Flash Memory Summit

大
普
微

Introducing DPU:
Data-storage Processing Unit

Placing Intelligence in Storage

Qing Yang 杨庆

CSO, Shenzhen Dapu Microelectronics Ltd Co.

深圳大普微电子科技有限公司董事长

Distinguished Engineering Professor

University of Rhode Island

Data Growth + Media Tech

大
普
微

✧ **Big Data, Cloud：Data Explosion**

✧ **Applications Demand**

High Performance, Secure, Reliable, Recoverable

✧ **Emerging Device Tech, more Cost-Effective：**

- **Flash、PCM、MRAM、RRAM**

$\Delta_{18m} = \Sigma_{-\infty}$

Ali-Cloud last year:

5 Times!

These placed great challenges to the storage control and management

Existing storage controllers are far behind！

History of Computing

大普微

- Decades ago，Displays were controlled by CPU/MCU
  - Resolution, color, pixels increased greatly
  - CPU/MCU could no longer control modern displays
  - As a result：GPU was born and developed
- Today: GPU Plays a revolutionary role in computing！
- We Believe

Storage control of big data has come to a historical point!

CPU/MCU can no longer manage exponential growth of data and a variety of storage media technologies：

Therefore, We introduce：

DPU Data-storage Processing Unit

Major Functions in DPU

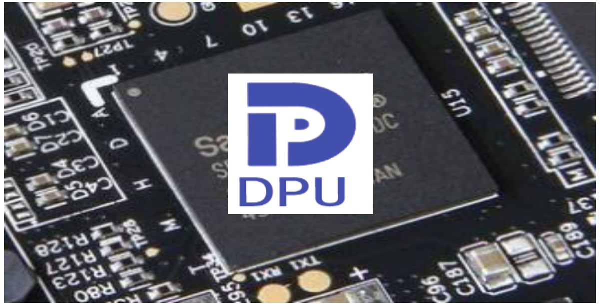

- **Media Managements**
  - **Flash Control**
  - **Machine Learning of I/O Patterns**
  - **Minimizing Erasures & Adaptive RAID**
- **Advanced Data Analytics**
  - **Processing in Storage: PIS**
  - **Placing data intensive computation closest to where data is stored**
- **Storage Architecture**
  - **Hierarchy and Tiering:** ***Cache***
  - **Dedupe, Snapshot, Replication, and Failure Tolerance**
  - **Distributed SAN, E-W connectivity, NVMe over the fabric**

Storage Media Management

- **Physical Properties of Flash Memory**
  - **Reads are faster than writes**
  - **Limited erase cycles**
  - **No in place write → GC, WL**
- **Write Amplification Problem**
  - **Slow down I/Os, Increase wearing, and hogging resources**

