

IT TAKES ADVANCED MANUFACTURING TECHNOLOGY TO MAKE A NATIVE MULTI-CORE PROCESSOR

Part III of a Series of AMD White Papers on 64-bit Computing

Since the dawn of human history, our species' ability to create and deploy advanced technology has differentiated us from earlier species. More than 8,000 years ago, humankind began its progression through the Stone, Bronze, and Iron Ages. During the last three centuries we have evolved through the Industrial, Space, and Information Ages, each characterized by increasingly sophisticated manufacturing technologies. Some regard rocket science as the most challenging technology of our age. Rocket scientists, to be sure, face enormous pressures to get it right—service calls in outer space can be prohibitively expensive. Nevertheless, AMD argues that modern semiconductor manufacturing technology, not rocket science, represents the true pinnacle of humankind's tool-making achievements.

For more than 37 years, AMD has contributed to advancements in semiconductor manufacturing technology. In this white paper we will share some of the approaches that position AMD among the leaders in the industry.¹ First, we will explore how its patented

APM (Automated Precision Manufacturing) technologies enhance the speed, accuracy, and agility of production facilities (known in the industry as “fabs”). Then we will examine how AMD stays at the forefront of semiconductor technology development by leveraging its intellectual and financial resources through collaboration with other industry leaders. Next, we will discuss how AMD introduces new technology into its fabs using evolutionary, rather than revolutionary tactics, and how it refines that technology using a unique “lab in fab” approach. Finally we will outline the production capacity already in place to manufacture award-winning AMD Opteron™ processors, AMD Turion™ 64 mobile technology, and AMD Athlon™ 64 processors, and how it intends to augment that capacity over the next few years. We have targeted this document at technology users, rather than technologists, and we have tried to keep our descriptions at a high enough level so you should not need a PhD in semiconductor technology (or even rocket science) to follow what we have to say.



¹ At the risk of appearing immodest, we note that results of independent, third-party benchmarking by a leading industry manufacturing consortium founded to improve semiconductor manufacturing effectiveness consistently ranks AMD among the industry's top performers on key metrics.

APM Forms the Nervous System of AMD's Fabs

It takes almost 800 discrete manufacturing steps to produce the silicon in AMD's advanced microprocessors. Even small mistakes during the manufacturing process can have big impacts on the performance and durability of the processors. Small variations in the settings for each piece of manufacturing equipment in the fab can also influence the number of marketable processors made from each wafer (known in the industry as "yield"), which in turn affects product costs and the prices at which AMD can profitably market these processors. With so much riding on effective fab management, it made sense to AMD to invest heavily in tools to manage its fabs. AMD has refined these tools for almost two decades; it started back in the days when the industry still relied on 150mm wafers, half the diameter of the 300mm wafers used in the newest fab. Today's APM (version 3.0) addresses five major aspects of fab operation:

- Equipment Performance Optimization (EPO) monitors each machine in the fab and takes remedial action whenever one of them misbehaves or drifts out of spec. Even machines operating within their specified limits exhibit small variations in their behavior, but APM dynamically adapts the instructions it gives to each to compensate for these minor variations.
- Product Performance Targeting (PPT) allows AMD's fabs to "tweak" their product recipes to enhance product performance (for the fastest chips) or to maximize product yield (for mid-range and value offerings). Sometimes, this requires shrinking the size of the gates in transistors by just one or two nanometers (billionths of a meter), or to adjust the thickness of those gates by just a few angstroms (tenths of a nanometer). APM's database contains all of these parameters, and transmits just the right values to each machine that touches a given wafer as it passes through the fab.
- Advanced Process Control (APC) lets APM handle many of the routine tasks that keep a fab running at full speed. APM provides the detailed recipes that each machine uses for a specific lot of wafers. In the old days, operators would tote cassettes filled with wafers from one production station to the next, and manually enter the production parameters for that cassette. This left room for human error. Today, APM completely automates the inputting of recipe parameters. It tracks each wafer in each lot individually as it passes through each of the hundreds of manufacturing steps needed to turn raw silicon into high performance microprocessors. This helps eliminate all those human errors, and allows AMD's fabs to post some of the highest yields in the industry.
- Integrated Production Scheduling (IPS) gives APM the responsibility to manage the schedules for each machine in the fab and each lot of wafers flowing through the fab. This involves matching thousands of wafer lots with almost 400 individual machines, some of which may be out for service at any given time. Not all machines can handle all tasks; there may be 40 steppers (the complex devices that transfer images of a processor's features onto the surface of the wafer), but only a few have the accuracy needed for certain "critical dimensions." To make things even more complicated, APM knows that wafers passing through stepper A may yield better if they are subsequently routed to etcher X, even though other etchers (Y and Z) can also handle the task. Think about the problems you had the last time you tried to prepare a five course meal for 20 guests, using ovens, stovetop burners, microwaves, and maybe a mixer or blender. Then multiply by about 100,000 and deduct 50 points for each overcooked or undercooked entree. APM does this 24 hours a day, 365 days per year, and rarely "overcooks" a chip.



