Navid Malek
Project Phase 1

Operating Systems
Spring 2017

Part 2

finding the important functions which get information about some hardware

list of hardwares:

- CPU
- Memory
- network interface
- sound card
- power supplier

inorder to implement our information on system we should find :

1. what are the type of variables we want?
2. What is implementing them?
3. Where are they defined?
4. How to implement from number 2, to number 1?

so in this part we have to answer these 4 questions.

IF YOU WANT ALL OF THE HARDWARES YOU CAN EMAIL ME

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

1.CPU information

first we start grepping on "cpuinfo" thats the first clue!

console:

grep -rnw . -e "cpuinfo"

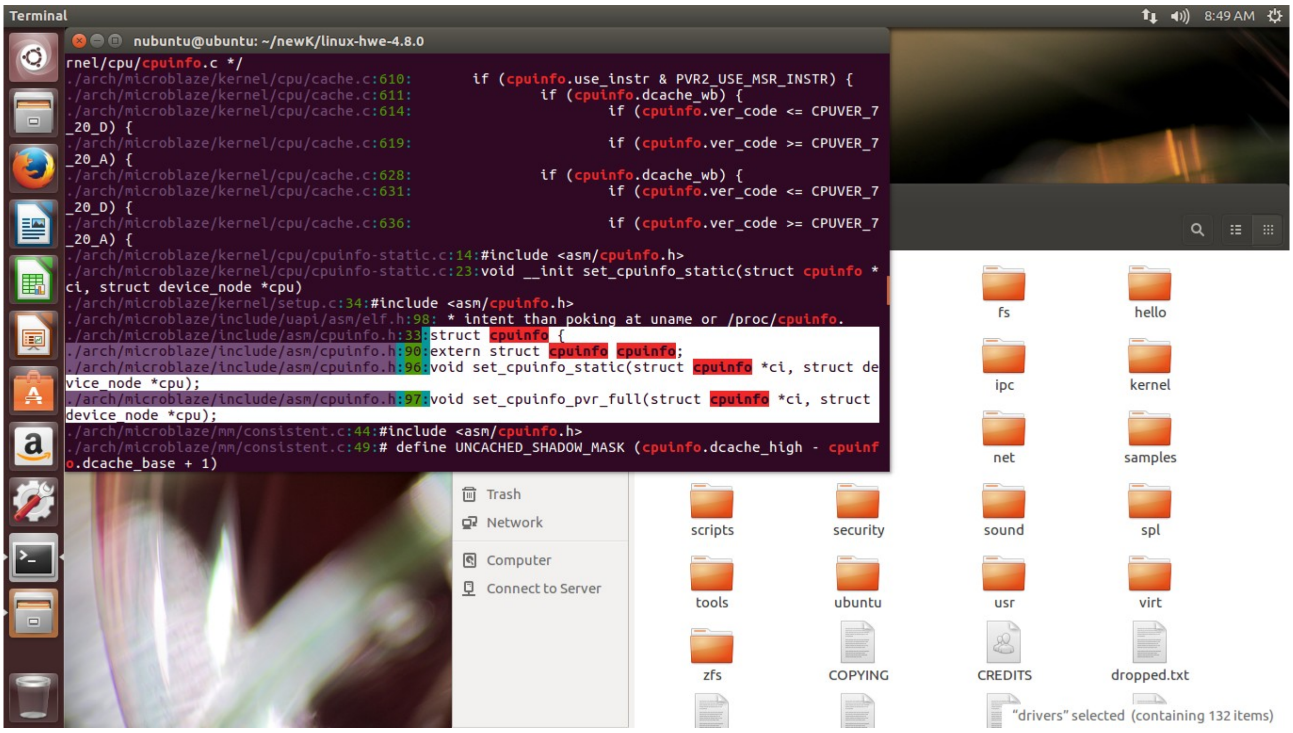

```
nubuntu@ubuntu: ~/newK/linux-hwe-4.8.0
rnel/cpu/cpuinfo.c */
./arch/microblaze/kernel/cpu/cache.c:610:       if (cpuinfo.use_instr & PVR2_USE_MSR_INSTR) {
./arch/microblaze/kernel/cpu/cache.c:611:               if (cpuinfo.dcache_wb) {
./arch/microblaze/kernel/cpu/cache.c:614:                       if (cpuinfo.ver_code <= CPUVER_7
_20_D) {
./arch/microblaze/kernel/cpu/cache.c:619:                       if (cpuinfo.ver_code >= CPUVER_7
_20_A) {
./arch/microblaze/kernel/cpu/cache.c:628:               if (cpuinfo.dcache_wb) {
./arch/microblaze/kernel/cpu/cache.c:631:                       if (cpuinfo.ver_code <= CPUVER_7
_20_D) {
./arch/microblaze/kernel/cpu/cache.c:636:                       if (cpuinfo.ver_code >= CPUVER_7
_20_A) {
./arch/microblaze/kernel/cpu/cpuinfo-static.c:14:#include <asm/cpuinfo.h>
./arch/microblaze/kernel/cpu/cpuinfo-static.c:23:void __init set_cpuinfo_static(struct cpuinfo *
ci, struct device_node *cpu)
./arch/microblaze/kernel/setup.c:34:#include <asm/cpuinfo.h>
./arch/microblaze/include/uapi/asm/elf.h:98: * intent than poking at uname or /proc/cpuinfo.
./arch/microblaze/include/asm/cpuinfo.h:33:struct cpuinfo {
./arch/microblaze/include/asm/cpuinfo.h:90:extern struct cpuinfo cpuinfo;
./arch/microblaze/include/asm/cpuinfo.h:96:void set_cpuinfo_static(struct cpuinfo *ci, struct de
vice_node *cpu);
./arch/microblaze/include/asm/cpuinfo.h:97:void set_cpuinfo_pvr_full(struct cpuinfo *ci, struct
device_node *cpu);
./arch/microblaze/mm/consistent.c:44:#include <asm/cpuinfo.h>
./arch/microblaze/mm/consistent.c:49:# define UNCACHED_SHADOW_MASK (cpuinfo.dcache_high - cpuinf
o.dcache_base + 1)
```

bingo!

As we see, there is a struct named cpu info that instance of it, has been passed by refrence to other functions. But this is for mips ! We want x86(intel) So we now grep only on struct cpuinfo.

Console:

grep -rnw . -e "struct cpuinfo"

and after that :

grep -rnw . -e "struct cpuinfo_x86"

