ELF1 7D Virtual Memory

Young W. Lim

2021-01-12 Tue

Outline

Based on

"Study of ELF loading and relocs", 1999
http://netwinder.osuosl.org/users/p/patb/public_html/elf_relocs.html

Compling 32-bit program on 64-bit gcc

- `gcc -v`
- `gcc -m32 t.c`
- `sudo apt-get install gcc-multilib`
- `sudo apt-get install g++-multilib`
- `gcc-multilib`
- `g++-multilib`
- `gcc -m32`
- `objdump -m i386`

Logical addresses

- logical address
  - generated by CPU while a program is running
  - since it does not exist physically,
    it is also known as virtual address
  - used as a reference to access
    the physical memory location by CPU

- logical address space
  - the set of all logical addresses
    generated by a program's perspective.

https://www.geeksforgeeks.org/logical-and-physical-address-in-operating-system/

Physical addresses

- physical address
  - identifies a physical location in a memory
  - the user <u>never</u> directly uses the physical address
    but can access by the corresponding logical address.

- physical address space
  - <u>all</u> physical addresses corresponding
    to the logical addresses in a Logical address space

https://www.geeksforgeeks.org/logical-and-physical-address-in-operating-system/

Virtual addresses

- virtual addresses
  - the address you use in your programs,
  - the address that your CPU use to fetch data,
    is not real and gets translated via MMU
    to its corresponding physical address

- virtual address space
  - Linux running 32-bit has 4GB address space
  - each <u>process</u> has its <u>own</u> virtual address space

`https://stackoverflow.com/questions/15851225/difference-between-physical-logical-`

- MMU (memory-management unit) hardware
  - maps logical address to its corresponding physical address

- OS along with MMU
  - the user program generates the logical address and
  - thinks that the program is running in this logical address
  - but to access physical memory for its execution, this logical address must be mapped to the physical address by MMU

`https://www.geeksforgeeks.org/logical-and-physical-address-in-operating-system/`

Logical vs virtual addresses (1)

Whenever your program is under execution CPU generates logical address for instructions which contains (16 bit Segment Selector and 32 bit offset).Basically Virtual(Linear address) is generated using logical address fields. Segment selector is 16 bit field out of which first 13bit is index (Which is a pointer to the segment descriptor resides in GDT,described below) , 1 bit TI field (TI = 1, Refer LDT , TI=0 Refer GDT)

https://stackoverflow.com/questions/15851225/difference-between-physical-logical-v

high / low memory (1)

- Low memory
  Memory for which logical addresses exist in kernel space. On almost every system you will likely encounter, all memory is low memory.

- High memory
  Memory for which logical addresses do not exist, because it is beyond the address range set aside for kernel virtual addresses.

https://www.oreilly.com/library/view/linux-device-drivers/0596005903/ch15.html

kmap

-kmap returns a kernel virtual address for any page in the system. For low-memory pages, it just returns the logical address of the page; for high-memory pages, kmap creates a special mapping in a dedicated part of the kernel address space. Mappings created with kmap should always be freed with kunmap; a limited number of such mappings is available, so it is better not to hold on to them for too long. kmap calls maintain a counter, so if two or more functions both call kmap on the same page, the right thing happens. Note also that kmap can sleep if no mappings are available.

https://www.oreilly.com/library/view/linux-device-drivers/0596005903/ch15.html

Virtual address and physical address (1)

- Physical addresses are provided directly <u>by the machine</u>
    - one physical address space <u>per machine</u>
    - addresses typically range
      from some minumum (sometimes0) to some maximum,
    - though some portions of this range are usually used
      by the <u>OS</u> and/or <u>devices</u>,
      but not available for <u>user processes</u>

`https://www.student.cs.uwaterloo.ca/~cs350/F07/notes/mem.pdf`

Virtual address and physical address (2)

- Virtual addresses (or logical addresses) are addresses provided by the OS
    - one virtual address space per process
    - addresses typically start at zero, but not necessarily
    - space may consist of several segments

`https://www.student.cs.uwaterloo.ca/~cs350/F07/notes/mem.pdf`

Virtual address and physical address (3)

