InterOffice Memo

To: Bill Gates

From: Nathan P. Myhrvold

Date: May 11, 1992

Subject: Topics for System Strategy

The personal computer industry is at a very interesting stage in its evolution, and this has many implications for Microsoft. I've collected a number of these thoughts in this memo. I'll state from the beginning that this is *not* a complete exposition of systems strategy, but rather a bunch of opinions, predictions and guesses about what will happen in the near future - primarily the next 24 months. It is intended as food for thought.

Whither RISC?

In the last 9 months or so we have seen some real changes in the processor equation.

MIPS was late with the R4000. In a certain sense this did not matter very much for us since our software was not ready, but it did hurt their public perception. The financial condition of MIPS the company became a bit of an embarrassment for MIPS the architecture until the point where SGI felt they had to stabilize the situation with an acquisition. Recently MIPS/SGI announced both low end and high end follow on chips to the R4000, (which are being financed largely by NEC) so the architecture is still healthy.

Meanwhile Intel has been very aggressive at claiming that they have closed the gap with RISC, and this perception has been quite pervasive in the PC industry, even among Microsoft people.

Is there a future for RISC? Will it be important in the PC business?

I believe that the issue of RISC in the PC world will now be driven by Apple and IBM. This does not mean that we are totally out of the picture, but for a variety of reasons (see below) the ACE initiative is not going to *drive* anything - among other things, nobody is in the driver's seat.

Meanwhile, Apple and IBM will aggressively move to create RISC based systems. They already have a 50 Mhz RIOS chip set that is getting between 40 and 90 SPECmarks - the only technical issues that face them are getting this into Motorola's process on a single chip and perhaps doing some re implementation to reduce cost. I believe that they will certainly succeed, and that they will be *at least 2X faster* than the fastest Intel machine available at the time.

Can this claim be believed? Hasn't RISC disappointed us so far? Didn't Intel catch up?

The notion that Intel has "caught up" with RISC is simply false - it is based on two things. First, Intel is being aggressive about *saying* that they have caught up. Second, they have supported this claim by numbers they are often quite misleading - such as using machines with a lot of static RAM cache (256K and higher) to report numbers that are often compared to machines without caches, or doing SPEC specific compiler hacks.

One amusing fact is that the supposed "catch up" did not result from any new technology - it occurred because of new marketing! The 50 MHz 486 is out, but that is

Plaintiff's Exhibit
5476
Comes V. Microsoft

hardly a major renewal of their technology base. If fact, the *very designation of the 486/50 is a typical Intel move to be aggressive with the "truth"!* The so-called 486DX2 50 MHz part has an external clock frequency of 25 MHz, and an internal frequency of 50 MHz. By this standard, the R4000 is a 100 MHz processor and the DEC Alpha is a 200 MHz processor. Although DEC and MIPS mention the internal frequencies, they follow the usual convention of naming the part by the external frequency because that is what the rest of the system sees. Intel isn't really lying, but they are certainly taking every opportunity to push things and create a favorable impression which is not borne out by the actual performance numbers.

The 586 will probably be better than the 486, and it is true that they have finally gotten around to moving in the direction of modern processor design methodology. Nevertheless, they have a long way before they have anything unique or innovation. On the basis of any real change at Intel, any notion that it will close the gap with RISC is unproven and unlikely.

One way to see this is to look at some actual data. Here is a chart of the integer SPEC benchmarks for various processors. Note that this is based on the latest version of the SPEC benchmarks, known as SPEC92, which has integer and floating point components. We'll look at integer first.

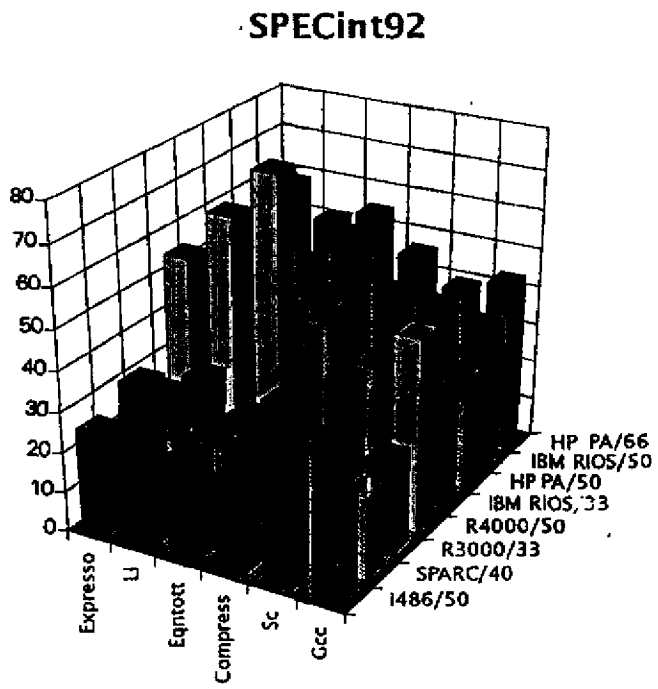

MS7059951
CONFIDENTIAL

	i486/50	SPARC /40	R3000 /33	R4000 /50	IBM RIOS/33	HP PA/50	IBM RIOS/50	HP PA/66
Expresso	25.8	21.5	25.4	55.1	27.8	40.8	41.7	58.5
Li	40.2	21.6	31.6	67.6	28.9	36.2	43.5	50.6
Eqntott	25.2	22.8	24.8	79.7	35.4	37.5	53.1	54.5
Compress	24	17.3	18.6		27.2	25.8	40.2	45.9
Sc	44.6	28.1	27		26.5	25.2	39.2	39.1
Gcc	26.6	20.9	23.8	48	24.6	27.7	36.4	42.9
SPECint92	30.1	21.8	24.9		28.2	31.6	42.1	48.1

