COMMENTS ON THE PROPOSAL FOR ATLAS VIRTUAL MEMORY MILESTONE

by
Peter J. Denning
Naval Postgraduate School
Monterey, CA 93943 USA
pjd@nps.edu

Peter Denning
Digitally signed by Peter Denning
Date: 2021.03.07 15:47:07 -08'00'

Full Endorsement of the Proposal

Virtual memory is indisputably one of the major hallmarks of computer system architecture since its public presentation in 1962. The authors of the proposal have documented its history and influence exceptionally well. There is hardly a computer chip and operating system that does not incorporate virtual memory. This technology is well deserving as a milestone of computing and Tom Kilburn deserves mention for his key role in its invention and design. I wholeheartedly endorse the proposal.

Was it really an invention?

One of the reviewers has questioned whether "invention" is the right word for the Atlas virtual memory, on the grounds that others were concerned with the same problem (managing page transfers between main and secondary memory) and would have come up with virtual memory on their own without knowledge of Kilburn's work. It is hard to take this objection seriously because Kilburn was issued five patents covering all aspects of the virtual memory system. Through its diligent investigation of the submitted patents and their claims, the patent office determined that this was an invention and there was no prior art. When a patent issues, the patent office is certifying that it was indeed an invention.

Patent law does not recognize later work that might duplicate the invention even if discovered independently. In the case of virtual memory, the historical record contains no instance of later, independent discovery.

Someone mentioned the PhD thesis of Fritz Güntsch in 1957 as a possible instance of prior art. A close look at his proposal (it was not implemented) reveals that he envisioned a cache not a virtual memory. His cache was an interface that speeded up access to a rotating drum. The cache did not hide the drum; instead the cache was invisible and made the drum seem faster. In contrast, a virtual memory gives the CPU the illusion it is accessing a large main memory and completely hides the disk or drum. Two other important differences are that cache does not support multiprogramming or attempt to minimize system data transfers, whereas virtual memory does support multiprogramming and through its replacement algorithm seeks to minimize data movement. Others picked up the cache idea later after virtual memory was invented. A notable example is Maurice Wilkes, who published a paper "Slave memories and dynamic storage allocation" in 1965 *IRE Transactions on Electronic Computers* (pp 270-271) discussing how a cache could speed up a main memory. It is noteworthy that Wilkes used some of the techniques of the virtual memory to describe how to make a cache efficient; he credited [Kilburn 1962] for these

insights. I had several conversations with Wilkes at Princeton in 1970, and he clearly liked the Atlas form of virtual memory much more than the slave memory idea.

There are only a handful of people that I am aware of who were thinking about the memory management issue at the time and were capable of inventing it themselves. These include Jack Dennis of MIT, David Sayre and Les Belady of IBM, and Maurice Wilkes of Cambridge. All their writings on the subject refer to [Kilburn 1962] and none suggests they had already thought of virtual memory.

Kilburn's invention was unique. There were no competitors.

Should Kilburn's name be mentioned in the citation?

I support including Kilburn's name in the citation. He masterminded the project, led the team that worked out the design, led the implementation, and applied for all the patents that covered the system. In fact, he masterminded the Atlas operating system, which was much more ambitious than virtual memory and was the host for the virtual memory.

I rate Kilburn's accomplishments with the design of early computing and operating systems right up there on par with Wilkes and Turing.

My involvement in the field

I was a graduate student at MIT in 1965 when I first learned of virtual memory from Jack Dennis. The Multics designers (led by F. Corbato) planned to include a virtual memory. By that time they were aware of serious performance problems including poor performance and susceptibility to thrashing (sudden collapse of system throughput when one too many jobs was loaded into main memory). I devoted my PhD work to understanding what caused these problems and I found solutions. I invented the working set model and articulated the principle of locality. When applied to the design of virtual memory, these new concepts prevented thrashing and made the system throughput very close to optimal. I was part of a large research community, with several hundred researchers in the US and Europe, investigating performance issues of memory systems through the decade of the 1970s. I have had the good fortune to watch the good works of these people grow up and mature into really useful concepts for memory management in caches, operating systems, and the Internet. Working sets have been a universal reference model for page replacement algorithms and the principle of locality has been accepted as a fundamental principle of computing, seen in every aspect of computing from individual programs, to the Internet, and even as a fundamental requirement for a procedure to be called an algorithm. I have kept my ever-interested eye on this field for the past half century. I recently wrote a tutorial paper for ACM *Computing Surveys*, "Working set analytics", which gives a current perspective and its use in evaluating the performance of modern cache memory systems.