



All I/O is Random I/O

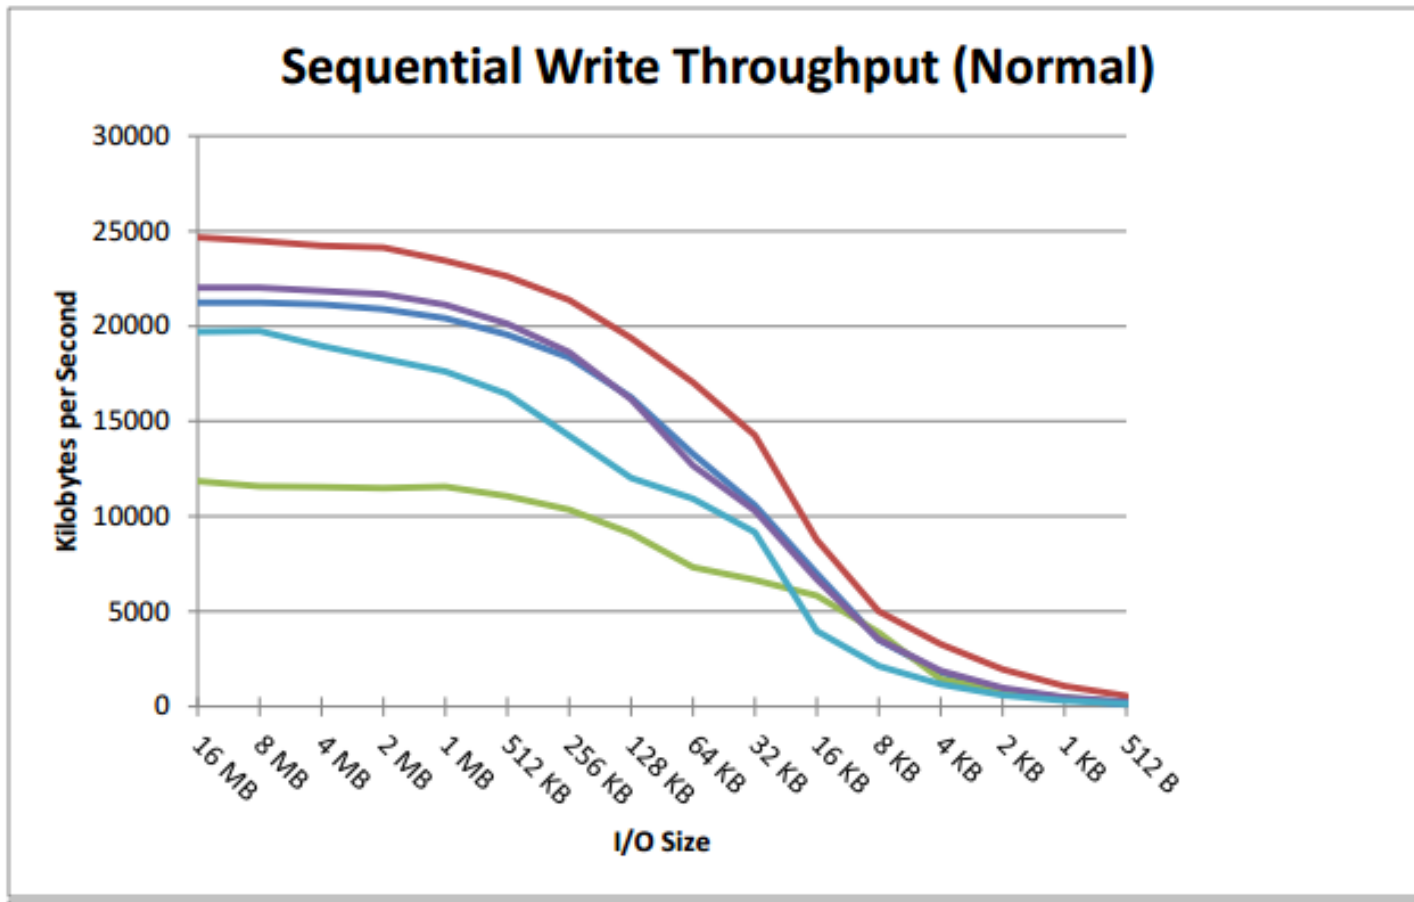
Flash Memory Summit 2013

Thom Denholm
Technical Product Manager

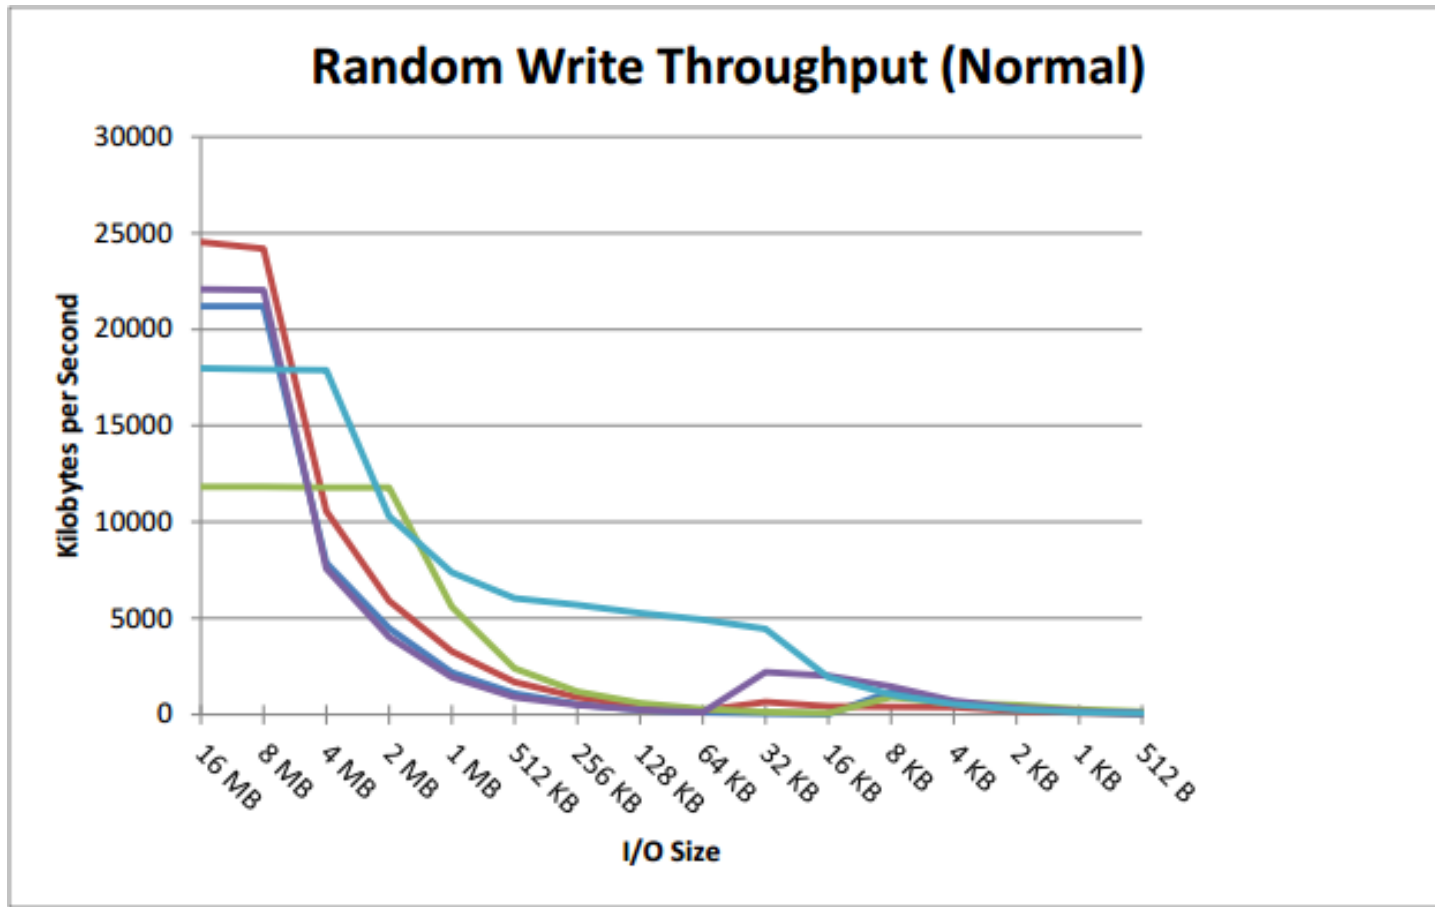
Talking Points

- Performance diffs – sequential/random
 - Contribution of Database on Android
- Benchmarks (measuring RAM)
- Write Amplification
- Correcting with Software
- Other expectations (ACID, lifetime)

Sequential I/O looks great




Random I/O not as good




Android and SQLite

AndroBench – Erase Counts




Each erase represents 2k of data written


Measuring Cache and Buffers



AndroBench




Write Amplification



- 4k file, minimum 512 byte write size
- Flash system Metadata
- File system Metadata & Directory Info

Correcting with Software



- Align writes to match Flash media
- Write sequentially where possible
- Collect writes into Atomic groups
 - Best match to underlying firmware

Flash Media Lifetime




Figure 1 | A life cycle and ECC comparison of NAND flash by process node shows how an increase in correction capability is not enough to maintain endurance of the memory cell.

a

Atomicity
Transactions
are all or
nothing

c

Consistency
Only valid
data is
committed

i

Isolation
Transactions
do not affect
each other

d

Durability
Written data
will not be
lost



Summary

- Performance Benchmarks need to account for RAM and Use Case
 - Ideally on the embedded target
- Software can improve access to the media and reduce overhead & Write Amplification
- Added benefits to Flash Media Lifetime and Reliability