The processor and machine details are given in the following chart. Note that I do not have all of the SPEC92 results for an R4000 machine, which is unfortunate - this is why there are gaps in the table and chart. The various designs are new enough that this data hasn't been published yet. In addition, no Alpha figures have been released. Judging from its clock rate and other features it should probably have a SPECmark above 100 in a reasonably designed system.

	i486/50	SPARC /40	R3000 /33	R4000 /50	IBM RIOS/33	HP PA/50	IBM RIOS/50	HP PA/66
Ext. MHz	25	40	33	50	33	50	50	66
Int. MHz	50	40	33	100	33	50	50	66
Superscalar	No	No	No	No	Yes	Yes	Yes	Yes
Internal Cache	8K	0	0	16K	0	0	0	0
2nd Level Cache	256K	64K	128K	1M	72K	96K	72K	512K

The first thing to note is that there are basically three generations of chips in this chart. The SPARC and R3000 chips are quite old - in each case there is a new implementation that has either just been released or is expected soon (i.e. Viking and R4000). The IBM and HP PA chips are quite recent, but are built with an intermediate generation of semiconductor technology. They are multi-chip implementations with each individual chip having a much smaller size limit than the 486 or R4000 (roughly 1M transistors). One result is that have no on chip caches. By this standard, the 486 and R4000 are the only really large chips in the group , and the *only single chip implementations on the chart.*

The overall integer SPECint92 for the 486/50 is 30.1, but most of this is due to just two programs, Li and SC. The other benchmarks are very consistently at the 25 level. This is very similar to the R3000 at 33 MHz. The R3000 machine has a cache that is one half the size of the Intel machine - and that does not even count the fact that the R3000 has no on chip cache. The cache size clearly hurts "compress" which uses large memory. The two benchmarks that are anomalously fast for the 486 might be due to cache effects. The HP machine is the only other large cache machine (with full SPEC92 results) and it gets comparable effects. The IBM machine also does as well without having a

large cache, but RIOS is superscalar and uses more internal parallelism than the others. They may have the ability to have their compiler do a lot of load scheduling to hide the memory latency. Another possibility is that there are compiler optimizations which happen to do well on key loops in those two programs - either as a happy accident, or a deliberate speed hack for SPEC.

The HP and IBM compilers would appear to have their own "favorites," although none are as dramatic as the two Intel winners. The "real world" question is how often do the factor behind the speed up in these programs come up? If the ratio is really that 2 out of every 6 typical integer programs for will get the same kind of speed up, then it is fair to average them in. Otherwise, it is quite deceptive. The *real* answer (from our perspective) would be to get some benchmarks of a reasonably large suite of Windows apps, but those are still not available on any architecture, much less in a portable form for chip evaluation.

My overall conclusion is that for many problems the 486/50 is about the same speed as an R3000, and that on systems with the same size cache the R3000 would probably beat it slightly. The implication is that an R4000 with a reasonable 2nd level cache would be 2X faster than the 486/50, which is borne out by the existing R4000 data (albeit numbers collected with a larger cache). The initial experience with the "Fusion" machine (Jazz with 2nd level cache) and NT performance bears this out. The data for the other 50 Mhz RISC chips shows a very similar result. The 486 is outclassed by about 2X or more on every benchmark, with the exception of the two "miracle" benchmarks. This is really not very surprising because the Intel chip is really a 25 MHz processor which has been souped up to run internally at 50 MHz, whereas the R4000 is a 50 MHz processor souped up to run at 100 MHz internally. HP-PA and IBM RIOS all execute at one clock rate, but they use a superscalar approach (multiple instructions per clock) which is quite analogous to running faster internally - it is the question of superpipelined versus superscalar. In any event it is easy to see why the 486 is still lagging behind.

That is the integer situation. Now we can look at floating point performance.

SPECfp92

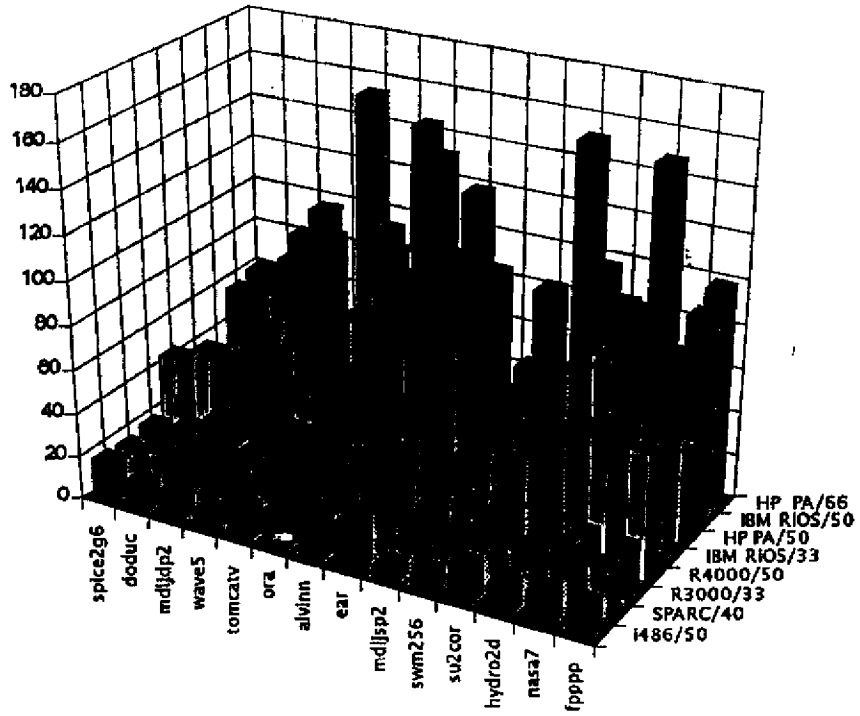

We had better hope that FP performance does not become an issue! The 486 is dead last in the set of chips above - by almost 10X when you compare to the IBM or HP. SPARC and R3000 are not much better than the 486 because they each have quite wimpy, off chip FP coprocessors. The R4000 is not quite as hard core as the IBM and HP machines, but it is in the same basic league.

