accumulation is a huge performance hit that carries a 4X memory requirement with a 4X to 5X memory bandwidth requirement, and therefore is only practical for low resolution. "With our approach," he points out, "there's no performance gain if [anti-aliasing is] turned off."

The other approach, edge-only anti-aliasing, produces low-quality polygon intersections that are aliased. He cites the compatibility problems of edge anti-aliasing, and points out that it may require polygon sorting or an increase in memory requirements. But most damning of all is that applications need to be rewritten.

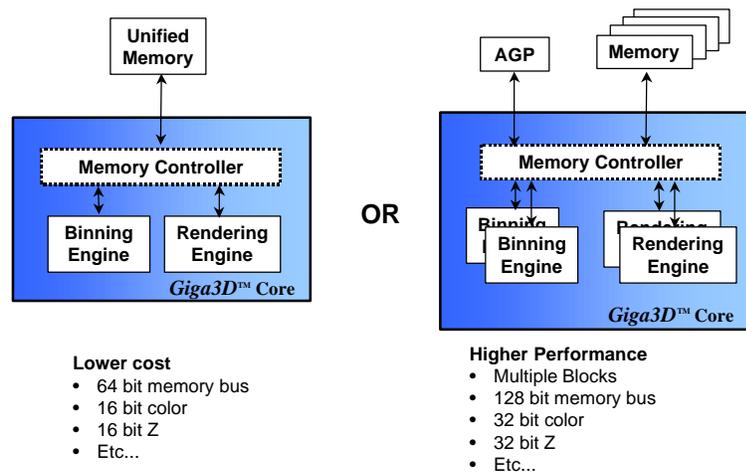
Well, he was preaching to the choir. I like tiling solutions — always have — they just make so much sense. Only visible pixels and polygons are setup and rendered, texture cache is extremely efficient, and

The Peddie Report is published weekly by Jon Peddie Associates, a technical marketing and management consulting firm for the computer industry at 4 Saint Gabrielle Court, Tiburon, CA 94920, ☎+1.415.435.1775, 📠+1.415.435.1599, 🌐 <http://www.jpa.com>, e-mail info@jpa.com, all rights reserved. Copyright ©1999 by Jon Peddie Associates: Jon Peddie, Publisher; Kathleen Maher; Editor-in-Chief, Bob McQuillan; Executive Editor, Analyst; Cindy Krause, Editorial Assistant; Tommy McClain, Online Editor; Bob Raikes, Displays Editor Guy Wright, Contributing Editor; Christine Arrington, Contributing Editor. **No part of this document may be used or reproduced by any party, government or person, in any form** or by any means, mechanical or electronic, without prior written permission of Jon Peddie Associates. The material in this report is not guaranteed as to its accuracy. It represents our interpretation of information generally available to the public and does not contain any information given to us in confidence by our clients. This report contains a "review" of various products. It is not an endorsement or attempt to sell any products. Under the rules of the "Fair Use Doctrine," JPA assumes no responsibility for the correct or incorrect usage of any trade marks or service marks.

there's higher texture coherence. With a tiling engine you can prefetch textures through AGP or UMA (and yes, of course you can with other solutions, too), all intermediate data is stored locally, and if designed properly the memory system uses highly efficient bursts which means no more read-modify-write issues.

So I asked George what makes it work so well in his chip. He showed me the following diagram.

Scalable & Flexible Architecture



Source: GigaPixel

GigaPixel's Giga3D core block diagram

Pretty boring, isn't it? I asked George what's in the rendering engine and with that devilish smile of his, he said, "It's so simple and beautiful; if I tell you, you'll be able to implement it."

So we went and looked at some demos, and of course that meant looking at a game. And that meant looking at *Quake III*. And it looked good, at 1024 x 768 x 32 @ 80 Hz with full-scene anti-aliasing and frame rates approaching 200 fps in some scenes. They had a G400 running next to it, same scenes, and of course the G400's scenes were riddled with edge jaggies.

I asked if the Giga3D core could do ray-casting like Imagination Technologies' PowerVR. "No," Henry Choy, GigaPixel's Director of Marketing, answered, "we don't do ray casting like the PowerVR. We can support things like bump mapping and shadows like other traditional 3D architectures. Our architecture has a traditional rendering pipeline after the visibility sub-system. The biggest difference is all the manipulation of pixels occurs on chip (lower bandwidth). There is no need for going off chip to write out intermediate results."

So I asked if they could do bump or environmental mapping in a single pass, and they said no. "But," George quickly added, "even though we take two passes, our system is so fast it doesn't

matter.” George had obviously been stung by that before and was probably thinking what I was thinking — that the PowerVR has a special hardware accelerator in it to handle single-pass bump.

“Something we can do that they [PowerVR] can’t, however, is provide Z data at any time, even within a frame once rendering has been started.” He then cited the corona effect around lights in *Quake*. “Without Z you can’t render that effect,” he said. I nodded my head appreciatively, and made a note.

The next day I called Imagination Technologies (formerly VideoLogic) and asked them to respond to the point. Trevor Wing told me (and later John Metcalf did also) that the PowerVR250 does not provide Z information during a frame. However, the next generation (series 3) will have that capability. “There’s nothing inherent in the architecture to stop it,” Metcalf added.

“However,” he added, “I don’t understand why they cite corona around light sources as a specific problem — we can certainly do those without access to Z information during a frame, and do so.”

And so the not-so-great corona war ended with a snooze.

But what about other features in the Giga3D core? Well it does bilinear, trilinear, and anisotropic filtering, can generate up to 5 Gpixels/sec., 32-bits/pixel fully turned on, at any depth complexity, and almost any resolution, and it’s scalable. In the current development stage the engine is generating 32 x 32 pixel tiles but they can be any size and don’t have to be symmetric. The memory interface is the main I/O port and other sources (i.e., 2D, video) can feed the device via it. It’s a multi-port controller like a crossbar but it’s not a crossbar, and it can be adapted to UMA or a frame buffer.

The proof of concept part I looked at was built at TSMC in 0.25-micron gate array and had about 600K gates. (www.gigapixel.com)

So GigaPixel has proven the part... something they had to do after the delays they suffered. But what’s that old expression... worth waiting for. Their customer (whom they still won’t name publicly) will probably introduce the device (in a larger system) in the second half of 2000. A lot of tweaking can happen between now and then so this design may surprise you even more.

GigaPixel was founded in late 1997 by three engineers from SGI’s Advanced Graphics Division and the ex-CEO of CompCore. Their plan was to be THE IP provider of 3D technology. The company is privately held with investment from US Venture Partners (Rounds A and B) and Synopsys, CMEA (Round B) and, as a result, it has a strong cash position. And it will need it to keep its 33 employees (29 engineers) fed and happy.

They targeted market leaders in existing and emerging market segments such as workstation, desktop / laptop / set-top / consumer boxes, and game consoles. Given that the Giga3D engine is completely API neutral and compatible, they have a good argument. The part should be economical to build, doesn’t require exotic memory, and has some real performance advantages. They’ll probably get some resistance from OEMs who think they know more about 3D than they

October 4, 1999

The Peddie Report

really do, and will object to the part's non-traditional approach. But like you've heard me say before, you've gotta kiss a lot of toads in this business... — JP