



**SDC<sup>18</sup>**

September 24-27, 2018  
Santa Clara, CA

[www.storagedeveloper.org](http://www.storagedeveloper.org)


# **Memory Class Storage and its Impact**

**Bill Gervasi**  
**Principal Systems Architect**  
**Nantero, Inc.**


# Agenda

- ❑ Describe carbon nanotube basics
- ❑ Apply CNT to memory architectures
- ❑ Compare efficiency of CNT memory to DRAM
- ❑ Application: main memory
- ❑ Application: storage devices
- ❑ Application: cell phones
- ❑ Application: artificial intelligence
- ❑ Embedding CNT into an ASIC
- ❑ Data encryption in a persistent memory world



# Carbon Nanotube Basics



Attractive Electrostatic Interactions



Repulsive Electrostatic Interactions




Van der Waals barrier keeps CNTs apart or together




Data retention >300 years (more likely >12,000 years)

No wear-out mechanism detected in >10<sup>12</sup> cycles

# CNT as a Memory Cell



- RESET: CNTs are not in physical contact = high  $\Omega$
- SET: CNTs are in physical contact = low  $\Omega$
- Delivers a 5ns core with no thermal sensitivity




# NRAM Scaling

4 layers CNT  
→  
16Gb




Add layers  
of CNTs

8 die stack  
→  
128Gb




Die stacking using  
standard TSV

Process agnostic  
Can be built using  
memory or logic  
processes




Example:  
4Gb/layer  
w/ 28nm logic

Logic/Memory process  
Substrate





Multi-level cell  
as a function  
of timing, voltage



Process scaling is a  
function of #CNTs  
per bit...  
well understood  
< 5nm

7nm process  
→  
512Gb/die  
1Tb/stack

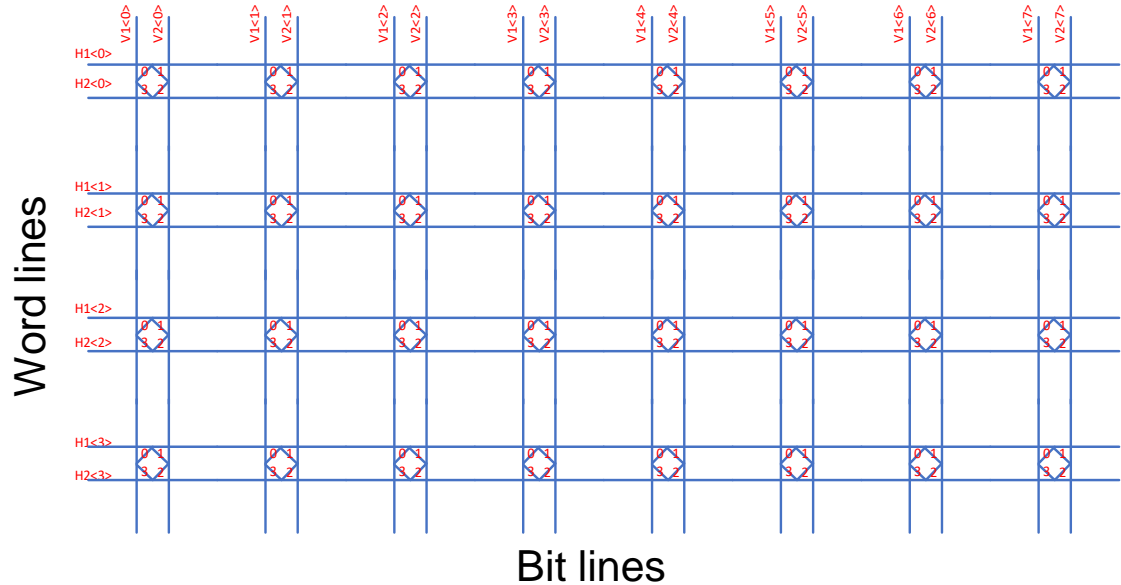

# NRAM Implemented in 1T-1R



1T-1R approach is good for smaller arrays (Mbits to low Gbits)

