Architectural Support for Single Address Space Operating Systems

Utpal Kumar

Conventional Architectures

- Store virtual-to-physical mappings in linear page tables maintained separately for each protection domain.

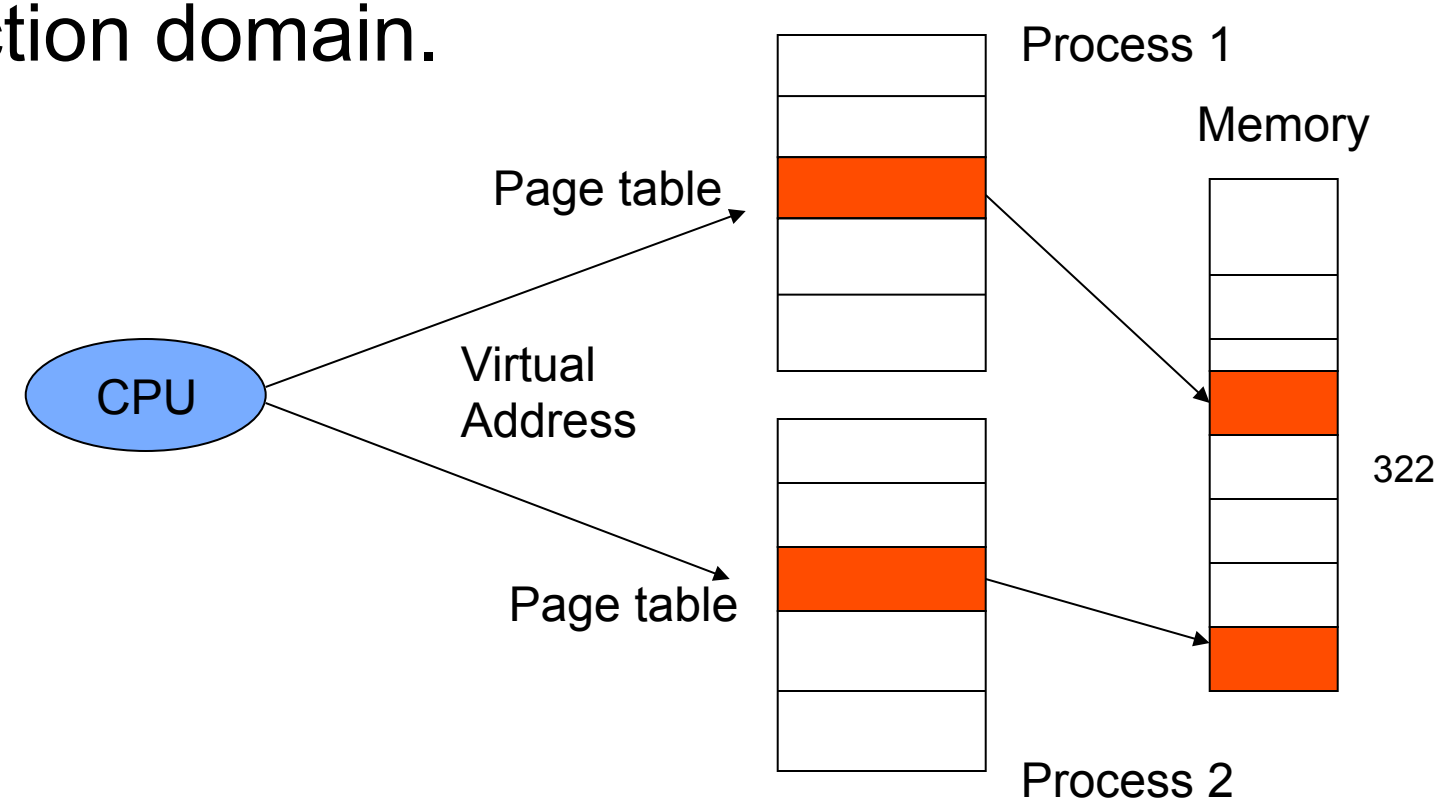

Shared Pages : Problems

- Forces kernel to keep the mappings consistent.
- In case of Tagged TLB there duplication of translation Information.