The interesting thing here is that it *is totally in IBM and Apple's interest to attempt to find mass market uses for floating point.* If they can find something of this sort, they will be able to use it to cream the Intel world. One obvious example is graphics and visualization, which is quite FP intensive. Another one is multimedia signal processing - MPEG video compression etc. It is not yet clear whether these things will become vital to end users, but if FP does become an issue, the Intel based world is going to have a very tough road ahead of it.

So, Intel has not caught up, and RISC has indeed met earlier claims. It is true that some specific chip projects have been late (notably the R4000 and the Sun/TI Viking SPARC project), but in the meantime others have been early - witness the 50 Mhz RIOS and the 66 Mhz HP-PA and the 100 Mhz DEC Alpha. The real competition in the PC market will

occur when IBM/Apple/Motorola get their act together and make a modern, large chip version of the "PowerPC" architecture with decent on chip caches.

The political situation with RISC is now at a stage where I believe that this event is going to be the primary driving force - IBM and Apple will now become the prime movers for RISC in the PC industry. This does not mean that we are out of the game, but:

- They are very committed to pushing this as a PC market machine, rather than as a workstation. It is their "last best hope" for evading the phenomenon of Windows which both companies see as anathema to their proprietary strategies.
- Apple will have an opportunity to use a port of System 7 to bring applications to the machine quickly. There is no dependency on Pink.
- The primary thing which will motivate OEMs in our camp (and us for that matter) will be the threat from Apple/IBM. Until this machine ships and the threat is concrete, there will be little RISC activity in our world.

The basic observation is that Apple and IBM will be aggressively pursuing a RISC based strategy, and this is going to be a *lot* more focused than our MIPS activities to date. We have been hampered by many things - including the fact many of the hardware companies that we were counting on have self-destructed (Compaq, DEC, MIPS...) and finally, our software hasn't been ready.

The biggest question is who will push the MIPS based Windows world? To date, we have not been very proactive in doing this. Admittedly there is only so much we can do before we have a product to ship, but even so we have not been laying the groundwork for a big push in the near future. We are going through the motions of supporting MIPS, but that is all.

The fact is that simply developing NT is a quite a task and I am not sure that we have had or will have sufficient bandwidth to focus a lot of energy on the unique challenges which face the MIPS world. In my view our attitude toward the MIPS based world could be summed up by saying that we *privately* want it to be a success, but we are not sure enough about what we are involved or what we should do to take more action than we have done. We'd think it was cool if somebody else would swoop in from nowhere and made it take off, but in lieu of that we're too busy to get really involved.

Here is what I think will happen:

1. IBM/Apple/Motorola will introduce a hot chip. Apple will base a hot new Mac on it.
2. They will promote this strategy to other OEMs. I do not think that they will create a truly open market, but they will attempt to get some others involved. Apple will license a ported version of System 7 for the machine, and also dangle the specter of IBM/Apple PowerOpen UNIX and Pink as well. It is entirely conceivable that they would enlist somebody like Dell or Zenith (through the IBM/Bull deal) and perhaps some Japanese company, or even what it left of Compaq.
3. Any of our OEMs who do not decide to go with IBM/Apple are likely to want a response. We will need to point them all toward a RISC strategy.

In the past I have argued that the best strategy for us would be to have a healthy RISC market going, or at least launched, before steps 1 & 2 occur. I still think that this is optimal, but I do not think that it is possible at present for a variety of factors - some due to us, but many due to external events as will be discussed below.

Failing that, we need to have a RISC alternative for 3 above. MIPS is still the only really viable choice and as a result we should support it within whatever niches we can. The

high end graphics niche is an obvious point of strength. SGI utterly dominates this market, and if we work with them effectively, it could be a safe harbor for NT. Sony has recently claimed to us that they plan to have a 75 Mhz (150 Mhz internal) version of the R4000 in a $4000 PC by the middle of 1993. An aggressive machine like this could initiate "grass roots" or "bottom up" support for RISC in the user community. We should certainly support such efforts - with NT itself, and also with Excel and other key applications.

Alpha is quite unlikely as the choice for step 3 above. The fact that it is 64 bit only makes it quite tough to justify for a mass market PC machine. In addition, DEC is a dead company (see below) and we do not have sufficient architecture rights to be able to have Alpha survive DEC (unlike the MIPS case). Note that I say that Alpha is unlikely - it doesn't mean that it is utterly impossible.

There is also the possibility that MicroUnity or some other dark horse candidate will pop up and save the day. We should be open to this sort of possibility, but we should also realize that it is not anything to count on.