Can be tuned for much faster access times, array efficiency

# NRAM as Crosspoint




**Transmission characteristics of lines defines access time**

**Better solution for large arrays e.g., many Gbits**

# A Core is a Core is a Core



Any memory PHY  
can be inserted

All applications  
benefit from non-  
volatility






# “Memory Class Storage”



Storage Class  
Memory

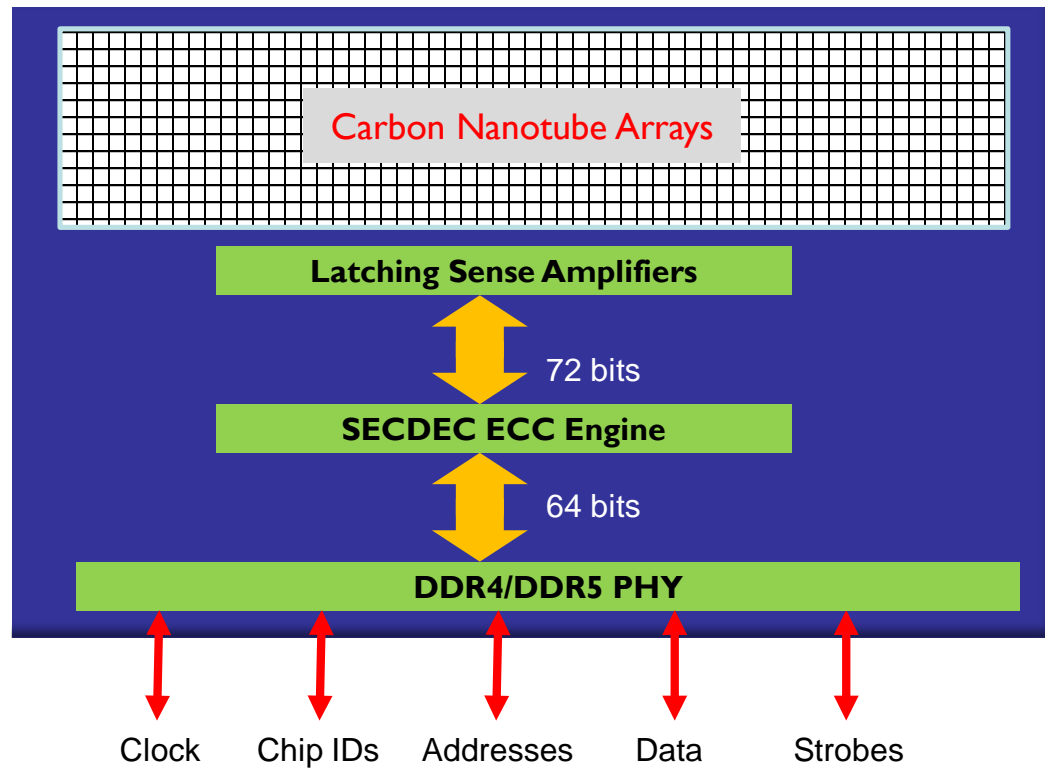


## Memory Class Storage

> DRAM performance  
= DRAM endurance  
> DRAM capacity  
< DRAM price



# DDR4/DDR5 NRAM With On-the-fly ECC



**Meets all DDR4/DDR5  
SDRAM timing**

**Translates DRAM  
protocol to NRAM  
accesses**

**Non-Volatile**

- **No refresh needed**
- **Non-destructive read**

# Power States

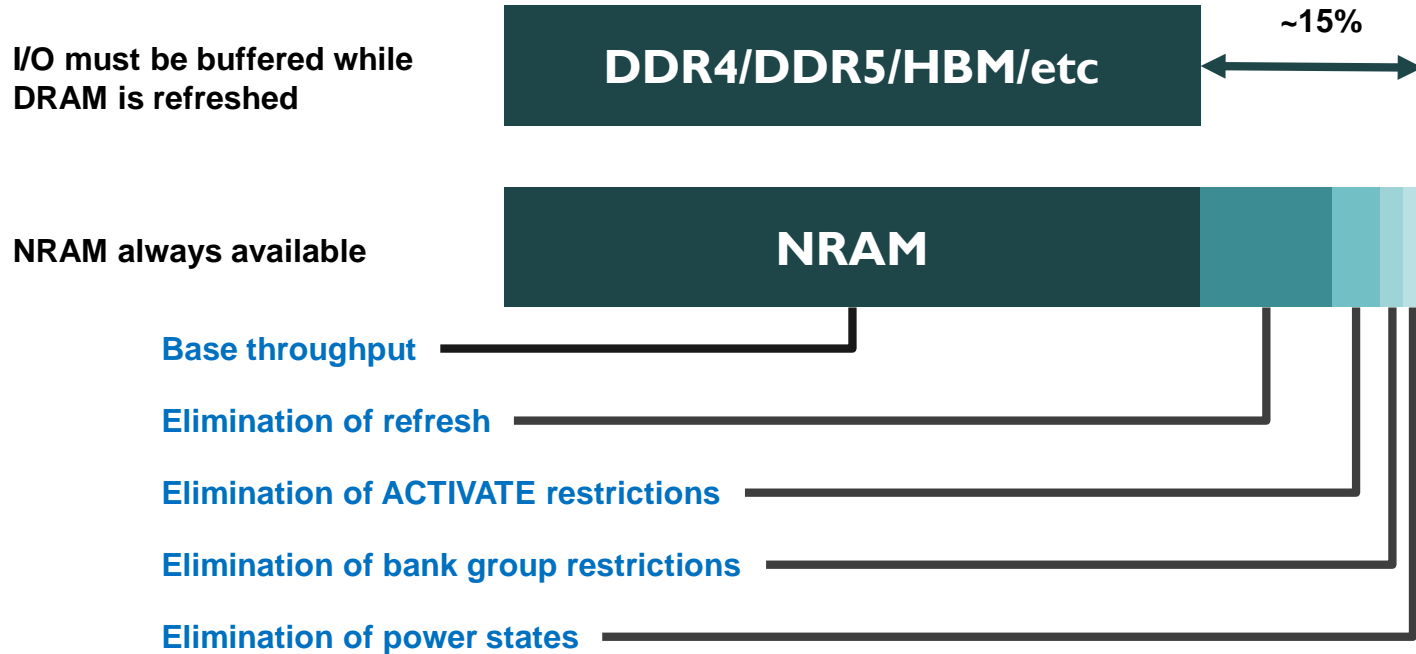
## DRAM

Active Power	Normal runtime
Active/Precharge Standby	Standby with banks open/closed
Must exit standby to refresh	
Refresh	High power read/write cycle
Self Refresh	Self refresh, banks closed, data read/write cycle
Power always consumed	

