Gernot Hoyler

Spansion Inc.

Pre-silicon Software Development of Linux MTD Drivers

October 30, 2009

SPANSION®

Product Development Process

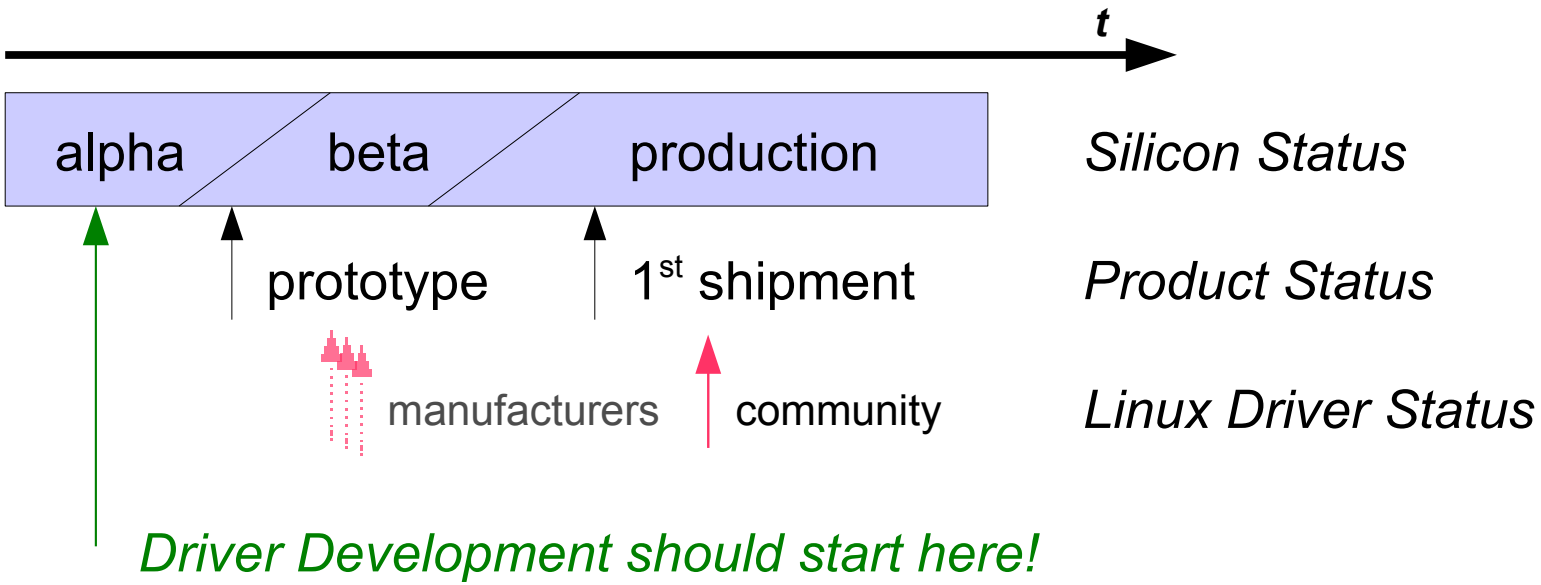

Ideally, device drivers are already available when customers get their hands on first samples. A simulator based approach can help work around hardware issues that might exist during this early phase (alpha phase).

Flash Basics

- Electrical charge is moved onto an isolated storage element in order to ***program*** data (done on a bit basis for NOR flash or on a page basis for NAND flash)

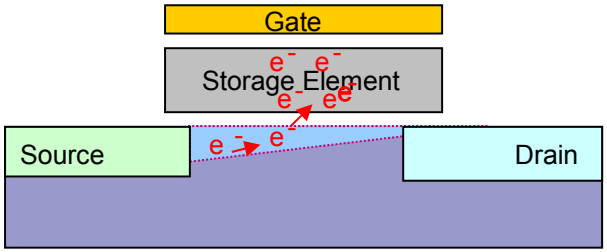

- Electrical charge is removed from the isolated area in order to ***erase*** data (done on a sector/block basis, e.g. 128 kB)
- To ***read*** data, a reference voltage is applied between source and drain and the current is measured (if current flows, the cell has been programmed)
- A special command set is being used to issue and to control the program and erase operations

