Leading-Edge Logic LSI Strategy
Construction of 300mm fab for 90nm/65nm production

Remarks by

Toshihiko Ono
Corporate Senior Vice President and
Group President, Electronic Devices Business Group
Fujitsu Limited

March 19, 2004

[Title slide]

Good morning. As you know, we have called you here today to discuss our plans to construct a 300mm wafer fab, but rather than talk about methods and means I'd like to explain our strategy in logic LSI devices. Perhaps it's just my imagination, but it seems to me that there are all kinds of rumors, suppositions, and projections floating around out there, so I thought it best that I set the record straight, and that is the reason for today's meeting.

First, as I said, this talk concerns the construction of a 300mm fab, which we see as simply a "tool" for ensuring production capacity. In fact, we are already engaged in 90-nanometer production at our Akiruno facility. Although it's with 8-inch wafers, the point is that we already have this production at Akiruno. Today I would like to talk about our plans going forward.

To describe our current business portfolio in simple terms, where we really make our money is in logic chips in general, including 0.13-micron microcontrollers and analog devices. So if we don't have a sufficient business foundation in this segment, there's not much we can do. In this regard, we've increased productivity at our existing plants, improved cost competitiveness, and reallocated resources, including, for example, transferring over 300 software designers from another group to my group. We've also consolidated four assembly operations into one and are continuing with other streamlining efforts. All of these moves have helped to secure our competitiveness. We have also spun off our flash memory and gallium-arsenide compound semiconductor operations. I should say that we've structured them so that they are no longer consolidated operations. To be frank, there is still much that we need to do, but in terms of really responding to our customers' needs, I believe we are taking an important step forward with the construction of this plant. Later, I will touch on the strategic nature of this decision again.

[Slide 2]

So, to get back to my original point about thinking of the production facility as a tool, let me explain what I mean by that. In Japan and abroad, Fujitsu is highly respected for our capabilities, especially with customers overseas. In terms of evaluating our actual capabilities, though I don't have any hard numbers to show you today, let's just say we're a few levels beyond certain other companies and we've received some extremely positive evaluations. To build on that, we are investing in technology for mass production using 300mm wafers with 90-nanometer and 65-nanometer process technologies, and the reason for this investment is that it will enable us to achieve good

cost-competitiveness and generate profits. In short, we're building a 300mm fab because we want to improve cost-competitiveness and expand the overall competitiveness of our business.

I'll touch on this later, but what we're doing is creating various partnerships with customers who share our vision, what I call the "New IDM" (integrated device manufacturer) model. We hope to rapidly recoup our R&D investment and of course increase profitability, centering on application specific processors.

The new fab will be built in Mie Prefecture on land adjacent to our current Mie facility. It will focus on 90-nanometer designs and beyond using 300mm wafers. It will have 12,000 square meters of clean-room space and a production capacity of 13,000 wafers per month. The production system will obviously be able to handle that kind of high volume, but the plant will also be equipped to carry out multi-part processing with small-batch control as low as a single wafer. This is a little different from single-wafer processing in the usual sense, but still, you should think of it in terms of wafer-level production management.

As to the clean room, we're designing it to accommodate 300mm wafers and for very high-volume customers. Production here will have ramifications throughout the world, so we need to take even unlikely contingencies into account, such as something like the major earthquake that we experienced at our Iwate plant last year. But here such an event would have even larger ramifications. From our customers' point of view, just building a plant in Japan at all is taking a risk. To cope with this, we're using a hybrid seismic isolation structure. Our schedule is to commence operations in April 2005 and to begin input immediately. We plan to begin volume shipments by September 2005 at the latest.

[Slide 3]

I mentioned that we'll be using land at our Mie plant site that we already have. As you can see in the photo on the left, the white buildings are, from the right, buildings number 1, number 2, and number 3, and so logically the new plant will be building number 4. This is land we had planned on using all along. The blank space to the left is still there for expansion or another building. And so that's where we're building the 300mm facility. We're taking advantage of shared resources that are already in place at the site, such as the existing power supply.