## NRAM

Active Power	Normal runtime
Standby	Standby, banks always open
Never need to exit standby	
No refresh needed; treated as NOP	
Low power	Banks open (no access), data retained without energy
Power off	Data retained, banks close, interface reinitialize at power restore

# More Deterministic than DRAM



# Implications of MCS on Applications

**Elimination of power fail concerns & backup energy**

**Elimination of data reload times**

**Elimination of recalculation times**


**Elimination of checkpointing**

**Reduction in data buffering**

**Data security concerns**

# NRAM is Inherently an NVDIMM-N


NVDIMM



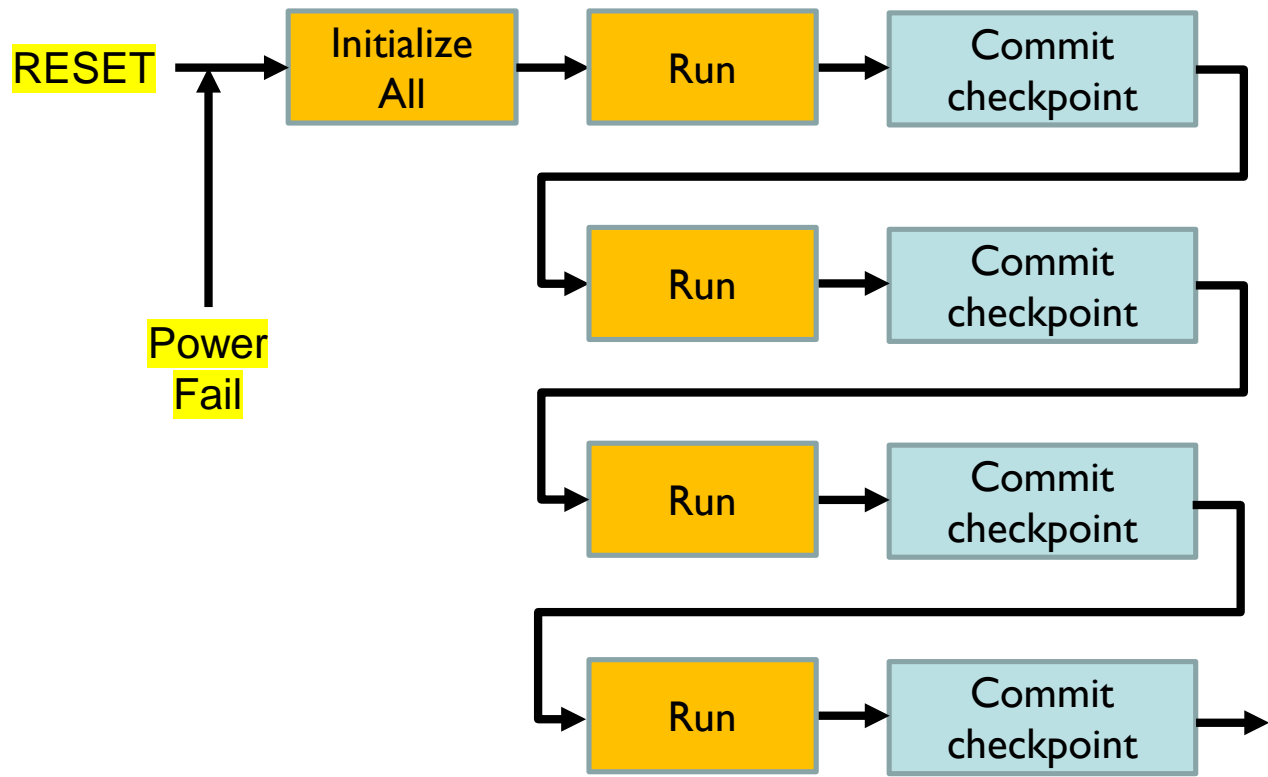
Supercapacitor array, 10W for 2 minutes

NRAM doubles capacity  
Eliminates support circuits  
Eliminates energy source  
At much lower cost