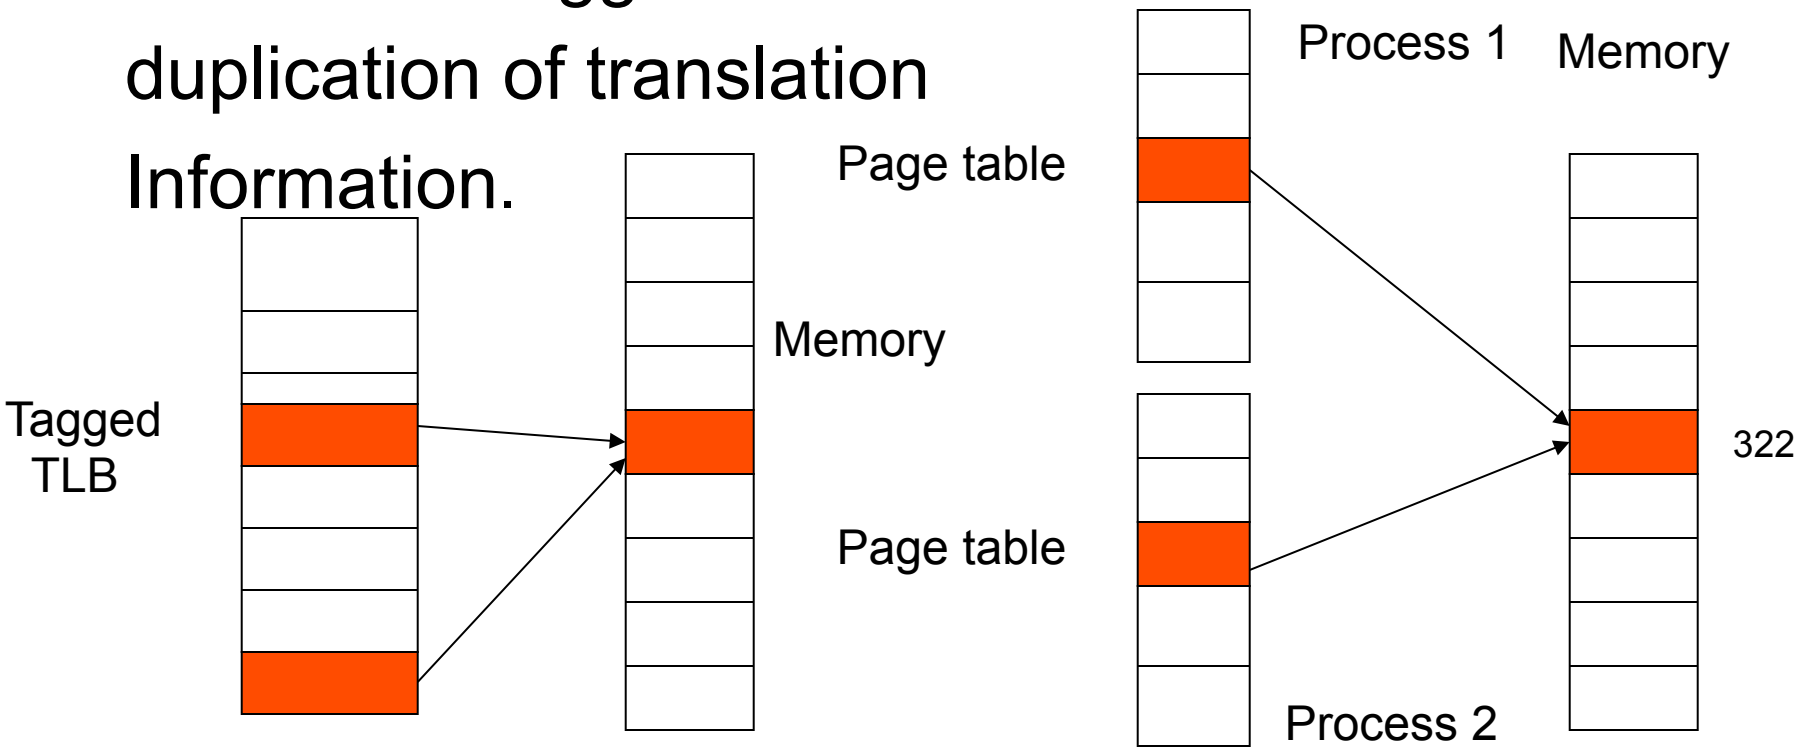

Inefficient Sharing

- Cannot pass pointers(virtual address) between domains.

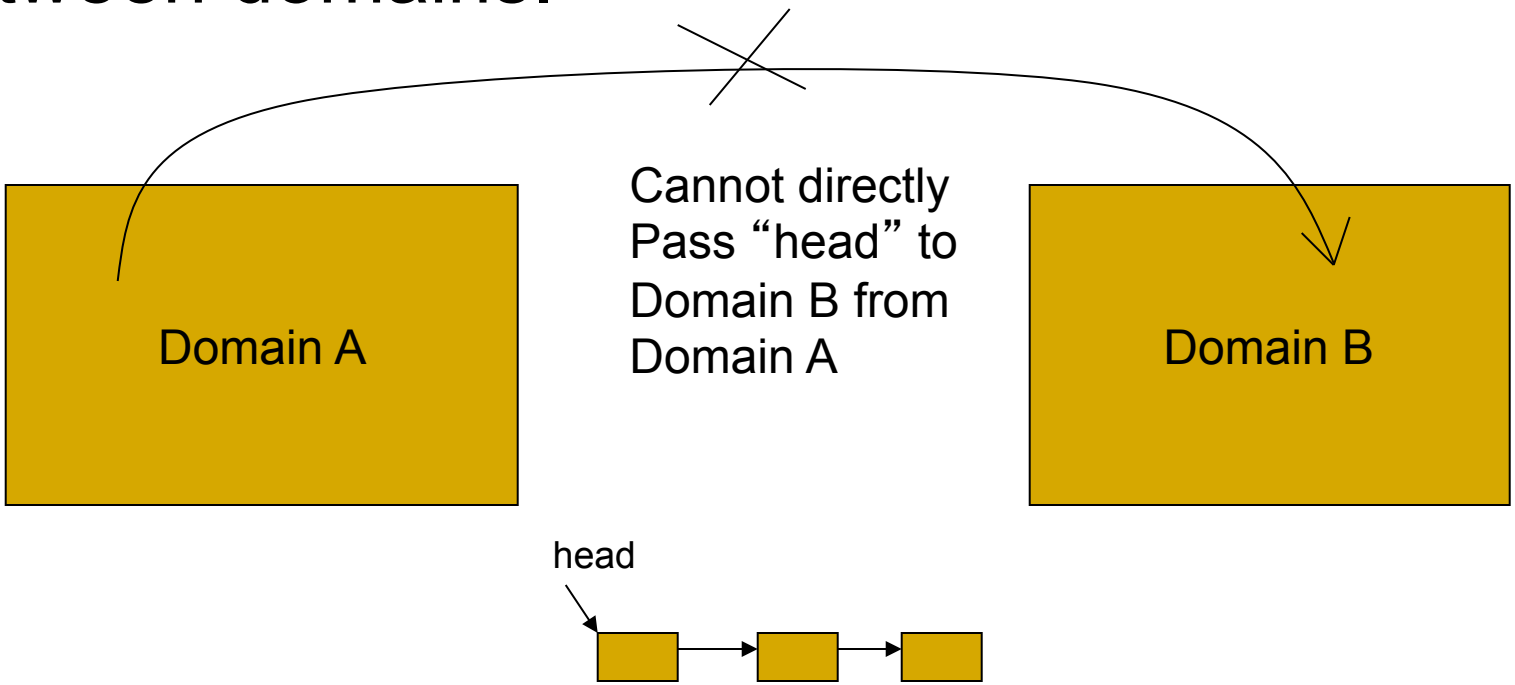

Virtually Indexed Caches - Problems

- Synonyms – When a physical page is mapped into two or more virtual pages. This causes the item to occur in multiple cache lines.