- address translation (or address binding) means
  mapping virtual addresses to physical addresses

`https://www.student.cs.uwaterloo.ca/~cs350/F07/notes/mem.pdf`

Virtual address and physical address (4)

- <u>size</u> of each <u>section</u> except stack is <u>specified</u> in ELF file
- sections which are initialized from the ELF file
    - code (i.e., `.text`)
    - read-only data
    - initialized data segments
- other remaining sections are initially <u>zero-filled</u>
- <u>sections</u> have their own <u>specified alignment</u>
- <u>segments</u> are <u>page aligned</u>
- 3 segments = (`.text` + `.rodata`), (`.data` + `.sbss` + `.bss`), (stack)
- not all programs contain this many segments and sections

https://www.student.cs.uwaterloo.ca/~cs350/F07/notes/mem.pdf

Single address space (1)

- simple systems
- sharing the same memory space
    - memory and peripherals
    - all processes and OS
- no memory proctection

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Single address space (2)

- CPUs with single address space
    - 8086 - 80286
    - ARM Cortex-M
    - 8 / 16-bit PIC
    - AVR
    - most 8- and 16-bit systems

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Single address space (3)

- portable c programs expect flat memory
  - multiple memory access methods limit portability
- management is tricky
  - need to know / detect total RAM
  - need to keep processes separated
- no protection

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual memory (1)

- a system that uses an address mapping
- maps virtual address space to physical address space
    - to physical RAM
    - to hardware devices
        - PCI devices
        - GPU RAM
        - On-SOC IP blocks

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual memory (2)

- Advantages
  - each process can have a different memory mapping
    one process' RAM is invisible to other processes
    built in memory protection
    kernel RAM is invisiable to user space processes
  - memory can be moved
  - memory can be swapped to disk

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual memory (3)

- Advantages (continued)
    - hardware device memory can be mapped
      into process' address space
      requires the kernel to perform the mapping
    - physical RAM can be mapped
      into multiple processes at once
      shared memory
    - memory regions can have access permissions
      read / write / execute

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual memory (4)

- Physical addresses
  addresses used by the hardware (DMA, peripherals)
- Virtual addresses
  addresses used by software
    - RISC: load/store instructions
    - CISC: any instruction accessing memory

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual memory (5)

- mapping is performed in hardware
- no performance penalty
  for accessing already mapped RAM regions
- permissions are handled without penalty
- the same instructions are used
  to access RAM and mapped hardware
- software will only use virtual addreses
  in its normal operation

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

MMU (Memory Management Unit) (1)

- MMU is the hardware responsible for implementing virtual memory
- sits between the CPU core and memory
- usually the part of the physical CPU on ARM, it's part of the licensed core
- separate from the RAM controller DDR controller is a separate IP block

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

MMU (Memory Management Unit) (2)

- transparently handels <u>all</u> <u>memory accesses</u>
  from load / store instructions
- maps <u>memory acceses</u> using virtual addresses
  to <u>system RAM</u> and <u>peripheral hardware</u>
- handles <u>permissions</u>
- generates an exception (page fault)
  on an invalid access

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

TLB (Translation Lookaside Buffer)

- TLB is consulted by the MMU
  when CPU accesses a virtual address
- if the virtual address is in the TLB,
  the MMU can look up the physical address
- if the virtual address is not in the TLB,
  the MMU will generate a page fault exception
  and interrupt the CPU
- if the virtual address is in the TLB,
  but the permissions are insufficient,
  the MMU will generate a page fault

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Page faults

- a page fault is a CPU exception
  generated when software attempts
  to use an invalid virtual address

  - the virtual address is not mapped for the process requesting it
  - the processes has insufficient permissions for the address
  - the virtual address is valid, but swapped out

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel virtual address (1)

- in linux, the kernel uses virtual addresses
  as user space processes do
  this is not true of all OS's

- virtual address space is split
  1. the upper part is used for the kernel
  2. the lower part is used for user space
  3. 32-bit linux have the split address 0xc0000000

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel virtual address (2)