NRAM RDIMM

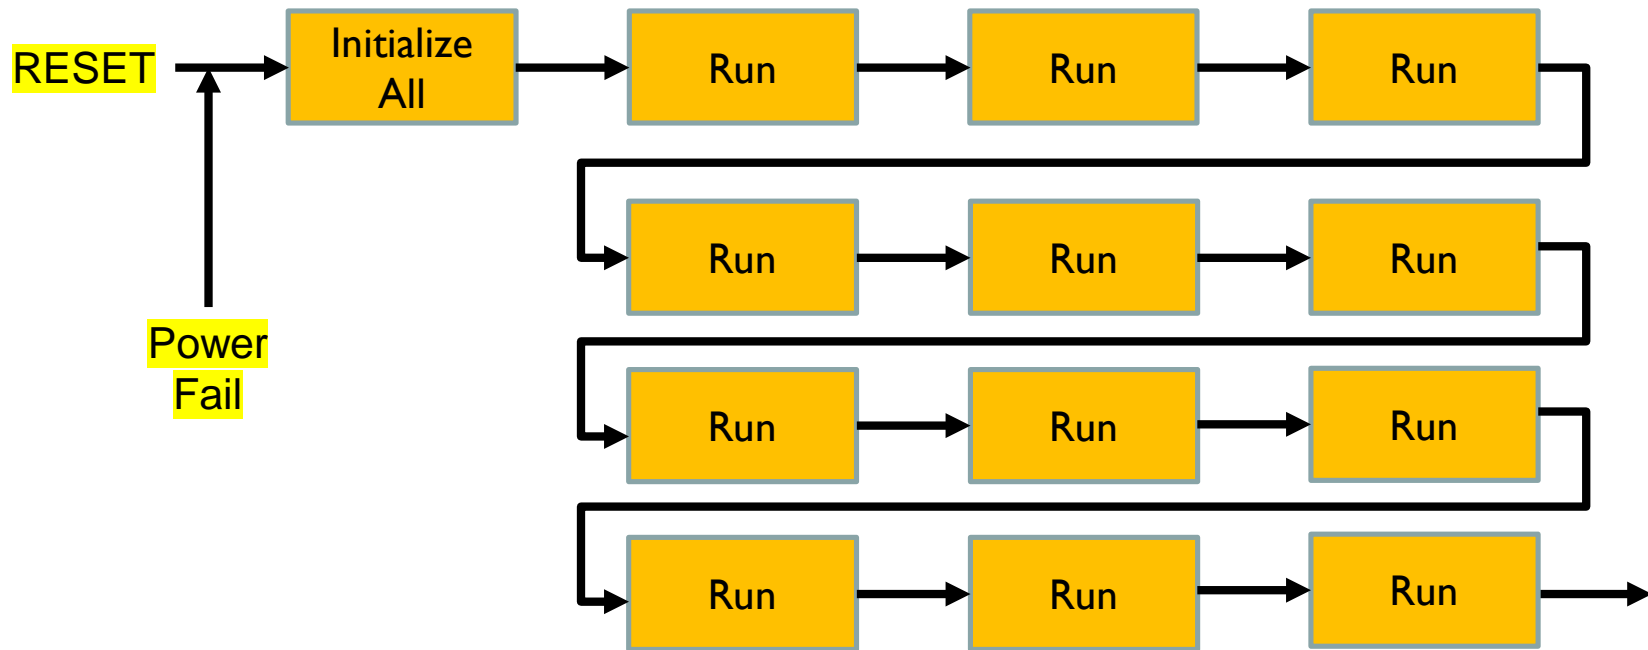


# Databases Without Persistence



Without persistence, databases must save periodically to ensure immunity from power loss

