



Flash Memory Summit

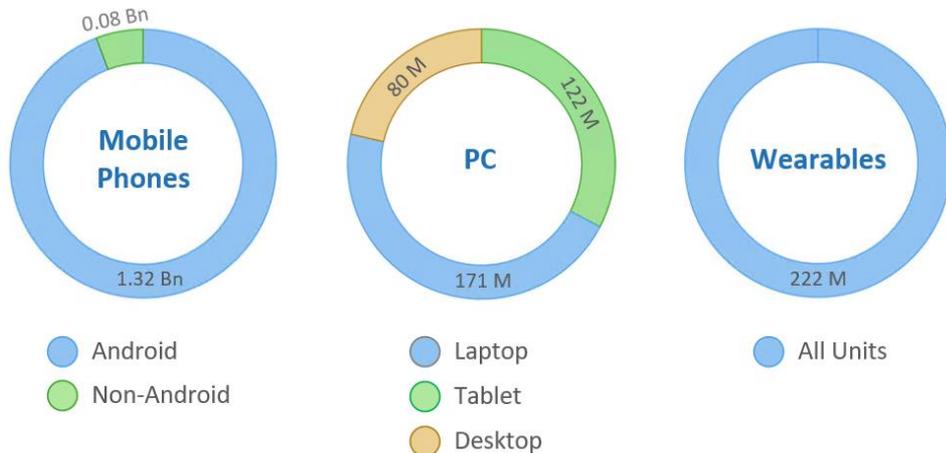
Applying NVMe in Embedded and Mobile Applications

Horace Chen, Phison Electronics



BGA Storage Opportunity

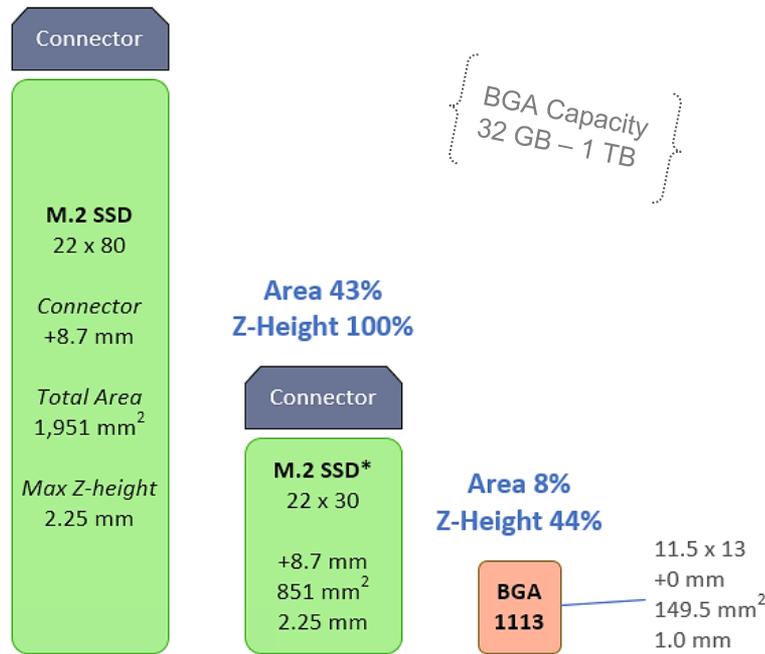
- Problem
 - By 2023, annual shipments for mobile phones are projected at 2.4 Billion and wearables will outnumber laptops
 - Global demand for small formfactor storage is enormous
- Solution
 - NVMe BGA is aligned with the rapid evolution of PCIe Gen-4 and Gen-5
 - Very small formfactor, low power, excellent performance





BGA Area Reduction

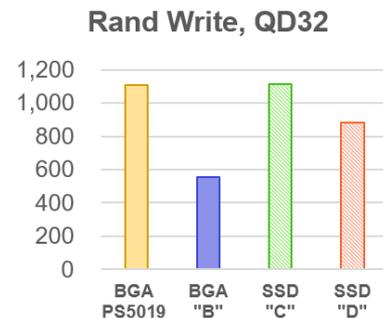
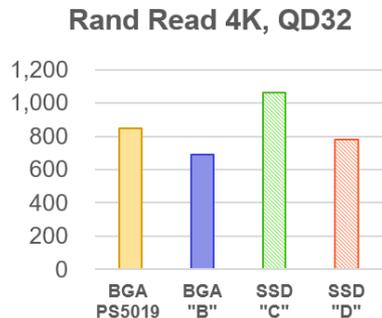
- The M.2 2280 formfactor is the most common module size today.
- On the next slide we will show that a properly designed DRAMless solution, with HMB enabled, can provide equivalent performance
- The BGA 1113 is thinner and smaller, allowing more room for sensors or a larger battery
- Only need to qualify one solution. The BGA can be mounted on a 2230 PCB for enterprise environments that want removable storage
- BGA supports co-layout with eMMC and UFS for added design flexibility