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

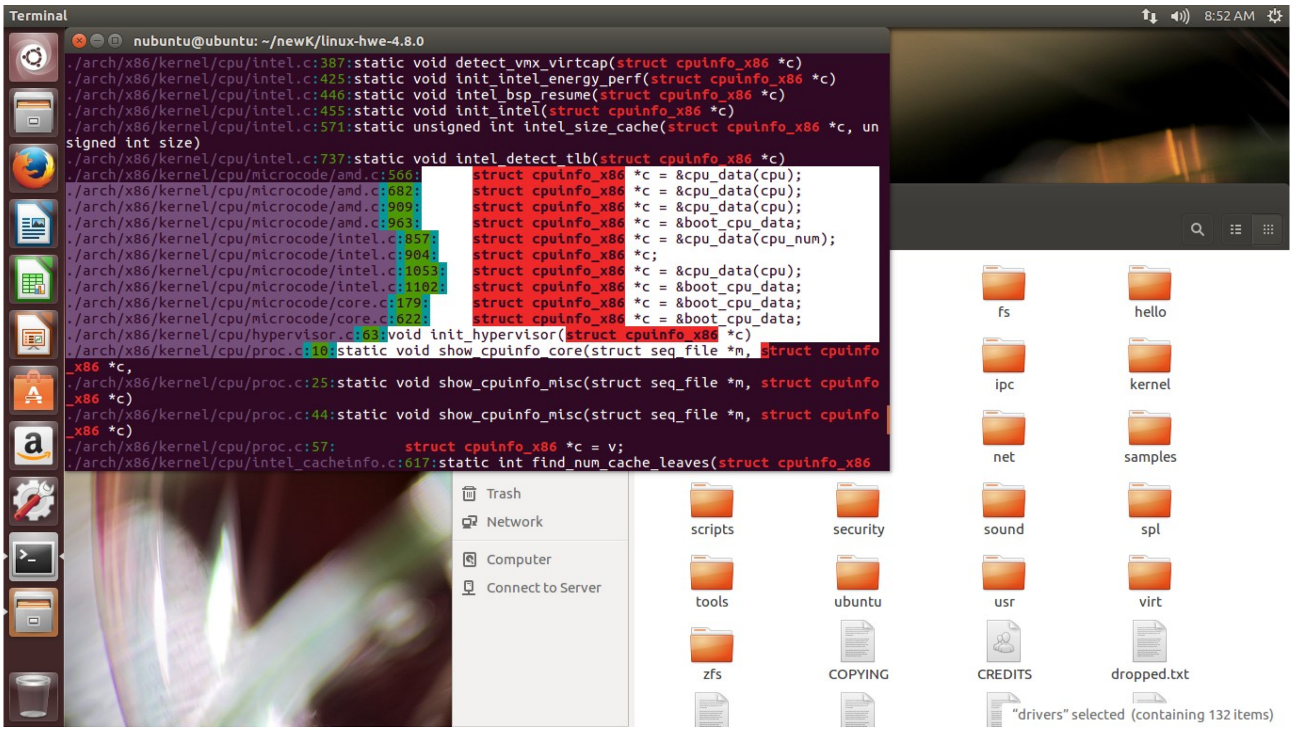

```
./arch/x86/kernel/cpu/intel.c:387:static void detect_vmx_virtcap(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/intel.c:425:static void init_intel_energy_perf(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/intel.c:446:static void intel_bsp_resume(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/intel.c:455:static void init_intel(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/intel.c:571:static unsigned int intel_size_cache(struct cpuinfo_x86 *c, un
signed int size)
./arch/x86/kernel/cpu/intel.c:737:static void intel_detect_tlb(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/microcode/amd.c:566:        struct cpuinfo_x86 *c = &cpu_data(cpu);
./arch/x86/kernel/cpu/microcode/amd.c:682:        struct cpuinfo_x86 *c = &cpu_data(cpu);
./arch/x86/kernel/cpu/microcode/amd.c:909:        struct cpuinfo_x86 *c = &cpu_data(cpu);
./arch/x86/kernel/cpu/microcode/amd.c:963:        struct cpuinfo_x86 *c = &boot_cpu_data;
./arch/x86/kernel/cpu/microcode/intel.c:857:      struct cpuinfo_x86 *c = &cpu_data(cpu_num);
./arch/x86/kernel/cpu/microcode/intel.c:904:      struct cpuinfo_x86 *c;
./arch/x86/kernel/cpu/microcode/intel.c:1053:     struct cpuinfo_x86 *c = &cpu_data(cpu);
./arch/x86/kernel/cpu/microcode/intel.c:1102:     struct cpuinfo_x86 *c = &boot_cpu_data;
./arch/x86/kernel/cpu/microcode/core.c:179:       struct cpuinfo_x86 *c = &boot_cpu_data;
./arch/x86/kernel/cpu/microcode/core.c:622:       struct cpuinfo_x86 *c = &boot_cpu_data;
./arch/x86/kernel/cpu/hypervisor.c:63:void init_hypervisor(struct cpuinfo_x86 *c)
./arch/x86/kernel/cpu/proc.c:10:static void show_cpuinfo_core(struct seq_file *m, struct cpuinfo
_x86 *c,
./arch/x86/kernel/cpu/proc.c:25:static void show_cpuinfo_misc(struct seq_file *m, struct cpuinfo
_x86 *c)
./arch/x86/kernel/cpu/proc.c:44:static void show_cpuinfo_misc(struct seq_file *m, struct cpuinfo
_x86 *c)
./arch/x86/kernel/cpu/proc.c:57:        struct cpuinfo_x86 *c = v;
./arch/x86/kernel/cpu/intel_cacheinfo.c:617:static int find_num_cache_leaves(struct cpuinfo_x86
```

okay! So cpuinfo_x86 has been passed by reference and cpu_data is implementing it.

The argument for cpu_data is cpu,cpu_num,0,1….

So the atgument type is int,and its number of core.

Okay first question! Type of variables

Go to http://lxr.free-electrons.com/

and search about cpuinfo_x86

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

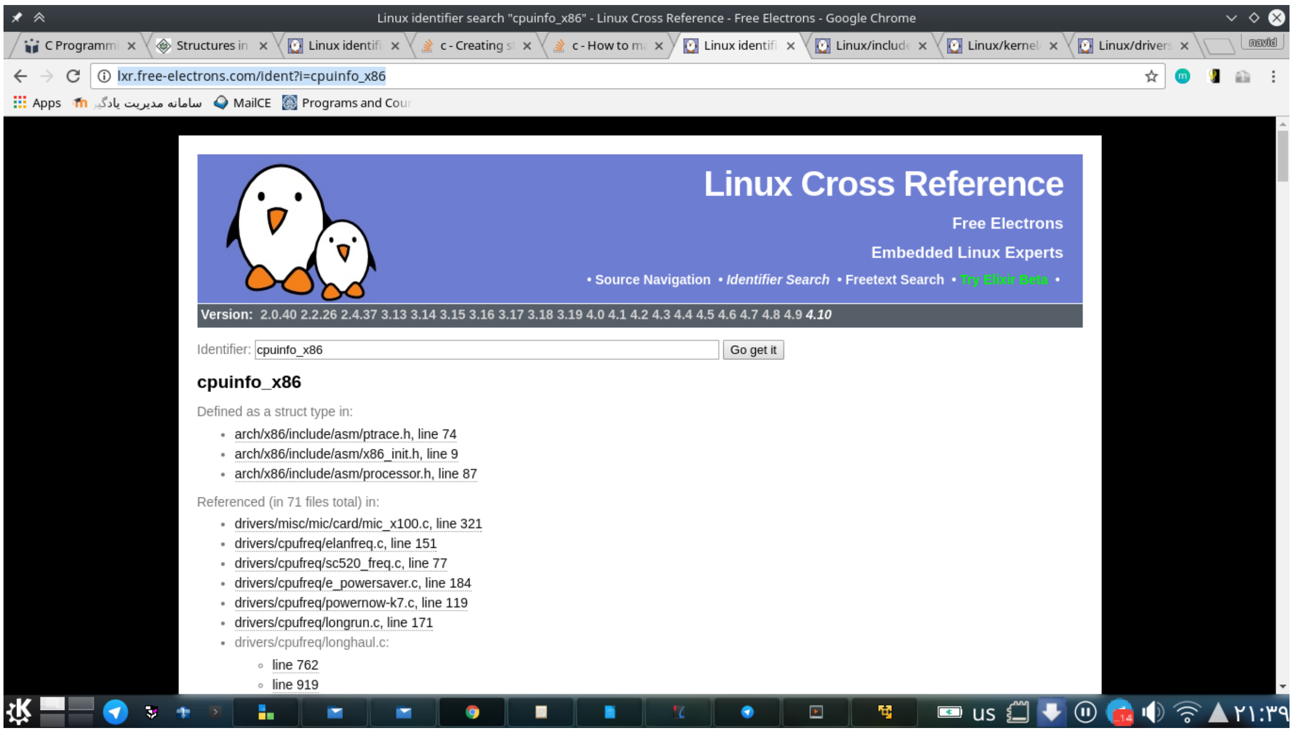
Linux identifier search "cpuinfo_x86" - Linux Cross Reference - Free Electrons - Google Chrome
lxr.free-electrons.com/ident?i=cpuinfo_x86
Linux Cross Reference
Free Electrons
Embedded Linux Experts
• Source Navigation • Identifier Search • Freetext Search • Try Elixir Beta •
Version: 2.0.40 2.2.26 2.4.37 3.13 3.14 3.15 3.16 3.17 3.18 3.19 4.0 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10
Identifier: cpuinfo_x86
Go get it

cpuinfo_x86

Defined as a struct type in:

- arch/x86/include/asm/ptrace.h, line 74
- arch/x86/include/asm/x86_init.h, line 9
- arch/x86/include/asm/processor.h, line 87

Referenced (in 71 files total) in:

- drivers/misc/mic/card/mic_x100.c, line 321
- drivers/cpufreq/elanfreq.c, line 151
- drivers/cpufreq/sc520_freq.c, line 77
- drivers/cpufreq/e_powersaver.c, line 184
- drivers/cpufreq/powernow-k7.c, line 119
- drivers/cpufreq/longrun.c, line 171
- drivers/cpufreq/longhaul.c:
  - line 762
  - line 919

okay so we know where it’s defined. Look into it

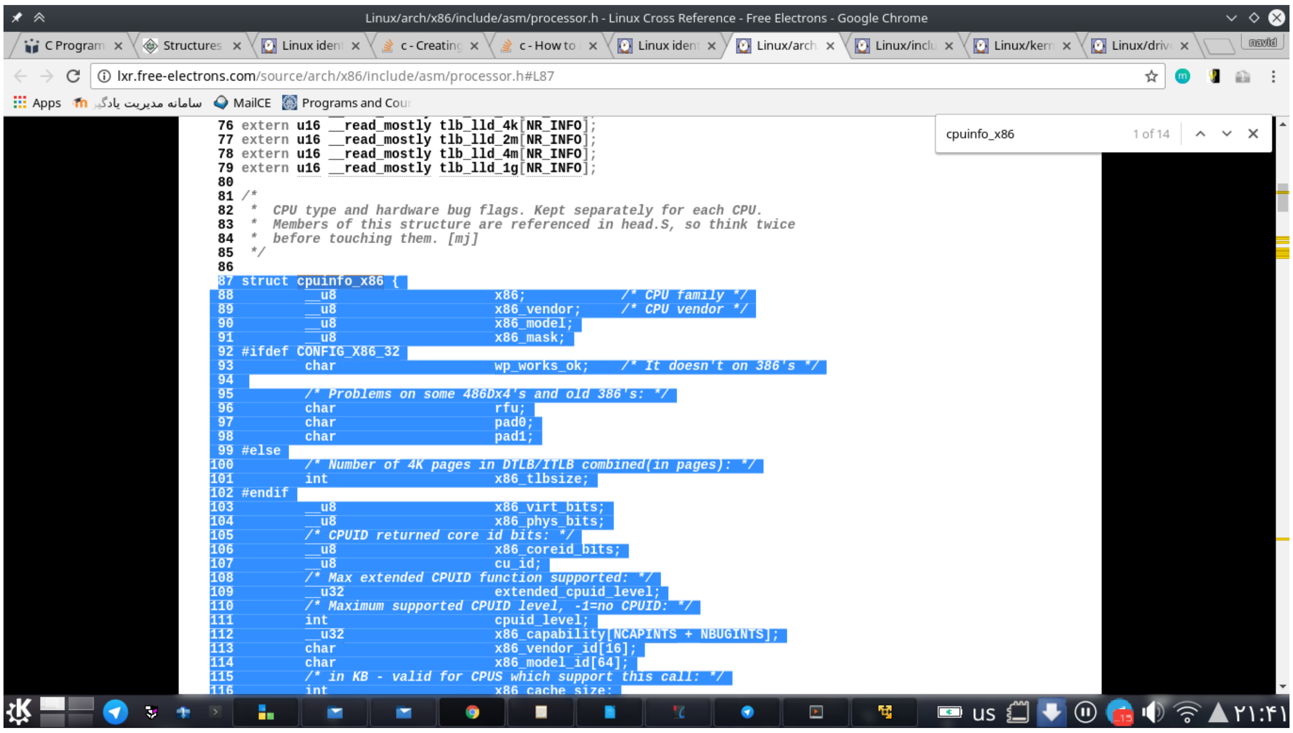
Linux/arch/x86/include/asm/processor.h - Linux Cross Reference - Free Electrons - Google Chrome
lxr.free-electrons.com/source/arch/x86/include/asm/processor.h#L87
cpuinfo_x86
1 of 14

```
extern u16 __read_mostly tlb_lld_4k[NR_INFO];
extern u16 __read_mostly tlb_lld_2m[NR_INFO];
extern u16 __read_mostly tlb_lld_4m[NR_INFO];
extern u16 __read_mostly tlb_lld_1g[NR_INFO];

/*
 *  CPU type and hardware bug flags. Kept separately for each CPU.
 *  Members of this structure are referenced in head.S, so think twice
 *  before touching them. [mj]
 */

struct cpuinfo_x86 {
        __u8                    x86;            /* CPU family */
        __u8                    x86_vendor;     /* CPU vendor */
        __u8                    x86_model;
        __u8                    x86_mask;
#ifdef CONFIG_X86_32
        char                    wp_works_ok;    /* It doesn't on 386's */

        /* Problems on some 486Dx4's and old 386's: */
        char                    rfu;
        char                    pad0;
        char                    pad1;
#else
        /* Number of 4K pages in DTLB/ITLB combined(in pages): */
        int                     x86_tlbsize;
#endif
        __u8                    x86_virt_bits;
        __u8                    x86_phys_bits;
        /* CPUID returned core id bits: */
        __u8                    x86_coreid_bits;
        __u8                    cu_id;
        /* Max extended CPUID function supported: */
        __u32                   extended_cpuid_level;
        /* Maximum supported CPUID level, -1=no CPUID: */
        int                     cpuid_level;
        __u32                   x86_capability[NCAPINTS + NBUGINTS];
        char                    x86_vendor_id[16];
        char                    x86_model_id[64];
        /* in KB - valid for CPUS which support this call: */
        int                     x86_cache_size;
```

here is the defenitions and types!

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

Now question 2!

we found cpu data is implementing it.