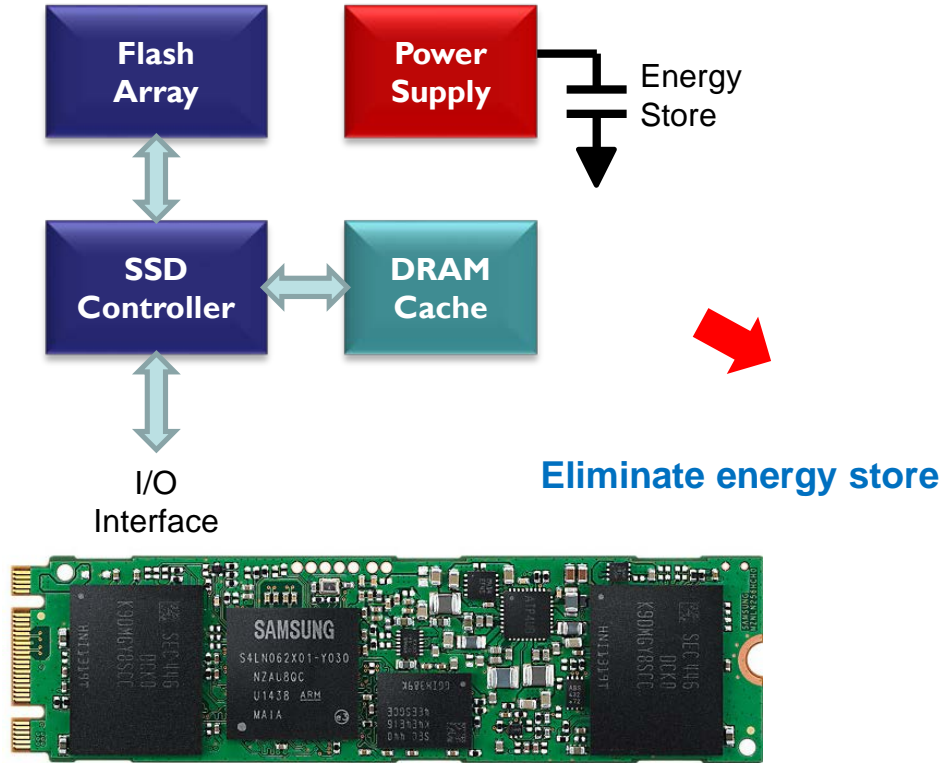
# Databases With Persistence



No power loss risk → no need for checkpointing




# Death to the Battery!




Cache size decoupled  
from energy requirements

More room for storage  
chips

Can add features such as  
hibernate partitions