- By default, the kernel uses the <u>top 1GB</u>
  of virtual address space

- each user space process gets the <u>lower 3GB</u>
  of virtual address space

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Virtual addresses - linux (1)

- kernel address space is the area
  above `CONFIG_PAGE_OFFSET`

- for 32-bit, this is configurable at kernel build time
  - the kernel can be given a different amount
    of address space as desired

- for 64-bit, the split varies by architecture
  but it is high enough

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Virtual addresses - linux (2)

- three kinds of virtual addresses in Linux
- Kernel
  - Kernel Logical Address
  - Kernel Virtual Address

- User Space
  - User Virtual Address

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Kernel virtual / logical addresses (1)

- the kernel maps <u>most</u> of the *kernel virtual address space*
  to perform <u>1:1 mapping</u> with an <u>offset</u> of
  the top part of physical memory (3GB - 4GB)
  - slightly *less then* for 1Gb for 32bit x86
  - can be different for other processors or configurations
- for kernel code on x86 address 0xc00000001
  is mapped to physical address 0x1.
- This is called logical mapping
  - a 1:1 mapping (with an offset) that allows the kernel
    to access most of the physical memory of the machine.

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

- in the following cases, the kernel keeps a region
  at the <u>top</u> of its <u>virtual address space</u>
  where it maps a "random" page
    - when we have *more then* 1Gb physical memory on a 32bit machine,
    - when we want to reference <u>non-contiguous</u>
      physical memory blocks as contiguous
    - when we want to map <u>memory mapped IO</u> regions

- this mapping does <u>not</u> follow the 1:1 pattern of
  the <u>logical mapping</u> area.

- This is called the virtual mapping.

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

Kernel virtual / logical addresses (3)

- on many platforms (x86 is an example),
  both the logical and virtual mapping are done
  using the same hardware mechanism
  (TLB controlling virtual memory).

- In many cases, the logical mapping is actually done
  using virtual memory facility of the processor,
  (this can be a little confusing)

- The difference is in which mapping scheme is used:
  - 1:1 for logical
  - random for virtual (paging)

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

- 3 kinds of addressing

1. Logical Addressing : Address is formed by base and offset
   This is nothing but segmented addressing,
   where the address (or offset) in the program is always used
   with the base value in the segment descriptor

1. Linear Addressing : also called virtual address
   Here virtual adresses are contigous,
   but the physical address are not contiguous
   Paging is used to implement this.

1. Physical Addressing : the actual address on the Main Memory

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-add

Kernel virtual / logical addresses (5)

- in linux, the kernel memory (in address space) is beyond 3 GB, i.e. 0xc000000.
- the addresses used by Kernel are <u>not</u> physical addresses
  - to map the virtual address from 3GB to 4GB it uses `PAGE_OFFSET`.
    - no page translation is involved.
    - contiguous address
    - except 896 MB on x86.
  - beyond the address space from 3GB to 4GB, paging is used for translation.
    - `vmalloc` returns these addresses

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

Kernel virtual / logical addresses (6)

- when virtual memory is referred
  in context of user space,
  then it is through paging

- if kernel memory is mentioned
  then it is the address mapped
  - by `PAGE_OFFSET` (kernel logical address)
  - by `vmalloc` (kernel vitual address)

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

- `PAGE_OFFSET` in x86 is `0XC0000000`,
  (1100_0000_0000_0000_0000_0000_0000_0000)
  or 3 gigabytes ($3 * 2^{30}$)

- this is where the 3G/1G split is defined.

- every address <u>above</u> `PAGE_OFFSET`
  is the kernel virtual address

- any address <u>below</u> `PAGE_OFFSET`
  is a user space address

`https://linux-mm.org/VirtualMemory`

- to get <u>kernel memory</u> in byte-sized chunks.
    - `kmalloc()`
      <u>virtually</u> <u>contiguous</u>
      <u>physically</u> <u>contiguous</u>
    - `vmalloc()`
      <u>virtually</u> <u>contiguous</u>
      <u>not</u> necessarily <u>physically</u> <u>contiguous</u>

`https://stackoverflow.com/questions/116343/what-is-the-difference-between-vmalloc-`

