The Chip Insider

January 10, 2008 – The Cook's Tour: *IBM's R&D Partnering Model*.

The Cook's Tour: **IBM's** R&D Partnering Model — up close and personal. I traveled to New York, camera in hand, to get a deep immersion course in what they are doing at both **Albany Nanotech** and at IBM's Fishkill facility. It was intentionally a one-day whirlwind tour of both facilities, so that I could get a real feel for what they are accomplishing. The hard part about these sorts of things is that when operations are so far apart and with so many overlapping partnership models they can all become disjoint. Moreover, 'partnering' is one of the most abused words in this industry. At the low end of the spectrum, it is just a guise to get lower prices from suppliers or higher prices from customers. At the highest end lies IBM, where there have been significant and, as I'll show later, measurable benefits to all parties in this company's innovative approach.

IBM's R&D partnering model is the most significant new private R&D model to emerge since the last World War. It started in 1988, almost 20 years ago with **Siemens** and **Toshiba** in an effort to develop DRAMs. Since then it has blossomed as one of the most important ways to defray rapidly escalating R&D costs. Back in the late eighties, it cost about a billion dollars to develop a DRAM process. Since these were commodities anyway, it made sense to compress what would have been a $3B tab into one. The only real issue was could they really work together: three different companies and three dramatically different cultures. They did make it work and today, IBM's R&D partnering model has become the gold standard. They've managed to attract some of the best and even some of the strongest companies to participate in applied research that now stretches all the way to manufacturing with their Common Platform alliance. Its members include AMD, Chartered Semiconductor, Freescale, Infineon, Samsung, Sony, not to mention the State of New York. One fact speaks volumes about the breadth of this alliance: that the cumulative sum of capital expenditures for all the partners accounts for almost half of the world's total.

If you're lucky enough, you might just get the full guided tour of a research corridor that stretches from Albany to Fishkill. I won the fab rat lottery and managed to pack it all into a single day, which is detailed here.

Albany Nanotech:

Starting at Albany, this place has a very unique feel to it. It has the excitement of a college town and the vibrancy of a university. Yet it has the focus of a results oriented research team. I reckon this comes from the fact that working here as a student is like being on one of the best school ball teams: there are always scouts in the stands for whom your performance and their opinion can make or break your future career.

Inside Albany Nanotech: A worker lugs a 300mm FOUP to its next station. The yellow lighting tells you it is the lithography bay. Even though today's photo-resists are no longer sensitive to visible light, old traditions die hard anywhere where this is manufacturing intent. AlbnyNT_0703_011

Another reason for Albany's vibrancy is that wafer movement is not highly automated and it is a place where lots of people from many different organizations are always coming and going. Yet don't think that the place is not automated at all. Most of the research tools are built on real production platforms with full FOUP to FOUP automation. It's only in the original clean room where you see truly manual prototypes.

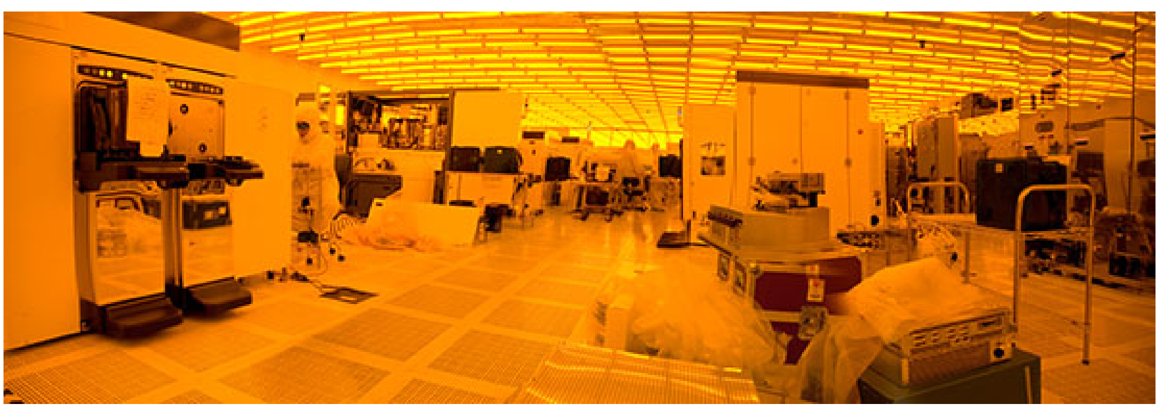

Inside Albany Nanotech: This panorama made from 3 images shows a modern research cleanroom. AlbnyNT_0703_008

