Flash Memory: Key Signal Processing Issues

Jaekyun Moon

EE Dept
KAIST

April 19, 2010

NVRAMOS Forum

Communications & Storage Lab

HDD vs SSD: Price Trend

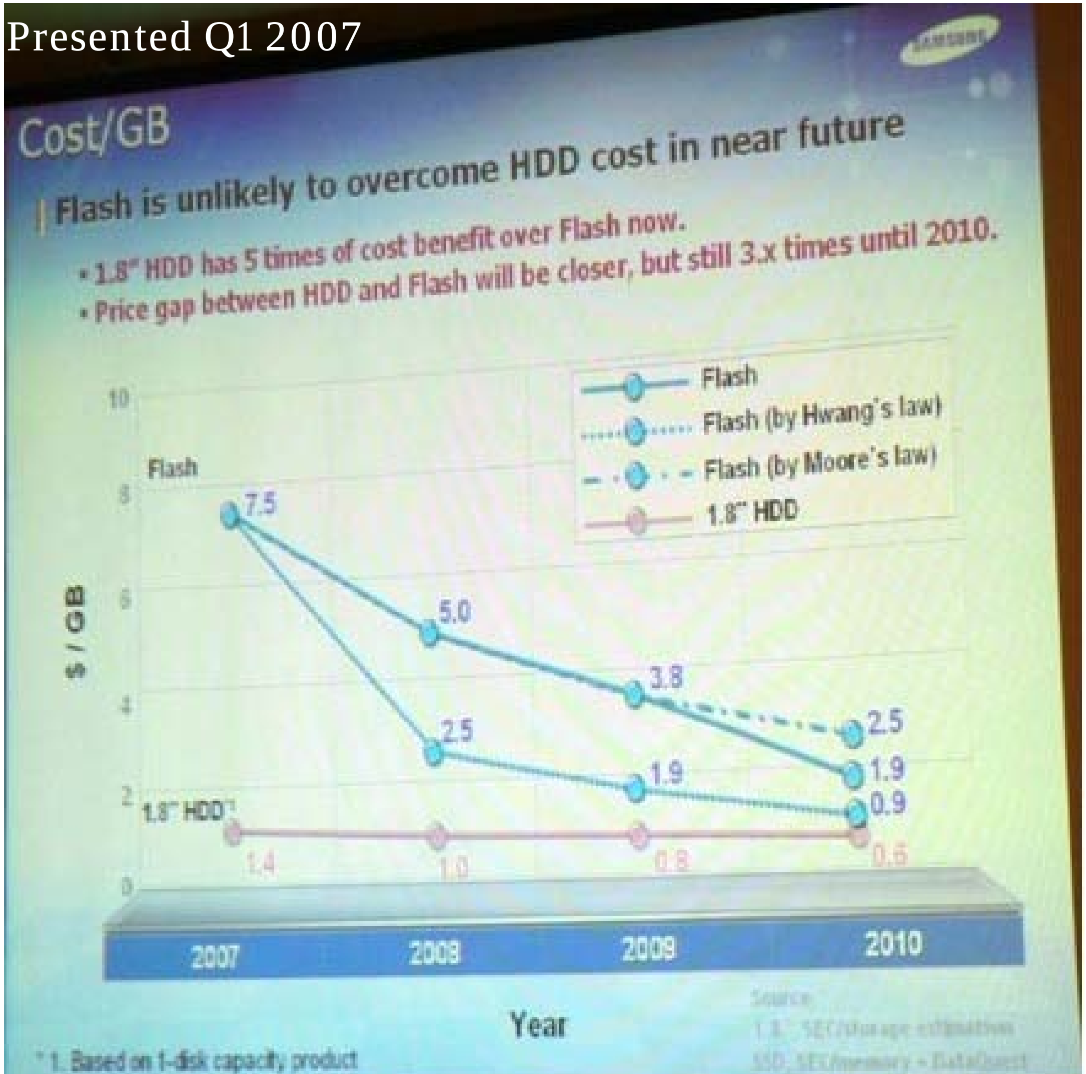

SSD prices were assumed to fall by about <u>60% annually</u>.