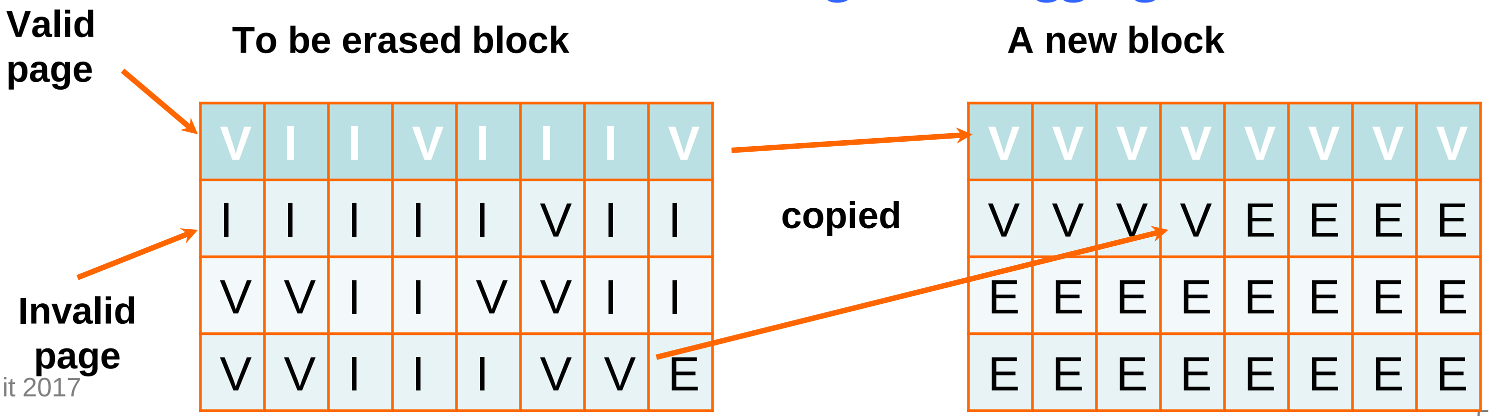

Reinforcement Learning

- Classify I/Os into groups of similar or same rewrite intervals
  - Features and attributes
    - {access history, recently accessed data, time stamps, R/W LBAs, Re-reference interval, Recency}
  - Pages of the same class will be written in one block
    - High performance, minimal WA
- Recognize I/O Patterns at Production Site
  - Train and learn I/O behaviors after deployment
  - Optimization kicks in after a week or so
  - Adapt to any environment and applications

Reinforcement Learning

- Objective function

$$J=\sum_{n=1}^{N}\sum_{k=1}^{K} r_{nk}(x_n-\mu_k)^2 \quad (1)$$

where N is total number of data points, K is number of clusters, $||x_n-\mu_k||$ is the Euclidean distance between data point x_n and the center of set k μ_k,
If data point n belongs cluster k , $r_{nk}=1$, else $r_{nk}=0$.

- If we can put x_n in its closest cluster, J will be minimized.

Dynamic Clustering at Run Time

- Let us consider the derivative of J by μ_k and set it to 0, we have

$$\mu_k = \frac{\Sigma_n r_{nk} x_n}{\Sigma_n r_{nk}} \quad (2)$$

It means μ_k, which is the center of cluster k, is the average of all points in the cluster.

- Based on Equations (1) and (2), we put current page to the cluster with minimum variance. The center of each cluster can be determined by the average distance of that cluster.