Question 3:

in order to use it we have to know where it is defined, again seach in the freeelectrons.com and go to where it is defined:

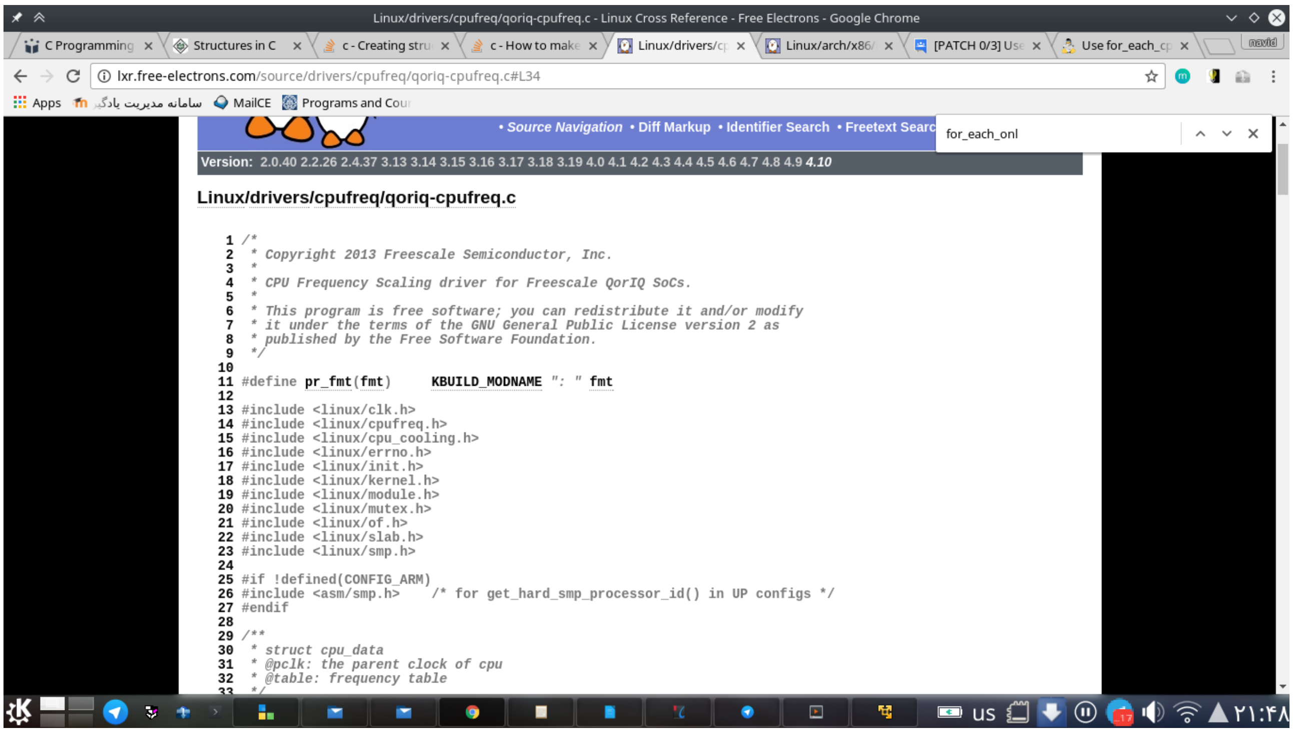
Linux/drivers/cpufreq/qoriq-cpufreq.c - Linux Cross Reference - Free Electrons - Google Chrome
lxr.free-electrons.com/source/drivers/cpufreq/qoriq-cpufreq.c#L34
• Source Navigation • Diff Markup • Identifier Search • Freetext Search
for_each_onl
Version: 2.0.40 2.2.26 2.4.37 3.13 3.14 3.15 3.16 3.17 3.18 3.19 4.0 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 4.10

Linux/drivers/cpufreq/qoriq-cpufreq.c

```
/*
 * Copyright 2013 Freescale Semiconductor, Inc.
 *
 * CPU Frequency Scaling driver for Freescale QorIQ SoCs.
 *
 * This program is free software; you can redistribute it and/or modify
 * it under the terms of the GNU General Public License version 2 as
 * published by the Free Software Foundation.
 */

#define pr_fmt(fmt)     KBUILD_MODNAME ": " fmt

#include <linux/clk.h>
#include <linux/cpufreq.h>
#include <linux/cpu_cooling.h>
#include <linux/errno.h>
#include <linux/init.h>
#include <linux/kernel.h>
#include <linux/module.h>
#include <linux/mutex.h>
#include <linux/of.h>
#include <linux/slab.h>
#include <linux/smp.h>

#if !defined(CONFIG_ARM)
#include <asm/smp.h>    /* for get_hard_smp_processor_id() in UP configs */
#endif

/**
 * struct cpu_data
 * @pclk: the parent clock of cpu
 * @table: frequency table
```

we don’t know what to include so we include all of them :D

Question 4:

how to implement it??

there is a function :   for_each_online_cpu

for knowing the functionaly of it we grep on it and see where it is used and how.

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

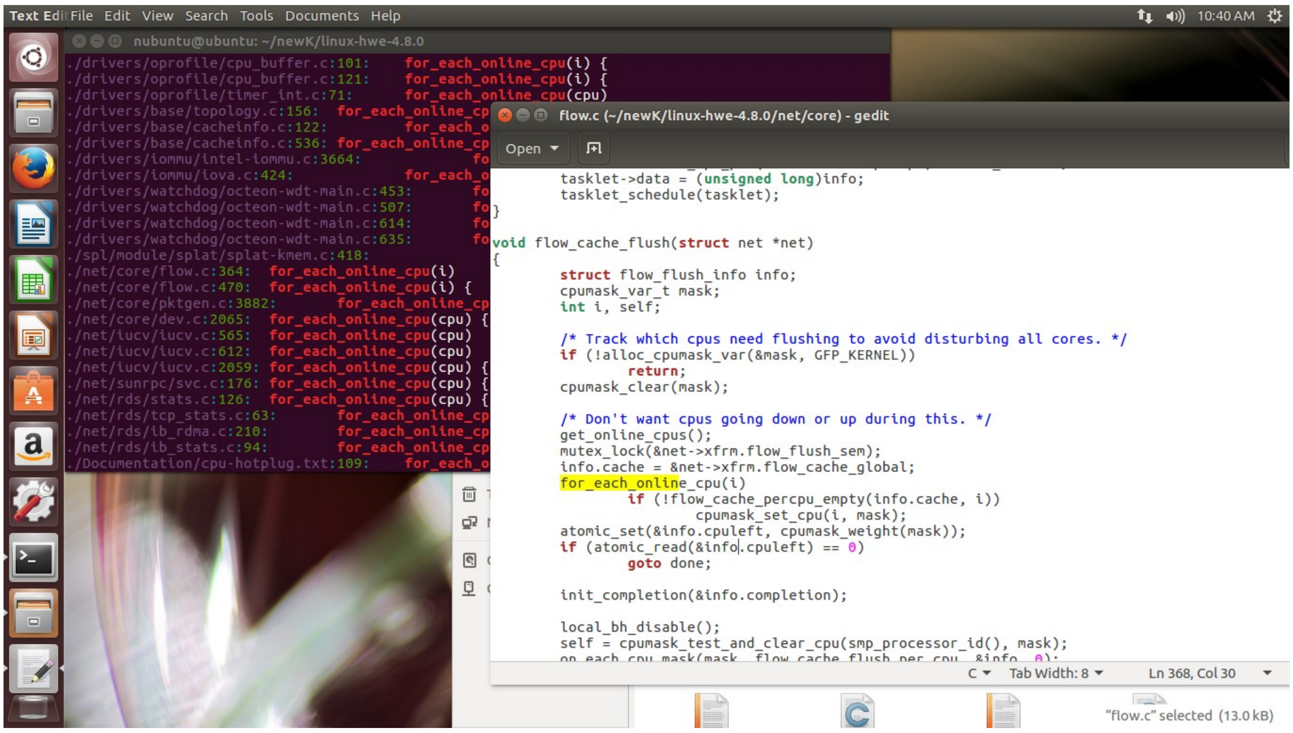
Text Edi File Edit View Search Tools Documents Help
10:40 AM
nubuntu@ubuntu: ~/newK/linux-hwe-4.8.0
flow.c (~/newK/linux-hwe-4.8.0/net/core) - gedit
Open
C
Tab Width: 8
Ln 368, Col 30
"flow.c" selected (13.0 kB)

```
./drivers/oprofile/cpu_buffer.c:101:     for_each_online_cpu(i) {
./drivers/oprofile/cpu_buffer.c:121:     for_each_online_cpu(i) {
./drivers/oprofile/timer_int.c:71:       for_each_online_cpu(cpu)
./drivers/base/topology.c:156:  for_each_online_cp
./drivers/base/cacheinfo.c:122:          for_each_o
./drivers/base/cacheinfo.c:536: for_each_online_cp
./drivers/iommu/intel-iommu.c:3664:               fo
./drivers/iommu/iova.c:424:              for_each_o
./drivers/watchdog/octeon-wdt-main.c:453:         fo
./drivers/watchdog/octeon-wdt-main.c:507:         fo
./drivers/watchdog/octeon-wdt-main.c:614:         fo
./drivers/watchdog/octeon-wdt-main.c:635:         fo
./spl/module/splat/splat-kmem.c:418:
./net/core/flow.c:364:  for_each_online_cpu(i)
./net/core/flow.c:470:  for_each_online_cpu(i) {
./net/core/pktgen.c:3882:       for_each_online_cp
./net/core/dev.c:2065:  for_each_online_cpu(cpu) {
./net/iucv/iucv.c:565:  for_each_online_cpu(cpu)
./net/iucv/iucv.c:612:  for_each_online_cpu(cpu)
./net/iucv/iucv.c:2059: for_each_online_cpu(cpu) {
./net/sunrpc/svc.c:176: for_each_online_cpu(cpu) {
./net/rds/stats.c:126:  for_each_online_cpu(cpu) {
./net/rds/tcp_stats.c:63:       for_each_online_cp
./net/rds/ib_rdma.c:210:        for_each_online_cp
./net/rds/ib_stats.c:94:        for_each_online_cp
./Documentation/cpu-hotplug.txt:109:     for_each_o
```

```c
        tasklet->data = (unsigned long)info;
        tasklet_schedule(tasklet);
}

void flow_cache_flush(struct net *net)
{
        struct flow_flush_info info;
        cpumask_var_t mask;
        int i, self;

        /* Track which cpus need flushing to avoid disturbing all cores. */
        if (!alloc_cpumask_var(&mask, GFP_KERNEL))
                return;
        cpumask_clear(mask);

        /* Don't want cpus going down or up during this. */
        get_online_cpus();
        mutex_lock(&net->xfrm.flow_flush_sem);
        info.cache = &net->xfrm.flow_cache_global;
        for_each_online_cpu(i)
                if (!flow_cache_percpu_empty(info.cache, i))
                        cpumask_set_cpu(i, mask);
        atomic_set(&info.cpuleft, cpumask_weight(mask));
        if (atomic_read(&info.cpuleft) == 0)
                goto done;

        init_completion(&info.completion);

        local_bh_disable();
        self = cpumask_test_and_clear_cpu(smp_processor_id(), mask);
```