BGA Performance is Competitive

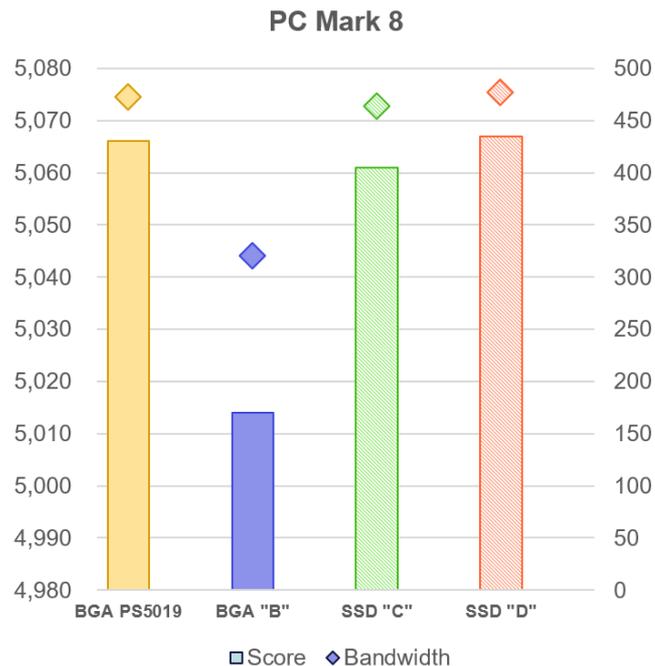
- BGA are DRAM-less + Host Memory Buffer (HMB)
- SSD have DRAM
- BGA 1113 limited to 2x PCIe Lanes
- Additional tuning and careful hardware design allows BGA to be competitive
 - Design with smart accelerators instead of CPU to save power and optimize simple lookup and DMA operations





BGA Real World Applications

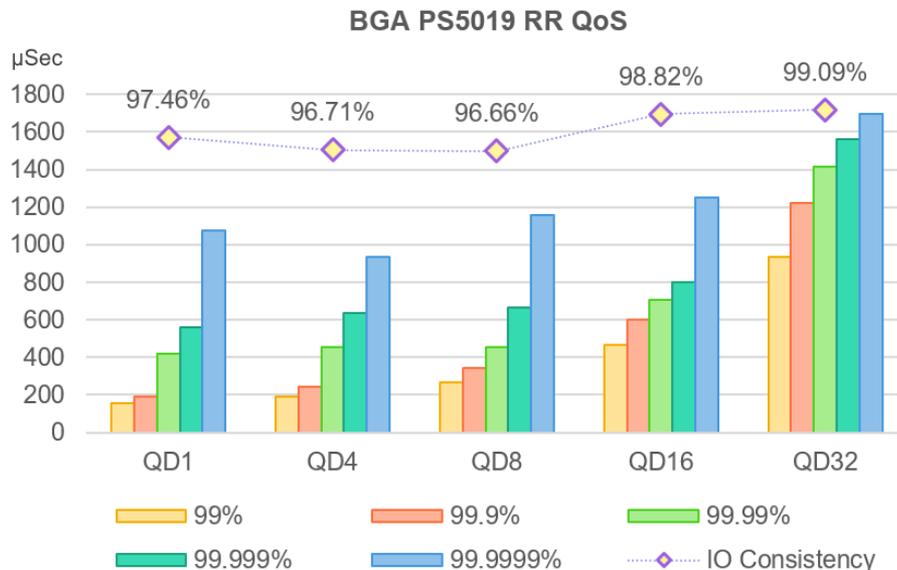
- No significant difference on properly tuned BGA solutions
- BGA supports additional embedded features
 - Replay Protected Memory Buffer (RPMB)
 - Boot Partition
 - Namespace (ie: Logical Unit)
 - Host Memory Buffer (up to 64MB on Windows host)
 - Deallocate (ie: Trim, Unmap)
 - TCG Opal, TCG Pyrite





BGA Latency & IO Consistency

- With proper firmware tuning, a BGA solution can compete with desktop SSD
- Smooth progression from QD1 from QoS-1 to QoS-6
- Optimal performance is achieved at QD16 which takes advantage of NAND independent plane operations
- Latency jumps up at QD32 due to having more commands than points of parallelism, though it's still smooth
- Excellent IO Consistency