There is an obvious strategy of porting NT to other architectures, including the PowerPC machines. I think that it is just as obvious that it is a terrible idea *as a strategic alternative* (the technical issues are minimal). Having it on PowerPC might be interesting as a "spoiler" gambit in addition to having it on a real alternative, but I think it is the height of folly to *rely* on this as our strategic defense against Apple and IBM. The whole reason that these guys are pushing their stuff is to try shut us out, so the chance that we will get a fair shot at the OS market for this machine seems unlikely at best. Porting to other random machines (HP-PA etc.) amounts to further confusion in the market before we come to do 3 above.

If a chip is a serious candidate for creating a mass market binary standard, then and only then is it suitable as a defense against IBM/Apple/Motorola. In that case we can consider supporting it. If it *isn't* a candidate for this purpose (which includes whether it is open, do we have the necessary architecture rights to keep it open etc.) then it is a real *drawback* to support it. The first thing we will need to do when we start *seriously* promoting one of these puppies as the new top dog is to drown the others. We can't get consensus in the industry or a binary standard to be based on *two* chips (at least I don't see how) so once we pick our favorite RISC we will have drop the others or position them safely in backwaters. The fewer of these we need to drop, the easier we will have it.

How will the RISC world play out? My guess is that when 1-2 above actually start to happen in a major, publicly visible way, we will get excited and finally become hard core about RISC. With a little luck we will have at least some ongoing level of MIPS activity at that point and we can fan the smoldering embers into some kind of backfire to stop IBM and Apple from getting major momentum. With a *lot* of luck, MIPS will be reasonably strong, and/or a dark horse candidate will have come in from nowhere. If luck isn't with us, then we could face some very serious competition in a long, drawn out battle.

Without the work that we have done in addressing RISC up to this point, I think that we would be extremely vulnerable. Originally it was an insurance policy against SPARC and other near term RISC worries. It served that purpose admirably. The new role of RISC in our strategic line up is as a weapon to counter the IBM/Apple threat. Ideally it would have shut them out by getting popular before they came out. Now it seems more likely that it will be in reserve for when they do get their act together.

The Hardware Hospice

Some very fundamental changes have occurred in the industry and we are beginning to see the effect this is having on firms that dominated and built the computer industry in earlier days. All around us there are hardware companies going belly up. Microsoft is the primary beneficiary of much of this change, but ironically I do not think that we have really come to grips with the full implications of what this means. A roll call of the dead and dying would comprise the largest and proudest companies in our industry, and this will cause some enormous changes for us.

DEC is as good as dead. I think it is well within the realm of possibility that our recent strategic deal with them will rank alongside the Ashton Tate and 3Com strategic relationships, although I admit this is a somewhat radical view. Note that "dead" doesn't mean that they will instantly disappear. A company with the name "Digital Equipment Corporation" will probably exist for some time, but they will never again enjoy the market size or influence that they have had in the past. The very fact that Wang and Unisys still exist (in name only) shows that it is very hard to actually disband the final remnants. Nevertheless, I think that it is fair to say that a hypothetical DEC of several years from now that trades on its former glory as its twist on the mail-order PC business is not the DEC we have known over the years.

Alpha is in some sense the epitome of their downfall. They had something very much like this *six years ago*, and repeatedly screwed the project. Now they have made a huge public show out of the stuff (with some updated technology), but by all appearances they lack a delivery vehicle. Their software stc., is amazingly weak (except for the system written here by the guy they spurned!), and they may not be quick enough at delivering Alpha in hardware systems. Alpha contains good technology, but that only highlights the other problems they have - they didn't really take it seriously enough soon enough to get the supporting components in place. Now they have decided it is their crown jewel, but will not be able to apply it. This ultimately goes back to their serious weakness in top management. Convulsions like the one that toppled Strecker will probably occur repeatedly until they have shaken off the last remaining people who could lead them out of the present mess.

Compaq is also finished. They may continue in business for some time by being a slightly upscale version of Dell, but their days as an influential trendsetter in the PC business are over. There is just no way that they could take the lead again now that they have fallen off, and there are a thousand ways that they could drop further in the market.

IBM is not dead yet, but they have been diagnosed as HIV positive and it is just a matter of time. Maybe the miracle drug will be found before they totally succumb, but the odds are against it. In the meantime they can be very dangerous, but they will never again be the IBM that once ruled supreme over the computer industry. Within the PC world they insisted on putting their power and prestige explicitly on the line with OS/2. The rhetoric behind the OS/2 2.0 versus Windows jihad they instigated makes this very clear to all segments of the PC industry. As users vote for Win 3.1 with their purchase orders, and ISVs vote with their application "design win" decisions, it will hasten the demise of IBM's credibility.

Their mainframe business is in slightly better shape, but unless they pull a miracle out of the hat they will be in real trouble. The trouble that they face here is that microprocessors are far and away the dominant technology for *any* computer large or small. MP machines of varying sizes from 2 to 1000 processors will provide much

better performance at lower prices. The number of people gunning for them here is enormous and they will be hard pressed to compete for long.

I think that the next 12 months will be quite critical to IBM within the PC business. First, they will have to face the fact that OS/2 is a failure - denial, repositioning and rationalization can only go on so long. One theory says that they will carve out a modest business in some select niches - big corporate customers etc. This can buoy them for only so long because their cost structure and company wide level of commitment to OS/2 can't be supported with this level of success. They would amount to being a bloated version of Quarterdeck, trying to sell a bloated version of DeskView.