Measured Erase Counts

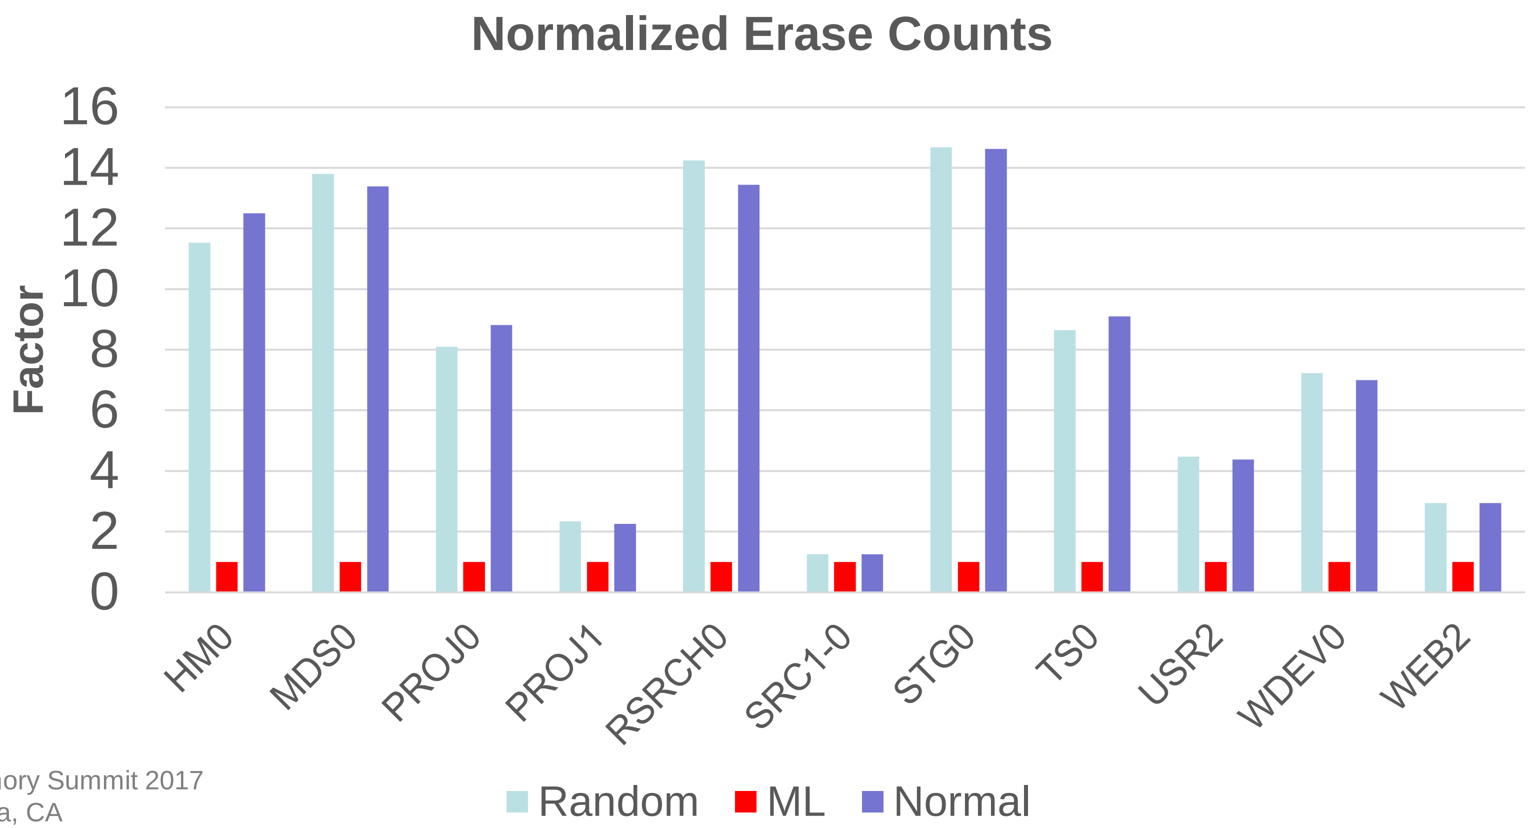

Flash Memory Summit

Major Functions in DPU

大
普
微

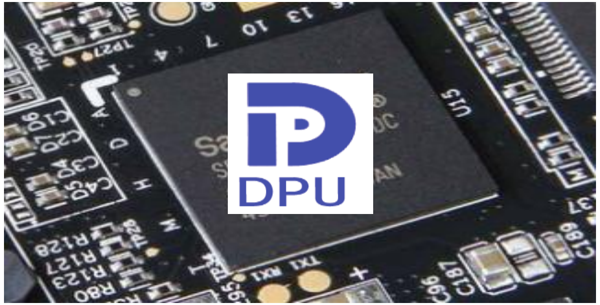

- Media Managements
  - Flash Control
  - Machine Learning of I/O Patterns
  - Minimizing Erasures & Adaptive RAID
- **Advanced Data Analytics**
  - **Processing in Storage: PIS**
  - **Placing data intensive computation closest to where data is stored**
- Storage Architecture
  - Hierarchy and Tiering
  - Dedupe, Snapshot, Replication, and Failure Tolerance
  - Distributed SAN, E-W connectivity, NVMe over the fabric

Cache

ADA1: HW Search in DPU

- **Over 80% of Data are Unstructured**
  - **Process of text data is critical**
  - **Software scanning is slow**
- **Research on Accelerators for Text Search**
  - **Maximizing DRAM bandwidth**
  - **I/O is still a burden**
- **Secure Data Search**
  - **Keyword search in encrypted data**

ADA1: REGISTOR in DPU

- **Regular Expression Grabbing Inside STORage**
  - **HW search in SSD where data is stored**
  - **Only results or related files are sent to the host**