[Slide 4]

I mentioned before that we'll be introducing a hybrid seismic isolation structure, but since that may not be very clear, allow me to explain it in a bit more detail. In the case of minor tremors, what you're concerned about is the equipment, such as aligners. The hybrid design will be able to absorb micro-vibrations, say, near the aligners, and macro vibrations will be absorbed as usual. If we suffer a serious earthquake, the structure will be able to absorb the seismic impact and release it. To put it another way, the impact of a magnitude-7 quake (on the Japanese scale) would be reduced to that of a magnitude-3 quake. At magnitude-3, there would be virtually no damage to the fab. At the bottom of the slide, you can see structural equipment to handle vertical vibrations on the left, horizontal vibrations in the center, and at the very right is a hydraulic setup that releases seismic acceleration. These pillar structures will be used throughout the building, making the entire structure seismically isolated.

[Slide 5]

Now, let me return to the central question of why we're building this facility. In terms of our investment strategy, this is part of a major shift in our advanced logic LSI strategy. For instance, years ago we were making investments in DRAM and other memory chips. But these investments were always predicated on there being a market for these products. For example—and here I'd like you to look at the left-hand side of this figure—if you ask why we have been able to maintain leading-edge technological capabilities, or why we have continued to invest in development, it's because these capabilities form the foundation that supports our software and services offerings. This includes the quality of our offerings as well as our ability to develop new services, such as in the field of security solutions, an area in which there is growing attention. Of course it also includes SPARC processors for our Unix servers, and what we call Pleiades, as an internal code-name, for large-scale Linux servers in open systems. Another example is our media processor for communications. We can come up with new concepts, but it is the products in which they are used that ultimately determine if they will be brought into reality. To put it another way, we feel that if you don't have advanced products, you aren't going to maintain leading-edge technological capabilities. So when we're talking about advancing that technology one step further, that is the background and rationale for why we are doing it.

If you look at the upper-left here – this is what we mean by development for in-house use – in the past, this is where we were investing our development funds. Now, let's say that you invest at the same time you expect the market to become commoditized, so, for example, it roughly coincides with the entry of Taiwanese manufacturers. Now, with that kind of investment, the products are already becoming commodities, so it's very difficult to recoup your investment costs. So if you're in the midst of development for in-house use, or if you enter the market and aren't able to secure a sufficient volume, it doesn't bode well for recouping your investment costs.

Now, I don't want to create any misunderstandings, but it's often said that once a system-on-chip device becomes a commodity, it is no longer profitable. In some sense that is true. You can see the effects of commoditization, for example, in digital audio-visual field. One-time development expenses may be recouped for 2 or 3 years. Against that, you can see that product life cycles are getting much more compressed. Take mobile phones, for example, and I'm sure you'll see what I mean. Development is moving on incredibly short three- and six-month cycles. And enormous development resources are required to create the system-on-chip devices for those products within those time-spans. In such cases, even before volumes have ramped up enough to generate returns on the existing product, you have to start investing in development of the next-generation product. So, with this model, it becomes very difficult indeed to recoup development costs. One possible way to resolve this issue would be to try to keep manufacturing going for a very long time, or to get the volumes needed to recoup the development costs.

That's easy to say, but if you try to simply get out of low-volume items, that certainly doesn't help your customer. So, as you see in the lower left here, we've moved into structured ASICs and, as we announced the other day, dynamically reconfigurable chips, and the key point in these products is the software. We offer FPGA and ASIC concurrent design services, and that's a way of providing customers with continuing support. At the same time, we are also presenting new solutions to our customers.

Still, these measures alone are not enough. Given the pressure to recover development expenses as quickly as possible, we are changing our approach to investment and our dealings with customers. As you can see on the right, we have switched to a model where we are developing technology and recouping development costs at the same time. More specifically, we have a basic IP library that we have developed for internal use. For example, we have IP for SRAM, ROM, basic cells for logic, high-speed I/O, PLL, and so on. When a customer tests a product, we simultaneously provide the minimum necessary IP library. On top of this, we'll work with partners that value our technology and continue the required development work. Then, when partners want to quickly ramp up production volumes, it makes it easier to recoup development costs quickly. Even when the product eventually becomes a commodity, we already have a cost advantage. I'll speak later about our partner strategy, but in terms of strengthening our strategic emphasis, it is similar to what companies like IBM have been saying. We are positioning ourselves as a leading edge technology partner, and we already have more than ten partners lined up.