The second major event is that the IBM/Apple/Motorola zaibatsu will bear fruit in the form of a hot chip and (in fairly short order) a hot new Mac. *What will IBM do with it?* In particular, what system software could they use? They won't have a portable OS/2, and even if they did what good would it do? They can use PowerOpen (a.k.a. AIX), which will let the Austin workstation division sell a cheaper version of the RS/6000 and potentially give Sun some grief, but it is hardly going to make a dent in the PC market. Of course, Apple will be right there with PowerOpen competing with them so they can't run too far with it. They could wind up with System 7 licensed from Apple, but how on earth would they rationalize *that* and maintain any prestige or credibility?' Pink won't be ready in time, but my guess is that either this, or some equally insane plan to mix AIX and OS/2 in a new system is likely to be their current Plan Of Record. I think that it is quite significant that *system software is the critical element in this situation.* They have shown an amazing degree of unreality in how their executives make assessments of OS/2 software progress and schedule and the *same* people are the ones making this decision.

Both of these challenges carry twin threats - they are damaging to their business *and to their pride.* In the first case they explicitly hung their credibility on achieving goals that are certainly impossible. In the second case, they will look like utter fools because their bold new plan to work with Apple will wind up shooting them in the foot. How will they deal with this? There are many rational but painful ways to deal with the situation, but the cure will be hard for them to swallow. There is a substantial chance that they will not turn to reason until it is much too late. The pattern they have been following would be to do what blackjack players call *double down* - i.e. stick with what they have and make another Big Macho Bet.

Across the industry you find the same pattern time and time again. The fundamental issue is the message of the *computerless computer company.* A more accurate statement is that *hardware manufacturers are no longer the architecture supplier.* The "architecture" is what end users really buy. It comprises the world of compatibility for application software and it is a very valuable intellectual property asset. It used to be that you could extract a fee for this in the form of a nice fat margin. Applying the "architecture vigorish" or premium of X margin points against a large hardware sale was certainly a nice business, but technology eventually turned against them. In today's world the "architecture vig" is unbundled from the hardware and instead is sold as a pure intellectual property asset which is a fixed fee per system. No more proportional uplift, and hence no more big high margin hardware companies. The loss of architecture agenda and the multiplicative premium that goes with it is the comet that will kill this particular set of dinosaurs.

There are two factors which is causing the ownership of architecture to slip from their grasp. The first is VLSI microprocessor technology. You just can't build a fast machine *of any size* without using microprocessors. Even Seymour Cray's company has wound

MS7059958
CONFIDENTIAL

up on the endangered species list, in large part[1] because massively parallel machines from a wide variety of vendors are creating computers which are faster than his monolithic computational leviathans, at lower cost and more flexibility. If you can scale from 1 to a 100 to 1000 processors then you have a big advantage. The reason that microprocessors hurt the big hardware companies is that it put the CPU architecture in the hands of chip makers - which is practice means nearly anybody. Poor DEC found that any bunch of grad students could make a chip faster than VAX - and not only could they, *but they did*, licensing the resulting designs to everybody who owned a chip fab. The economies of scale are quite different than the discrete CPU business - a small group can *design* a chip on a shoestring, so nearly anybody can. Once you go to *manufacture* the chips it is just the opposite - you need a fortune for your fab, so it is hard to justify keeping the chips to yourself - you are much better off selling them as a standard part on a wide scale. Neither scale fits with proprietary architectures, so CPUs with wide availability became important.

Software has been the other real killer. Economics are once again the critical element - it makes far more sense (i.e. *more money*) if you unbundle software from hardware and treat it as a separate business. Once you do this, the effect of installed base and standardization drives the market. Third party software is better than what you can write for yourself. No sooner have you made that leap, than they the next effect kicks in - third party software written in a large competitive market tends to be better, and have more variety, and fit more user needs more responsively than software you write yourself, or which comes from smaller markets. This is because there is more business drawing the software developers, which makes the stakes higher, which eventually translates into more resource being put on the problem - either at a single developer, or across N competitors shooting for the same market.

The inevitable evolution is to the point where mass market binary software (and therefore software compatibility) calls the shots. If it hadn't been for standard, mass market microprocessors, mass market software wouldn't have been possible so this clearly *started* the trend. Once the software compatibility snowball gets rolling down the hill and an installed base forms, it soon dictates the fate of microprocessor architecture rather than the reverse.

Apple is an exception to the mass extinction, and will be discussed below. Sun is another exception, but interestingly enough they have used a software oriented view to achieve this. They get a small architecture premium on the hardware precisely because of their software. Even in purely hardware oriented issues like the SPARC chip itself, they have taken a software oriented approach. They invented the concept of an open chip architecture and have used it to good effect.

For the first time, the computer industry will be primarily *a manufacturing oriented business* - I assert that it hasn't really been in this situation up to this point. You can make cool hardware just as long as you compete on price and features rather than on lock in or brand name. As I have mentioned in the past, this is *not* a "commodity" business in the normal sense because the diversity of hardware will increase enormously. Without the big fat industry leaders, the "clone" market will ultimately go away - who are they cloning after all? Instead we will see a very diverse market with no clear leadership. OEMs will copy each other, and will individually innovate in their own preferred areas, but there will be no clear cut direction. Windows is the unifying feature that makes this all possible because the user is buying a Windows Machine, not a brand X.

---

[1] World peace and shrinking government budgets are partly to blame, but in general the folks who pulled out of buying the Cray 3 are still buying MPP machines.