- Yield Management Systems (YMS) give AMD's fab managers the ability to adjust the output of their fabs to meet customer requirements. Few outside the industry realize that fab operators must often trade processor performance for volume; that a given wafer can yield a large number of medium-performance processors, or a smaller number of high-performance devices, but not both. Faster processors often command premium prices, but many customers prefer to sacrifice a little performance to get a more affordable system. Thus the fab wants to produce just enough of its fastest chips to satisfy market demand, but not more. If a fab makes too many fast chips, it must sell its excess high-speed chips at the lower prices normally reserved for slower chips, a tactic the industry refers to as "down-binning." Then it gains neither the premium price nor the volume advantage that it might have attained by targeting moderate performance in the first place. APM lets AMD's fab managers favor either performance or volume on a lot-by-lot basis. It even offers some opportunities for mid-course corrections that swing the balance after the wafers have been partially processed.

Alliances Leverage AMD's Intellectual and Financial Resources

These days, if you want to be in the semiconductor manufacturing business, you need more than just a two or three billion dollar fab. The semiconductor industry has progressed through three distinct eras during the 37 years AMD has been in business. Back in the 1960s and 1970s, AMD and other semiconductor pioneers developed their own manufacturing equipment and process technology. Fortunately, the processes were simpler and the fabs were smaller and more affordable. As the industry grew, equipment suppliers like Applied Materials, ASM Lithography, and Nikon emerged and ushered in the industry's second phase. The rise of these independent equipment suppliers allowed any company that could afford a few hundred million dollars for a state-of-the-art fab to enter the industry. This



facilitated the emergence of "semiconductor foundries" like TSMC and UMC that manufacture processors designed by others, and "fabless" semiconductor companies that design and market processors, but leave the actual manufacturing to others. In those days, companies could improve processor performance merely by deploying new, more advanced equipment, and had little need to carry out basic research into material science and semiconductor physics.

By the end of the 1990s, the industry had harvested most of the low-hanging CMOS (complementary metal oxide semiconductor) fruit, and entered its third phase of development. The technology needed to keep companies on the periodic 2X transistor density improvement mandated by Moore's Law could no longer be purchased from equipment suppliers. Rather, companies needed to employ small armies of physicists and chemists to invent new transistor structures that relied on advanced materials like copper, silicon on insulator (SOI), and strained silicon to drive improved transistor performance. AMD realized its requirements closely matched those IBM was pursuing with Sony and Toshiba in their Advanced Semiconductor Technology Alliance (ASTA) and it joined that group in early 2003.

Since that time, many of AMD's semiconductor development and process technology experts have worked side-by-side with colleagues from IBM and other major chip makers in IBM's East Fishkill (NY) fab. This group has already met its goal of developing a 65nm SOI/strained silicon manufacturing process that can be deployed by IBM in New York, by AMD in Germany, by other Japan-based OEMs, and by Chartered Semiconductor in Singapore. AMD will soon begin manufacturing products based on this technology in its recently commissioned 300mm Dresden fab.

The collaboration between AMD and IBM has been so successful that AMD and IBM have extended their relationship through 2011. They added research on 32nm and 22nm processes, as well as development for 45nm process technology, to their shared research agenda. AMD scientists and engineers work side-by-side with colleagues at IBM's Watson Research Center and East Fishkill fab. AMD and IBM have presented co-authored papers on key aspects of our research at major industry conferences around the globe. They recognize that the challenges inherent in these advanced technologies are far too big for any one company to tackle successfully. AMD's ability to cooperate successfully with a company like IBM leverages financial resources, but more importantly, strengthens the industry and helps maintain effective competition within this very important industrial sector.