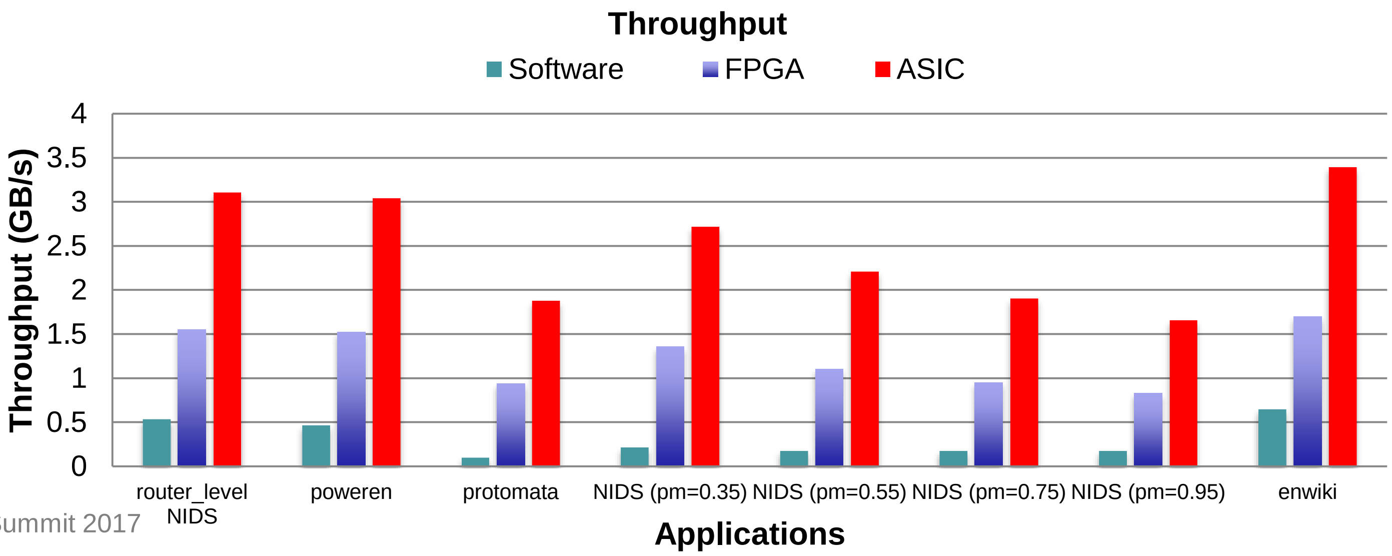

ADA2: HW Deserialization

- **Future world will be sensor driven world**
  - **Huge amount of sensing data files**
  - **Numbers are stored in readable and exchangeable formats: ASCII, Unicode etc.**
- **To Process Data Using Computers**
  - **Readable data need to be converted to binary**
  - **Host CPU is very inefficient**
  - **Time Consuming, up to 60% of Total Processing Time**