So here is the syntax in one random file.

Note: if you want to understand better and know how to implement this info (show it in user space) skip to the part 3 and then return here again.

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

2.memory info

ok grepping on meminfo does not return anything special, so we search the web.

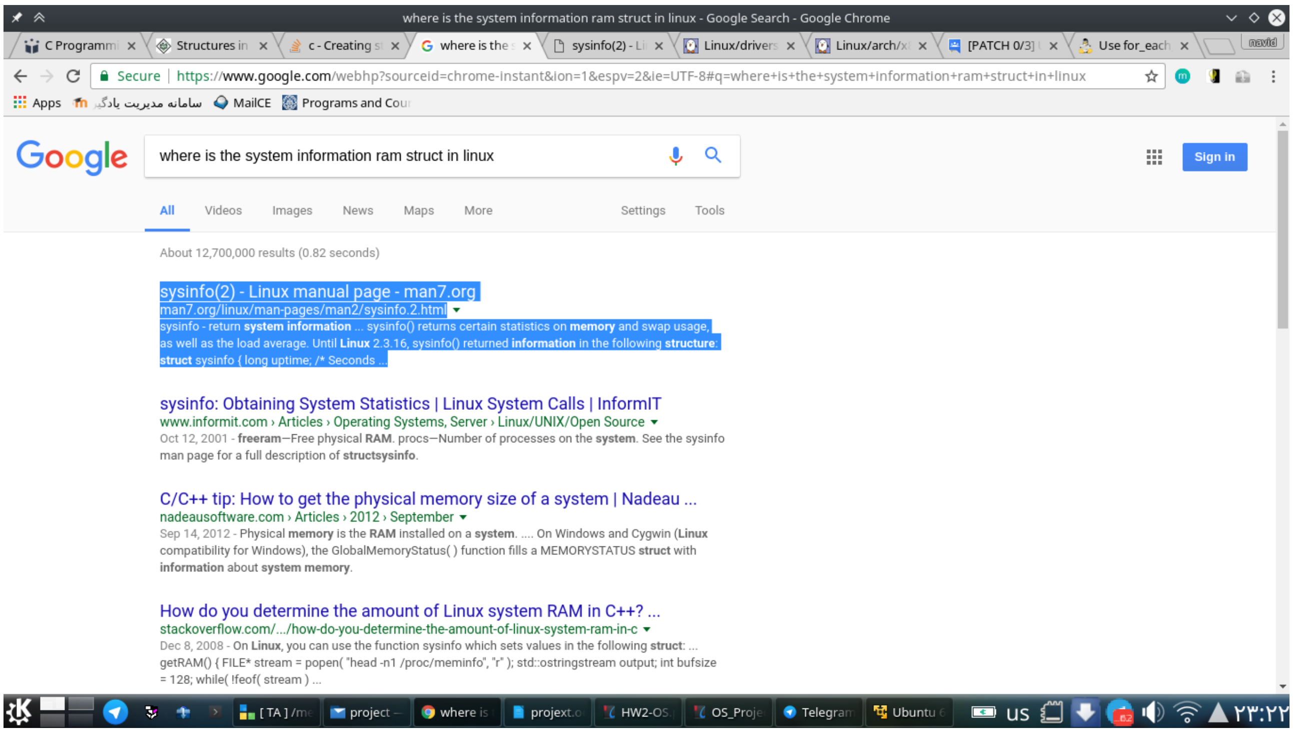
where is the system information ram struct in linux - Google Search - Google Chrome
C Programmi
Structures in
c - Creating s
where is the s
sysinfo(2) - Li
Linux/drivers
Linux/arch/x
[PATCH 0/3]
Use for_each
Secure
https://www.google.com/webhp?sourceid=chrome-instant&ion=1&espv=2&ie=UTF-8#q=where+is+the+system+information+ram+struct+in+linux
Apps
سامانه مدیریت یادگیری
MailCE
Programs and Cou
where is the system information ram struct in linux
Sign in
All
Videos
Images
News
Maps
More
Settings
Tools

About 12,700,000 results (0.82 seconds)

sysinfo(2) - Linux manual page - man7.org
man7.org/linux/man-pages/man2/sysinfo.2.html
sysinfo - return **system information** ... sysinfo() returns certain statistics on **memory** and swap usage, as well as the load average. Until **Linux** 2.3.16, sysinfo() returned **information** in the following **structure**: **struct** sysinfo { long uptime; /* Seconds ...

sysinfo: Obtaining System Statistics | Linux System Calls | InformIT
www.informit.com › Articles › Operating Systems, Server › Linux/UNIX/Open Source
Oct 12, 2001 - **freeram**—Free physical **RAM**. procs—Number of processes on the **system**. See the sysinfo man page for a full description of **structsysinfo**.