The impression that this facility makes on you is one of its overwhelming size for a research cleanroom — much less, one that is a university cleanroom. The beauty of mini-environments is that lots of messy and dirty things can happen in a cleanroom without affecting yields. This messiness is an inherent part of every good research environment I've been in, as a good researcher learns as much, if not more, from the failures as he or she does from the successes. But unlike most semiconductor research environments, this one is spacious and well equipped. Like Formula 1 racing, you can't win without the most advanced machines, and a common mistake is to try to win on a Malibu Grand Prix budget. Not so here, as you can see.

Inside Albany Nanotech: Technology moves much faster than fabs age. So new tools always arrive for installation. It wasn't long after I left that the bagged tool to the left was up and running wafers.   AlbnyNT_0703_016

Another interesting thing is the degree to which the different suppliers work together. This is fairly unprecedented. I walked freely into and out of **Applied Materials'** research area and over to **TEL's**. I talked to the people in each and they were very open about how good this environment was for both. They do stay out of each other's areas on an honor system. Yet they are not shy about working together informally to solve a common integration problem that affects the customer or just to lend the occasional hand. It's amazingly collegial for two companies who compete fiercely against each other and yet can fiercely work together to help a common customer. I'd call it collegial competition . . . the bottom line is that the people here really work together for the common good. Adam Smith's invisible hand works here daily.

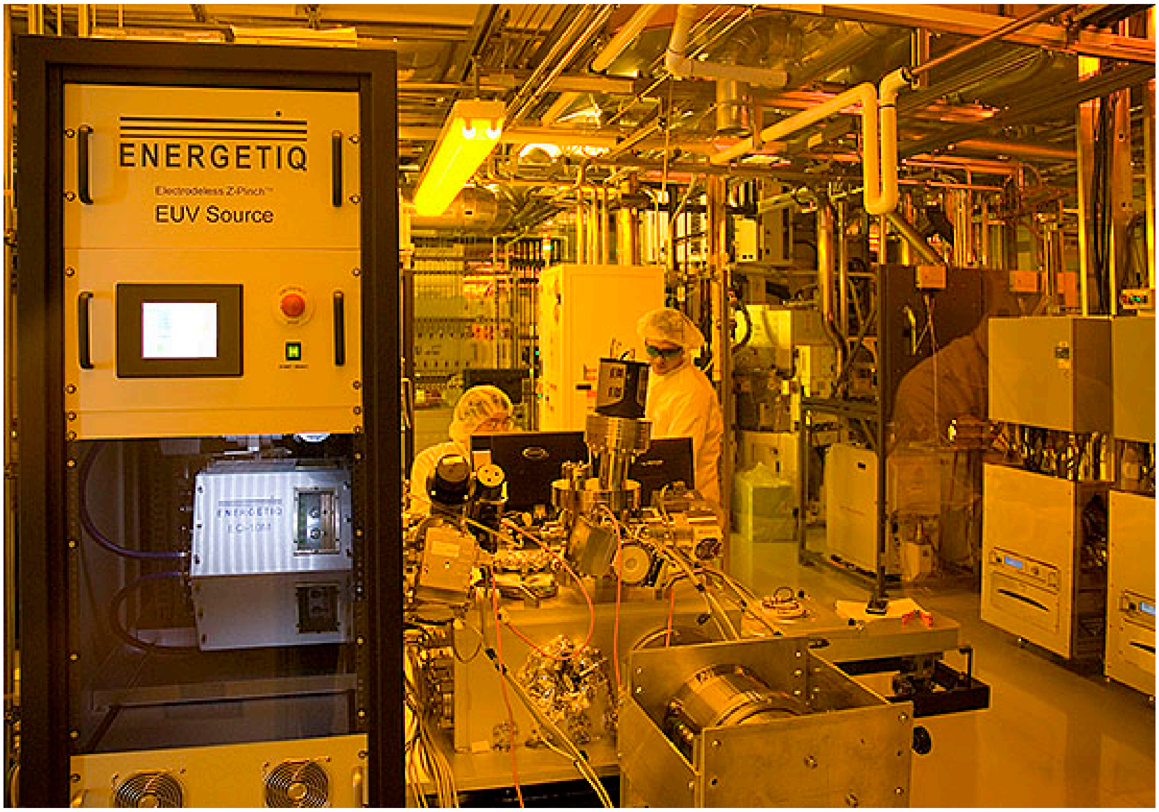

Inside Albany Nanotech's first cleanroom: here you see one of the world's first alpha EUV tools. AlbnyNT_0703_021