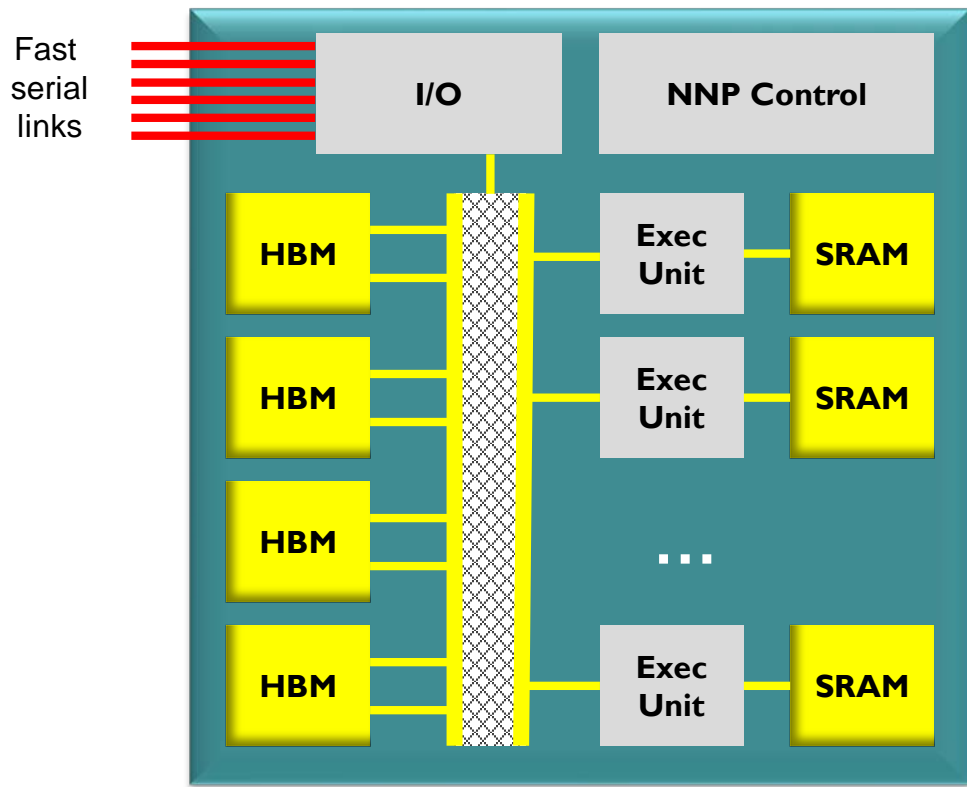


# NRAM in a Cell Phone




**True zero power standby**

# Artificial Intelligence Architectures



Hypercube/Toroid



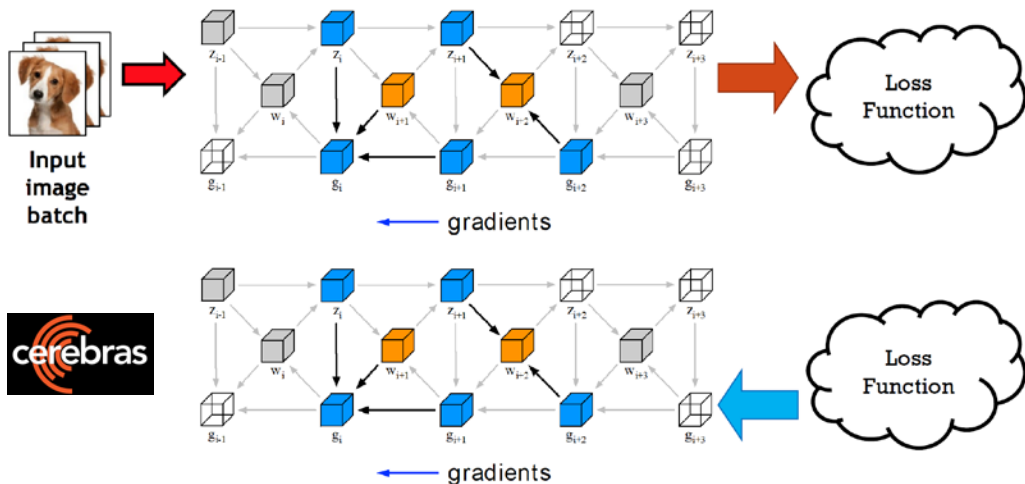
SIMD architectures

Matrix interconnections

Serial spigots limit load/save time