Data Conversion Performance of DPU vs. Server CPU

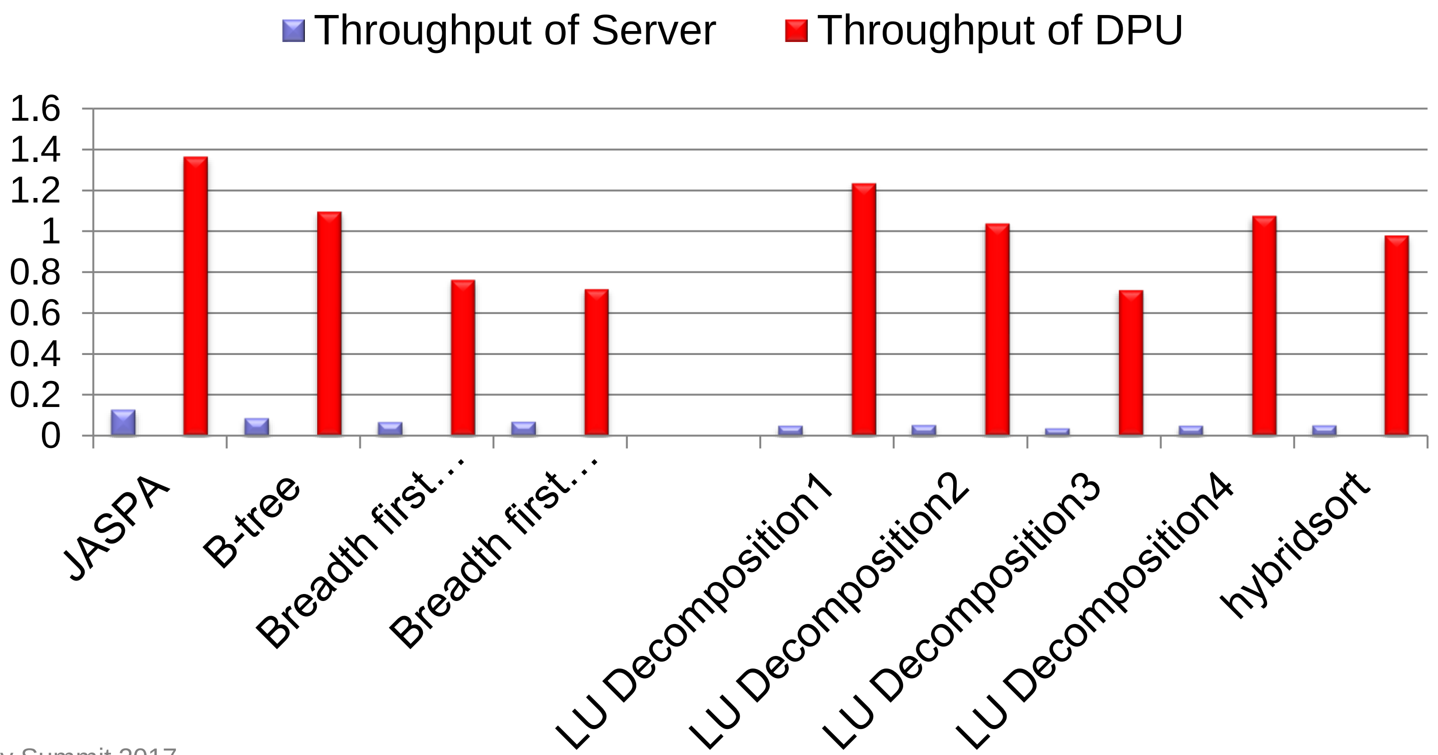

Major Functions in DPU

大
普
微

- Media Managements
  - Flash Control
  - Machine Learning of I/O Patterns
  - Minimizing Erasures & Adaptive RAID
- Advanced Data Analytics
  - Processing in Storage: PIS
  - Placing data intensive computation closest to where data is stored
- **Storage Architecture**
  - **Hierarchy and Tiering**
  - **Dedupe, Snapshot, Replication, and Failure Tolerance**
  - **Distributed SAN, E-W connectivity, NVMe over the fabric**

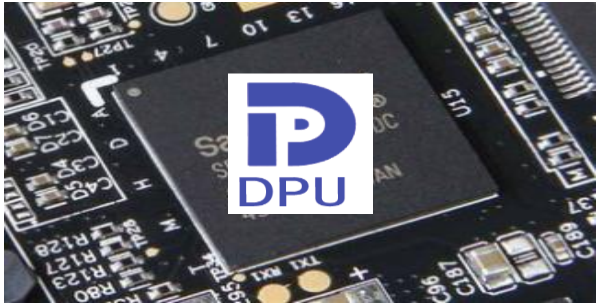

Cache

IST: Intelligent Storage Tiering

- ➢ **Media: Flash, PCM, MRAM, Memristor etc:**
  - ✓ Different Speed
  - ✓ Different Cost
- ➢ **What Do Users want:**
  - ❖ **$↓ & Speed↑ & Power↓ & Ease of use & Reliability↑**

