



What I didn't know about coding in college

Getting Code into the Linux Kernel
and
the Power of Peer Review

Dave Hansen
haveblue@us.ibm.com

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, with horizontal lines through the letters, positioned in the bottom left corner of the slide.

IBM



Legal Statement

Yes, I work for a big company...

This work represents the views of the author(s) and does not necessarily reflect the views of IBM Corporation.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries: IBM (logo). A full list of U.S. trademarks owned by IBM may be found at <http://www.ibm.com/legal/copytrade.shtml>.

Linux is a registered trademark of Linus Torvalds.

Other company, product, and service names may be trademarks or service marks of others.

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, with horizontal stripes through the letters. It is positioned in the bottom left corner of the slide.

IBM



Classes: Coding in Isolation

- Discourages Cheating
- Develop peculiar habits
- No fear of humiliation
- No shame
- No one to catch your mistakes
- Projects have to “just work”
- Selfish behavior encouraged



Real Word: the Bazaar

- Cheating and stealing encouraged!
- Saves time and money
- Bad code means shame and humiliation
- Selfless behavior encouraged



Additude in the Bazaar

From Greg Kroah-Hartman:

Linux USB maintainer

(gets lots of code into the kernel)

- I want to read your code
- I want you to read my code
- I want you to fix my bugs
- I want you to change my code
- I want to build on your code

IBM



How do I trick you?

(to do work for me)

- I write code that is interesting
- I write code that **looks** good
 - Follow conventions
 - Be consistent
 - Self-document
- I write code that works well
 - use language features
 - use thorough debugging techniques



I Write Dull Code

- If you wrote it as cleverly as possible, then it is by definition impossible to debug
- Things are harder to debug than to write in the first place



Coding Style

- Every (sub)project different
- Far too lax in school
- Each academic class should have different requirements



I need more than 80 columns!

- Narrow means more windows at once
- Low levels of nested control
- Encourages good functions
 - small
 - simple
 - named loops
- Going past 80 means DANGER!

IBM



C

I didn't know C!

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, with horizontal lines through the letters, positioned in the bottom left corner of the slide.



static means private

- local
 - to a function
 - to a .c file
- “non-global” global variables
- combine with inline for type-safe “macros”



`extern` – the lazy keyword

- prototype – I'll define this later
- variables or functions
- used to access a something in another .c file
 - declare in .c
 - extern in .h
- can not be used with static



Macros vs. `inline`

- Use `inline` whenever possible
 - type safety
 - just as fast
 - make the compiler do the work
- Use Macros
 - structure initialization
 - assembly
 - `#include` hell



#ifdefs

- Evil in .c files
- OK in .h files
 - They're ugly anyway
 - Free Polymorphism

include/linux/preempt.h:

```
#ifndef CONFIG_PREEMPT
#define preempt_enable() \
do { \
    preempt_enable_no_resched(); \
    preempt_check_resched(); \
} while (0)
#else
#define preempt_enable() do { } while (0)
#endif
```



Types

- **unsigned long**
 - “natural type”
 - `sizeof(void *)`
 - good for array indexes
- **short**
 - rarely faster
 - storage space only



How to Write Good Loops

- Use small xterms
- Don't embed loops
- **break** and **continue** are OK
- Brackets are bad



Loop without continue

```
int sum_positive_numbers(int *array,
                        unsigned long size)
{
    for (i=0; i<size; i++) {
        if (array[i] >= 0) {
            if (i >= ARRAY_SIZE) {
                break;
            } else {
                total += array[i];
            }
        }
    }
}
```



Loop with continue

```
int sum_positive_numbers(int *array,
                        unsigned long size)
{
    for (i=0; i<size; i++) {
        if (i >= ARRAY_SIZE)
            break;

        if (array[i] <= 0)
            continue;

        total += array[i];
    }
}
```



Real Code

```
/*calculate node_start_pfn/node_end_pfn arrays*/
for (nid = 0; nid < numnodes; nid++) {
    int been_here_before = 0;

    for (j = 0; j < num_memory_chunks; j++){
        if (node_memory_chunk[j].nid == nid) {
            if (been_here_before == 0) {
                node_start_pfn[nid] = node_memory_chunk[j].start_pfn;
                node_end_pfn[nid] = node_memory_chunk[j].end_pfn;
                been_here_before = 1;
            } else { /* We've found another chunk of memory for the node */
                if (node_start_pfn[nid] < node_memory_chunk[j].start_pfn) {
                    node_end_pfn[nid] = node_memory_chunk[j].end_pfn;
                }
            }
        }
    }
}

for (nid = 0; nid < numnodes; nid++) {
    for (j = 0; j < num_memory_chunks; j++) {
        if (node_memory_chunk[j].nid != nid)
            continue;

        max_pages_per_node - node_present_pages;
        node_present_pages += chunk_size(j);

        if (node_start_pfn[nid] == -1)
            node_start_pfn[nid] = chunk_start(j);

        node_end_pfn[nid] = chunk_end(j);
    }
}
```

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, with horizontal lines through the letters.



Working with People

- Engineering is nothing without Marketing
- Technical arguments alone don't win
- Politics often *more* important than code



Working With People

- Taking feedback
 - have thick skin
 - don't take anything too seriously
- Be open to
 - rewrites
 - throwing your entire approach away
 - **starting over**
 - having a beer and laughing about it



Many have fallen...

- Anyone incapable of taking and incorporating feedback is incapable of writing code in an open environment

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned at the bottom left of the slide. The background of the slide features a vertical blue gradient on the left side, with a faint image of a hand holding a computer mouse and a document with a globe and code-like symbols.



Giving Feedback

- Do unto others...
- But
 - Be critical
 - Be firm
 - Be flexible
 - Know when enough is enough
- Play Stupid
 - If you don't understand something, 100 other people don't either



Open Source

- Open does not just mean source
- *Design* in the open
- Beneficial, not burdensome
- Remember Greg's rules
 - Open helps *you*

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, is positioned at the bottom left of the slide. The background of the slide features a vertical blue gradient with a faint, stylized globe and a hand holding a computer mouse.



Questions?

Dave Hansen

haveblue@us.ibm.com

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, positioned at the bottom left of the slide. The background of the slide features a vertical strip on the left side with a globe, a hand holding a mouse, and the IBM logo.



Tools - Scripting

- Script *everything*
- When do I write a script?
 - When you do something often
 - When you do something rarely
- Common things save you time
- Rare things save you having to look it up again



Scripting Examples

- Lunch
- To make your Mom happy
- Preempt something you'll have to do anyway
- Subvert company bureaucracy
- Perl sucks



Email

- Find a client that you like
- .procmailrc
- Who/what should I listen to?
 - Linus
 - Richard B. Johnson
 - People at my company
 - People at a competitor
 - SCO



Really Done

Dave Hansen

haveblue@us.ibm.com

The IBM logo, consisting of the letters 'IBM' in a bold, sans-serif font, positioned at the bottom left of the slide. The background of the slide features a vertical strip on the left side with a globe, a hand holding a mouse, and the IBM logo.