BGA Configurable Power

- Uses 0.9, 1.2, 2.5V supply rail to reduce active power by 14% over same ASIC as M.2
- On DRAM-less BGA, power scales with read bandwidth ~1W / GB/s
- Power evenly split between IO, NAND and ASIC; approximately 0.3/0.3/0.3
- Silicon process node not a primary factor; can choose to optimize for cost
- BGA molding compound uses advanced heat transfer materials
- Platform should provide passive cooling through direct contact with metal chassis
- Using NVMe Power States, the BGA can be set to limit power 0.5W – Max, if performance is not a concern
 - One part can be used in many solutions ranging from wearables to gaming tablets and content creation laptops
 - One solution simplifies the supply chain



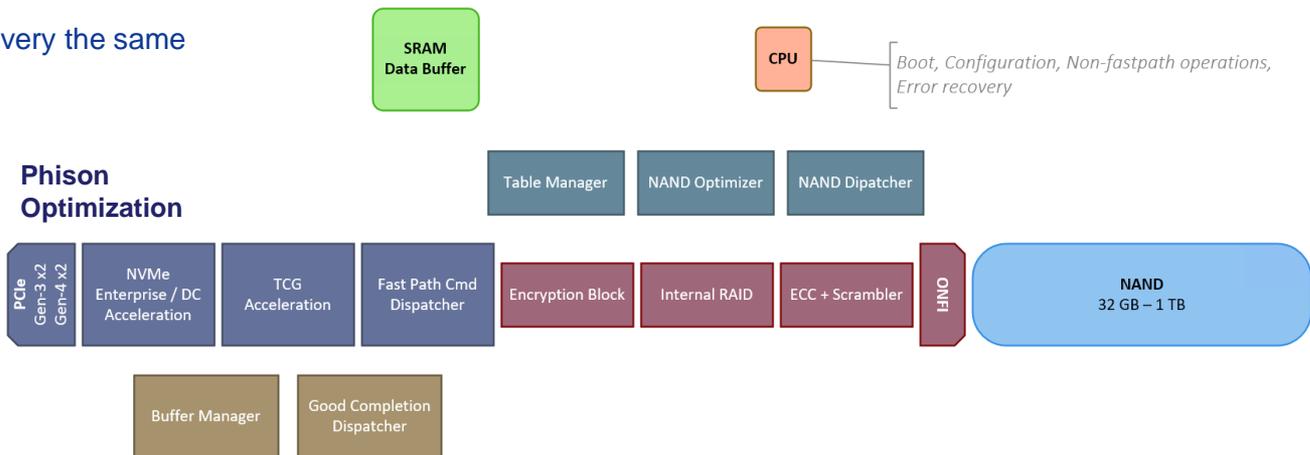
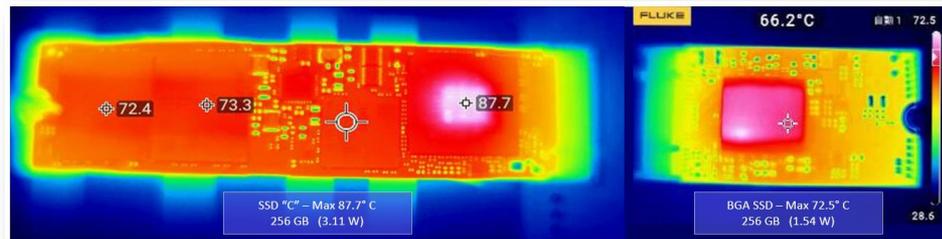
Current maximum BGA Gen-4 x2

- Read Speed = 3,500 MB/s
- Write Speed = 1,800 MB/s



BGA Competitive Power

- PS5019 BGA and SSD “C” are both connecting as Gen-3 x2
- Both are running a sequential workload at the same throughput
- PS5019 BGA is on an M.2 carrier
- Advanced automation allows the design to run substantially cooler
- Uses less energy while delivery the same performance and latency





Key Takeaway

1. NVMe BGA has excellent performance which means no more lag
2. Tightly aligned to the rapid evolution of PCIe Gen-4 and Gen-5
3. Small formfactor leaves more room for battery
4. Low thermal allows the device to drive cutting edge applications
5. Low power means the battery lasts longer



Flash Memory Summit

PHISON
Knows What You Need

“PUSHING BOUNDARIES”
YOUR PCIe GEN-4 SSD LEADER

VISIT US AT BOOTH #219!