Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
1. it is reasonable to assume that a first-tier manufacturer would run only 20 million units of a product iteration (Geilhufe, Tr. 9562:10-9563:4; 9725:1-9726:23)	Geilhufe Report at 10. ("The cost estimates in Tables 1-4 below are my best estimates based on my 30 years of integrated circuit manufacturing experience and are based on a mature product generation at a volume of 20,000,000 units per product produced by a leading manufacturer.") Geilhufe Depo. At 72:13-16 ("Q. . . . You are talking about a mature product generation at a volume of 20 million units per product produced by a leading manufacturer? A. Correct.") Geilhufe Depo. at 75:2-6 ("Q. Why did you use the number 20 million bits per product? A. I was trying to estimate what volume a top-tier manufacturer at this stage in production would be outputting. It's a best estimate that I took.")
2. use of fixed CAS latency parts is difficult and costly because (a) based on all options contained in the JEDEC standard as adopted (and not on industry usage or practice), 3 separate parts would be required (Geilhufe, Tr. 9578:10-23, Tr. 9682:20-9683:2)	Geilhufe Report at 11 (first column of table relating to the fixed CAS latency alternative shows that Geilhufe was basing his cost estimates on three separate parts) Geilhufe Depo. at 122:7-17 ("Q. Now, you concluded in this table, first column under Inventory, that there would be three parts -- A. Right. . . . Q. What did you base your conclusion on that there would be three parts manufactured? A. The JEDEC spec provides for the option for three parts.")
(b) it would cost approximately $100,000 more than programmable CAS latency in design costs (Geilhufe, Tr. 9575:9-21)	Geilhufe Report at 11 (first column of table relating to the fixed CAS latency alternative shows that Geilhufe was estimating approximately $100,000 more in design costs) Geilhufe Depo. at 80:25 – 81:21 ("Q. . . . When I look at the cells relating to the product design cost element, it says add, and that squiggly line in front of the dollar sign, does that mean approximate? A. It is approximate. Q. Add approximately $100,000 per CAS latency part. Is that an accurate reading of the sentence? A. That's accurate. Q. Does that mean this cell -- sorry. Does that mean that the design, the product design for the fixed CAS latency part is $100,000 more than it would cost to design an SDRAM that had programmable CAS latency using

EXHIBIT A

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
	the mode register, or does that refer to the cost of the product design itself? A. Good question. First your example has very little design effort per se. It just has the effort related to a simulation being able to create the tools. Given that for fixed CAS latency you now have more than one product, each product has to have its own photo tools. So to do each one of those products in my estimation is about 100,000, given that I have a baseline and I have a product that's done.")
(c) it would require assumptions about the speed grade of the parts (Soderman, Tr. 9347:8-9348:11)	Soderman Report at 24 ("Moreover, manufacturing a number of different fixed latency parts would have a significant impact on yield, a key consideration for DRAM manufacturers. . . . As noted above, different DRAMs from the same wafer will have different speed ratings, and the ratings cannot be accurately determined until the part is packaged and tested. Some of the DRAMs in question will not be fast enough to operate with a latency of only 2 at the contemplated bus speed.") Soderman Depo. at 146:10-25 ("Q. What do you mean by "significant guesswork"? A. In other words, if you fixed the latency when you were manufacturing the die; in other words, you have no idea what the performance of that part is going to be; in other words, the performance of that particular die, because you can't measure it in the die form, you have to package it, but you had to fix the latency at that point, much like whatever technique you're using to fix it, much like maybe if you wanted to use lasers or whatever, if you fixed it at that point, you're never quite sure whether the parts are fast or slow. You're making an educated guess. And that is not something a very efficient and cost-competitive industry can afford. There is no guesswork in DRAM manufacturing. It's really a fine-tuned art, engineering.")
(d) it would interfere with a manufacturer's ability to speed grade parts (Soderman, Tr, 9348:12-9349:15)	Soderman Report at 24-25. ("Moreover, manufacturing a number of fixed latency parts will lead to a loss of much of the price premium that manufacturers currently obtain from sales of faster DRAMs. Currently DRAM manufacturers sort each product into 2-3 speed-grade bins (after packaging and testing) depending on their performance and then distribute them at different prices. The faster performance products command a premium price that can be as much as 50% [Denali 2002] higher than slower parts. Under Professor Jacob's alternative,