- On a 32-bit system, `kmalloc()`
    - returns the kernel logical address (it is a virtual address)
    - the direct mapping (constant offset)
    - a contiguous physical chunk of RAM.
    - suitable for DMA where we give only

- `vmalloc()`
    - returns the kernel virtual address
    - paging (not direct mapping)
    - not necessarily a contiguous chunk of RAM
    - Useful for large memory allocation and
      in cases where non-contiguous physicl memory is allowed

https://stackoverflow.com/questions/116343/what-is-the-difference-between-vmalloc-

- kernel logical addresses use
  normal CPU memory access functions.
- On 32-bit systems,
  only 4GB of kernel logical address space exists,
  even if more physical memory than that is in use.
- logical address space supported by physical memory
  can be allocated with `kmalloc()`

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-add

- kernel virtual addresses do not necessarily have corresponding logical addresses.

- You can allocate physical memory with `vmalloc` and get back a virtual address that has no corresponding logical address (on 32-bit systems with PAE, for example).

- use `kmap()` to assign a logical address to that virtual address.

https://stackoverflow.com/questions/8708463/difference-between-kernel-virtual-addr

Kernel logical addresses (1)

- normal address space of the kernel
  `kmalloc()`
- virtual addresses are a <u>fixed offset</u>
  from their physical addresses
  virtual 0xc0000000 → physical 0x00000000
- easy conversion between physical and virtual addresses

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel logical addresses (2)

- kernel logical addreses can be converted to and from
  physical addresses using these macros
  `__pa(x)`
  `__va(x)`
- for small memory systems (less than 1G of RAM)
  kernel logical address space starts at `PAGE_OFFSET`
  and goes through the end of physical memory

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

- kernel logical address space includes
    - memory allocated with `kmalloc()` and most other allocation methods
    - kernel stacks per process

- kernel logical memory can <u>never</u> be <u>swapped out</u>

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Kernel logical addresses (4)

- kernel logical addresses use a fixed mapping
  between physical and virtual address space
- this means virtually contiguous regions
  are by nature also physically contiguous
- this combined with inability to be swapped out,
  makes them suitable for DMA transfers

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel logical addresses (5)

- for 32-bit <u>large</u> memory systems (> 1GB RAM)
  <u>not all</u> of the physical RAM can be mapped
  into the kernel's address space
- kernel address space is the <u>top 1GB</u> of
  virtual address space, by default
- <u>upto 104 MB</u> is reserved at the <u>top</u> of
  the kernel memory space
  for <u>non-contiguous</u> allocation
  vmalloc()

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

- in a large memory case,
  only the bottom part of physical RAM
  is mapped directly into
  kernel logical address space

- only the bottom part of physical RAM has
  a kernel logical address

- this case is never applied to 64-bit systems
  - there is always enough kernel address space
    to accommodate all the RAM

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Low and High Memory

- low memory
  - physical memory which has a kernel logical address
  - physically contiguous

- high memory
  - physical memory beyond ~896MB
  - has no logical address
  - not physically contiguous when used in the kernel
  - only on 32-bit

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel virtual addresses (1)

- kernel virtual addresses are
  above the kernel logical address mapping

- kernel virtual addresses - `vmalloc()`
- kernel logical addresses - `kmalloc()`

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Kernel virtual addresses (2)

- kernel virtual addresses are used for
  - non-contiguous memory mappings
    - often for large buffers which could potentially be too large to find contiguous memory
    - vmalloc()
  - memory-mapped I/O
    - map peripheral devices into kernel
    - PCI, SoC IP blocks
    - ioremap(), kmap()

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Kernel virtual addresses (3)

- the important difference is that memory
  in the kernel virtual address area (`vmalloc()` area)
  is <u>non-contiguous</u> <u>physically</u>
- this makes it easier to allocate, especially
  for large buffers on small memory systems
- this makes it <u>unsuitable</u> for DMA

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Kernel virtual addresses (4)

- in a <u>large</u> memory situation,
  the kernel virtual address area is smaller,
  because there is more physical memory

- an interesting case, where <u>more</u> memory means
  <u>less</u> space for kernel virtual addresses