Moving New Technology from the Lab to the Fab

Years ago, AMD developed new process technologies in a fab dedicated exclusively to research and development. When the engineers at its Submicron Development Center (SDC) in Silicon Valley thought a process was ready for prime time, they handed it off to the engineers in production fabs, and they made it work on a larger scale. AMD realized it could simplify the process and save time if it developed the new processes

inside the actual fab where they would be deployed. AMD could add the new tools needed for its next-generation process to the equipment already in the production fab, so operators would be familiar with these devices by the time it was ready to roll out the new process. It called this approach "lab in fab," and it has served well over the past five years.

AMD's APM capability greatly simplifies the implementation of the lab in fab concept; indeed it would have been nearly impossible to make the concept work without APM. It's easy to understand why fab managers responsible for manufacturing the products on which AMD's customers depend might worry that new machinery or process flows could upset ongoing production. To help ensure that the lab doesn't disrupt the fab, AMD uses APM to schedule new machinery and development wafers along with production lots. Operators in the fab see no difference between wafers designated for customers and those used for development, unless they happen to glimpse a cassette moving to a new tool not used in the established process flow. Once the developers and the fab managers agree a new process is ready for deployment, all the equipment and training needed has already been put in place in the fab that will use it. The handoff entails nothing more than a few signatures on a management form. Given the benefits this lab-in-fab concept delivers, it is a wonder that other companies prefer to continue massive investments in development-only fabs.

AMD's lab-in-fab approach provides another big advantage with regard to the introduction of enhanced process technologies. Other semiconductor companies develop their next-generation processes while they ship products based on the current generation; then (every two years or so) they throw a switch, introduce chips designed (or redesigned) for the new process, and begin to ramp the new process up in volume as they ramp



the old one down. Problems with the design of the new processors or the high-volume deployment of the new process can (and sometimes do) derail this volume ramp and disrupt the production schedules of customers awaiting these next-generation chips.

AMD takes a more evolutionary approach to process development and phases in small improvements to its processes on a quarterly basis. Over the course of eight quarters, these small changes add up to a whole generation of improvements, but this approach brings the benefits of the process improvements to customers more rapidly. For example, the 90nm process deployed in AMD's fabs in early 2006 uses the same transistors as the 65nm process it plans to deploy later in the year. This approach also helps mitigate the risks inherent in the more revolutionary process transitions undertaken by other suppliers.

Today's Capacity Investments Will Support the Growing Demand for AMD's Chips

Until recently, AMD relied on the output from just one megafab for all its microprocessor shipments. Its 20,000 wafer per month Fab 30 in Dresden, Germany produced over 50 million processors in 2005, but its market share has been growing and it realized a few years ago that an exciting product roadmap would generate more sales than Fab 30 could accommodate by itself. AMD is pursuing two complementary strategies to help ensure that manufacturing capacity will never be a factor that limits its growth.

AMD's first move was to add a second megafab, known as Fab 36, adjacent to its existing fab in Dresden. It formally commissioned that fab, its first to use 300mm wafer technology, in October 2005, and it began shipping chips to customers in early 2006. By early 2007, this fab will have the capability to process over 13,000 wafers per month. It will more than double AMD's capacity and allow it to ship over 100 million

chips per year.² The clean room in Fab 36 has room to hold enough equipment to handle up to 20,000 wafers per month as capacity requirements grow in future years.

AMD's second move was to enter into a production agreement with Chartered Semiconductor, a leading foundry company based in Singapore. Chartered's new 300mm Fab 7 in Singapore uses the same process technology IBM uses in East Fishkill, which is remarkably similar to the process technology AMD uses in Fab 36. (The 65nm processes at IBM, Chartered, and AMD Dresden will be even more closely aligned.) This allows AMD to shift some production to Chartered without the need to redesign products to accommodate differences in processes. AMD's relationship with Chartered gives it the ability to augment capacity in a flexible and capital-efficient manner, and lets it respond quickly to seasonal shifts in demand or to the success of marketing programs at major new or existing OEM customers. AMD has even licensed many components of its APM technology to Chartered, so they can attain the same precise control of their Fab 7 environment that AMD has in its Dresden fabs.