# Artificial Intelligence Applications

## BACKPROPAGATION



Learning algorithms often update the model

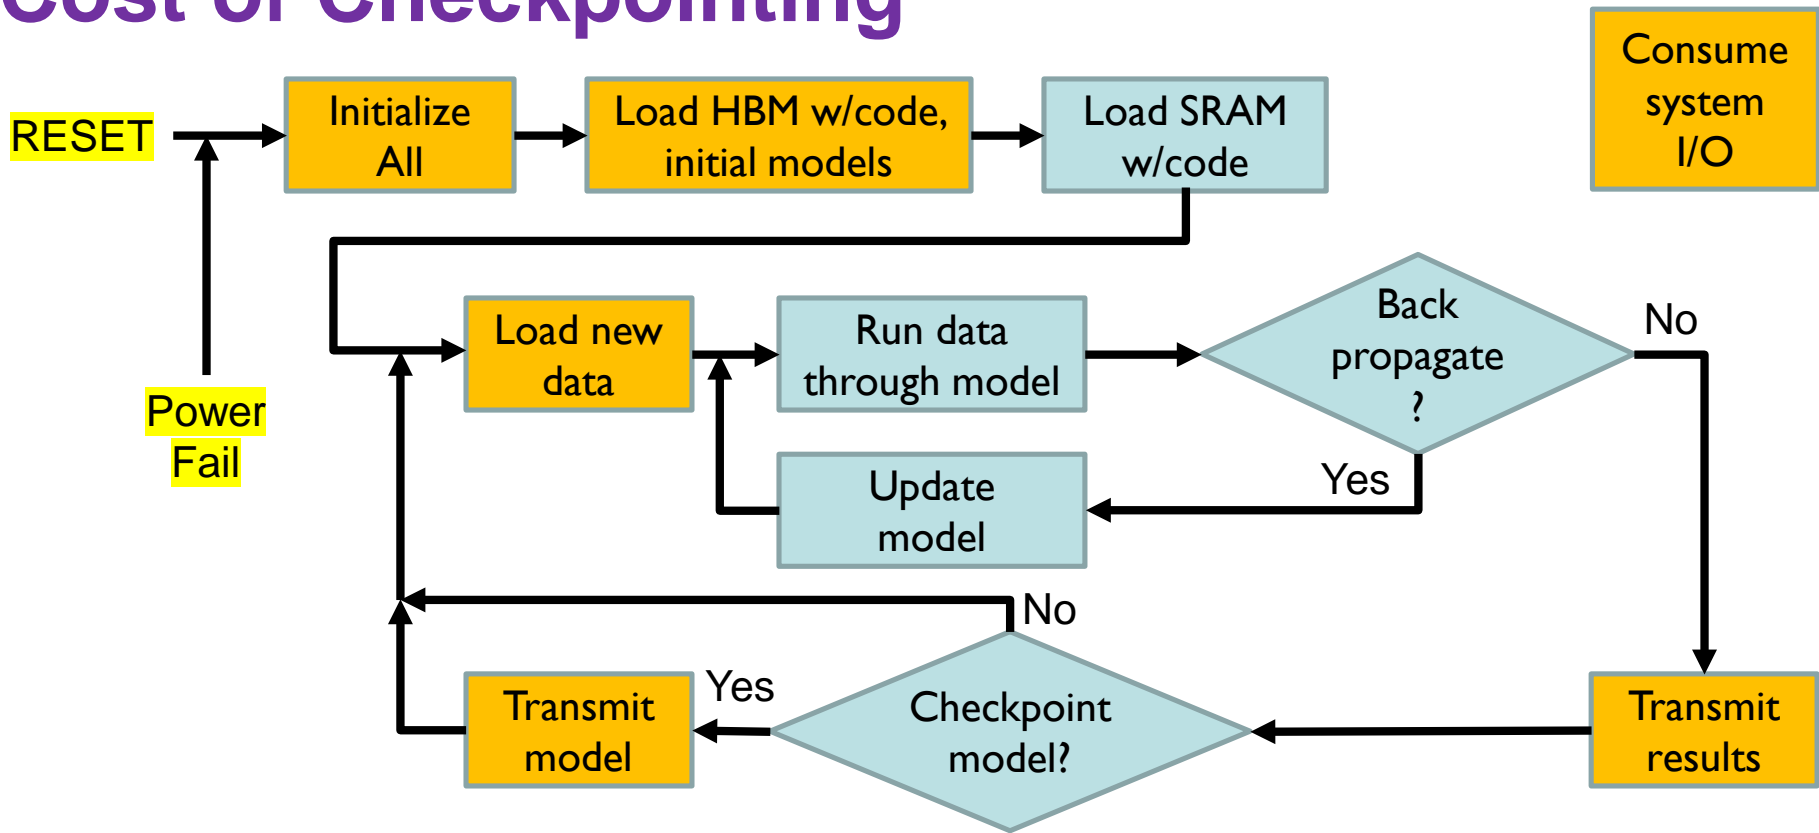
On power fail, the updates are lost

Potentially forever

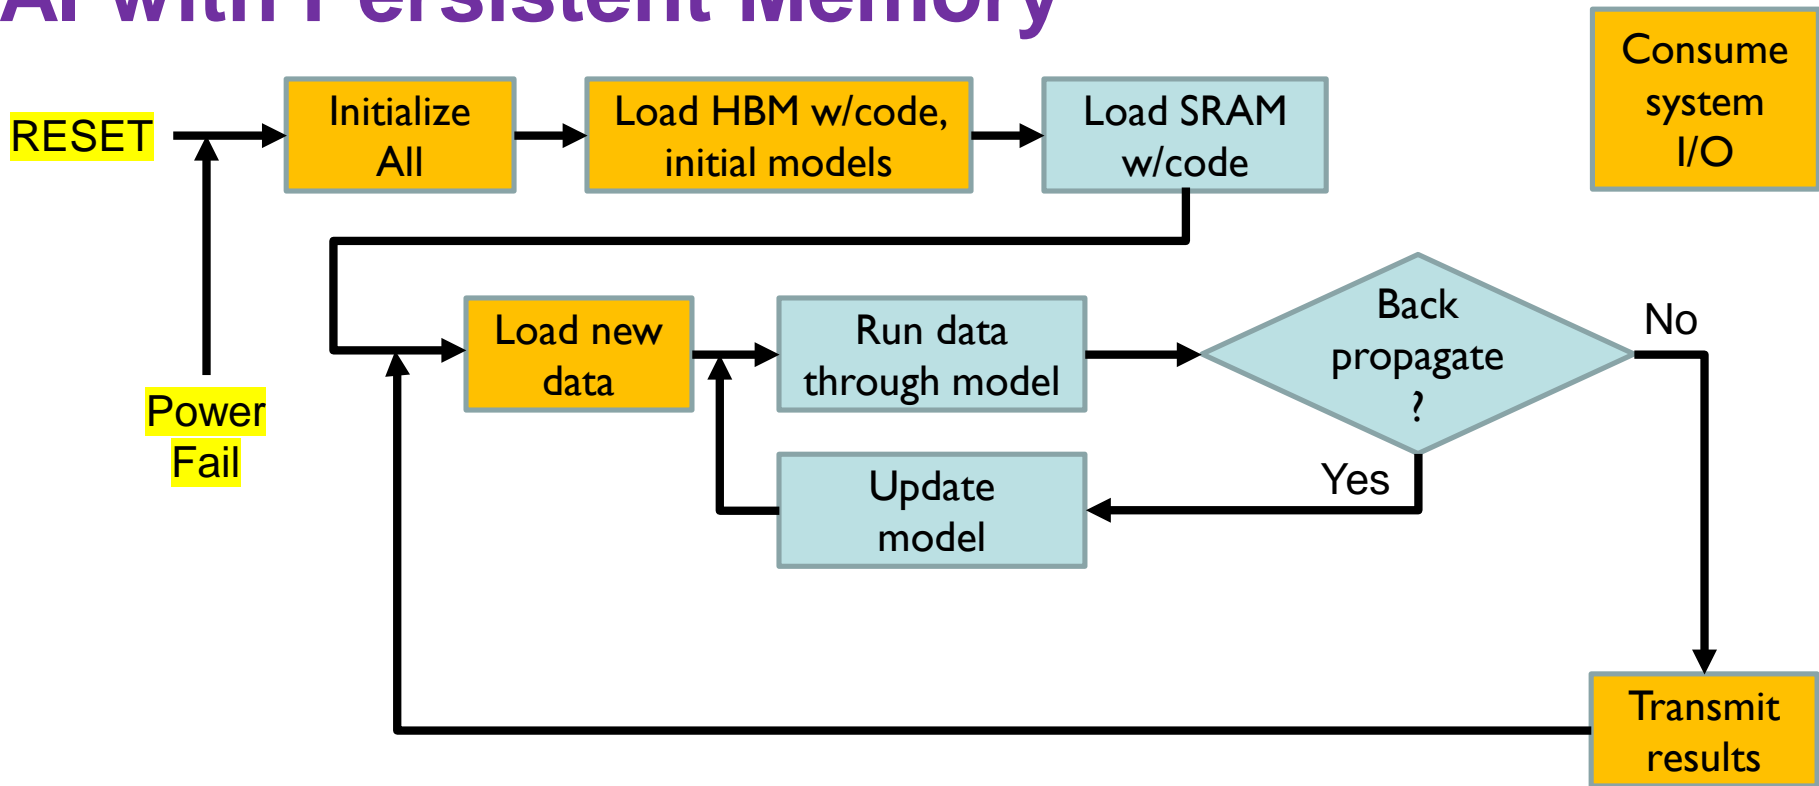
Backing up modified models is slow,  
consumes channel bandwidth