HDD vs SSD: Price Trend

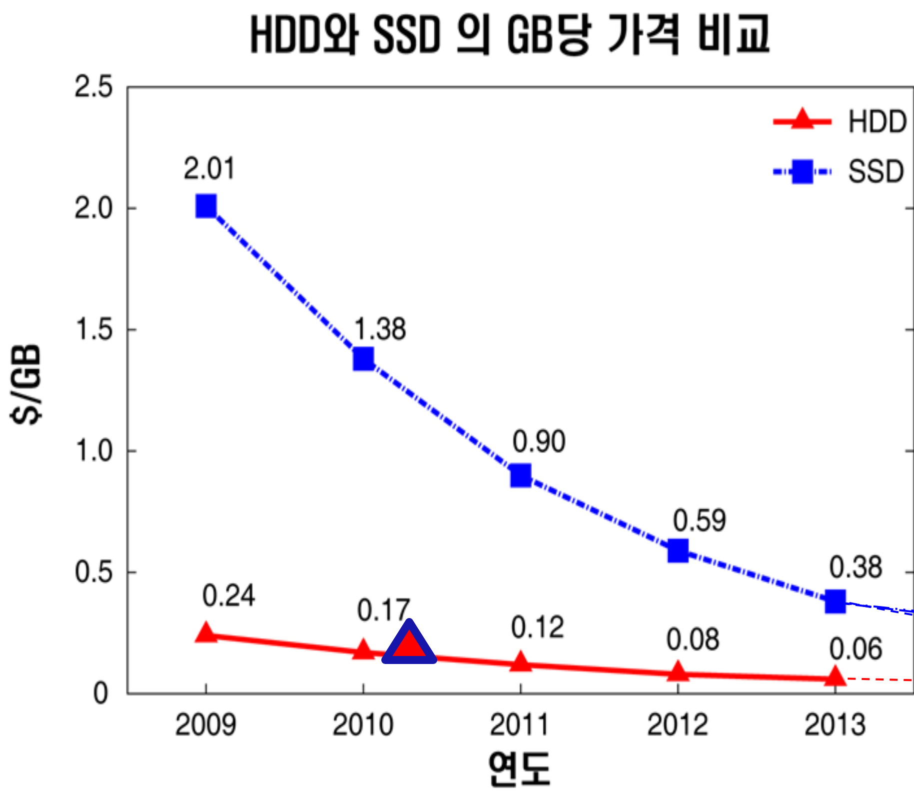

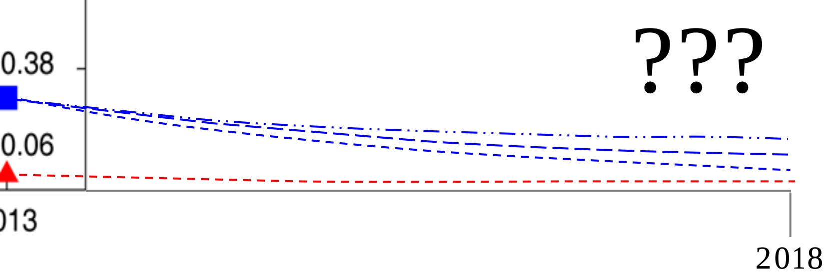

As of March 2010:

G3: $400/120GB=3.33

Vertex: $230/60GB=3.83

Samsung 256: $695/256GB=2.71

Intel X-25: $499/160GB=3.12

Seagate Barracuda XT: $340/2TB=0.17

WD Caviar Blue: $60/250=0.24

연도별 **HDD**와 **SSD**의 **GB**당 가격 비교 **(**출처**: iSuppli)**

Home Storage Server

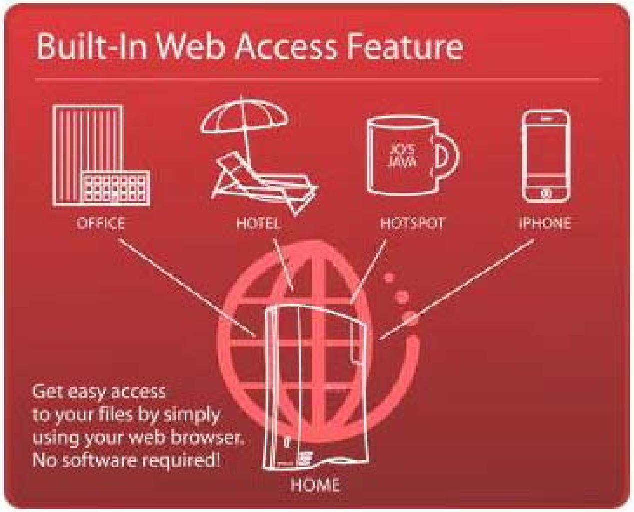

Home Storage Server

1TB NAS Server, 480/12 Mbps
$138.00

Price: $0.14/GB

Flash Memory in IPad

$499 for 16G iPad versus $699 for 64G iPad

KAIST

Sustaining Storage Density Growth…

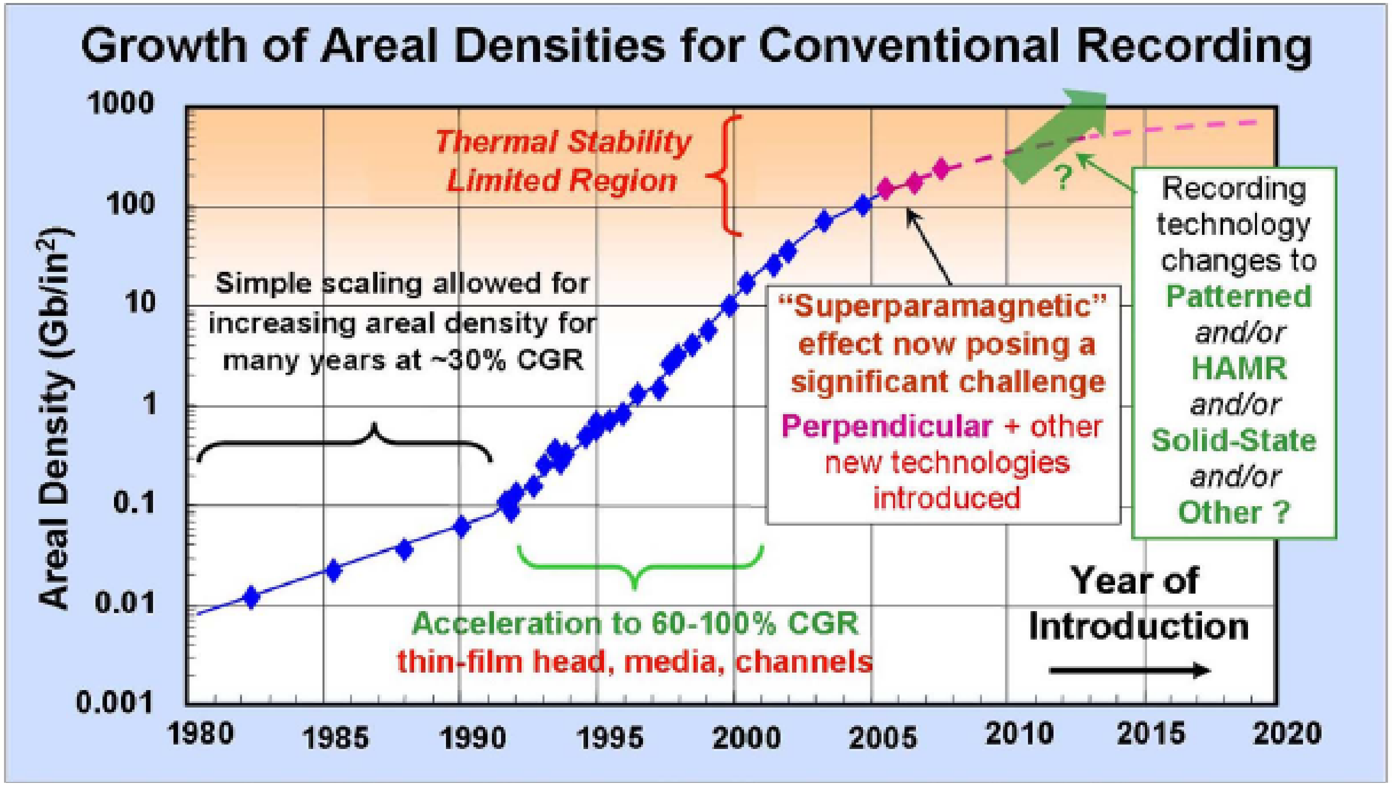