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
	however, in order to maintain acceptable yields for a wafer many of the fastest parts that currently command premium prices would have to be fixed with a latency value that would result in the full performance capability of the part not being realized and, consequently, could no longer be sold at a premium price.")
(e) it would add expense due to decreased die yield (Geilhufe, Tr. 9577:1-9578:9)	Geilhufe Report at 11 (first column of table relating to the fixed CAS latency alternative shows that Geilhufe was estimating an additional expense of approximately $.03 per unit due to decreased die yield)
3. use of fixed burst length parts is difficult and costly because (a) based on all options in the JEDEC standard as adopted (and not on industry usage or practice), it would require 4 separate parts (Geilhufe, Tr. 9594:25-9595:3)	Geilhufe Depo. at 154:20 – 155:11 ("THE WITNESS: I needed to refresh my memory with a JEDEC spec. I have the JEDEC spec in front of me. And the JEDEC spec calls for four SDRAM burst lengths and four -- 3 DDR DRAM burst lengths. Q. So where did the six parts for the fixed burst length column come from? A. I would have to look at my notes if I have that. My first answer is the number very likely came from the burst length and burst type combination, but I can't be certain right now. So I think a better approach is to assume -- not assume but to change the report to four burst length elements or parts at this time and change -- as you can tell, I used roughly one cent of inventory cost per two part types within -- and I would amend the inventory cost to three cents as opposed to four cents to make sure the analysis stays reasonably consistent.")
(b) it would involve extra photo tool costs of $50,000	Geilhufe Report at 12 (first column of table relating to the fixed burst length alternative shows that Geilhufe was estimating approximately $50,000 in extra photo tool costs)
(c) it would cost approximately $100,000 more than programmable burst length in design costs (Geilhufe, Tr. 9594:5-12)	Geilhufe Report at 12 (first column of table relating to the fixed burst length alternative shows that Geilhufe was estimating approximately $100,000 in extra design costs)

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
4. based on all options in the JEDEC standard as adopted (and not on industry usage or practice), use of both fixed CAS latency and fixed burst length would require 12-15 separate parts (Geilhufe, Tr. 9601:7-16)	Geilhufe Report at 15 ("JEDEC Standard 21-C specifies three different CAS Latency values for the SDR SDRAM, and five different values for the DDR SDRAM. Four different Burst Length values are specified for the SDR SDRAM and three different Burst Length values are specified for the DDR SDRAM. To fulfill the JEDEC Standard with fixed CL and B.L. products could require 12 or more different SDR DRAM products and 15 or more different DDR DRAM products for each current product – a major increase in manufacturing complexity, in inventory management, in distribution channel complexity, and in systems manufacturing complexity.") Geilhufe Depo. at 152:9-21 ("Q. So I understand, I had asked you if you made a determination of the number of DRAM products that would be manufactured if JEDEC went to a fixed CAS latency and fixed burst length structure. As I understand it, you said you did. What was that determination based on? . . . A. It was based on the JEDEC specification. Q. Was it based on anything else besides the JEDEC specification? A. No, it was based on the JEDEC specification.")
5. use of fixed CAS latency would not permit the mode register to be removed from the DRAM (Geilhufe, Tr. 9736:24-9737:19)	Geilhufe Depo. at 152:2-4 ("The mode register, in addition to CAS latency and burst length, has another element called burst type . . . .")
6. (a) electrically blown fuses and anti-fuses are not reliable (Soderman, Tr. 9356:18-9357:2)	Soderman Depo. at 149:18 – 150:6 ("Q. Are you saying it's no longer technically feasible to use electrically-fused DRAMs? A. It depends on what type you are talking about. If you're talking about an anti-fuse, everything is possible. The question is what is economically profitable. And that's what drives the DRAM industry. The anti-fuse technology does not have the same reliability as required for the DRAMs. They use it for programmable logic. And there is some percentage of those that we'll call regrade; in other words, have a defect either change with time or just don't get programmed correctly initially. So it's not a -- the electrically-programmed anti-fuse is not 100 percent