[Slide 6]

As a result – and what I'd like to emphasize here is – we will no longer make investments based on speculation about what we think will be the next big market area. We won't invest unless customers are clear about their needs. What we're doing is switching over to investment that is centered on the customer. Obviously, we won't go ahead unless we can make money, and that means that we have to rely on self-generated cash flows. So we need to avoid getting into a vicious cycle with investments in items that have a high development-cost ratio. Again, our basic stance is to pursue a business model that is based customer responsiveness.

This graph should give you a more concrete understanding of what I've been talking about. It shows overseas partners, domestic partners, and internal partners.

I'll get around to a more complete explanation of internal partners later, but first look at overseas partners, shown here at the very top in dark blue, and domestic partners—I'm talking about companies to which we are providing application specific processor solutions—shown in pale blue. Internally – this includes areas such as server chipsets – this is shown in white. The vertical axis on the left shows wafer demand based on 8-inch format. This takes into account continuing production at Akiruno, so to avoid the confusion of converting to 12-inch, we've just calculated everything in this graph in terms of 8-inch. These figures do not refer to capacity. These are our demand forecasts, and we believe that 2004, 2005 and the first half of 2006 are pretty firm. To be perfectly frank, if we took customer demand strictly at face value, we'd have an overflow. But looking carefully at this business and the market trends, and by market I mean customer trends, we have to account for a certain amount of risk. So we've cut the stated requirements by roughly half and are planning our production accordingly. To be honest, we've had various requests from customers, including requests to accelerate production, and we will continue to make adjustments, but this gives you a basic idea of what our plans are.

Below that you'll see our corresponding planned investment amounts. These are raw investment figures, and actually the equipment will be leased. So, as shown, we are planning roughly 25 billion yen in fiscal 2004, 50 billion yen in fiscal 2005, and 85 billion for fiscal years 2006 and 2007 combined – the majority of that will be in fiscal 2006, so you could figure perhaps 55 billion yen of that in 2006. We've tried to account for the fact that demand could obviously increase or slip from one year to the next. So we're

going to re-evaluate this in fiscal 2005 fiscal year; the first phase will be up through fiscal 2005, and the next phase – 2006 and 2007 combined – should be 85 billion yen.

Below that – and we debated about whether we should disclose this amount or not – you can see that we have 30 billion yen in capital infusion from partners. Now, if you're wondering what this represents, from the customers' standpoint, they want guaranteed capacity; and from our standpoint, we want a commitment that those revenues will materialize. So we've tried to resolve these respective needs in the form of a deposit.

This isn't just one company, it involves several, and they will be investing in the form of a deposit. We're always taking a portion of our current profits and reinvesting them. For example, with earnings from everything from ASICs to microcontrollers to analog devices, a portion of that may be recognized as profits in fiscal 2003, but another portion will be invested in future products. For example, the expenses associated with development costs at Akiruno. There are also the development costs of prototypes used in Fujitsu's next generation of products.

We also believe that, in a few years, chips and applications will inevitably merge, so we are making forward-looking investments in that regard. Now, obviously we don't receive any current return from such forward-looking investments. When you put all of those costs together and you look at the reported profitability figures of our semiconductor operations, they may look pretty flat. But if you take out those investments in future technologies, the truth is that our real profitability now is higher than ever before. In that sense, I'm confident that we're in a position that no competitor can match. I think we need to further increase our forward-looking investments. And if we combine the deposits from partners with our own internal sources, I'm confident we can fund these investments ourselves. So I have no intention of asking our CFO to issue additional corporate bonds or to secure equity financing.

[Slide 7]

Over the past year or two, I've been discussing our "New IDM" model with various partners. Suppose for example that a customer becomes a partner in this model. As you know, with 90-nanometer technology a single chip may contain many different functions, and the customer might ask for a single chip that contains all kinds of functions. To put it in extreme terms, you can literally put the whole system inside the chip, entailing a very complex design. Obviously there would also be software involved, so we would need to work together with the partner from the start to properly test that software. In this case, in contrast to the foundry model, domestic and foreign customers obviously want to make use of the software and design skills and environment that we as an IDM can provide. And on the production end, they want to have flexibility—that's the direction things are going. In that sense, you have software, chip design, and system design, and if the technologies underlying that are not consistent, you're not going to be able to produce a high-performance chip or get good yields. That's why there are limits to the old horizontal organizational model.

