iVDIMM – World First Intelligent Memory Module

Handle Big Data Faster by Combining Memory and Compute

By: Mike Amidi

CEO of Xitore, Inc.

How Does Data Get Stored Today?

- Historically, data comes from the outside world (i.e. Cloud, IoT) into local system via network connections.
- The incoming data will reside momentarily in volatile memory (DRAM) prior to getting stored in to traditional storage (HDD, SSD, etc.)

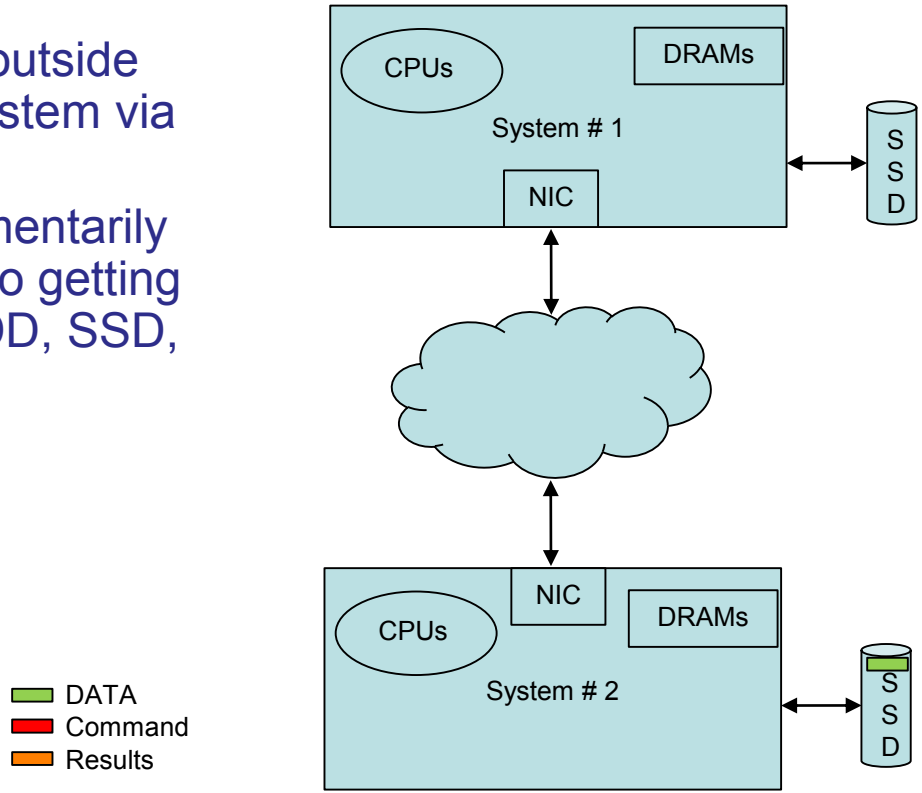

How does Application Access Stored Data today?

- First application will fetch required data segments into local volatile memory
- The processor will perform required task, analysis on the raw data and save resultant outcome into temporary local volatile memory
- Finally the OS will save a copy of final results into traditional storage units while a copy will be passed to requested application.

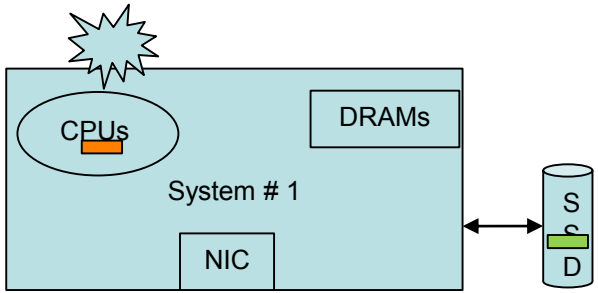

DATA
Command
Results

Requires Many steps:

- 1) Save:
  - Store external data into volatile memory (DRAM)
  - Move data into non-volatile Memory (SSD)
- 2) Retrieve:
  - Fetch stored data into volatile memory (SSD -> DRAM)
- 3) Compute:
  - Perform computation (CPU)
- 4) Move:
  - Store the results back into non-volatile memory (DRAM-> SSD)