S29WS-R Command Set

Command Sequence		Cycles	Bus Cycles (Notes 1–5) First Addr Byte Word	First Data	Second Addr Byte Word	Second Data	Third Addr Byte Word	Third Data	Fourth Addr Byte Word	Fourth Data
Read		1	RA	RD						
Reset		1	XX	F0						
Write Buffer Load (9)		4-35	(SA) AAA (SA) 555	25	(SA) 555 (SA) AAA	WC	(SA) PA (12)	PD	(SA) PA (13)	PD
Buffer to Flash		1	(SA) AAA (SA) 555	29						
Chip Erase		2	(SA) AAA (SA) 555	80	(SA) 555 (SA) AAA	10				
Sector Erase		2	(SA) AAA (SA) 555	80	(SA) 555 (SA) AAA	30				
Status Register Read		2	(SA) AAA (SA) 555	70	(SA)	RR				
Status Register Clear		1	(SA) AAA (SA) 555	71						
ID/CFI Command Definitions										
ID/CFI	ID/CFI Entry (8) (11)	1	(SA) XAA (SA) X55	90 or 98						
	ID/CFI Read	1	(SA) RA	data						
	ID/CFI Exit	1	XXX	F0						

...

Simulator Based Approach for MTD Drivers

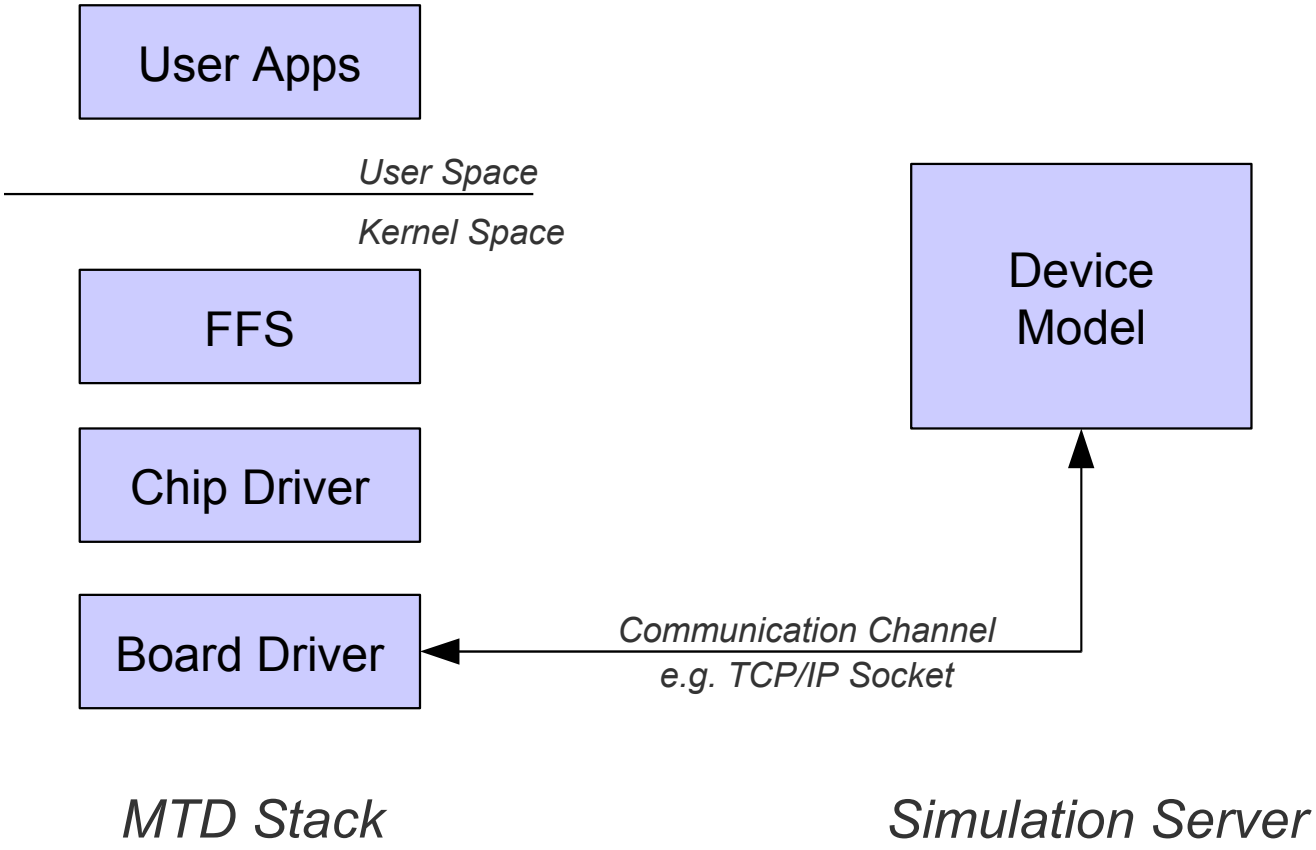

Device Model / Simulation Server