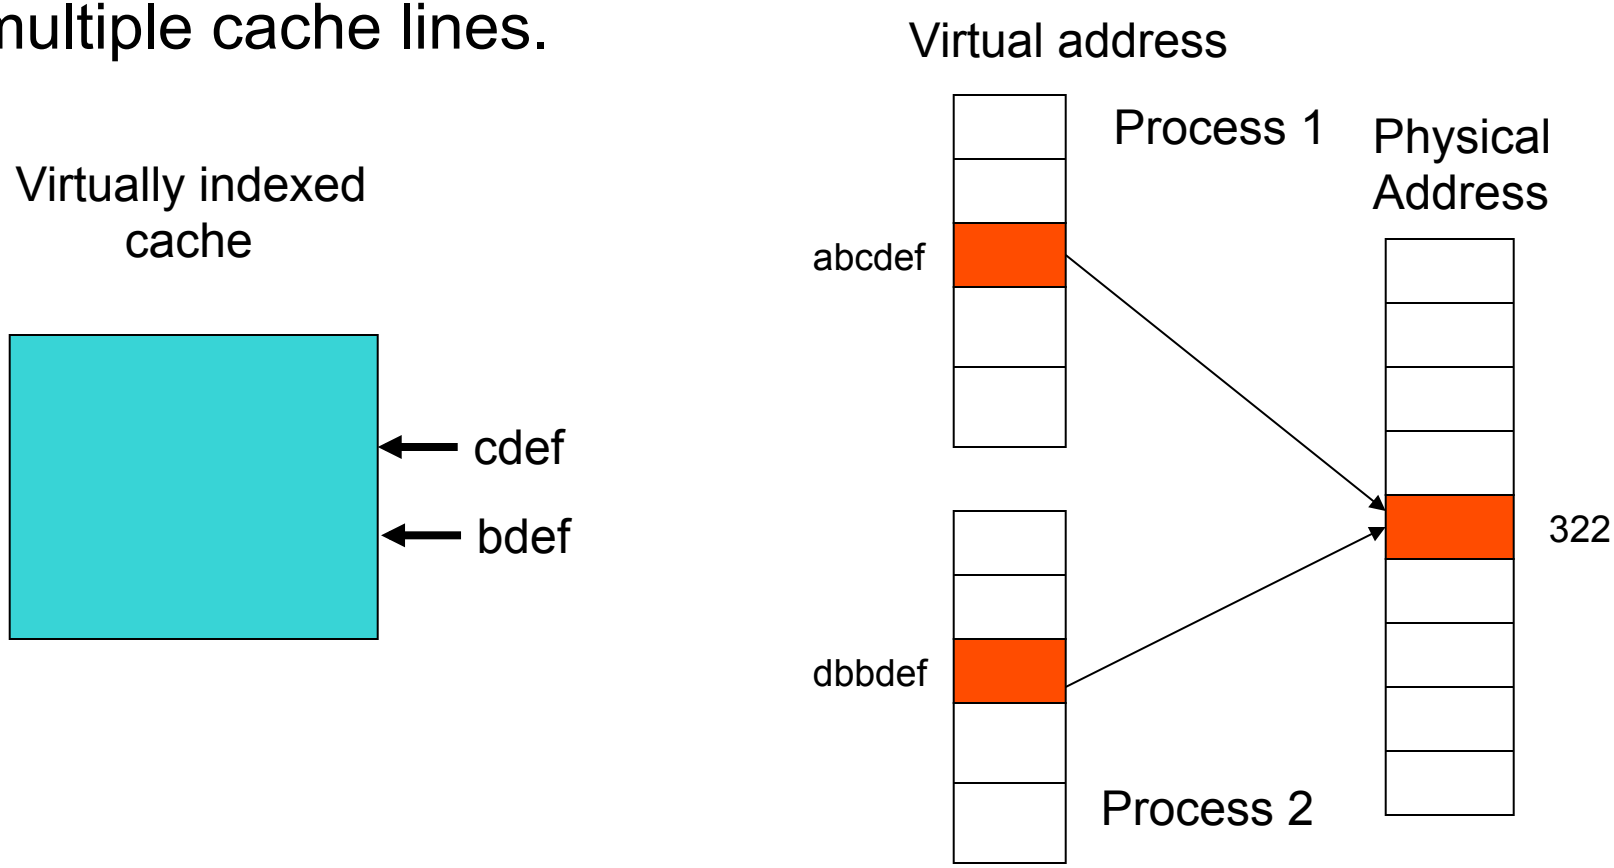

Virtually Indexed Caches - Problems

- Homonyms – When each address space has a different translation of the same virtual address.
  - Flushing the cache on process switches.
  - Extending virtual address with address space identifier.

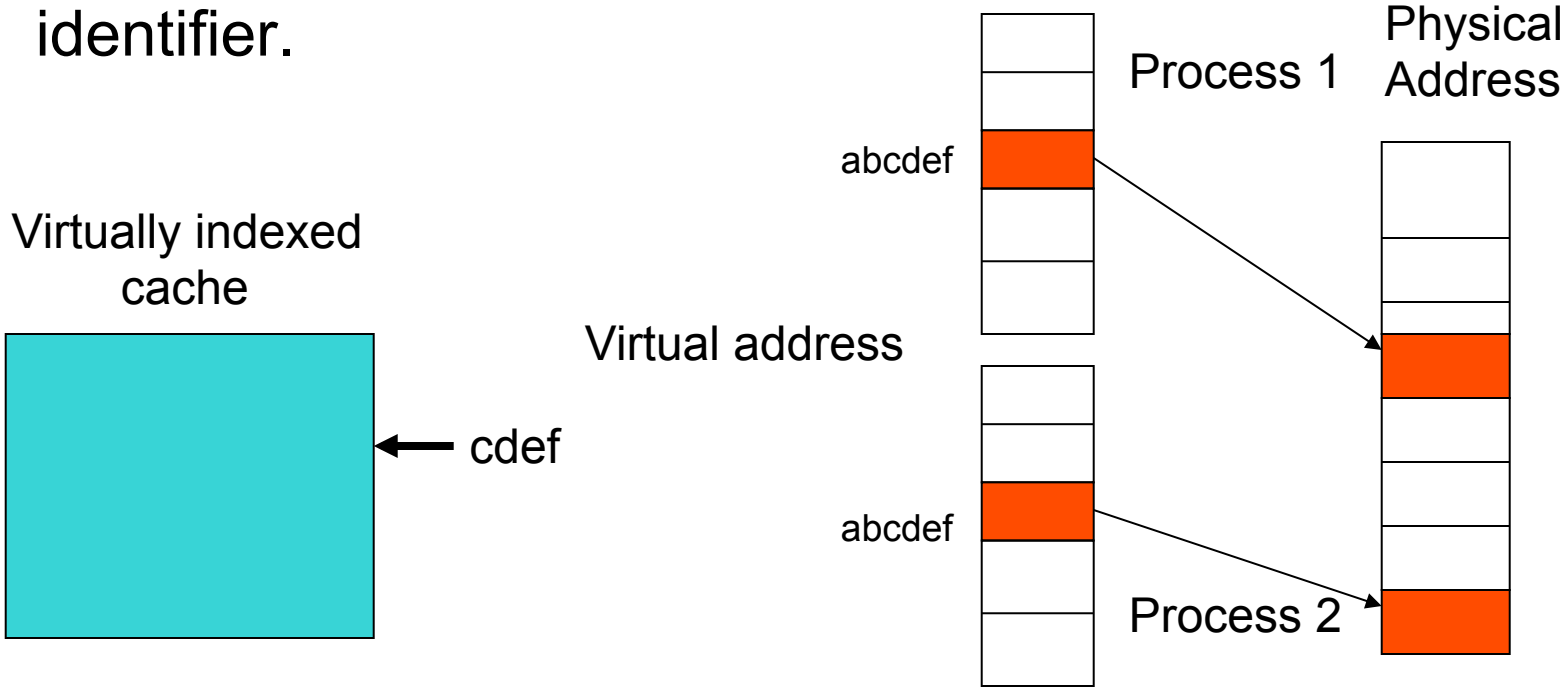

Single Address Space Operating System

- All processes run within a single global virtual address space. Possible because of transition from 32 bit to 64 bit virtual address space.
- Protection provided through protection domains that dictate which pages of the global address space a process can reference.

