A LOOK AHEAD PAGE REPLACEMENT STRATEGY FOR PAGING ON PARALLEL PROCESSOR ARCHITECTURE

Maya S. Rainford
Open University of Sri Lanka.

A paged system is one of a number of techniques which make a distinction between address space, the set of identifiers used by a program, and memory space, the set of physical main memory locations that are available. There is no a priori correspondence between address space and memory space and it is only at execution time that such correspondence can take place.

The price paid for such machine independence is an increased complexity in the addressing mechanism. A function f is required, where

f(a) = b if item a is in memory at location b, or
0 if item a is missing

Such address translation function translates a virtual address into a real address at execution time. When the process refers to an address which is within a page not currently in memory, a page fault occurs. The address transformation cannot be completed until the referenced page has been transferred from secondary storage into memory. The process is, therefore, blocked whilst this is going on. In a demand paging environment, no action is taken until a demand is made for a particular page and when such demand is made, the question facing the Operating System is which page to move out, a factor which can seriously affect the machines throughput.

This paper proposes an anticipatory paging environment as an alternative to demand paging, suitable for parallel processor architecture. A knowledge based virtual machine is described which supports a strategy of 'least expectation to be required next'. This is computed by means of an aggregate function of least worth (Rainford, 1989,1990) implicitly defined by mean of a production system.

References:

Rainford, M.S.(1989) - 'An Adaptive Knowledge Base for Learning Systems, Proc.9th Nat.Comp.Conf.CSSL
Rainford, M.S.(1990) - 'An A.I.Perspective of Human Cognition Towards an Architecture for the Thinking Machine', Proc.10th Nat.Comp.Conf.CSSL.