DevSim is a cycle accurate SystemC model of Spansion Flash. It supports the following formats:

- **DevSim object for Cadence Design Systems' NC-Verilog on Linux**
- **DevSim object for Synopsys' VCS on Linux**
- **DevSim object for Microsoft Visual Studio**
- **DevSim executable/server and command line client**

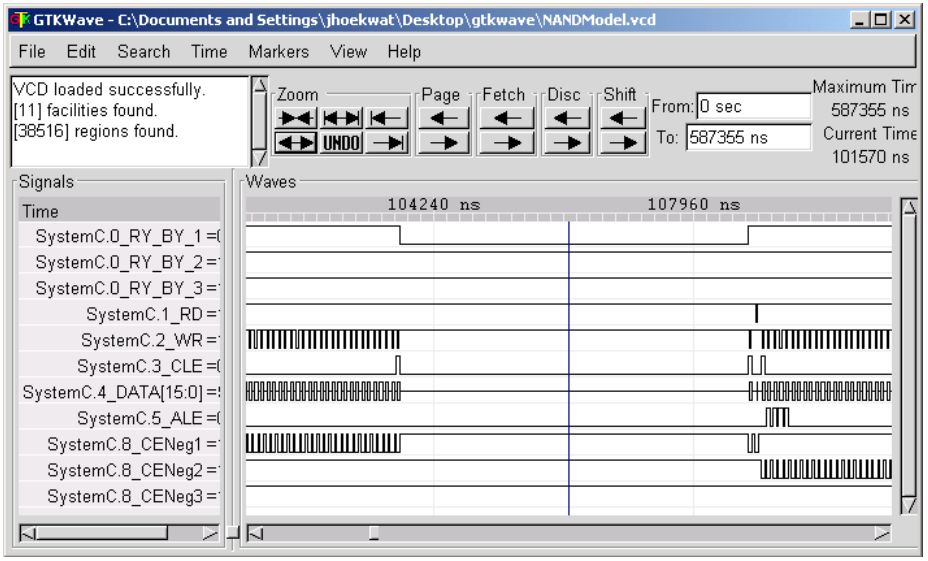

Benefits:

- **Developed by Spansion for internal verification**
- **Cycle accurate**
- **Detects and reports timing violations**
- **Creates a waveform of the session**
- **Resource sensitive - Simulates as much of the array as you use**
- **Truly non-volatile from session to session**
- **View the array outside of a session**
- **Array pre-loading supported**
- **Flash architecture controlled by CFI input file**
- **Works great for software development**
  - Program and erase times can be altered to reduce real-time execution