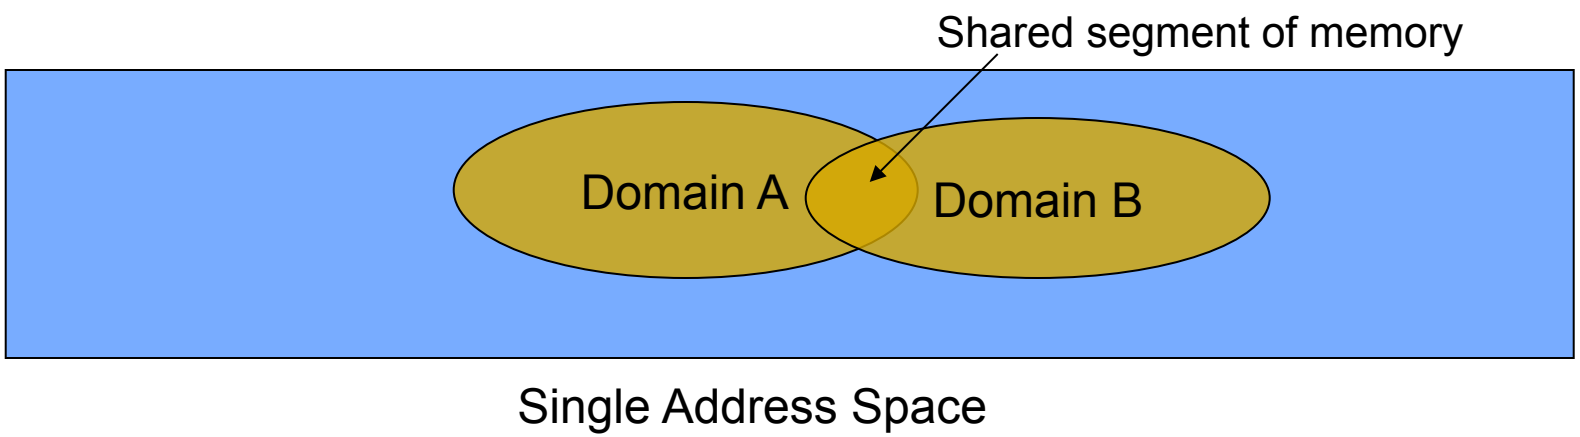

Advantages

- Encourages sharing of memory between protection domains.
  - Pointers (virtual addresses) can be passed between domains.

- Define unique mapping between virtual and physical address leading to the introduction of new constraint which can be exploited in the architecture.
  - Simplification of virtually indexed caches.
  - Address translation can be removed from critical path of process memory references.

Domain Page model

- Specifies access rights explicitly for each (domain, page) pair.
- It can be implemented in single address space architectures by removing protection domain tags from the TLB and placing them in separate Protection Lookaside Buffer.

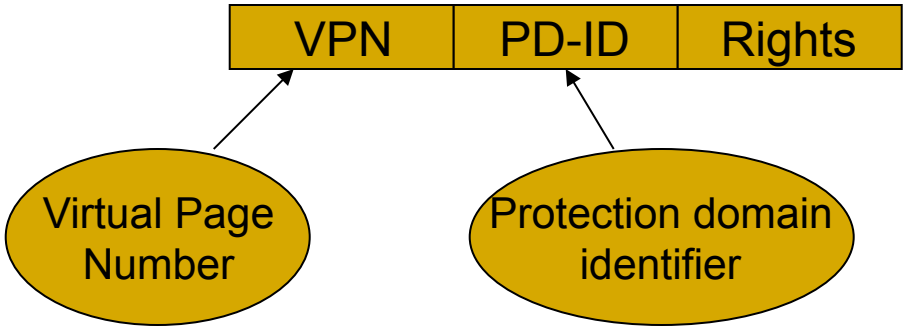

Advantages

- VPN bits are used for both the cache and PLB lookups. Thus cache and PLB searches can occur in parallel.
- Protection domain switches do not require purge of either the PLB or the TLB.
- TLB can be moved off processor chip.
  - Address translation required only for a small percentage of accesses that either miss in the cache.
  - Permits larger TLB.

Page Group Model

- Defines logical grouping of pages called page-groups.
- Each page is a member of single page group.
- Protection domain is defined by the set of page-groups that it can access.
- Each page within a group has access rights that are used by all domains with access to the group.

Evaluation of Protection Model

Attaching Segments

- Domain page model – The individual PLB entries for each domain-page pair are lazily faulted into the PLB.

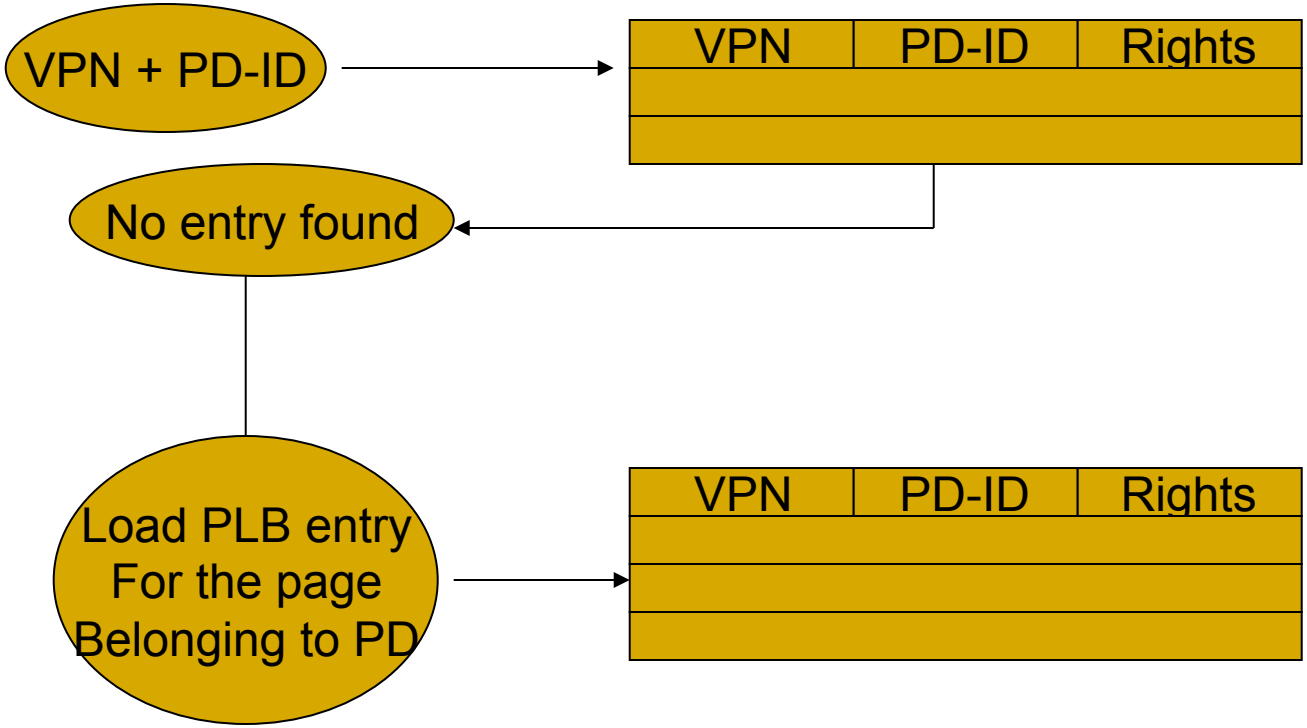

Attaching Segment

- Page group Model – OS adds the page group representing the segment to the set of groups accessible to the current domain (Add entry in page group cache)