Inside the older part of the research facility, everything is much more crowded and looks more like most R&D environments that are not in a modern cleanroom. This is actually one of the first EUV exposure tools ever built and it is still being used. Like Eskimos, part of a good semiconductor research environment is the ability to extract every last bit of value from the carcasses of early generation tools.

IBM's 323 Fab:

Around 10 years ago I wrote that they were no longer practicing catch-and-release at IBM . . . they were killing the fish at Fishkill. IBM had gone through a rough spot in the early nineties as they moved from a company that only made chips for internal consumption to one that sold them to all takers. This meant they had to have real sales and marketing and they had to have cost- effective manufacturing. I wrote that comment, because I was both surprised and pleased to see them so focused and doing so well. I have been fortunate enough to see them continue to build on that excellence over the years, as you'll see in the following.

One of the most important aspects of the common platform is that processes and products are easily transferable from one company to another and across fabs. By doing this the partners can guarantee a virtual second source, which is essential to having the pricing control to get the value out of their efforts. It is also essential to keep IBM's customers happy with quick delivery. Thus, IBM's 323 fab is a fully functioning manufacturing line even though they also do development here. Doing both in one place ensures that processes transfer more easily from one location to another. It also insures that what's in development comes to the fab manufacturing-ready. At the same time, this fab has evolved quite a bit over recent years.

So, let's head over to the cleanroom and get suited up. First stop on any cleanroom tour is always mainstreet.