Inefficiency of current implementation

- Requires too many unnecessary raw data movements to and from volatile to non-volatile memory.
- Keeps both memory channel and storage channel busy for this data movement
- It's inefficient and has a cost associated with power usage of this data movement
- Mostly, CPU idles waiting for data movement between non-volatile and volatile memory

Another Approach Based upon NAND

- 1) Send the required computation command to storage unit
- 2) Perform computation local within non-volatile memory
- 3) Send the final outcome of computation to volatile memory for application to access the results

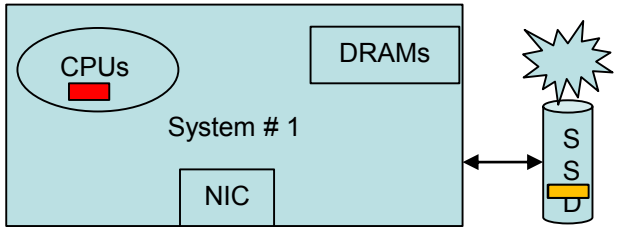

DATA
Application Command
Results

Requirements of NAND Approach

- This new method requires a storage with two components:
  1) A local operating system
  2) A local copy of Application to run required task (i.e. Search, RDBMS, SQL Server, etc.)

  And still requires moving back the final results between non-volatile to volatile memory for host application to access the results.

Benefits of NAND Approach

- Improve memory channel and storage channel traffic due to lower data movement
- Improve CPU utilization
- It's great for RDBMS, and search algorithms
- Excellent choice for sequential data access as flash is architected to do so.

New Approach – Based upon DRAM

- Using an intelligent volatile memory called (iVDIMM)
- Putting computation unit inside volatile memory module (i.e. DRAM)
- Application will send a function call to iVDIMM, which required task and algorithm to be executed on local data.

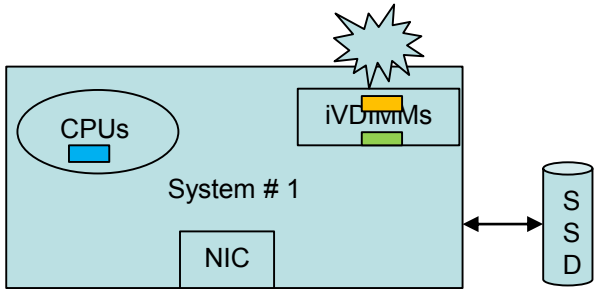

DATA & Application
Function () Call
Results

iVDIMM - DRAM Based Approach

- iVDIMM controller has all required computational algorithm within its local controller to be performed on its data
- Application will only make a function () call to select which algorithms needs to be executed by iVDIMM controller on local data.

iVDIMM - DRAM Based Approach

- iVDIMM's intelligent controller performs required task and computations on its local volatile memory data content.
- Application will fetch the results without any additional required steps or data movements between storage (SSD) unit and memory (DRAM).
- Host does not requires any hardware or operating system modifications to take advantage of this technology.

iVDIMM - DRAM Based Approach

- iVDIMM is the most efficient real time intelligent computation solution.
- Don’t move data or task back and forth from volatile to non-volatile memory when you can do the same more effectively and efficiently via in memory module itself.
- Release host CPU from unnecessary activities, free storage data movements and memory channel traffic. Reduce overall system power.
- iVDIMM will improve application latency, system power, and memory/storage channel bandwidth.
- iVDIMM is the only choice for in-memory computation of random access data due to DRAM architecture.

Summary

- iVDIMM is the world first intelligent volatile DIMM
- iVDIMM has local intelligent controller and required algorithm to execute required task by volatile DIMM itself
- iVDIMM does not require data movement between SSD and DRAM
- iVDIMM does not require hardware modifications to operate
- iVDIMM does not require any special system software to be running on the host system to function
- iVDIMM is the only single module intelligent memory solution of the future, capable of extremely fast byte addressable random access with backward compatibility.

iVDIMM - DRAM Based Approach

- For more information about this revolutionary disruptive patent protected technology contact us at:

info@xitore.com
marketing@xitore.com
sales@xitore.com