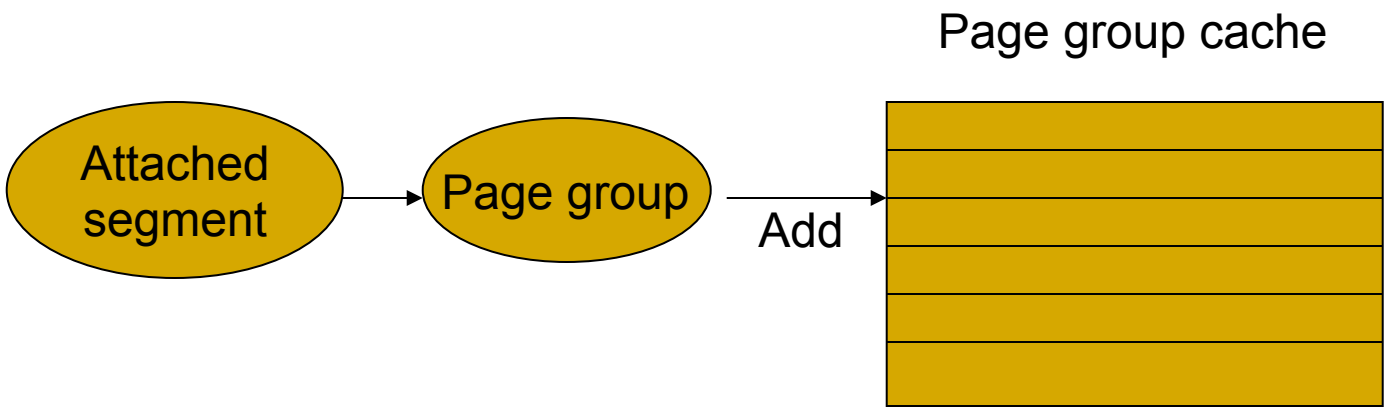

Evaluation

- Detaching Segment
  - Domain page model – OS must purge all entries from the PLB that map the detaching segment. In worst case, this could require inspecting all the entries in the PLB and eliminating those that match.

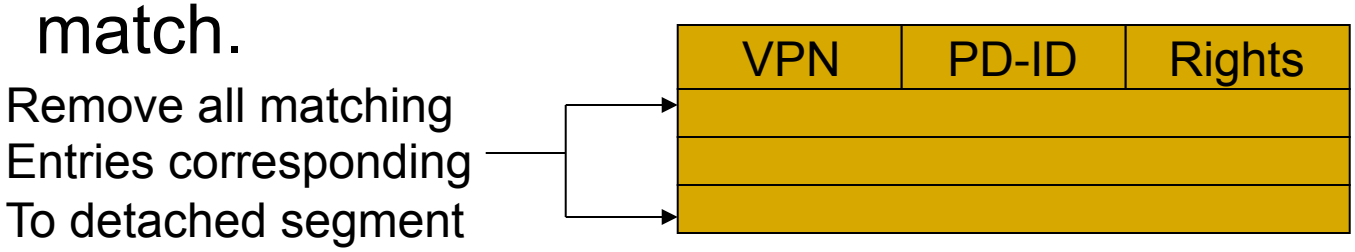

  - Page group model – Removing the appropriate page-group identifier from the set of page groups accessible to the current domain.

Manipulating Page Permissions

- Changing page protection on a per-domain, per-page basis
  - Domain page model – Straightforward.

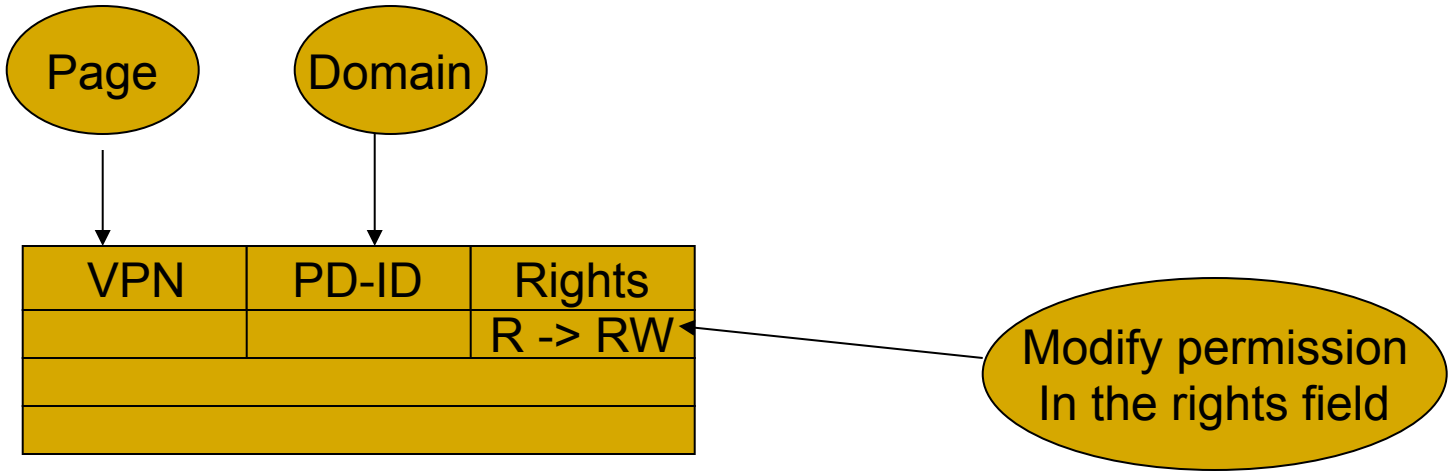

Manipulating Page Permissions

- Page group Model
  - Involves moving page between page group. E.g. if two domains require rw access to different pages in a readonly segment then create 2 additonal page groups each containing rw pages for its respective domain.