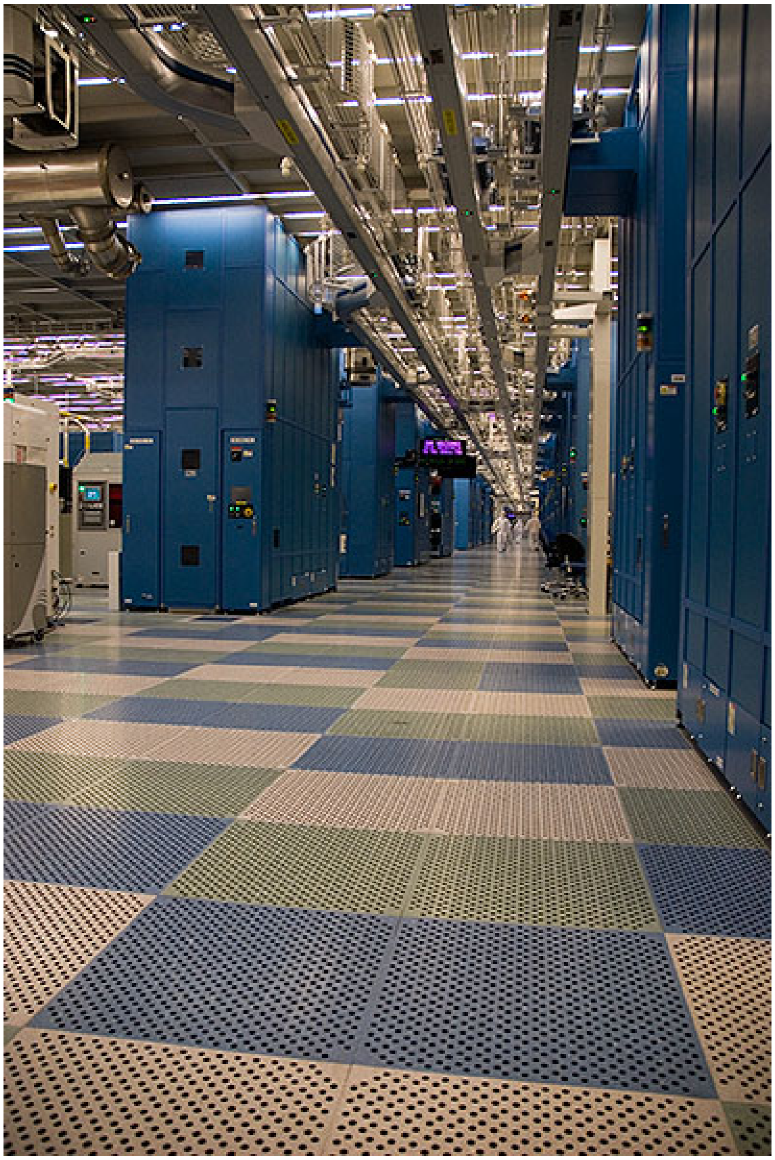

IBM's 323 facility: Mainstreet, Chip City, New York The big blue structures are stockers, which hold wafers in a clean environment while they wait for an available tool to process the next step. Note the purple plasma display screen in the center. One of the messages was welcoming me to the fab — thanks IBM, I never got such treatment. IBM323_0703_001

Mainstreet is the central corridor where all the product and people move to get to bays (where the actual tools that process the wafers are). Though 300mm fabs have gone back to the ballroom concept; automation and technical discipline still dictate a linear comb-like layout. IBM is very focused on automation of not just WIP flows, but also the flow and distribution of information. As one walks through they see information displays everywhere. The old ways of posting status printouts all over the walls are long gone. Engineers carry wireless notebooks that give them detailed access to real time information.

IBM's 323 facility: IBM323_0703_005

To understand how much product goes through this fab, all you have to do is look up. Multiple rows of monorail are there for trains carrying wafers to whisk by. It’s a real live miniature railroad. Also notice how spotlessly clean everything is. Special paints are used to repel rather than attract dust. Another big feature of this fab is that mini-environments allow it to use far less HEPA filtration and reduce air flows. This not only saves money, it also helps the environment by using less electricity.

IBM's 323 facility: this wafer sorter has direct access to the wafers in the stocker.   IBM323_0703_013

One very useful feature I was shown in this fab was this wafer sorter tied to stockers. Each sorter has direct access to the stocker, so engineers do not have to pull lots out to build special test runs. For example, suppose they want to run a chamber cleaning experiment and study degradation over, say 10K wafers. They don’t need to look at every wafer, just one per lot of 25. So this sorter can be programmed to automatically pull a wafer from every 25 wafers that comes off the same tool and chamber and then send them on to inspection. The engineer can fire-and-forget – coming back later to get results, not wasting time picking through a haystack of wafers. This makes them more productive in finding answers faster. It seems simple, but a lot of precious engineering time is wasted in fabs around the world as costly engineers wait for wafers to come through one at a time. The engineers love it, because they can get on with their analytical work rather than manage wafers.

IBM's 323 facility: It may look like everyone's gone to lunch, but in IBM's automated fab, you won't see operators – just freight cars of wafer filled FOUPs whizzing by along the ceiling.   IBM323_0703_021

What is so amazing about the level of automation here is that there are really no operators. Virtually everyone in the fab has a technical or maintenance job to perform. These are highly skilled people, which means that direct labor cost is not a significant factor between having a fab in New York or China. Another aspect of this is having tools and maintenance procedures that guarantee wafers are not sitting around waiting on broken equipment.

IBM's 323 facility: where all the engineers are above average and the light trees are all green. IBM323_0703_029

Here you see a bay of inspection tools on the border of the lithography bay. Note that the light stalks on every tool are green. That means that all are either running wafers or are waiting for them, like the tool on the right and the car coming to load it on the rail above. These tools run 24 hours a day and seven days a week at maximum speed. Your car won't do that. This is a very well run fab.

IBM's fab is also at the cutting edge of technology, as exemplified by the tools. Walk around and you'll see early generation tools everywhere. IBM was one of the first companies in the world to crack the defect density issue with immersion. Having this capability gives them higher yields and lower cost masks.

IBM's 323 facility: Too Cool for School Tool — the cutting edge of manufacturing lithography: Hyper-NA immersion
IBM323_0703_012

All in all, it's pretty much a mind-blower. But IBM is a master of blowing your mind on plant tours. My dad used to always come back dazzled thirty years or so ago. Of course, in IBM's favor is that they have been successfully adding partners and no one is dropping out like many other such ventures. On the other hand, maybe they've just done a great sales job on them as well. The real question is: are IBM and its partners getting financial traction from these efforts? After all, this looks very expensive. Do they really save any money?

