Optimizing Flash Controller Technology for Next-Gen Flash

Greg Huff, CTO, LSI

Agenda

- Don’t Be the Bottleneck
- All Markets Need Power Reduction
- Faster Wakeup Times
- Improve Flash Error Correction
- Extend Flash Endurance

Don’t Be the Bottleneck

Controllers need to keep both host and flash interfaces at max speed

Don't Be the Bottleneck

- SATA is limited to 500MB/s
  - Add native PCIe for >3x performance increase
- Maximize parallelism of flash channels
  - Make them independently controlled
- If using DRAM, running user data through it has a large cost
  - Full speed requires 2x host bandwidth
  - Increases costs and power consumption
  - Prevents true DevSleep ability

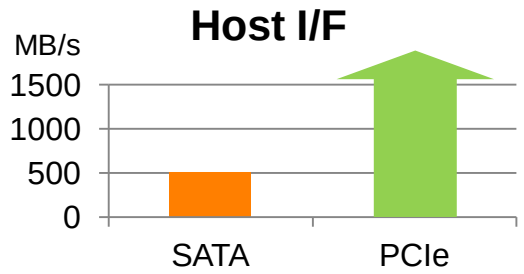

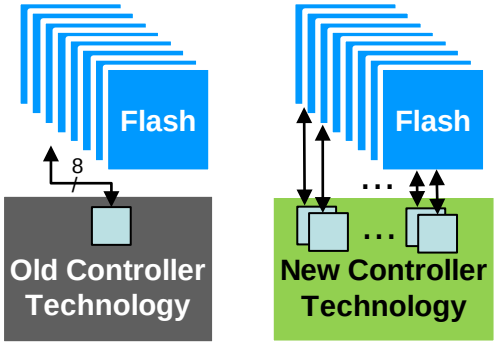

Push the bottleneck to the interfaces (host or flash)

All Markets Need Power Reduction

Future SSD performance must be “power efficient”, not “power wasteful!”

All Markets Need Power Reduction

- Need more performance without burning power
- Everything is going “green”
  - Ultrabooks and Connected Standby
  - Datacenter buildings are typically power-capped
- Keep active power as low as possible
  - Avoid extra power of DRAM
  - Reduce data transfer times (consider GB/watt and IOPS/watt)
- Leverage DevSleep and low power modes
  - LSI showed 400x idle power reduction at Computex
- Reduce leakage
  - Design hardware to minimize leakage
  - Consider low power nodes like 40LP

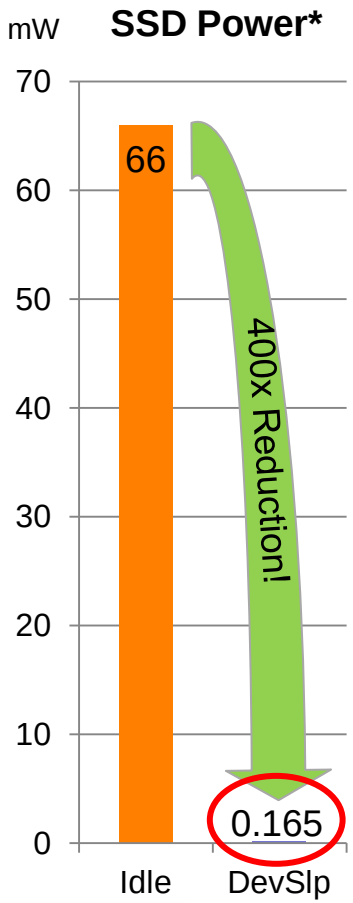

Power reduction is an ongoing industry challenge

* Power measurement tests conducted by LSI on standalone test bench on display at Computex, Taiwan, June 2013

Faster Wakeup Times

"Launch" from a dead stop as fast as possible

Faster Wakeup Times

- Transition between power states quickly
  - Minimize state information
  - Reduce data restore times after wake
  - Improve recovery time

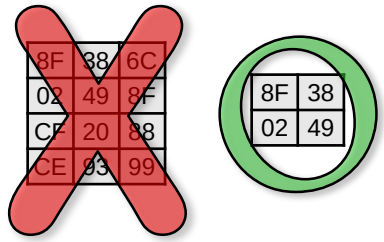

- Optimize Flash Translation Layer (FTL)
  - Reduce response time to host request immediately following wake up
  - Requires optimizing map architecture