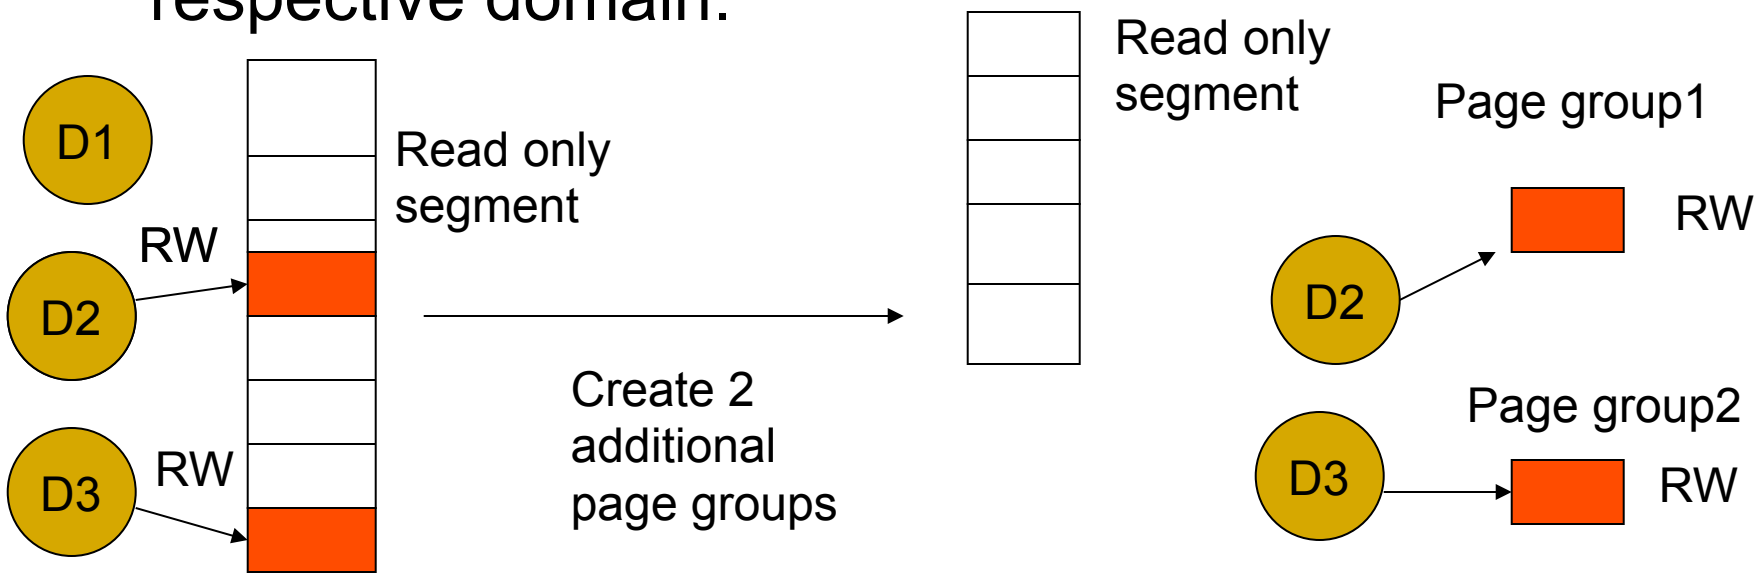

Domain Switches

- Domain page model - PLB based system requires changing only a single register, the processor control register in the processor. Rights for the new domain can be faulted in lazily.

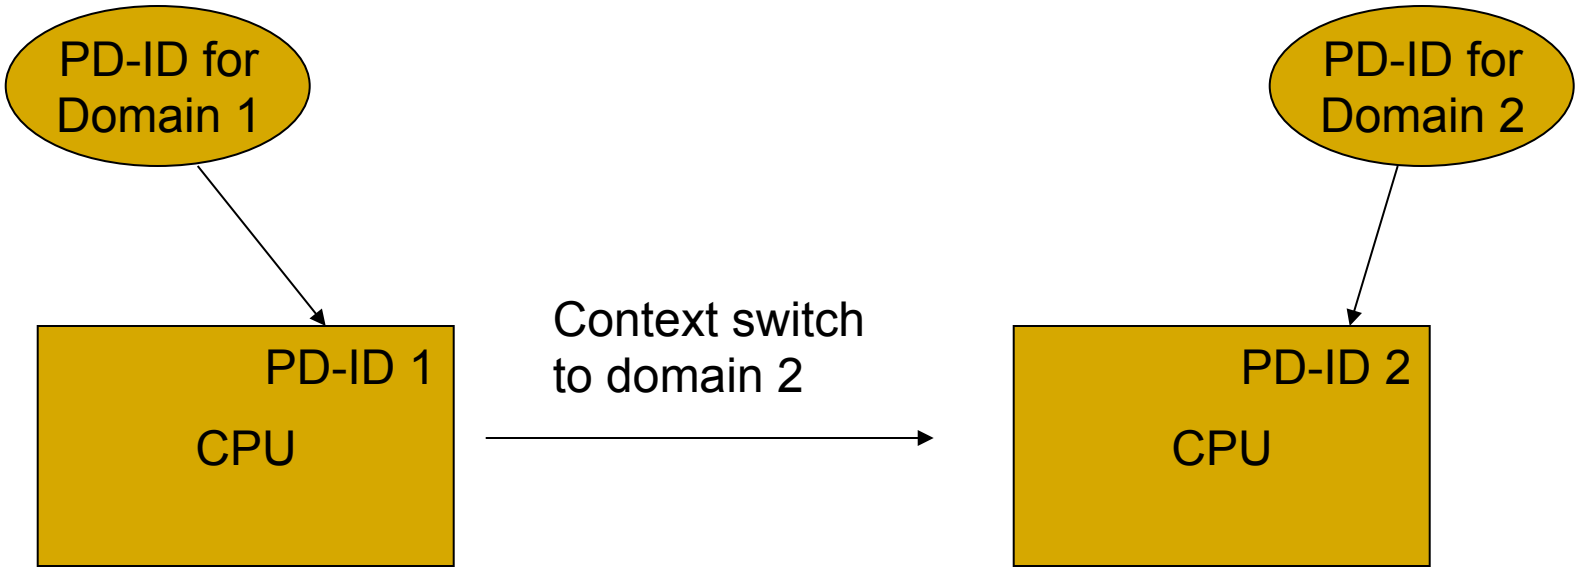

Domain Switches

- Page group model – Involves purging the active page group cache and loading in the page groups for the new domain.