Communication Channel Protocol

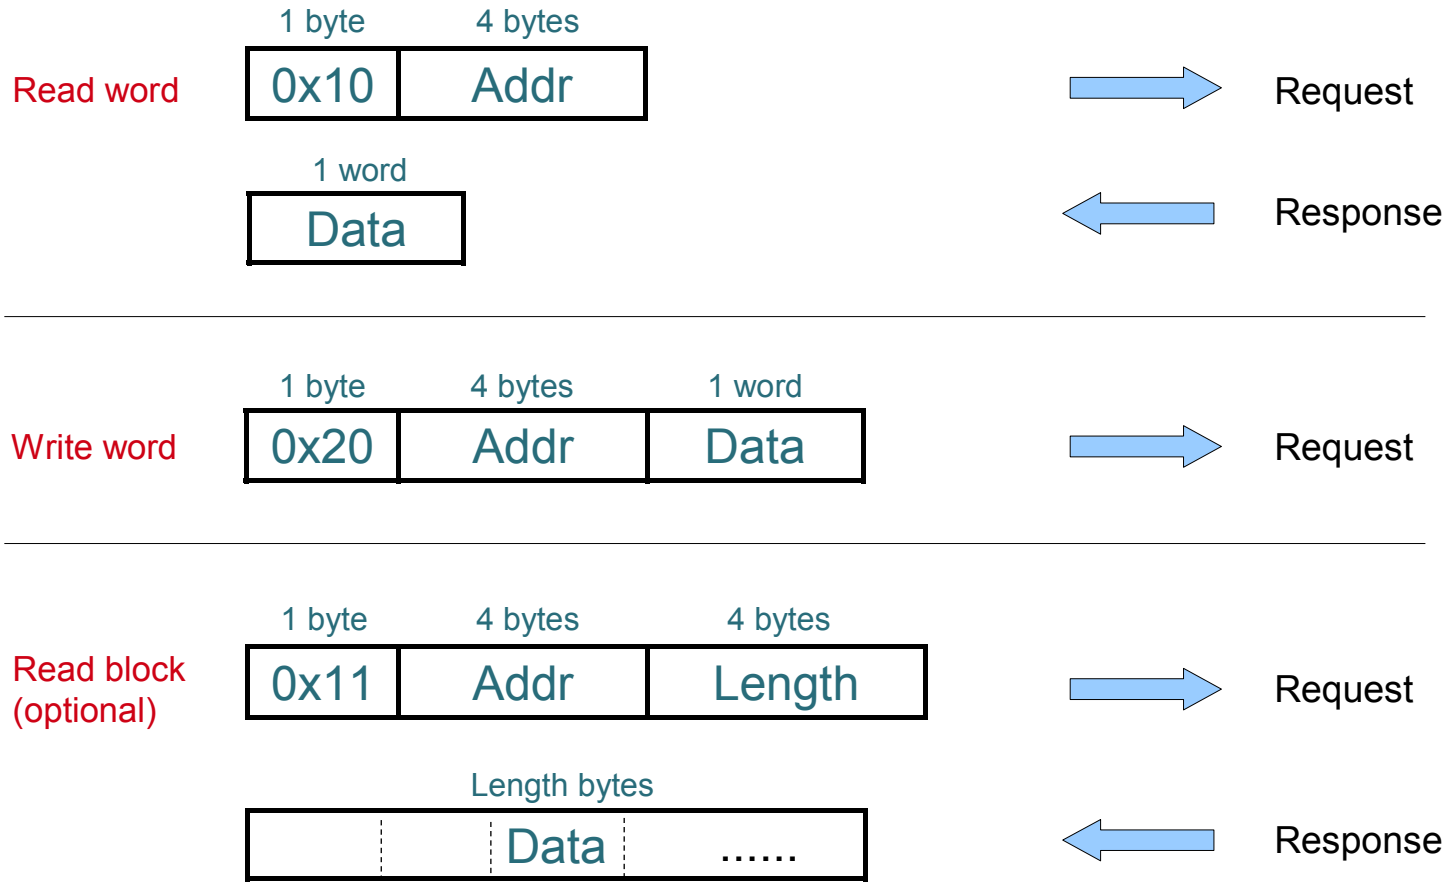

MTD Board or Map Driver

The pivotal element connecting a chip driver (NOR flash) to the hardware is a `map_info`, see `include/linux/mtd/map.h`

```
struct map_info {
    char *name;
    unsigned long size;
    resource_size_t phys;
    void __iomem *virt;
    void *cached;
    int bankwidth;

    map_word (*read)(struct map_info *, unsigned long);
    void (*write)(struct map_info *, const map_word, unsigned long);
    void (*copy_from)(struct map_info *, void *, unsigned long, ssize_t);
    void (*copy_to)(struct map_info *, unsigned long, const void *, ssize_t);

    void (*inval_cache)(struct map_info *, unsigned long, ssize_t);
    void (*set_vpp)(struct map_info *, int);
    ...
};
```

Function pointers to issue read & write cycles to the flash device

All the driver has to implement are functions that route the r/w requests over the communication channel to the model.

Proper Synchronization

- Replacing the standard I/O calls used for a physical device (`__raw_readw()`, `__raw_writew()`) by socket calls changes the characteristics of the I/O from non-blocking to blocking.
- This can cause issues if the kernel holds spin-locks during the I/O. Another thread trying to acquire the spin-lock at the same time can stall the entire system (deadlock).
- Some MTD chip drivers hold a spin-lock while issuing commands. This spin-lock needs to be released before the socket operations take place and reacquired afterwards.
- To prevent other threads from accessing the MTD while the socket operations are ongoing, the device has to be put into a special mode so that other threads have to wait, i.e. sleep and try again later.