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
	reliability like a laser fuse would be.") Geilhufe Depo. at 86:13-23 ("Q. What would be the cost of adding that module? A. The anti fuse? Q. Yes. A. It would be totally unrealistic. Anti fuse is so unreliable that you could not use it for this application. Q. Why do you say that? A. Actel, an anti-fuse company, they guaranteed 95 percent of their product to work then they ship it to you. Can you imagine a DRAM manufacturer saying I'm shipping you parts and 95 of them are not working?")
(b) based on a survey of "maybe 50" out of "hundreds" of data sheets, only about 2 out of 50 SDRAMs appear to incorporate electrically blown fuses (Soderman, Tr. 9357:3-9358:1)	Soderman Report at 27 ("While it is true that there are other types of fuses that could be employed post-packaging, these fuse technologies are not in common use in DRAM manufacturing . . . .") Soderman Depo. at 152:18-23 ("Q. So you're not aware of any company that does incorporate electrically-blown fuses in commodity products at this time? A. I have heard that Micron and some others did it for a while but, you know, it did not catch on with other companies.")
(c) anti-fuse technology is not generally available in DRAMs (Geilhufe, Tr. 9582:20-9583:19; Tr. 9732:11-9734:21)	Geilhufe Report at 15-16 ("The most common technology is to laser vaporize conductor links at Wafer Sort (electrical test of the die on the wafer). . . . Alternative fuse technologies that would offer programming after Wafer Sort and possibly by the Systems manufacturer are not compatible with DRAM wafer manufacturing processes of the mid 1990's.") Geilhufe Depo. at 124:9-24, 127:8-14 ("THE WITNESS: We looked at anti-fuse, I looked at nonvolatile fuses of all kinds, even considered silicon fuses. . . . Q. You say they are not compatible with DRAM wafer manufacturing processes in the mid 1990s, is that accurate, is that what you said? A. That's correct. Q. What's the basis for that? A. As we discussed before, these alternative fuse technologies require process modules that did not exist or do not exist in the DRAM manufacturing process. . . . Q. Okay. And your understanding that it does not exist in the DRAM process is based on what? A. It's based on my understanding, my experience of the DRAM processes that were active in that time period. Q. In 1991 to 1996? A. Up to 2000, roughly.")

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	**Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed**
(d) the use of laser blown fuses would lead to reduced yield due to speed distribution (Geilhufe, Tr. 9585:21-9586:9)	Geilhufe Report at 11 (third column of table relating to the alternative of programming CAS latency using fuses shows that Geilhufe estimated that there would be reduced yield due to speed distribution, leading to an additional cost of $.03 cents per unit) Geilhufe Depo. at 131:24 -132:18 (Q. The next column -- the next element I would like to talk about is the final test and good unit yield for programming CAS latency with a fuse. A. Yes. Q. Now, you note in here your belief that it would lead to a reduced yield due to speed distribution, and that you were adding three cents per unit, approximately? A. Approximately. Q. Why do you believe that? A. As we discussed earlier, manufacturing process results in a distribution of speed, and fixed CAS latency has the potential of not using the entire speed range that the process turns out. And, therefore, has, if you will, potentially fewer good parts, certainly for the fast parts. Q. Is it an assumption of this column -- and by "this column" I mean the programming CAS latency with fuse column -- that the fuses are being burned by a laser prior to packaging? A. Absolutely.")
7. (a) based on the number of bits provided for in the JEDEC standard as adopted (and not on industry usage or practice), setting CAS latency and burst length via pins each would require three bits of information (Geilhufe, Tr. 9589:22-9590:6; 9599:8-9600:1)	This follows logically from the undisputed fact that there are three bits corresponding to CAS latency and three bits corresponding to burst length in the mode register described in the JEDEC SDRAM standard.
(b) it would be necessary to add pins (Geilhufe, Tr. 9724:16-21;9741:8-974 2:1; Soderman, Tr. 9362:12-9363:3)	Geilhufe Depo. at 135:22 – 136:22 ("Q. Is there an assumption there that there would be an additional pin required for this alternative? A. I just realized it's a mistake. It should be two cents. You can't add one pin -- Q. Okay. A. -- in a package, you have to add them in pairs so you add pairs. . . . Q. What is your understanding that additional pins would be required to implement this alternative based on? A. It's the definition of

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed
	Professor Jacob's report. He suggested that you add a pin possibly to provide voltages with the CAS latency information. In other words, you are programming the DRAM with an additional pin.") Soderman Report at 30 ("To allow for several latency values without using multiple voltage levels on a pin, would require the addition of two or three pins to the DRAM.") Soderman Depo. at 140:14 – 141:12, 143:6-13 (Q. Let me direct your attention to paragraph 63, first sentence, The first implementation proposed by Professor Jacob would be expensive to implement because of the necessity of adding pins. What is your basis for the conclusion that it's necessary to add pins? A. If you look at the package options that are available, and we're talking about -- the common packages at the time were frequently 54 pin TSOP, there were not a lot of extra pins, the pins in which configuration you have chosen there. If you have a by four, there are a couple of extra pins; if you have it by eight, there are fewer. If you have it by 16, I think there's one. Obviously it would not work if you're trying to encode all of this information with one pin. Q. Okay. But even with a by 16 there is one no-connect pin; isn't that right? A. A no-connect pin means that JEDEC has not specified the function. It does not mean it's not connected to something. It just means that JEDEC has not defined it. Q. Okay. A. It could be used in testing. So you're not necessarily free to go and use it for something else. . . . A. Actually, some of these parts come in a 32-bit output where there are no extra pins. They've used them all. So it depends on which flavor of the number of -- the organization of the array, by four, eight, 10, 16, 32, and we're talking about the 16 does not leave many left, there's either one or two that are not defined, not necessarily not connected, they are just not defined.")