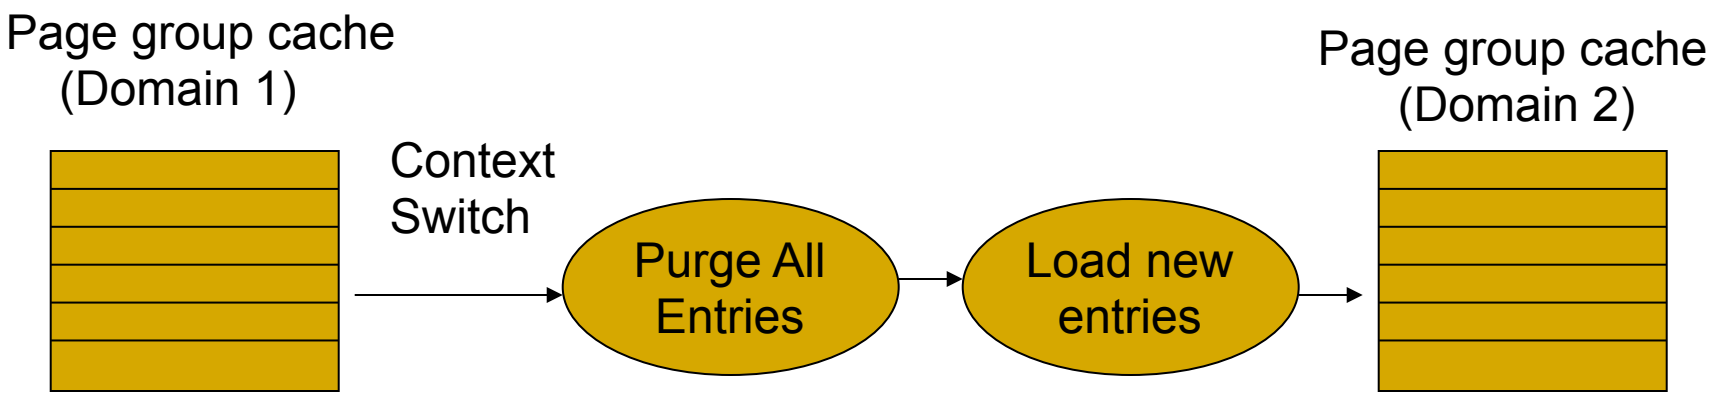

Granularity of Protection

- Domain Page Model – PLB explicitly separates protection and translation. Hence granularity of protection and translation can be different.
  - Protection control on sub-page units can lead to larger page size which improves TLB performance.
  - Protection control on larger than single translation page leads to less PLB entries.

Applications

- Well suited to the needs of “integrated software environments” – collection of software tools that work together to support users in complex tasks.
  - Protection is crucial because “tools” are separately authored programs.
  - Sharing is common among tools and hence should be done efficiently

Conclusion

- 64 bit address space facilitates a fundamentally new operation system organization, the single address space system. This led to reexamination of the conventional memory system architecture and in particular caching and protection structures:
  - Domain page model
  - Page group model

The model which will take best advantage of single address space characteristics depends upon how the system will be used.

