COMPUTER ARCHITECTURE AND DESIGN

B W LAMPSON

Rapporteurs: Richard P Hopkins and Michael J Elphick

DISCUSSION

Rapporteurs: Richard P Hopkins, Michael J Elphick

Lecture One

During the presentation itself Professor Rabin raised the question of the size of storage required for local view, to which Dr Lampson replied that the formalism was dealing with things at a very abstract level, for which such questions were not significant.

Much of the discussion concerned the question, initially raised by Professor Tanenbaum, as to whether the proposed approach was really a step forward, since it needed "wizards" to be able to handle it. Dr Lampson responded that wizards were in any case needed to implement abstractions such as critical regions even when the underlying primitives were more usable. He felt that so long as they had reliable methods for producing the magic this was perfectly satisfactory. Professor Rabin asked at what level did the wizard have to exercise the magic - in the operating system or in the program? Dr Lampson replied that what the wizard needed to produce was, for example, the code for entry and exit to a critical region, which would be in the implementation of a programming language providing that construct, not in the operating system.

Professor Ercoli then pointed out that one often needs shared memory interaction between independently produced programs, to which Dr Lampson replied that if they shared some area of memory then they were not independent. He amplified this by saying that the caching would work at the level of virtual addresses, not physical, and thus there would be some impact on the operating system, namely the need to, in effect, flush the cache when switching contexts.

Professor Anderson then commented that his reserve Chairman's remark would have been Professor Randell's standard question as to how the mechanism could be made recursive; but this had regretfully already been answered in the presentation itself. Professor Randell said that his actual question was something else. He observed that the proposed approach was analogous to that taken by Burroughs in the design of the B5500 where certain functionality was omitted from the hardware and left to the software on the assumption that the only software used would be that provided by Burroughs. He felt it might be seen as a rather perilous approach, and that system designs should minimise reliance on delegation upwards to an uncertain software environment. Dr Lampson agreed that there could be some dangers in that direction.

Finally Professor Knuth raised the question of whether the proposed system was able to recover properly from phenomena such as loss of power. Dr Lampson replied that this was not addressed directly - it is an independent issue which must be dealt with in any implementation, for which there were a variety of approaches.

Lecture Two

Professor Randell made the point that the steady 40% per year increase in performance of everything suggests that the balancing of processor/memory performance, which is usually considered the prime issue, is in fact a non-issue. In response, Dr Lampson said that to first order, these performance differences were not significant, but that there were some second order effects. Most advances were driven by improvements in silicon technology, but on the whole, magnetic technologies had also kept pace.

Professor Knuth queried whether his concept of 'mems' would continue to be a good measure if main memory technology did not keep pace with processor speeds. The speaker's view was that more levels would be added to the memory hierarchy (including two level caches on-chip), and these changes would require serious consideration of the

best way to handle locality. In reply to a question from Professor Rabin about the figures given for multiprocessor communications, Dr Lampson felt that they were not dramatically different from those realised in practice. Communications performance would certainly be reduced in overload situations, but the trend was to more general-purpose interconnection systems, rather than optimisation for special cases.

Professor Nygaard's query about backup media led the speaker to assert his confidence in the continuing importance of magnetic technologies; although the continual decreases in cost difference between magnetic and semiconductor technology would lead to the former only being used for backup in case of power loss. For example, a number of telephone answering machines now being sold used battery-backed RAM exclusively (for digital storage of speech).

Mr Thomas queried whether the figures quoted for communications bandwidth over LAN's or between processors could really be sustained for long periods. Dr Lampson pointed out that his figures were only concerned with the transport mechanisms; whether these rates could be matched by the producers or consumers of the traffic was another matter.

Professor Tanenbaum questioned whether the real distinction between shared memory multiprocessors and distributed systems with local memories had been brought out. The speaker did not agree that the organisation of memory was the essential difference between centralised and distributed systems, and felt that the important difference was in how one needed to deal with failures. In a centralised system it was acceptable to roll-back and re-process, whereas in a distributed system that would be unacceptable. Professor Randell expressed his support for that viewpoint.

Professor Rabin remarked that there were many intellectual challenges to be answered, e.g. lock management in shared databases. Dr Lampson agreed -- the need for updates (as opposed to read access only) leads to complications.

Mr Kerr asked whether the challenge of parallel programming should best be dealt with by basing one's efforts on a carefully designed library of standard algorithms? The speaker, however, had no particular views on this. Another query from the audience raised the question of the need for profiling in order to make reliable predictions. A fair amount of work had been done, and Dr Lampson believed that prediction was significantly improved with profiling data (which was not difficult to obtain). In reply to a query by Professor Rabin, he said that the effect of pre-scheduling on something like the Multiflow VLIW performance could be as much as a factor of 5.

There was some general discussion about the most effective way to exploit the increasing number of devices per chip. Dr Lampson commented that a watershed had now been reached. The strategy of putting on a single chip more powerful existing processor designs has been exhausted since it is now possible to put the whole of a CDC6600 on a single chip. He felt that inevitably the direction would be single-chip multi-processors. Professor Tanenbaum noted that the usual trade-offs between time and space still applied in the question of whether one put multiple processors on the chip or a single processor with a very large cache. Dr Lampson agreed, but commented that it would probably not be effective to have more than 90% of the chip area as cache.

Professor Rabin wondered whether "the ghost of wafer-scale integration" would re-appear, as it became more feasible to get a reasonable number of processors and memories on a wafer. The speaker noted that this particular ghost was constantly being revived! His own engineering judgement was that there were many better ways of improving inter-chip communications bandwidth, and that this approach was attempting to fight the laws of nature. However, people would undoubtedly try it yet again!

Professor Randell suggested that many of the major advances in computer architecture could be interpreted as the special hardware implementation of simple skew conditional

branches, quoting paging, interrupts, and overflow detection. He asked if the speaker saw any new approaches on these lines. Dr Lampson said that transactional memory was indeed one such.

Professor Ercoli (referring to an earlier question) felt that it was most important to know whether future parallel architectures would be general-purpose; could we continue using standard algorithms, or would each problem be best approached with a specialised organisation? Dr Lampson disagreed strongly: with so many levels between the hardware and the user interface, there could be lots of support for uniformity at higher levels. There was no need for uniformity at the architectural level, and we should stop having religious arguments about the lower levels of our systems and concentrate on how to build reliable higher-level constructs!