


Research Report: On the
Feasibility of Retrofitting
Operating Systems with
Generated Protocol Parsers

Wayne Wang, Peter C. Johnson
Middlebury College

Integrating
generated
protocol parsers
into existing OS



Process

Operating Systems

FreeBSD
illumos
Linux

Protocols

IP
USB 3.0 (xHCI)
SCSI

Research Process

- Ingress & Egress points

Protocol	Ingress Point	Egress Point
IP	hand-off from link layer (e.g., Ethernet)	hand-off to transport layer (e.g., TCP, UDP)
USB	bus controller interrupt handler	hand-off to device-specific driver
SCSI	HBA controller interrupt handler	hand-off to filesystem or block I/O layer

- traced execution of functions

- Labeled each line of code into 9 categories

Category: Parsing

(IP - FreeBSD)

```
if (m->m_flags & M_FASTFWD_OURS) {
    m->m_flags &= ~M_FASTFWD_OURS;
    /* Set up some basics that will be used later. */
    ip = mtod(m, struct ip *);
    hlen = ip->ip_hl << 2;
    ip_len = ntohs(ip->ip_len);
    goto ours;
}
```

Category - Multiplexing

(USB - Linux)

```
switch (trb_type) {
case TRB_COMPLETION:
    handle_cmd_completion(xhci, &event->event_cmd);
    break;
case TRB_PORT_STATUS:
    handle_port_status(xhci, event);
    update_ptrs = 0;
    break;
case TRB_TRANSFER:
    ret = handle_tx_event(xhci, &event->trans_event);
    if (ret >= 0)
        update_ptrs = 0;
    break;
```

Category - State management

(SCSI - illumos)

```
/*  
 * Only update state to IOCOMPQ if we were in the INTR state.  
 * Any other state (e.g. TIMED_OUT, ABORTED) needs to remain.  
 */  
if (pwrk->state == PMCS_WORK_STATE_INTR) {  
    pwrk->state = PMCS_WORK_STATE_IOCOMPQ;  
}
```

Category - Memory management

(USB - FreeBSD)

```
usb_d_get_page(&sc->sc_hw.root_pc, 0, &buf_res);
```


Category - Hardware manipulation

(SCSI - linux)

```
if (intstat & CMDCMPLT) {
    ahd_outb(ahd, CLRINT, CLR_CMDINT);

    /*
     * Ensure that the chip sees that we've cleared
     * this interrupt before we walk the output fifo.
     * Otherwise, we may, due to posted bus writes,
     * clear the interrupt after we finish the scan,
     * and after the sequencer has added new entries
     * and asserted the interrupt again.
     */
    if ((ahd->bugs & AHD_INTCOLLISION_BUG) != 0) {
        if (ahd_is_paused(ahd)) {
            /*
             * Potentially lost SEQINT.
             * If SEQINTCODE is non-zero,
             * simulate the SEQINT.
             */
            if (ahd_inb(ahd, SEQINTCODE) != NO_SEQINT)
                intstat |= SEQINT;
        }
    } else {
        ahd_flush_device_writes(ahd);
    }
}
```

Category - Synchronization

(SCSI - FreeBSD)

```
cam_periph_lock(periph);
```

Categories

- Queueing

```
/*  
 * If we are waiting on a queue, just remove the USB transfer  
 * from the queue, if any. We should have the required locks  
 * locked to do the remove when this function is called.  
 */  
usb_transfer_dequeue(xfer);
```

- Diagnostic

```
SCSI_LOG_HLCOMPLETE(1, scmd_printk(KERN_INFO, SCpnt,  
    "sd_done: completed %d of %d bytes\n",  
    good_bytes, scsi_bufflen(SCpnt)));
```

- Assertion

```
M_ASSERTPKTHDR(m);  
NET_EPOCH_ASSERT();
```







Fig. 1. Categorized code diagram for IP. Colors correspond to categorizations shown in Tables IV and V; line width corresponds to the length of the individual lines of code.



(a) FreeBSD



(b) Illumos



(c) Linux

Fig. 2. Categorized code diagram for USB. Colors correspond to categorizations shown in Tables IV and V; line width corresponds to the length of the individual lines of code.

Results - multiplexing

IP

	FreeBSD	Illumos	Linux
destination class		callback	
ip.protocol	jump table	switch	jump table

USB

	FreeBSD	Illumos	Linux
OS request handler		callback	callback
device driver	callback	callback	callback
storage layer			callback

SCSI

	FreeBSD	Illumos	Linux
XHCI event type	switch	switch	switch
transfer	callback		
endpoint type		switch	if / else
endpoint	callback		
transfer	callback	switch, callback	callback

Results - multiplexing

Netgraph in FreeBSD

&

NetFilter in Linux

What is done

- Survey
- Categorization
- Summarize

What is next

- enumerate general kernel functions used
- analyze their similarity within a single OS & across OS
- define an API for a universal shim