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	**Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed**
8. running a single edge clock at a higher frequency (a) would cause significant clock distribution problems (Soderman, Tr. 9393:20-9394:8)	Soderman Report at 42 ("The tolerance margins on the very critical circuits required to buffer and amplify this [higher frequency] clock signal then would be substantially reduced . . . .") Geilhufe Report at 18 ("To maintain signal and timing integrity at double clock frequencies an on-DIMM clock chip may be needed.")
(b) would require on-DIMM clock circuitry and possibly an on-DIMM PLL/DLL, which would cost $3.80 (Geilhufe, Tr. 9609:17-9610:5)	Geilhufe Report at 13 (fifth column of table relating to the alternative of running a single edge clock at double frequency indicates that an on-DIMM clock would be required and would cost approximately $3.80) Geilhufe Depo. at 193:18 – 194:1 ("Q. Why do you believe an on-DIMM clock is required in this alternative? A. I believe to accurately position the clocks, the positive edges of the clocks and have usable clock signals, that is not disappearing clock signals, the positioning of the edges has to be more accurate. And that requires then some form of synchronization, whether it's a DLL or PLL or something.")
9. moving the DLL to the module would cost $3.80 for the DLL (Geilhufe, Tr. 9613:13-25)	Geilhufe Report at 14 (fourth column of table relating to the alternative of moving the DLL to the module indicates that the cost would be approximately $3.80 for the DLL) Geilhufe Report at 19 ("From a cost point of view, removing the DLL from chip lowers the chip cost by about $.03 while it increases the DIMM cost by more than $3.80.") Geilhufe Depo. at 211:6-13 ("Q. What is your understanding of this alternative that you base your analysis on? A. Literally removing the DLL circuit from the DRAM and putting a DLL chip on the DIMM. Q. Okay. And the DLL on the DIMM you have listed under board complexity for this column and you say that it will cost $3.80 approximately? A. Correct.")

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	**Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed**
10. SLDRAM was unable to design a high speed DRAM using Vernier circuitry, without an on-chip DLL (Soderman, Tr. 9412:22-9415:9)	Soderman Report at 51 ("The inadequacy of the vernier alternative to on-chip DLL is illustrated by the failed efforts of a group of DRAM manufacturers to develop a DRAM without on-chip DLL. . . . In order to account for differences in the timing of various SLDRAMs in the system, the SyncLink design provided a vernier circuit on the controller that would analyze the timing difference and insert a programmable delay for data received from each SLDRAM. . . . Apparently, SyncLink determined that forgoing an on-chip DLL was not a workable solution. Ultimately, when an SLDRAM test chip was built, it incorporated an on-chip DLL.")
11. because the proposed alternatives didn't include circuit designs, they were poorly thought out (Geilhufe, Tr. 9673:17-9674:5)	Geilhufe Depo. at 206:16-24 (Vernier alternative): "Again, the description was inadequate for me to do a very clear analysis. . . . I don't know what specifically Professor Jacobs had proposed. Certainly there were no block diagrams or schematics.") Geilhufe Depo. at 232:25 – 233:7 ("THE WITNESS: The analysis of the alternatives is totally inadequate for a -- let's say if my design manager came to me in my general management role with these alternatives and said decide one, I would say go take another five engineers and go to work and do a better job and find serious alternatives analyzing carefully and give me the pros and cons of each one of them.")
12. DDR II (a) expands the use of programmable CAS latency (Soderman, Tr. 9351:7-9353:3)	Soderman Report at 23. ("The current proposal for the next generation 'DDR II' standard includes programmable latency values of 2, 4 5 and 6 (optional) plus a programmable additive latency of 0, 1, 2, 3, and 4, resulting in 8 possible read latency values . . . .")
(b) initially planned to use a single burst length, but subsequently reverted to programmable burst length (Soderman, Tr. 9369:12-23)	Soderman Report at 32 ("The JEDEC JC-42 Memory committee initially proposed fixing the burst length at 4 for the DDR2 specification . . ., but now intend to provide the variable burst length feature with burst lengths of 4 or 8.")

Complaint Counsel's Proposed List Of Rambus Expert Testimony To Be Covered In Rebuttal	**Pages In Rambus Expert Reports And Depositions Where Issues Were Addressed**
(c) limits the use of the burst terminate command because of timing difficulties (Soderman, Tr. 9376:19-9377:20)	Soderman Report at 34. ("The JEDEC JC-42 Memory committee is proposing to remove the burst stop command for the DDR2 specification . . . because the timing is difficult to control.")