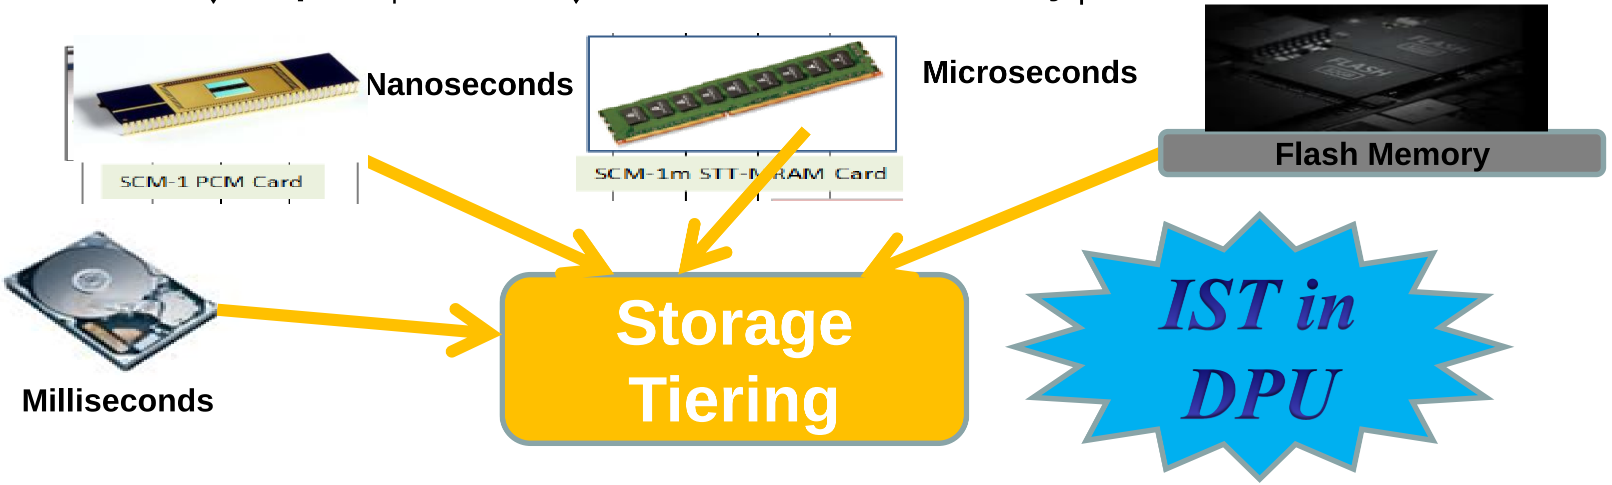

Summary and Conclusions

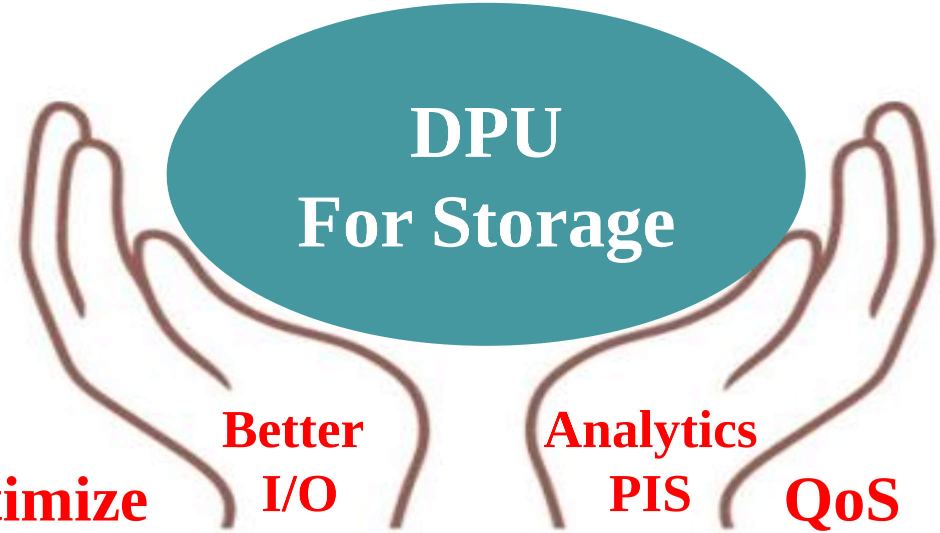

Optimize
Media

QoS
Apps

A New Concept for Next Generation Storage