The Protection Lookaside Buffer

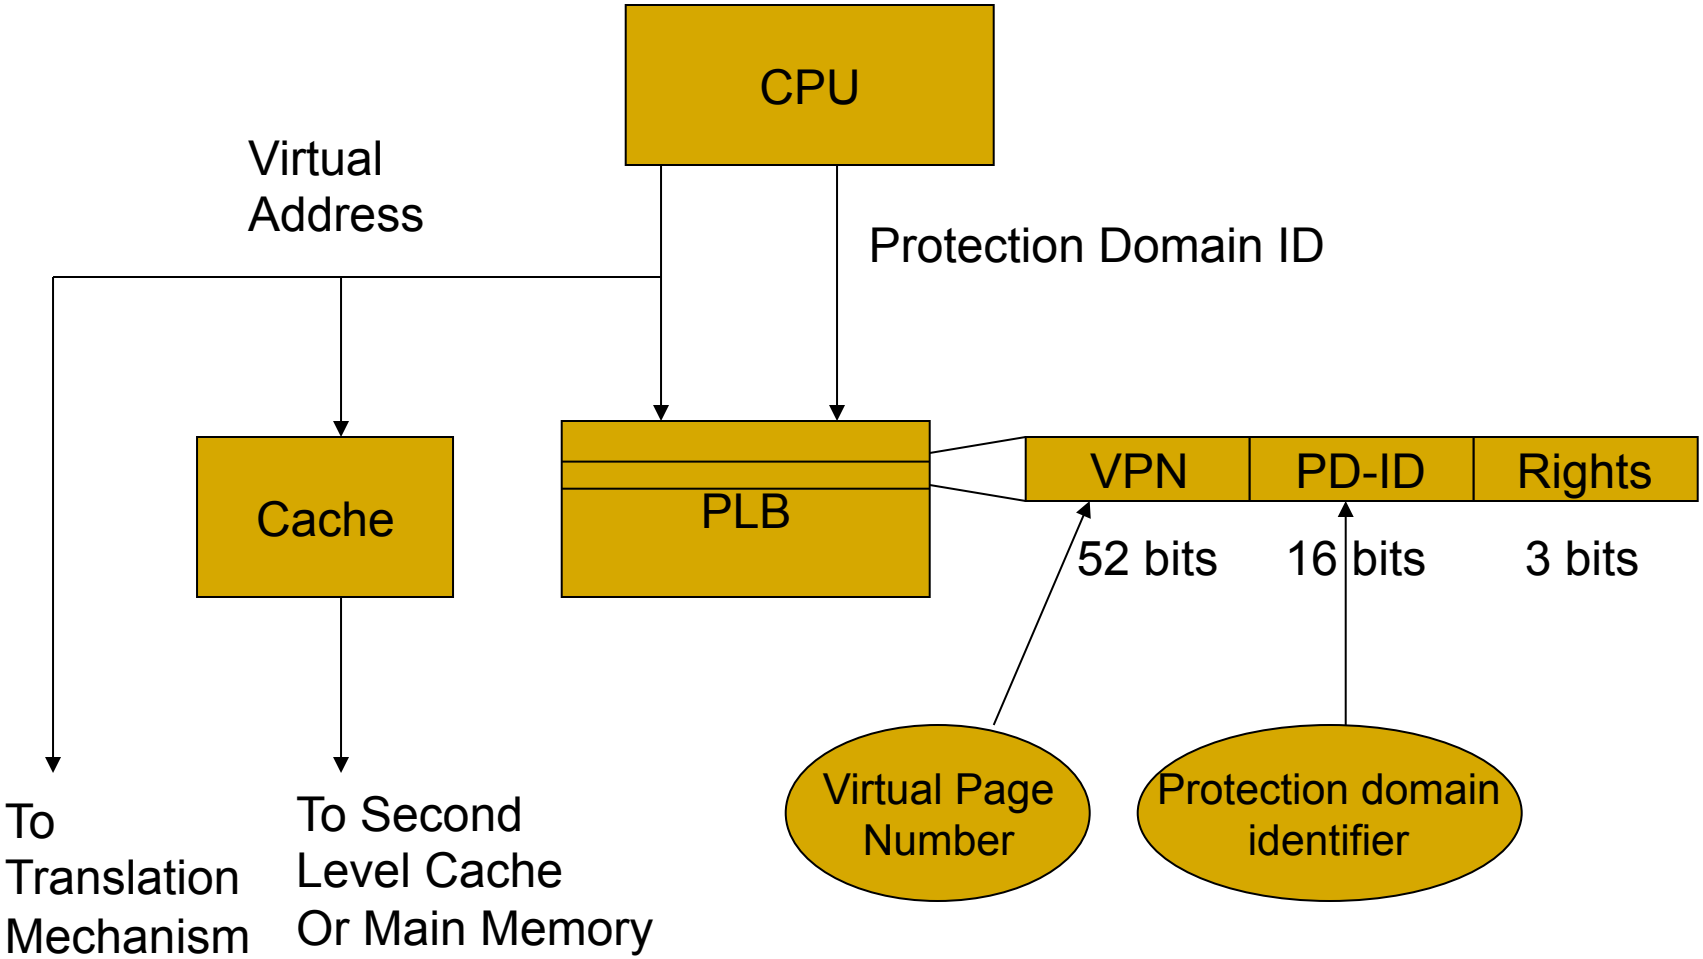

PA-RISC Protection Architecture

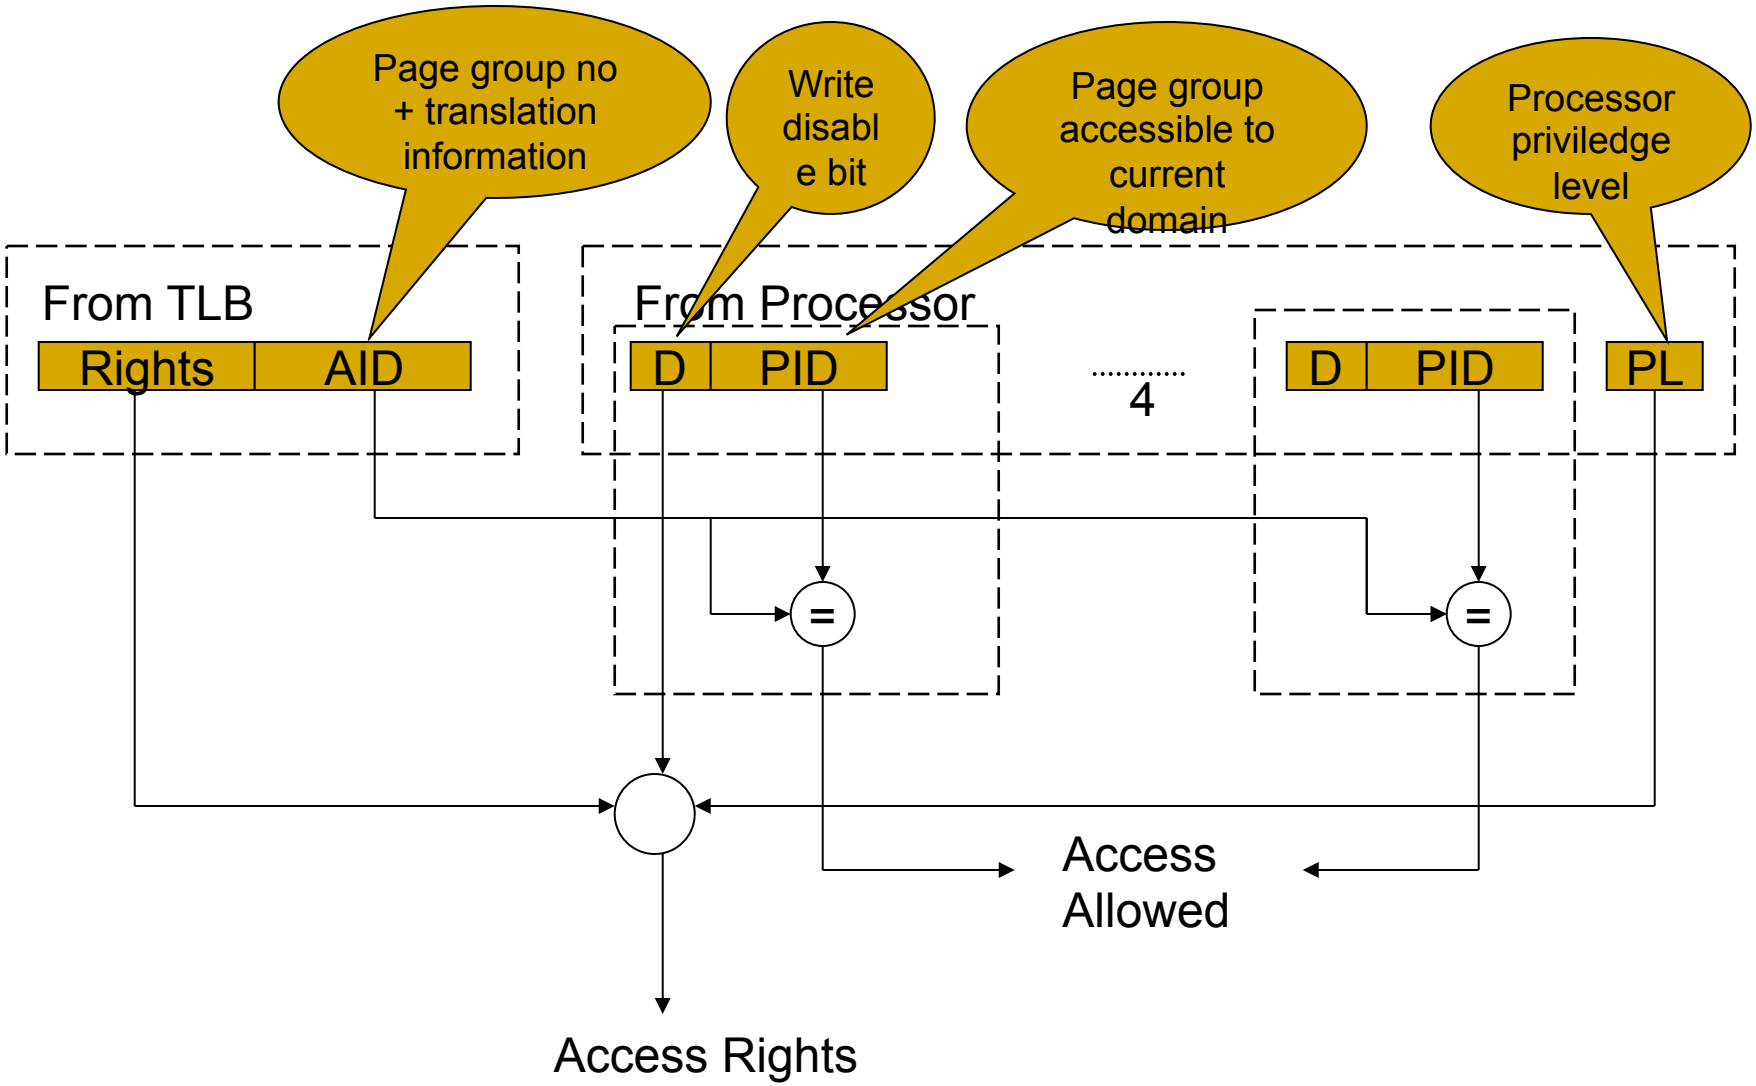

Paging Operations

- Pages must be protected from access by applications while page-in and page-out operations.
  - In a PLB system access rights are simply updated in the PLB. The number of entries changed depends on the number of domains that have access to the page.
  - In page group system a special page group is created to represent the paging server's access rights. Pages are moved to the paging server's group prior to paging operations.

Implementation Considerations

- Domain page model – Requires only a single cache lookup which provides the access rights.
- Page group model :
  - TLB is indexed to obtain the page group and access rights for the page.
  - Page group cache is checked to determine if the current protection domain has access to the page group.

  These cannot be done in parallel, since the second lookup depends on the result of the first.