To some degree, the current UNIX workstation market is actually an example of this sort of market in action. People within this world do not typically buy a proprietary machine in the old sense - instead they buy a "UNIX workstation". In doing so they sign up to the task of porting the software, which is a pain. If the pain is too great, you stay with Sun, even though they are not the best in *any* hardware category (price, performance, graphics, servers...). If you are a speed demon you switch every couple of years depending on who is hot at the moment. If you want a particular feature you might settle on a single favorite, such as high end graphics with SGI. Essentially every single company in this market has their own instruction set - both for the big boys (Sun, SGI, DEC, IBM, HP...) as well as the losers (Tektronix and Data General with 88K, Intergraph with Clipper, Oki and a few others with i860...). They each would use the word "architecture", but it practice it doesn't really mean the binary compatibility standard that I have used here. In fact, EACH of these companies has had a major change in instruction set *without* losing any significant number of customers - in fact, I suspect that they each actually *gained* at the time they switched (IBM to RIOS from ROMP, Sun from 68K to SPARC, HP from 68K and HP-PA88 to HP-PA90).

The only two companies which took a true "own an architecture" approach in the strong, binary compatible sense were Apollo and Sun. As is typical in such struggles, only one could take the category and Sun did. The amusing thing is that all of the companies are looking for the architecture win, yet this very fact has caused them to prevent it from occurring. Today we find lots of different instruction sets, but early in the UNIX market many people had the same chip - the 68K. In order to get the differentiation they craved they focused on incompatible proprietary versions of UNIX. A neat cross over has occurred - most of the player have thrown out their own random versions of UNIX to go with a standard version (usually OSF) at the same rate that they have thrown out standard architectures and invented their own random instruction set. There is probably some economic law here about the "conservation of net proprietaryness" - as one barrier drops another one is raised at the same time so as to maintain the same degree of market fragmentation. The nice thing about OSF from this point of view is that instead of having hardware companies differentiated by their UNIX version (which wasn't much fun for hardware guys) the switched to differentiating on instruction set.

The basic situation is that there is a unifying hunk of software (UNIX) which is common to everybody. OEMs create their own hardware bells and whistles to compete with each other, and this is tolerated precisely because the software allows users *to a certain degree*, to ignore the difference and have a common set of applications. Because the bells and whistles include the instruction set, this must occur at the source level, so it can never support a true mass market, which requires binary software distribution. The effort and time lag due to porting and the pain of having some software be unavailable represents a kind of "friction" for users moving between different *brands* of machine. The aggregate friction is too high to allow the PC software business model to take off - momentum becomes damped too quickly. Normally you would expect one company to pull ahead (as Sun has tried) but the parameters appear to be balanced well enough that this can't happen. There is a sort of equilibrium whereby the frictional losses generated by instruction set differences within the OSF camp are balanced by the fact that Sun is behind the performance curve of the hottest OSF company at any point in time[2].

The analogy to the future Windows world is that a similar thing *will happen with the non instruction set aspects of the machine* - video resolution and performance,

---

[2] Other effects also come into play. Too much of the workstation user base is willing to write their own programs, because those who really crave mass market apps defect to the PC world. The economics of the distribution channel has an effect as well.

motherboard, form factor, battery life etc. There will be a lot of different Windows machines, but they will share binary applications compatibility via Windows (and device drivers). This is a *far* stronger form of user indifference to hardware - there is a lot less brand mobility friction.

Is this good for us?

In the short term the answer is yes. It is also very good for our customers, who will get great prices and an unprecedented level of hardware innovation. No more waiting for IBM before the next video resolution or the next bundled motherboard feature.

In the long term we are going to have to deal with the reality that *in a world with no hardware leaders, OEMs will have no leadership other than that which we provide.* This is a responsibility that I do not believe we are prepared to shoulder at this point in time. We must develop this capability because nobody else is going to lead our OEMs to new architectural features, instead they will mill around making incremental steps while somebody else.

ACE in the Hole?

The ACE effort happened to straddle a period of time when the enrollment in the "hardware hospice" increased quite dramatically. At the onset of the negotiations which ultimately lead to ACE, there was a fundamental belief in the notion that a couple of key hardware companies were the critical "king makers" who would determine the success or failure of a new hardware architecture standard. We did everything we could to prevent Compaq from going with SPARC, and to assuage their concerns that Microsoft was getting too big for our britches in agitating for a standard. From today's vantage point we see that neither of the major companies involved (DEC and Compaq) is in a position to dictate standards to anybody. Each are hemorrhaging money and executives, and each has dropped out of ACE, either formally (Compaq) or in practice (DEC). None of us foresaw the magnitude of the problems that Compaq and DEC faced, and I am not going to focus on that aspect of the issue. The interesting thing to me is the state that this leaves things for the future evolution of hardware architectures.

I believe that the ACE effort marks a turning point in the industry. The perspective of the entire enterprise was that hardware companies, lead by a couple of powerful and influential companies, were going to actively create a new architectural standard. Microsoft directly caused this to occur, but after a certain point we were compelled to (in Hugh Barnes' words) *"stop driving the car, get in the back seat and shut up"*. We did this, even to the point of sharing the system software honors with SCO - of all people - primarily because Compaq wanted it this way.

My question to Hugh Barnes at this point is - *who's driving the car now?* Not only did the leaders of ACE enter the hospice, but all of the factors discussed above argue that they were the last of their generation, and we are at the end of an era. *We may never again see another large powerful PC hardware company which is able to set the architectural agenda.* In fact, you probably could replace "PC" with "general purpose computer".

There just isn't anybody who is going to play this role - because *we* are filling it. The technology won't allow it - at this stage the only things which determine application compatibility are the operating system and the CPU. There isn't room for a systems vendor to exert enough ownership of the architecture to extract a premium.