To answer that question I began to dig into the revenue and R&D spending data of the partners. I split the data for each company into pre- and post-partnering sets, focusing on companies that joined after 2003. As it turned out there was a significant drop in R&D spending as a percent of sales for most companies. I then modeled their pre-partnering spending data against industry averages to come up with an estimate of what they would have likely spent had they not partnered. I also did a similar trend modeling for the existing partners, assuming that they were benefiting as well from the pooling of knowledge. At the same time, using industry average spending as a driver keeps the estimates conservative, since it must include the savings of the partnership. Sum it all up, then take the difference and you should have a reasonable estimate of how much it saved all the partners so far.

Using this method, I came up with a cumulative savings since 2003 of almost $8B. That's well over $1B per partner! In 2006 alone, they saved $3.6B. That was up from $2.2B in 2005. It should be no wonder why IBM is attracting so many new members into their partnership platforms.

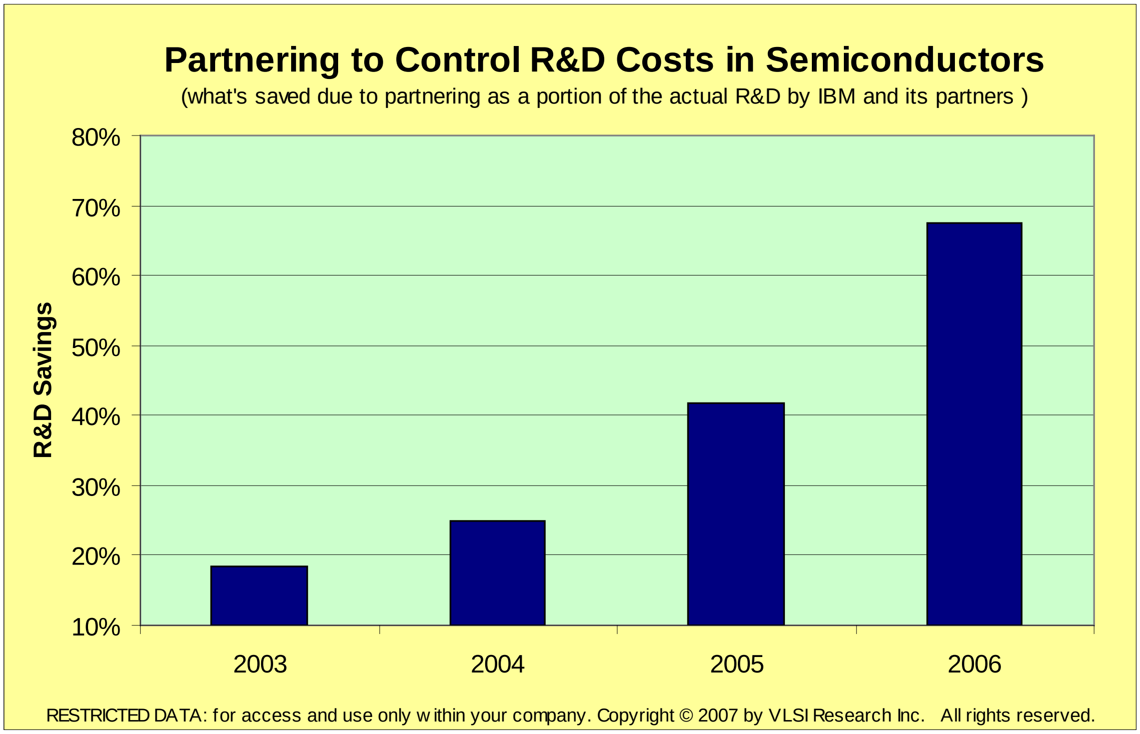

Had these companies tried to go it alone, they would have needed almost 70% more R&D funding than what they spent with the partnership. For a public company like **AMD**, who is a partner and spent $1.2B on R&D in 2006, this multiplier works out to a savings of roughly $800M, or about 14% of sales. These are save-the-company value levels.

One of the fun things about this industry is that it is not only innovative on technical things, it is also highly innovative in its business processes. IBM's partnering model represents a fundamental advance in R&D productivity.