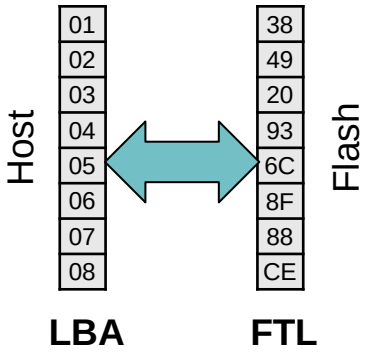

Nirvana would be instant on from zero power!

Improve Flash Error Correction

New flash requires more difficult “repairs”

Improve Flash Error Correction

- Lower NAND technology nodes increase error rates
- Stronger BCH requires too many spare bytes vs. LDPC
- LDPC & DSP technology becoming popular, but requires sophisticated soft-decision processing
  - Limit hard-decision processing to streamline most reads
  - Consider multiple code rates over flash life
  - Minimize number of read retries required to correct data
- Optimize time-to-data
- Add a higher level correction capability when ECC fails
- Create a NAND flash qualification team
  - Staff at double-digit manpower
  - Perform in-depth flash analysis across all vendors
  - Work closely with all flash manufacturers

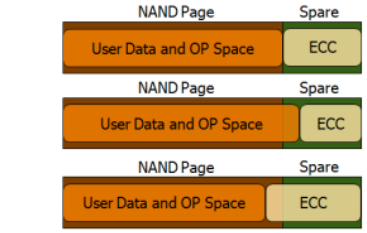

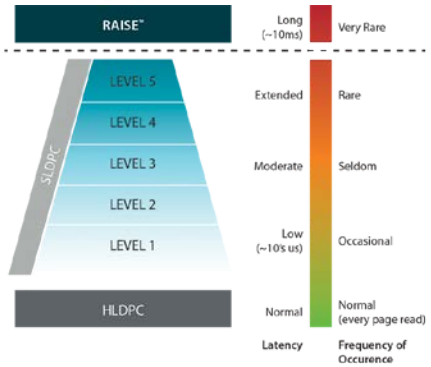

LSI SHIELD™ Error Correction Technology is an example of leading edge LDPC and DSP correction technology

Not all LDPC is created equal – the actual soft-decision processing is key

Extend Flash Endurance

Strengthen the underlying “roadway” and use it more “efficiently”

Extend Flash Endurance

NAND Trend

Higher error rate = lower endurance

- Increased ECC ability extends recoverable reads
- Write less data to flash, but keep the same info!
  - Use data reduction technology like DuraWrite™
  - Higher performance from writing less
  - Increases Dynamic OP which reduces Write Amp
  - Also reduces power draw from less prog. time
- Avoid background garbage collection
  - It increases Write Amp which accelerates wear

Limited Factory P/E Cycles

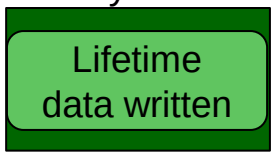

Increase P/E Cycles with Much Better ECC

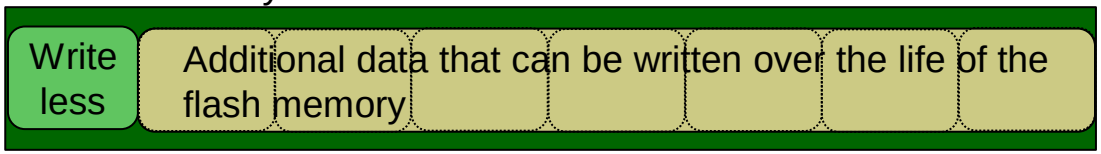

Read better to read longer, and write less to write more

Summary

- Push the bottleneck to the interfaces (host and flash)
- Power reduction is an ongoing industry challenge
- “Launch” from a dead stop as fast as possible
- LDPC is key to robust error correction, but not all LDPC solutions are created equal
- Increase flash endurance with better ECC and write more efficiently

Storage. Networking. Accelerated.™

Follow LSI SandForce