There will be *large* hardware companies, and there will be *successful, growing* hardware companies, but they will be based on a different model. If you don't own the architecture and charge a premium or "architectural vigorish" then you wind up with a

different mentality. The notion of investing a lot of money with a large potential loss associated with it is crazy if you are a low overhead operation like Dell or Gateway. Look at the effort that Apple had to expend to launch the Mac. The only way that this kind of Big Bet makes sense is if you are going to get enough head of the game that the payoff is large. This doesn't happen in the new style of PC marketing, because everybody is a low overhead box builder, and they are not structured for this kind of business. In the passage above I said *"is able to set the architectural agenda"*, but you could just as well add *"or would want to"*. The new model hardware company is based on aggressive marketing, low overhead channels and incremental improvements.

The new model hardware company is ideally suited for incremental evolution, and for the era of the *Windows Machine.* They can do plenty of value added engineering - in fact *even more so than in the old days of strict clones* - new graphics cards, sound support, multimedia gizmos like motion video, better power management and other things of this sort are the purview of Windows OEMs for the next five years at least. This will be a terrific era for incremental variation because many of the old constraints - like the bad old days when the PC world waited for IBM to move from EGA to VGA, simply do not exist.

Unfortunately, there is a fatal flaw that lurking behind this happy picture. *Who's driving the car when it comes to non-incremental innovation?* This is a very serious problem, because without some mechanism of this sort we will be vulnerable to competition from outside - particularly from Apple, but more generally to companies outside of the Windows world who make a cool new feature.

Microsoft is the only source of leadership in this new world and we will have to step up to the plate and accept this responsibility. We will have to lead the next major architectural innovation within the Windows world, especially the next change of instruction set to a RISC machine. This is going to require enormously more effort than any hardware evangelism we have done in the past. It most emphatically does *not* mean taking the sort of path we took with ACE - Microsoft will have to be out in front pushing the *architecture itself* very hard. We will be able to get hardware companies to make the machines and offer them in their mail order ads or superstores or other mid-90s channels but *we will not be able to count on hardware companies to do proactive work to stimulate demand for the machine and to sell people on the concept.* This is going to be a big change for us. In fact, I do not think that we even have a good grasp of the magnitude of the task because it is so beyond our experience.

The Windows Hardware Engineering Conference is an excellent example of a program which will help guide the *incremental* aspect of architecture innovation. It is terrific, but it isn't what I'm talking about here. The MPC consortium, and the work we did to stimulate pen machines are other examples of great hardware oriented projects, but the work in establishing a new major new architectural feature is going to make these pale by comparison.

Apple Strikes Back

The one glaring exception to the generalities about hardware companies in the discussion above is Apple. In many ways, Apple in as old style computer company - more like IBM and DEC than Compaq ever was. They have successfully maintained a lock on the architecture asset by controlling both the system software and the hardware, and they have extracted a huge premium as a result. They have been more true to the cause of proprietary control than any other company in the industry.

Unfortunately, for the last several years Apple has not been a very well run company from a technical perspective. Their record at pushing the edge of the envelope and

going into new areas is pretty pathetic. The success of the PowerBook machines only underscores how much money they left on the table for not having these machines three to five years ago - which was clearly possible if you look at what Toshiba, Grid and Compaq did in the PC world.

Even worse than the lack of innovation into obvious new areas, is the neglect for the core technology. The Macintosh line is in serious trouble from a *software* perspective because there does not seem to have been any focused effort to either build a unified successor or to inject new technology into the Mac itself. In hardware they also have a problem, but it isn't as severe. The PowerPC based Mac is likely to be nice, but it sure is late - Apple has been doing things like messing around with their own architecture, or planning on the 88K for a long time. They really should have been out there whipping us with a RISC Mac already.

Returning to software, the best things that have been added to the Mac system software in recent memory, Hypercard and Quicktime - were the spontaneous actions of individuals - not the product of any strategic initiative. The real cincher for me was the recent "Macintosh is not dead" announcement - when a company starts issuing press releases like that, you know *they* think they are in trouble.

Apple has done a lot of interesting work on Pink, but it is *not* a renewal of the Mac OS. It is not compatible with System 7 it in any strong sense, so it would make a quite clumsy successor - needlessly so in my view. In addition, Pink is out of Apple's hands, having been used as the dowry in the IBM deal.

Newton is a very similar story. It will undoubtedly be cool stuff when it comes out, but it too is incompatible with the Mac and does little to shore up the core business. It might get Apple launched into a new business, but from what we know today it does not appear to leverage the Mac any better than a Go machine or Pen Windows machine might. I suspect that most of the consumer electronics projects that they are working on fit in the same story - individually cool, but without any overall synergy. *Apple is a company that has some interesting technology, but no technical strategy to coordinate it, or to focus it in the areas where their business needs it most.* This is particularly true in software, which is an interesting situation considering that they keep declaring that they are "really a software company"

This makes Apple a very dangerous company for us - the technology they have gives them some technical weapons to use against us, while the lack of any coherent strategy means that they might wind up using them in all sorts of crazy ways. This could hurt us even if Apple winds up on a path to failure because they don't understand what business they are in, or want to be in.

Specifically, Apple is likely to try to move to something closer to our model by licensing their technology on a wide scale. One example is to make a RISC based Mac and then license both the hardware and the software. Another is to move into pen computing aggressively with Newton. The consumer electronics activities is yet another front on which they hope to seize the high ground. The good news is that they do not seem to have any discernible underpinnings to these strategies - they are almost in a panic mode. The bad news is that they could be very damaging - perhaps even more damaging - in a panic then they would if they knew what they were doing. If one of these wild experiments works out, they could become formidable competitors.