Another aspect of this is what I mentioned before about the difficulty of staying abreast of advanced technologies if you don't have your own advanced products in which they are employed. On the left side of this slide you can see a column for internal partners – that's Fujitsu – and we have internal groups carrying out set manufacture, set design, and chip design and software. All of these need to be thought through in conjunction with the semiconductor side. For example, with a Linux server, if you're talking about developing a large-scale, high-performance model and want to really maximize

performance, you really need to make the required chipsets on your own. Likewise with SPARC. The market for Unix servers is shrinking, but if you want top-level reliability to support software and services operations, that's what SPARC can deliver.

Let me go back to the subject of dynamically reconfigurable processors. This is a completely new approach to network processors. Trying to handle this function at the software level only is very burdensome, and for security you need to put a chip somewhere around a firewall. But if you can configure part of this function to reside at the hardware level, you can create a very fast, very reliable network processor. And that's just one example.

Next, regarding applications, we have a variety of them, such as our new Shunsaku application that we've made available to customers. I think some of you here may even have used it, but in any case, it's an application that enables extremely fast searches. When used in a grid configuration – we're using our own processors for parallel processing – it makes possible amazingly fast software and services. This kind of world is illustrated on the left. In the middle column, you can see domestic partners, and this would involve our solutions offerings in ASICs and ASPs, for which we handle set design, chip design and software together with our partners. This means that we've got better knowledge of the customer's previous chip design, and we know their software. Now this is becoming a very complex and advanced area, and, in terms of resources and skills, there are not many companies that can handle all of this under one roof. But Fujitsu has these software capabilities. We have embedded software and, as I mentioned earlier, we've transferred 300 software programmers into my group. So we have the resources in place to offer a superior level of service and skills. We're placing particular emphasis on the graphics and automotive fields.

Our overseas partners are looking hard at set manufacturers around the world. We can offer them graphics and MPU products, and for this it is essential that chip design and software be done together by Fujitsu and our overseas partners. In any case, quantities are increasing for the scenario on the right. On the left, this is mostly for industry, but volumes increase as we move from left to right. The largest volumes are from overseas partners.

The New IDM concept may be slightly difficult to grasp, but one point is that this business is increasingly moving to a model in which you can't do it alone. You hear often these days about design crisis. A design crisis, for example, means that we would burn out if we tried to do everything ourselves. Working with an overseas partner on a high-volume development project, you might think it would be difficult to allocate sufficient design resources. But if they have resources with extremely advanced design skills that they can allocate, this is where working together comes in. It allows us to avoid overtaxing our design resources. So think of it as a world geared for collaborative work in which we have these kinds of combinations.

[Slide 8]

Finally, I expect that you would like me to name our partners or the companies that will contribute the 30 billion yen. But this involves our customers' business – their business moves and shifting alliances – so I'm afraid it must remain confidential. There are cases, for instance, of companies that had been collaborating together and now one side has switched over to us. And as top management, we have agreed to keep such information confidential, so, I'm sorry that I am unable to disclose the names. Of

course, if the customers consent to it, I will be happy to disclose their names in the future.

However, Transmeta is one company that has already announced it themselves, so today I've included this example of Transmeta in the area of MPUs. Matt Perry, Transmeta's CEO, has kindly allowed us to use his comments. What he's saying here is that predicted performance, time, yields—all of these have gone according to plan and product is now into production. That's one point. The other point that I'd like to make is that, in my view, our work with Transmeta is really a successful model of how the New IDM can work. They're making a high-speed device like our SPARC processor. Typically, one of these devices would use a fair amount of electricity. But with our technology, we've developed a chip that has high performance, low power demands, and low leakage—seemingly contradictory properties. On top of that, Transmeta has achieved the lowest power demands yet. It's mentioned here, their LongRun2 technology. About a month or two ago there was some small coverage of it in the newspaper. I really don't understand why it didn't receive more coverage. This is really a spectacular method, and they've verified its effectiveness. I can't go into the details about the know-how that went into it because the design and software methods are secret. But combined with our technology, it is extremely effective, with no precedent. Great speed and incredibly low power—Transmeta can tell you more about the low power aspect. In any case, I wanted to introduce this to you as one example.

That ends my presentation.