C/C++ tip: How to get the physical memory size of a system | Nadeau ...
nadeausoftware.com › Articles › 2012 › September
Sep 14, 2012 - Physical **memory** is the **RAM** installed on a **system**. .... On Windows and Cygwin (**Linux** compatibility for Windows), the GlobalMemoryStatus( ) function fills a MEMORYSTATUS **struct** with **information** about **system memory**.

How do you determine the amount of Linux system RAM in C++? ...
stackoverflow.com/.../how-do-you-determine-the-amount-of-linux-system-ram-in-c
Dec 8, 2008 - On **Linux**, you can use the function sysinfo which sets values in the following **struct**: ... getRAM() { FILE* stream = popen("head -n1 /proc/meminfo", "r"); std::ostringstream output; int bufsize = 128; while(!feof(stream) ...

Seems like we have it

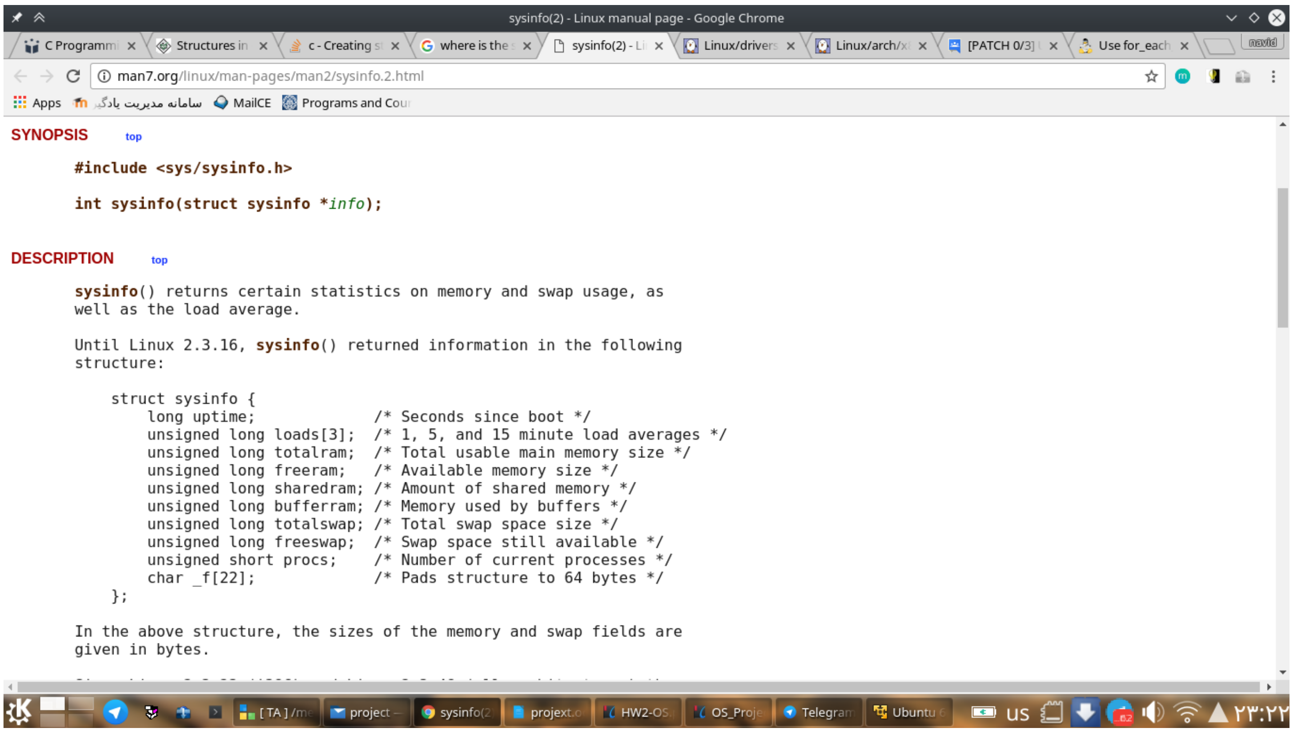
sysinfo(2) - Linux manual page - Google Chrome
man7.org/linux/man-pages/man2/sysinfo.2.html

SYNOPSIS top

```
       #include <sys/sysinfo.h>

       int sysinfo(struct sysinfo *info);
```

DESCRIPTION top

```
       sysinfo() returns certain statistics on memory and swap usage, as
       well as the load average.

       Until Linux 2.3.16, sysinfo() returned information in the following
       structure:

           struct sysinfo {
               long uptime;             /* Seconds since boot */
               unsigned long loads[3];  /* 1, 5, and 15 minute load averages */
               unsigned long totalram;  /* Total usable main memory size */
               unsigned long freeram;   /* Available memory size */
               unsigned long sharedram; /* Amount of shared memory */
               unsigned long bufferram; /* Memory used by buffers */
               unsigned long totalswap; /* Total swap space size */
               unsigned long freeswap;  /* Swap space still available */
               unsigned short procs;    /* Number of current processes */
               char _f[22];             /* Pads structure to 64 bytes */
           };

       In the above structure, the sizes of the memory and swap fields are
       given in bytes.
```

thats is all we need.

Now to know where is it we use free electrons site:

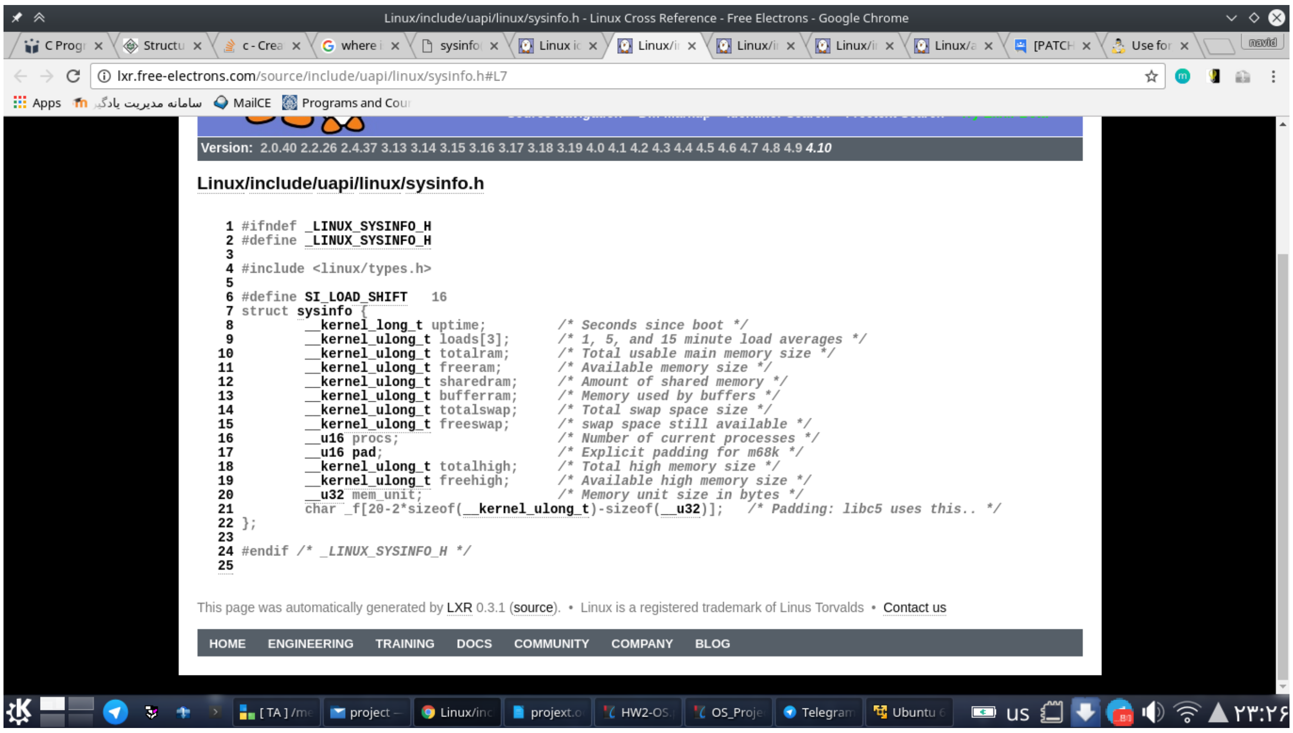

Linux/include/uapi/linux/sysinfo.h - Linux Cross Reference - Free Electrons - Google Chrome

lxr.free-electrons.com/source/include/uapi/linux/sysinfo.h#L7

Version: 2.0.40 2.2.26 2.4.37 3.13 3.14 3.15 3.16 3.17 3.18 3.19 4.0 4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9 **4.10**

Linux/include/uapi/linux/sysinfo.h

```
#ifndef _LINUX_SYSINFO_H
#define _LINUX_SYSINFO_H

#include <linux/types.h>

#define SI_LOAD_SHIFT   16
struct sysinfo {
        __kernel_long_t uptime;         /* Seconds since boot */
        __kernel_ulong_t loads[3];      /* 1, 5, and 15 minute load averages */
        __kernel_ulong_t totalram;      /* Total usable main memory size */
        __kernel_ulong_t freeram;       /* Available memory size */
        __kernel_ulong_t sharedram;     /* Amount of shared memory */
        __kernel_ulong_t bufferram;     /* Memory used by buffers */
        __kernel_ulong_t totalswap;     /* Total swap space size */
        __kernel_ulong_t freeswap;      /* swap space still available */
        __u16 procs;                    /* Number of current processes */
        __u16 pad;                      /* Explicit padding for m68k */
        __kernel_ulong_t totalhigh;     /* Total high memory size */
        __kernel_ulong_t freehigh;      /* Available high memory size */
        __u32 mem_unit;                 /* Memory unit size in bytes */
        char _f[20-2*sizeof(__kernel_ulong_t)-sizeof(__u32)];   /* Padding: libc5 uses this.. */
};

#endif /* _LINUX_SYSINFO_H */

```

This page was automatically generated by LXR 0.3.1 (source). • Linux is a registered trademark of Linus Torvalds • Contact us

now to question 2, what is implementing it?

We now grep on the :

grep -rnw . -e "struct sysinfo"

and review two random .c files to see how sysinfo is implemented

Navid Malek
Project Phase 1

Operating Systems
Spring 2017

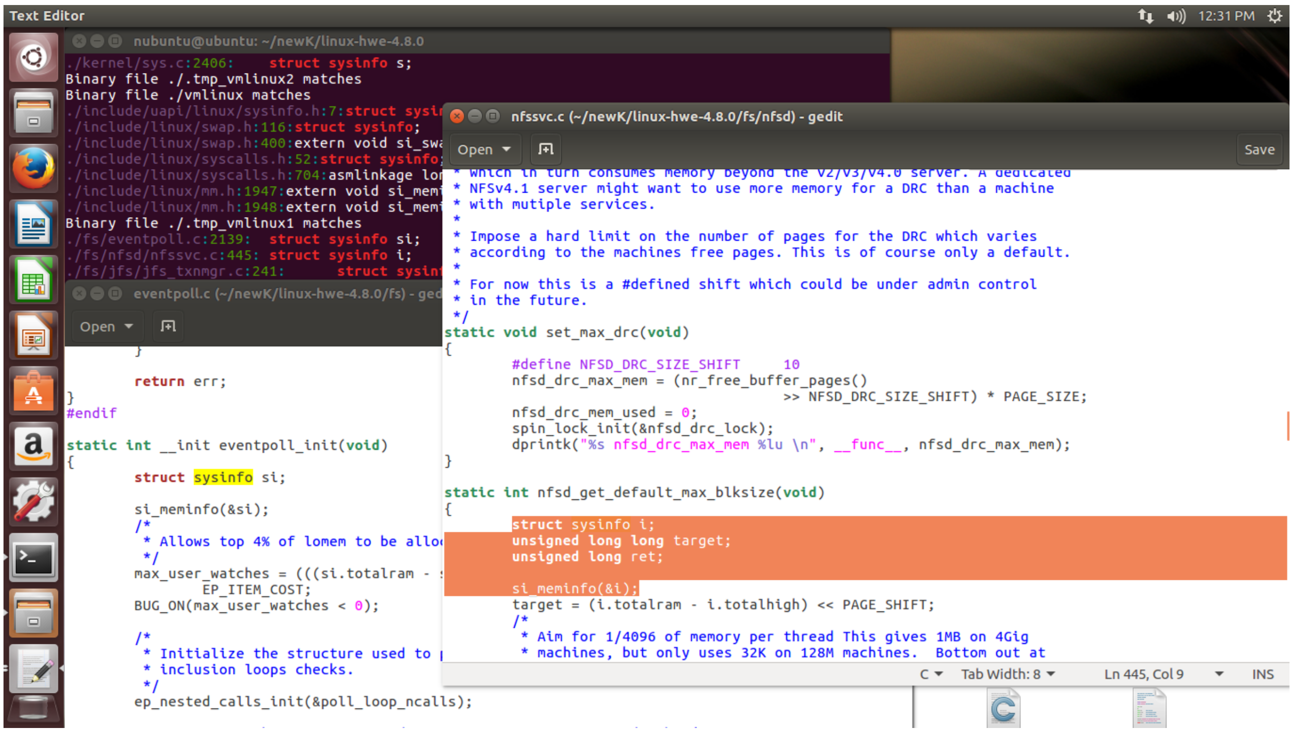

nubuntu@ubuntu: ~/newK/linux-hwe-4.8.0

```
./kernel/sys.c:2406:    struct sysinfo s;
Binary file ./.tmp_vmlinux2 matches
Binary file ./vmlinux matches
./include/uapi/linux/sysinfo.h:7:struct sysi
./include/linux/swap.h:116:struct sysinfo;
./include/linux/swap.h:400:extern void si_swa
./include/linux/syscalls.h:52:struct sysinfo;
./include/linux/syscalls.h:704:asmlinkage lo
./include/linux/mm.h:1947:extern void si_mem
./include/linux/mm.h:1948:extern void si_mem
Binary file ./.tmp_vmlinux1 matches
./fs/eventpoll.c:2139:  struct sysinfo si;
./fs/nfsd/nfssvc.c:445: struct sysinfo i;
./fs/jfs/jfs_txnmgr.c:241:      struct sysin
```

eventpoll.c (~/newK/linux-hwe-4.8.0/fs) - ged

```
        }

        return err;
}
#endif

static int __init eventpoll_init(void)
{
        struct sysinfo si;

        si_meminfo(&si);
        /*
         * Allows top 4% of lomem to be allo
         */
        max_user_watches = (((si.totalram - 
                EP_ITEM_COST;
        BUG_ON(max_user_watches < 0);

        /*
         * Initialize the structure used to 
         * inclusion loops checks.
         */
        ep_nested_calls_init(&poll_loop_ncalls);
```