- in 64-bit, of course, this doesn't happen,
  as `PAGE_OFFSET` is large, and
  there is <u>much more</u> virtual address space

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

User virtual addresses (1)

- represent memory used by user space programs
  - the most of the memory on most systems
  - where the most of the compilation is
- all addresses below `PAGE_OFFSET`
- each process has its own mapping
  - threads share a mapping
  - complex behavior with `clone(2)`

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

- kernel logical addresses use a fixed mapping
  user space processes make full use of the MMU
  - only the used portions of RAM are mapped
  - memory is not contiguous
  - memory may be swapped out
  - memory can be moved

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

User virtual addresses (3)

- since user virtual addresses are not guaranteed
  to be swapped in, or even allocated at all,

- user buffers are not suitable for use
  by the kernel (or for DMA), by default

- each process has its own memory map
  `struct mm` pointers in `task_struct`

- at context switch time, the memory map is changed
  this is part of the overhead

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Memory management unit (1)

- the MMU manages virtual address mappings
  - maps virtual addresses to physical addresses
- the MMU operates on basic units of memory : pages
  - page size varies by architecture
  - some architectures have configurable page sizes

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Memory management unit (2)

- common page sizes
    - ARM - 4k
    - ARM64 - 4k or 64k
    - MIPS - widely configurable
    - x86 - 4k

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

- a page is
    - a unit of memory size
    - aligned at the page size
    - abstract

- a page frame refers to
    - a physical memory block
      which is page sized and page aligned
    - physical

- the pfn (page frame number) is often
  used to refer to physical page frames
  in the kernel

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Memory management unit (4)

- the MMU operates on pages
- the MMU maps <u>physical</u> frames to <u>virtual</u> addresses
- a memory map for a <u>process</u> contains many mappings
- a mapping often covers multiple pages
- the TLB holds each mapping
    - virtual address
    - physical address
    - permissions

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Page faults

- when a process acceses a region of memory that is <u>not</u> <u>mapped</u>, the MMU will <u>generate</u> a page fault <u>exception</u>

- the kernel <u>handles</u> page fault <u>exceptions</u> regularly as part of its memory management design

- TLB can contain only the part of the required maps for a process
- page faults at context switch time
- lazy allocation

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Basic TLB mappings (1)

- user virtual address space
  - mapped pages unmapped space
- physical address space
  - allocated frames
- TLB mapings
  - TLB entries (page, page frame)
  - virtually contiguous regions
    not physically contiguous

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Basic TLB mappings (2)

- mappings to virtually <u>contiguous</u> regions
  do <u>not</u> have to be physically <u>contiguous</u>
- <u>easy</u> memory allocation
- almost all user space code does not need
  physically contiguous memory

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Multiple processes

- each *process* has its own set of mappings
- the same virtual addresses in two different processes
  will likely be used to map different physical addresses
    - (page, page frame1) for process 1
    - (page, page frame2) for process 2

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Shared memory (1)

- shared memory is easily implemented with an MMU
- simply map the same physical frame
  into two different processes
- the virtual addresses need not be the same
  - for pointers to values inside a shared memory region
    the virtual addresses must be the same

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Shared memory (2)

- the shared memory region can be mapped to
  different virtual addresses in each process

- the mmap() system call allows the user space process
  to request a specific virtual address
  to map the shared memory region
    - if the kernel cannot grant a mapping at this address,
      mmap() returns with failure

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Lazy allocation (1)

- the kernel does <u>not</u> <u>allocate</u> pages <u>immeidately</u> that are <u>requested</u> by a process
- the kernel will <u>wait</u> until those pages are <u>actually used</u>

- lazy allocation to optimize a performance
  - if the requested pages may not be actually used, then the allocation will never happen

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Lazy allocation (2)

- when memory is <u>requested</u> for allocation,
  the kernel simply creates
  a <u>record</u> of the *request* in its page tables
  and then <u>returns</u> (quickly) to the process,
  <u>without updating</u> the TLB

- when that newly-allocated memory is actually <u>accessed</u>,
  the CPU will generate a page fault,
  because the CPU doesn't know about the mapping
  (no entry in the TLB)