Countering Apple in these areas is easy in principle, but a lot of work in practice. The secret for us is to make damn sure that we *do* have a comprehensive strategy. *We cannot afford to ignore having synergy between projects*, or between markets such as consumer electronics and our core business. One reason for this is that it is the Right Thing - in the long run it is much easier to win if you have a master plan with synergistic components. The other reason is utterly pragmatic. We haven't had

platoons of people fooling around writing randomly incompatible-but-cool pen systems in Lisp dialects, or incompatible-but-cool object oriented operating systems. I'm not complaining, mind you: I think that we are spending our research and advanced development money a lot more wisely than they have. Nevertheless, if we want to dazzle somebody with hot new systems *we had better make good use of the stuff we have!* Fortunately this is precisely what our consumer strategy is all about.

So Long, Intel

The final company to discuss is Intel, our "partner" in the PC standard. Last year I started saying that within five years Intel would be out of the processor business. Afterwards I did wonder a bit about whether I was being too extreme, but those momentary waverings are now past - Intel is making excellent progress at abdicating their leadership position in the industry. This is the first step down the road to either leaving the processor business, or at any rate becoming a radically different kind of player than they have in the past.

Intel has already lost key portions of the market. AMD has taken over 80% of the laptop 386 market by some estimates. The Cyrix and C&T efforts are each quite dangerous, because TI might buy or deal with either one of them and enter the processor market with a bang. TI has a reputation for operating their chip divisions at very low margins and bombing the price of nearly everything.

The basic scenario that is unfolding was described in detail in the old memo Trends in the Microprocessor Industry. The power of an open architecture is amazing, and that is what the x86 world has become.

My guess is that the following things will happen to Intel:

- **They will dump a lot of money unwisely into R&D.** This has always been their pattern with other technologies that they lead in the past - they always try to spend their way out in the 11th hour, and quit amid heavy losses. I think that it is simply too late for them to pull a rabbit out the hat. Good R&D might help them in some areas, but what could they possibly do which would stop C&T, Cyrix, AMD, TI and a ton of others from grabbing the bulk of the market? The new R&D is only relevant to new high end versions, or new low power or integrated version etc. - *not* the mainstream core of the business.
- **The bulk of the market will go to x86 clones.** There is just no reason that 386 and 486 chips can't be made by a *lot* of people. Many companies have now the up front technical work, so it is only a question of turning the crank on execution. I think that it quite likely that Intel's market share in the mainstream middle of the road segments will drop to 30% or less. This is Intel's bread and butter, but they have no advantage in this area.
- **The clone x86 world will pioneer many valuable new niches.** The AMD low power fully static chips are an example. Integrated chipsets for handhelds and pocket machines are another example (the C&T PC on a chip is an example here). Growth in these new segments will expand the total market, but most of this will go to new x86 suppliers. This is terrific for the x86 world, but the net effect on Intel is that their market share of x86 (and hence influence and position in the industry) will drop even *faster* than their actual revenue will drop.
- **Intel will stick initially to the high end of the x86 market.** They do have an advantage in the fast chips - in part because they are good at high volume manufacturing of large chips (the 586 has 3X the transistors of the R4000 or 486), and in part because it is very hard to make the x86 architecture go fast so Intel's

resources and an advantage (just ask Nexgen). The trouble is that these are not fast chips on an absolute scale - they are only fast within the x86 world. Otherwise they are 2X slower than comparable RISC technology (or the same speed and 3X the transistors). As RISC becomes available in the PC market (either via NT or IBM/Apple or...), they will find the high end x86 world to be under price competition from chips with much lower manufacturing costs.

- **Competition from RISC and x86 competitors will trim margin from the high end.** In the market of the last few years there was little relationship between the manufacturing cost of an x86 chip (for x > 2) and the price to OEMs. Intel used to boast that they could afford to put 3X as many transistors as anybody else because COGS never was an issue for them. Competition from Cyrix, Nexgen, AMD and others at the high end will force Intel toward more realistic prices. Meanwhile RISC, such as the R4000 and the IBM/Apple/Motorola PowerPC chips, will put a cap from above. The 586 is likely to be caught in a squeeze between the two, which is not condusive to extracting Intel-style margins. The once lucrative high end will be a tough business for them.
- **They will fail to attract support for new architectural enhancements.** They currently are talking about things like adding a 64 bit RISC instruction set which they will position as the "64 bit x86" as part of a hybrid chip - say in the P6 or 686. By the time this is out, nobody will support it. Indeed *any radical new feature which is software visible will wind up languishing for want of software support, because Intel will soon no longer be able to dictate the definition of the x86 market.* At best they will find themselves in the position that IBM did with the PS/2 and Microchannel - they can try to go proprietary, but the rest of the x86 market will just thumb their nose at them because most users and ISVs will stick with the current x86 definition. At worst, Intel will find that their chips won't sell because there they have devoted millions of transistors to a feature than nobody uses.
- **Low power RISC will threaten the x86 at the extreme low end.** The MIPS-per-milliwatt issue will became crucial. It will put a lot of strain on the entire x86 market, but Intel will be the hardest pressed.
- **Intel will ignore the hand writing on the wall.** It will be very difficult for them to realize that the end is drawing near until it is too late to do anything about it. Even in the areas where they *do* realize that there is a problem, they will find in most cases that it is already too late.

Is this good for us?

The answer is *YES!* I think that this is an extremely good thing for anybody in the x86 world *except* Intel. Competition has dramatically effected the price performance ratio of machines which certainly helps Windows, Windows apps and Windows end users and the OEMs that supply them. If the scenario above makes sense, then Intel will driven to be ever more aggressive with pricing, especially in making the fast stuff get cheap quickly. As an example, the 586 will probably come out with quite reasonable prices because Intel will want as many people as possible to move up to it, because the 486 market is likely to be quite competitive by this time next year.