Proven Advantages on Current SSDs

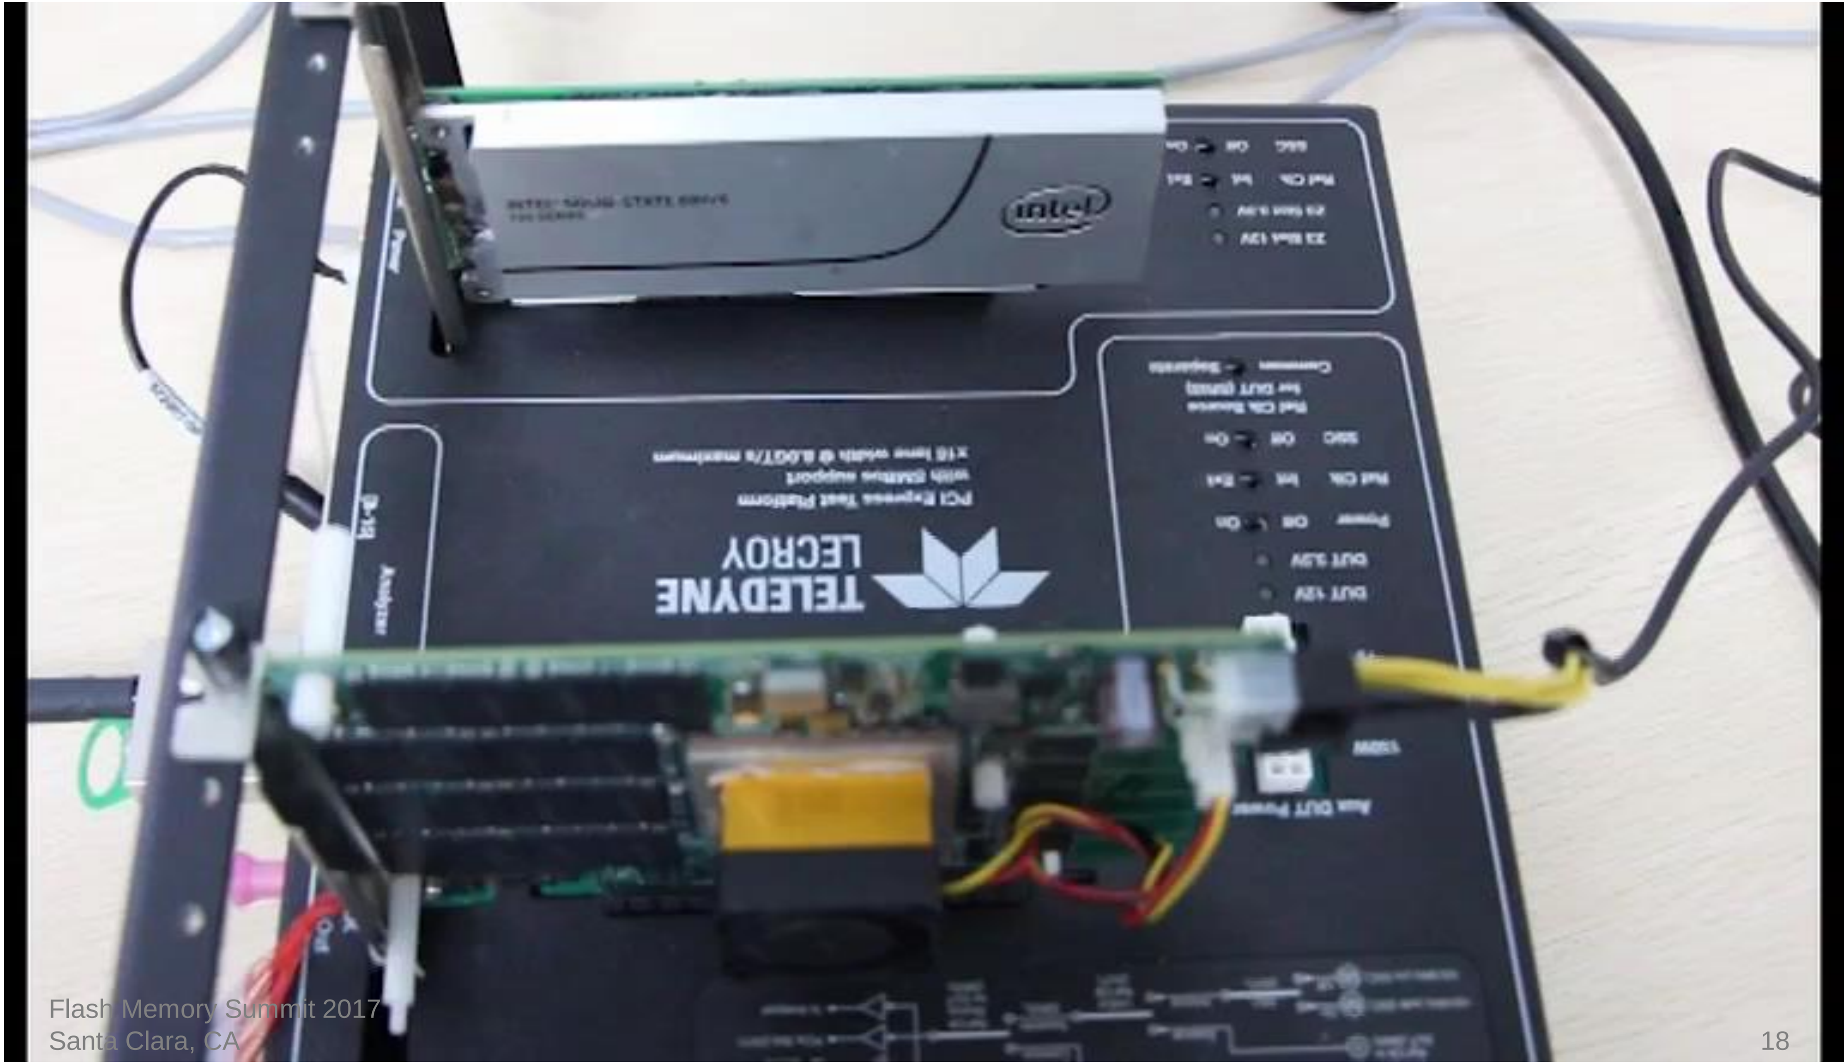
intel
TELEDYNE
LECROY
PCI Express Test Platform
with SMBus support
x16 lane width @ 8.0GT/s maximum