Storage vs Computing Networks

HiPINEB

February 5, 2017

NVIDIA

Bill Dally

Chief Scientist and SVP Research, NVIDIA Corporation

Professor, Stanford University

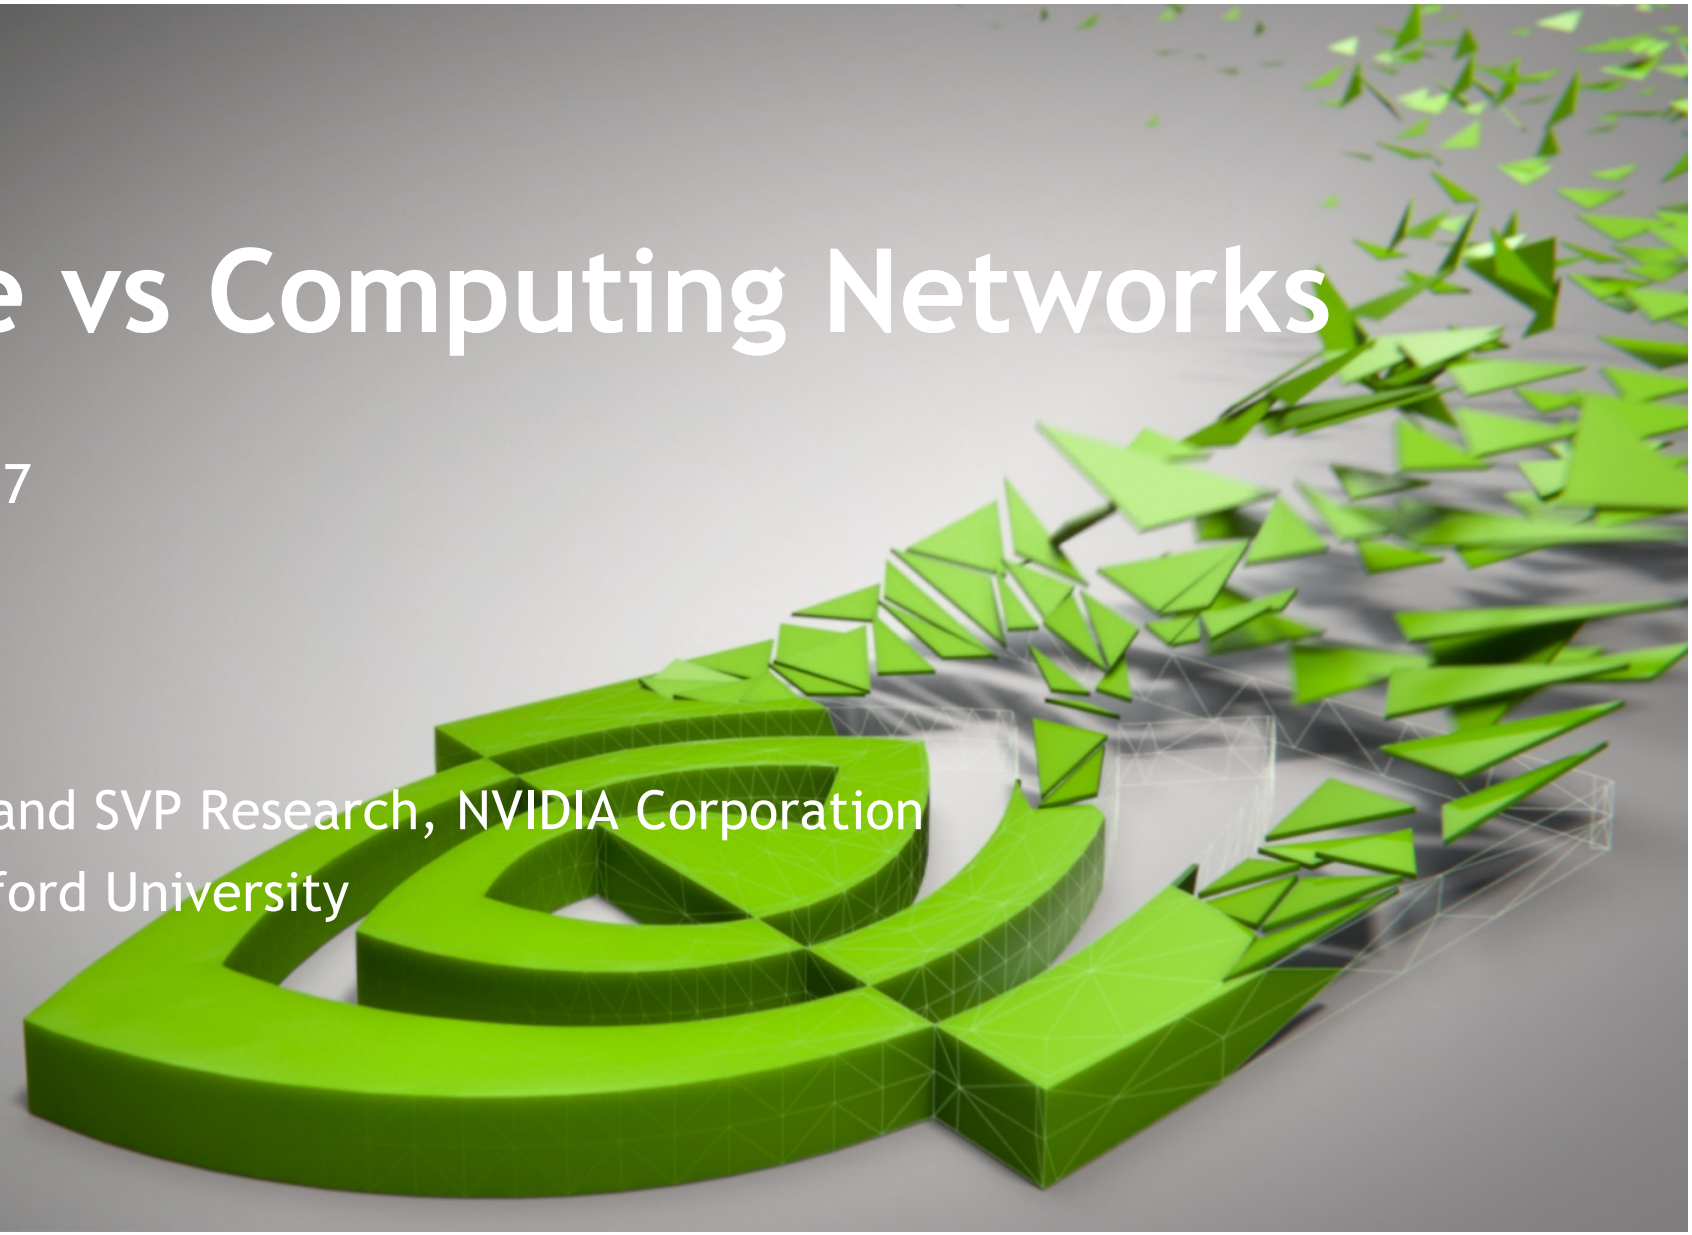

Bits are Bits

Bits are Bits

Networks Deliver Bits

Bits are Bits

Networks Deliver Bits

All Networks Want the Same Things

Bits are Bits

Networks Deliver Bits

All Networks Want the Same Things

Just in Different Amounts

Requirements: A Matter of Emphasis

- Storage Network
  - Cost-Effective Bandwidth
  - Latency not critical (rotating)
  - Scale to 10^4 endpoints
  - Large packets dominate
- Exascale Network
  - Cost-Effective Bandwidth
  - Low latency essential
  - Scale to 10^5 endpoints
  - Short packets important

Data Centers Want an Exascale Network

- Low-latency, fine-grained network access, scalability
- Latency of apps dominated by network
- Many data center apps could benefit from finer-grained parallelism (PGAS)
- And it shouldn't cost anything

$500

Bits are Bits

Networks are Networks