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Lazy allocation (3)

- in the page fault handler,
  the kernel uses its page tables
  to determine that the mapping is <u>valid</u>
  (from the kernel's point of view)
  yet <u>unmapped</u> in the TLB

- the kernel will <u>allocate</u> a physical page frame
  and <u>update</u> the TLB with the new mapping

- the kernel <u>returns</u> from the exception handler and
  user space program can <u>resume</u>

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Lazy allocation (4)

- in a lazy allocation case, the user space program
  is never aware that the page fault happened

- the page fault can only be detected
  at the time that was lost to handle it

- for processses that are time-sensitive
  pages can be pre-faulted, or simply touched,
  at the start of execution
    - see also `mlock()` and `mlockall()`

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

- the entries in the TLB are a limited resource
- far more mappings can be made than can exist
  in the TLB at one time
- the kernel must keep track of all of the mappings
  at all times
- the krenel stores all these informations
  in the page tables
  stuct_mm and vm_area_struct

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

- since the TLB can only hold a <u>limited subset</u> of
  the total mappings for a process,
  some <u>valid mappings</u> will <u>not</u> have TLB entries
- when these addresses are <u>touched</u>
  the CPU will generate a page fault
  because the CPU has no knowledge of the mapping
  only the kernel does

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Page tables (3)

- the page fault handler will
  - find the appropriate mapping for the offending addresses
    in the krenel's page tables
  - select and remove an existing TLB entry
  - create a TLB entry for the page
    containing the address
  - return to the user space process
    - observe the similarities to lazy allocation handling

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Swapping (1)

- when memory utilization is high,
  the kernel may swap some frames to disk
  to free up RAM

- the MMU makes this possible
  - the kernel may copy a frame to disk and
    remove its TLB entry
  - the frame may be reused by another process

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

Swapping (2)

- when the frame is needed again,
  the CPU will generate a page fault
  because the address is not in the TLB

- at a page fault time, the kernel can
  - put the process to sleep
  - copy the frame from the disk
    into an unused frame in RAM
  - fix the page table entry
  - wake the process

`https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf`

- note that when the page is restored to RAM,
  it is not necessarily restored to the same physical frame
  where it originally was located (before being swapped out)

- the MMU will use the same virtual address though,
  so the user space program will not know the difference
  - this is why user space memory cannot
    typically be used for DMA

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

User space

- there are several ways to allocate memory
  from user space
    - ignoring the familiar *alloc() functions,
      which sit on top of platform methods
- mmap() can be used directly to allocate
  and map pages
- brk() / sbrk() can be used to increase
  the heap size

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

mmap()

- mmap() is the standard way to allocate
  large amounts of memory from user space
- while mmap() is often used for files,
  the MAP_ANONYMOUS flag causes mmap()
  to allocate normal memory for the process
- the MAP_SHARED flag can make the allocated pages
  sharable with other processes

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

brk() / sbrk() (1)

- brk() sets the top of the program break
- this is the top of the data segment
  but inspecton of kernel/sys.c shows
  it separates from the data segment
- this in effect increases the size of the heap
- sbrk() increases the program break
  rather than setting it directly

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

brk() / sbrk() (2)

- lazy allocation
- see mm/mmap.c for do_brk()
- do_brk() is implemented similar to ~mmap()
- modify the page tables for the new area
- wait for the page fault
- optionally, do_brk() can pre-fault the new area
  and allocate it
  see mlock(2) to control this behavior

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

High level implementation

- malloc() and calloc() will use
  either brk() or mmap()
  depending on the requested allocation size

  - small allocations use brk()
  - large allocaion use mmap()
  - see mallopt(3) and the M_MMAP_THRESHOD parameter
    to control this behavio

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf

Stack

- Stack expansion
- if a process accesses memory beyond its stack, the CPU will trigger a page fault
- the page fault handler detects the address is just beyond the stack, and allocates a new page to extend the stack
- the new page will not be physically contiguous with the rest of the stack
- see __do_page_fault() in /arch/arm/mm/fault.c

https://elinux.org/images/b/b0/Introduction_to_Memory_Management_in_Linux.pdf