² The 300mm wafers used in Fab 36 contain roughly 2.25 times the area of the 200mm wafers used in Fab 30, so those 13,000 wafers per month represent the capacity of almost 30,000 200mm wafers.

Summary

We hope this document has demonstrated that AMD has the technological and financial wherewithal to meet the growing market demand for its industry-leading microprocessors. AMD invites comparisons between its technology and that of its competitors. AMD realizes that no vendor has a monopoly on good ideas, and has partnered with the best and the brightest semiconductor technologists to stay on the leading edge of transistor, and process development. As a result its AMD Opteron™ processors, AMD Turion™ 64 mobile technology, and AMD Athlon™ 64 processors deliver the capabilities end-users need to keep up with ever-increasing requirements, while at the same time being kind to both their pocketbooks and their air conditioners.

Of course, AMD knows that it takes more than a solid architecture, a solid ecosystem, and advanced manufacturing technology to make industry-leading multi-core processors. It also takes great people. Although this paper focuses primarily on the technologies that underlie AMD's manufacturing edge, those technologies are only as good as the people behind them, and its manufacturing organization includes some of the most talented contributors in the world.

AMD's approach to manufacturing mirrors that taken by the athletes and sports organizations it sponsors. AMD, like seven-time Tour de France Champion Lance Armstrong and the Discovery Channel Pro Cycling Team, works to improve its speed, accuracy, and agility, with the goal of consistently out-maneuvering, out-performing, and out-innovating the competition. And, just like Lance Armstrong and the Discovery Channel team, it takes great people, along with the most advanced technology, to create a winning team.



To Fab or Not to Fab, That is the Question

In the early days of the semiconductor industry, chip manufacturers of necessity built and operated their own fabs, and often differentiated themselves by unique aspects of their manufacturing technology. Over the past two decades, the emergence of a silicon foundry infrastructure has allowed many companies to enter the industry and rely on foundries for all their production. These “fabless” companies have far lower capital expenses, and (in theory) can respond more quickly to market demand than companies that continue to manufacture their own processors. This phenomenon causes some to ask why AMD doesn’t “go fabless?” Whenever AMD reviews this situation, it finds the benefits of an integrated manufacturing approach generally outweigh the costs, for the following reasons:

- **Economic:** Fabless suppliers must share some of the profits they derive from the sales of their chips with their foundry partners. If a company can finance a fab, and operate it at close to capacity, the profits it generates usually far outweigh the cost of capital to build the fab. Only when a company cannot finance a fab, or afford to keep up with technology advances, or absorb all a fab’s output, does it make sense to outsource the wafer manufacturing function. None of these constraints apply to AMD.

These economic issues take on increased importance when some competitors use integrated strategies, and others go fabless. In some market segments (e.g., discrete graphics chips), all the principal market participants use a fabless approach, and thus all face the same issues with regard to margins, time to market, and access to new technologies.

- **Business:** The semiconductor market is well known for short product cycles. Time-to-market plays a key role in determining a product’s overall profitability, since a product is usually most competitive, and commands its greatest price premium, at the time of its introduction. Factors that delay introduction are bad; factors that accelerate it are good. Integrated suppliers can move lots of prototype wafers through their fabs much more rapidly than fabless suppliers, as long as they are willing to place the prototypes ahead of revenue production, a move most foundries cannot tolerate. These so-called “rocket lots” can trim weeks from development schedules and make the difference between a competitive product introduction and a “me-too” product.

Some products require new and improved processes to reach their full potential. Reliance on a foundry’s ability to migrate to a new process generation, followed by the foundry’s inability to move that new process into volume, has scuttled many a product plan for fabless companies. AMD views the issue of “to fab or not to fab” as an economic and business issue, rather than an ideological one. It realizes that the use of foundries can make sense under certain circumstances, and has structured its relationship with Chartered Semiconductor for just those circumstances. But as long as the economic issues suggest that integrated manufacturing makes sense, AMD will continue to build and operate its own fabs.