Live Demo

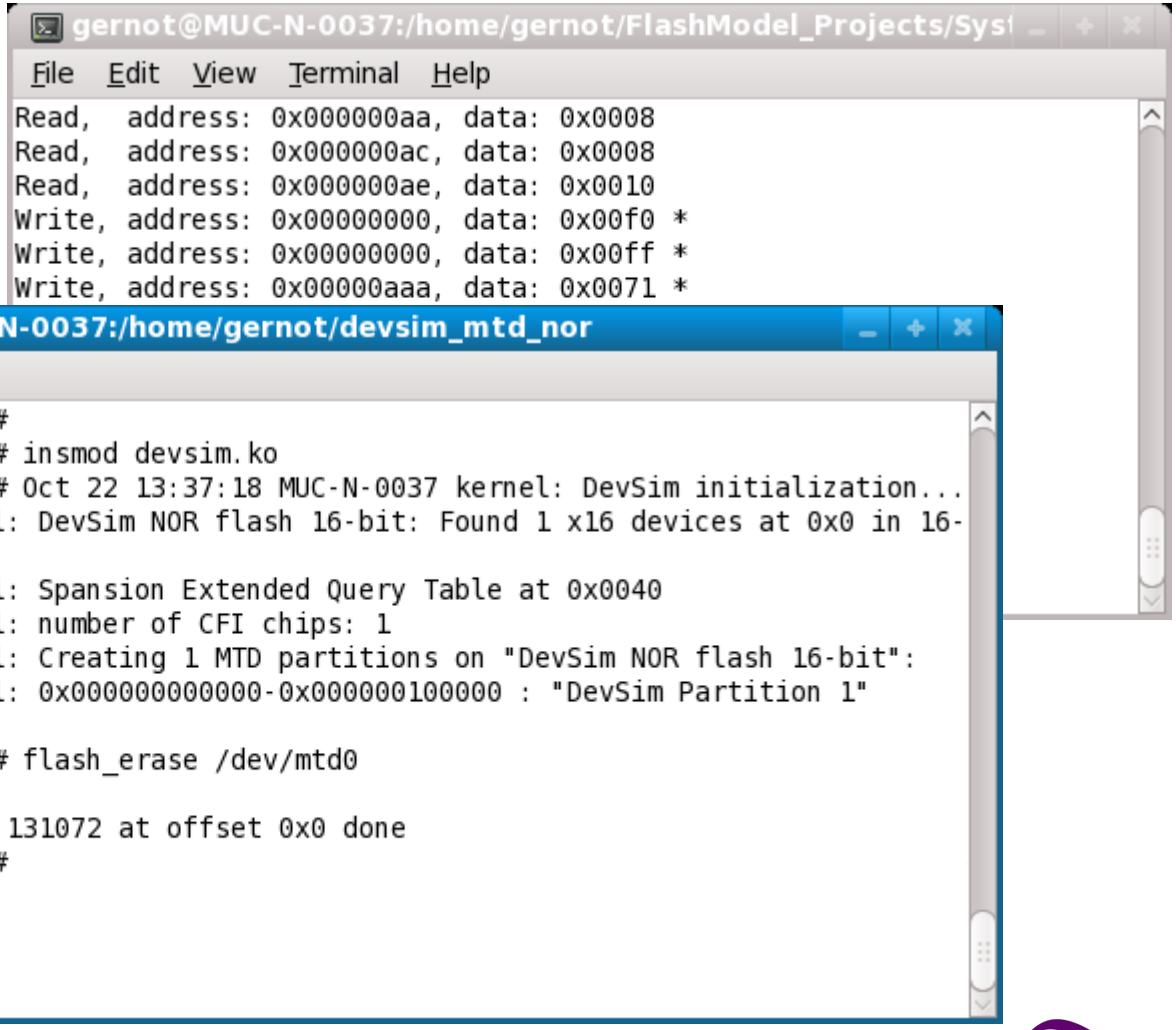

Further Advantages

Next to the early availability, a simulator based approach has further advantages:

- ✓ Simplified debugging
- ✓ Software trace support (virtual logic analyzer)
- ✓ Errors can easily be inserted
- ✓ Different device models can be loaded quickly
- ✓ Non-standard bus widths and setups can easily be tested

Sync	Type	Address	Data
00257585	Write	00000AAA	000000AA
00257586	Write	00000554	00000055
00257587	Write	00000AAA	00000090
00257588	Read	00000002	0000227E
00257589	Read	00000000	00002222
00257590	Read	00000000	00002201
00257591	Write	00000000	000000F0

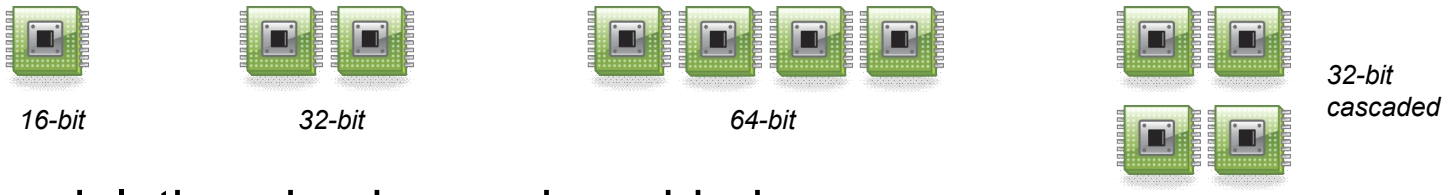

- ✓ Spec violation checks can be added

Summary

- A new approach for pre-silicon software development of MTD chip drivers has been proposed
- In the context of the development of a new NOR flash driver it proved to be stable and beneficial
- It can be used not only at an early stage where no Si is available yet, it also greatly simplifies testing and debugging
- A similar approach should be conceivable for other memory technologies such as NAND or SPI devices
- All it needs is to route the flash specific interface over a communication channel to a device model

Thank You!

Trademark Attribution

Spansion®, the Spansion logo, MirrorBit®, MirrorBit® Eclipse™, ORNAND™, ORNAND2™, HD-SIM™ and combinations thereof, are trademarks of Spansion LLC in the U.S. and other countries. Other names used are for informational purposes only and may be trademarks of their respective owners.

This document is for informational purposes only and subject to change without notice. Spansion does not represent that it is complete, accurate or up-to-date; it is provided “AS IS.” To the maximum extent permitted by law, Spansion disclaims any liability for loss or damages arising from use of or reliance on this document.