Persistent memory prevents data loss

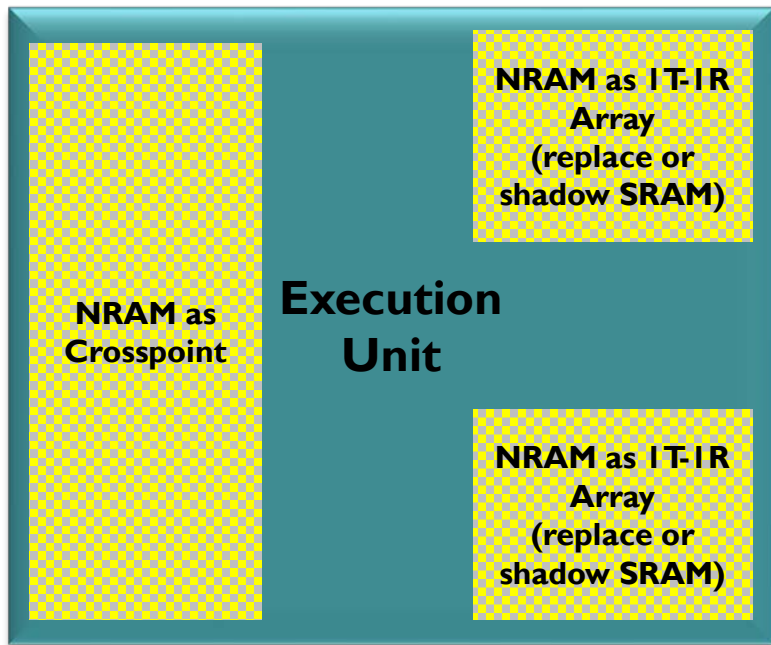
# Cost of Checkpointing



# AI with Persistent Memory



# Embedding NRAM in AI



CNT cells are built above the logic

Crosspoint and 1T-1R can be constructed at the same time

Just a matter of where the drivers and sense amps are placed


On power fail:

Registers copied to NRAM shadows

Data set saved automatically

Cache (SRAM) need not be reloaded


# Adding NRAM to the Mix



*Concept diagram...  
CNT areas don't  
actually stick out*

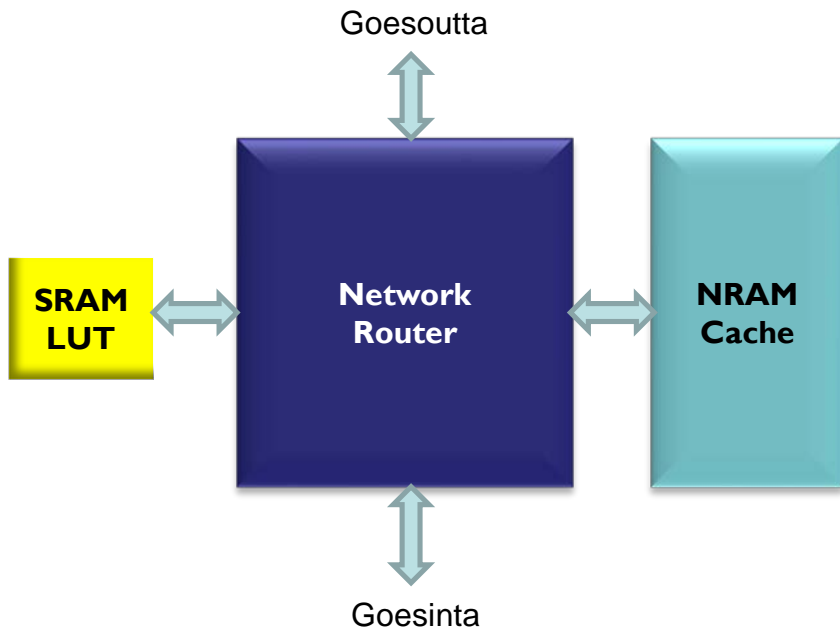
**CNT drivers & sense  
incorporated into  
customer design**

**CNT layers deposited  
on top of wafer**





# Exploiting Elimination of Refresh



**Data streaming through routers needs to be buffered during DRAM refresh cycles**

**350 ns refresh recovery =  
140K bit times @ 400 Gbps  
interface**


**Eliminating refresh greatly  
simplifies buffering**

# Data Security for Persistent Memory

Data persistence is a security challenge

Servers often perform encryption in CPU

Data on module is useless if moved



Some architectures moving to on-module encryption

More important if RDMA is supported

# Summary

- ❑ Carbon nanotubes (CNT) make a great memory core
- ❑ Flexible: any DRAM may be replaced by CNT
- ❑ Crosspoint or 1T-1R arrays have different uses
- ❑ May be incorporated into customer silicon
- ❑ Persistent memory is more efficient than DRAM
- ❑ Efficiency gains improve chip & system architectures
- ❑ Main memory & AI: eliminate checkpointing
- ❑ Storage devices: decouple cache from energy stores
- ❑ Communications: reduce need for buffering
- ❑ Data encryption choices include in CPU or on DIMM

***Thank you for your time***

Bill Gervasi  
bilge@nantero.com