nfssvc.c (~/newK/linux-hwe-4.8.0/fs/nfsd) - gedit

```
 * which in turn consumes memory beyond the v2/v3/v4.0 server. A dedicated
 * NFSv4.1 server might want to use more memory for a DRC than a machine
 * with mutiple services.
 *
 * Impose a hard limit on the number of pages for the DRC which varies
 * according to the machines free pages. This is of course only a default.
 *
 * For now this is a #defined shift which could be under admin control
 * in the future.
 */
static void set_max_drc(void)
{
        #define NFSD_DRC_SIZE_SHIFT     10
        nfsd_drc_max_mem = (nr_free_buffer_pages()
                                >> NFSD_DRC_SIZE_SHIFT) * PAGE_SIZE;
        nfsd_drc_mem_used = 0;
        spin_lock_init(&nfsd_drc_lock);
        dprintk("%s nfsd_drc_max_mem %lu \n", __func__, nfsd_drc_max_mem);
}

static int nfsd_get_default_max_blksize(void)
{
        struct sysinfo i;
        unsigned long long target;
        unsigned long ret;

        si_meminfo(&i);
        target = (i.totalram - i.totalhigh) << PAGE_SHIFT;
        /*
         * Aim for 1/4096 of memory per thread This gives 1MB on 4Gig
         * machines, but only uses 32K on 128M machines.  Bottom out at
```

as we see si_meminfo(&i) implements sysinfo.

now the 3rd question:

we now find where is meminfo implemented to know what should we include in kernel module.

Lets see meminfo definition.

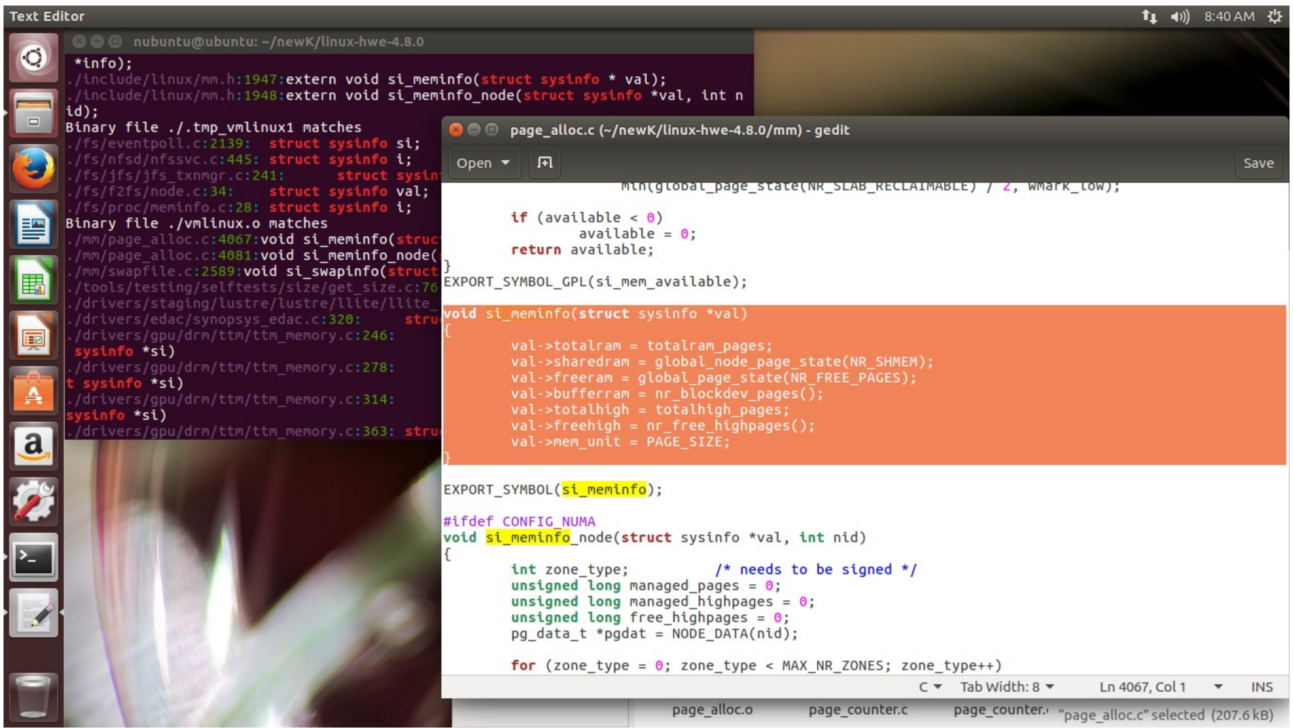

nubuntu@ubuntu: ~/newK/linux-hwe-4.8.0

```
 *info);
./include/linux/mm.h:1947:extern void si_meminfo(struct sysinfo * val);
./include/linux/mm.h:1948:extern void si_meminfo_node(struct sysinfo *val, int n
id);
Binary file ./.tmp_vmlinux1 matches
./fs/eventpoll.c:2139:  struct sysinfo si;
./fs/nfsd/nfssvc.c:445: struct sysinfo i;
./fs/jfs/jfs_txnmgr.c:241:      struct sysin
./fs/f2fs/node.c:34:    struct sysinfo val;
./fs/proc/meminfo.c:28: struct sysinfo i;
Binary file ./vmlinux.o matches
./mm/page_alloc.c:4067:void si_meminfo(struc
./mm/page_alloc.c:4081:void si_meminfo_node(
./mm/swapfile.c:2589:void si_swapinfo(struct
./tools/testing/selftests/size/get_size.c:76
./drivers/staging/lustre/lustre/llite/llite_
./drivers/edac/synopsys_edac.c:320:     stru
./drivers/gpu/drm/ttm/ttm_memory.c:246:
 sysinfo *si)
./drivers/gpu/drm/ttm/ttm_memory.c:278:
t sysinfo *si)
./drivers/gpu/drm/ttm/ttm_memory.c:314:
sysinfo *si)
./drivers/gpu/drm/ttm/ttm_memory.c:363: stru
```

page_alloc.c (~/newK/linux-hwe-4.8.0/mm) - gedit

```
                        min(global_page_state(NR_SLAB_RECLAIMABLE) / 2, wmark_low);

        if (available < 0)
                available = 0;
        return available;
}
EXPORT_SYMBOL_GPL(si_mem_available);

void si_meminfo(struct sysinfo *val)
{
        val->totalram = totalram_pages;
        val->sharedram = global_node_page_state(NR_SHMEM);
        val->freeram = global_page_state(NR_FREE_PAGES);
        val->bufferram = nr_blockdev_pages();
        val->totalhigh = totalhigh_pages;
        val->freehigh = nr_free_highpages();
        val->mem_unit = PAGE_SIZE;
}

EXPORT_SYMBOL(si_meminfo);

#ifdef CONFIG_NUMA
void si_meminfo_node(struct sysinfo *val, int nid)
{
        int zone_type;          /* needs to be signed */
        unsigned long managed_pages = 0;
        unsigned long managed_highpages = 0;
        unsigned long free_highpages = 0;
        pg_data_t *pgdat = NODE_DATA(nid);

        for (zone_type = 0; zone_type < MAX_NR_ZONES; zone_type++)
```

Operating Systems
Spring 2017

Question 4.

the implemention is easy its just calling meminfo function.

IF YOU WANT ALL OF THE HARDWARES YOU CAN EMAIL ME