Proceedings of the IEEE | Vol. 96, No. 11, November 2008

ComSto
Communications & Storage Lab

Outline

- How many raw errors can we realistically correct using advanced SSP?
  - LDPC code versus BCH
- Essential Ingredients
  - Channel characterization (J. Moon)
  - Capacity-achieving coding (J.S. Ha)

MLC NAND Flash Architecture

A 2-plane, 4K-page MLC architecture. The 4K page has 4,096 bytes of data and 218 bytes of spare area.

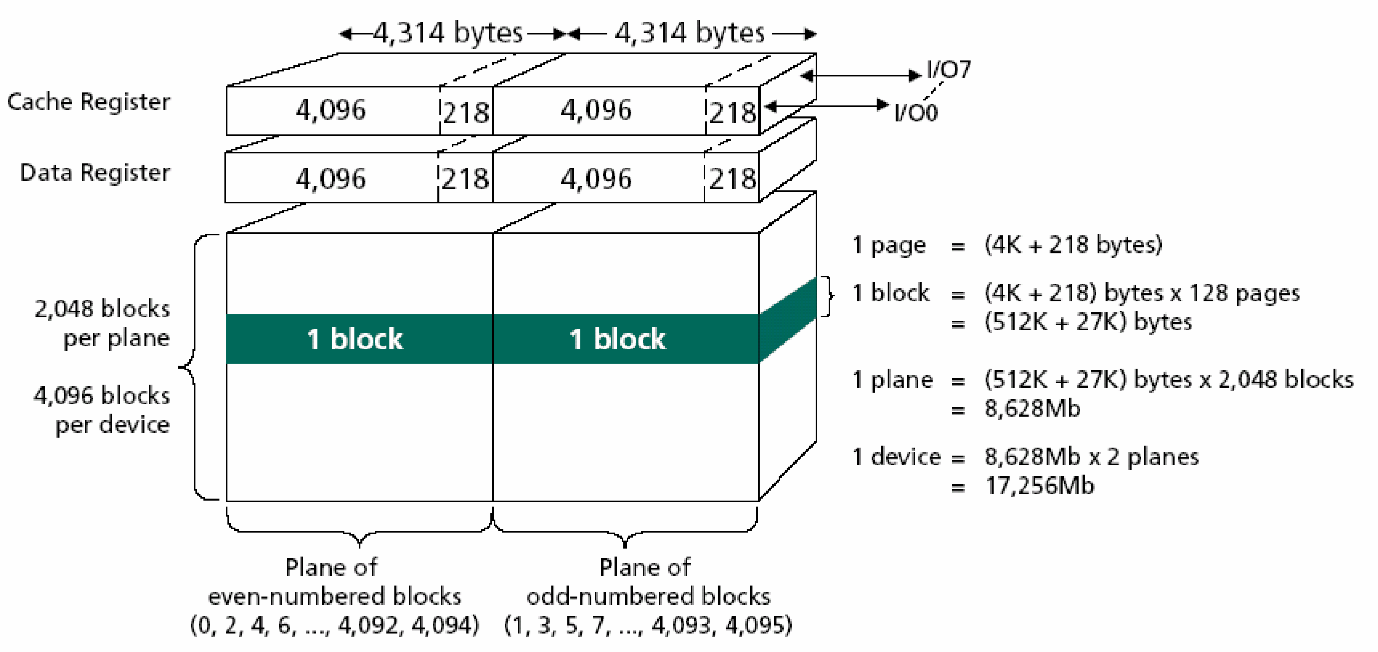

A 16Gb MLC NAND Flash Example

SSP Signal Flow

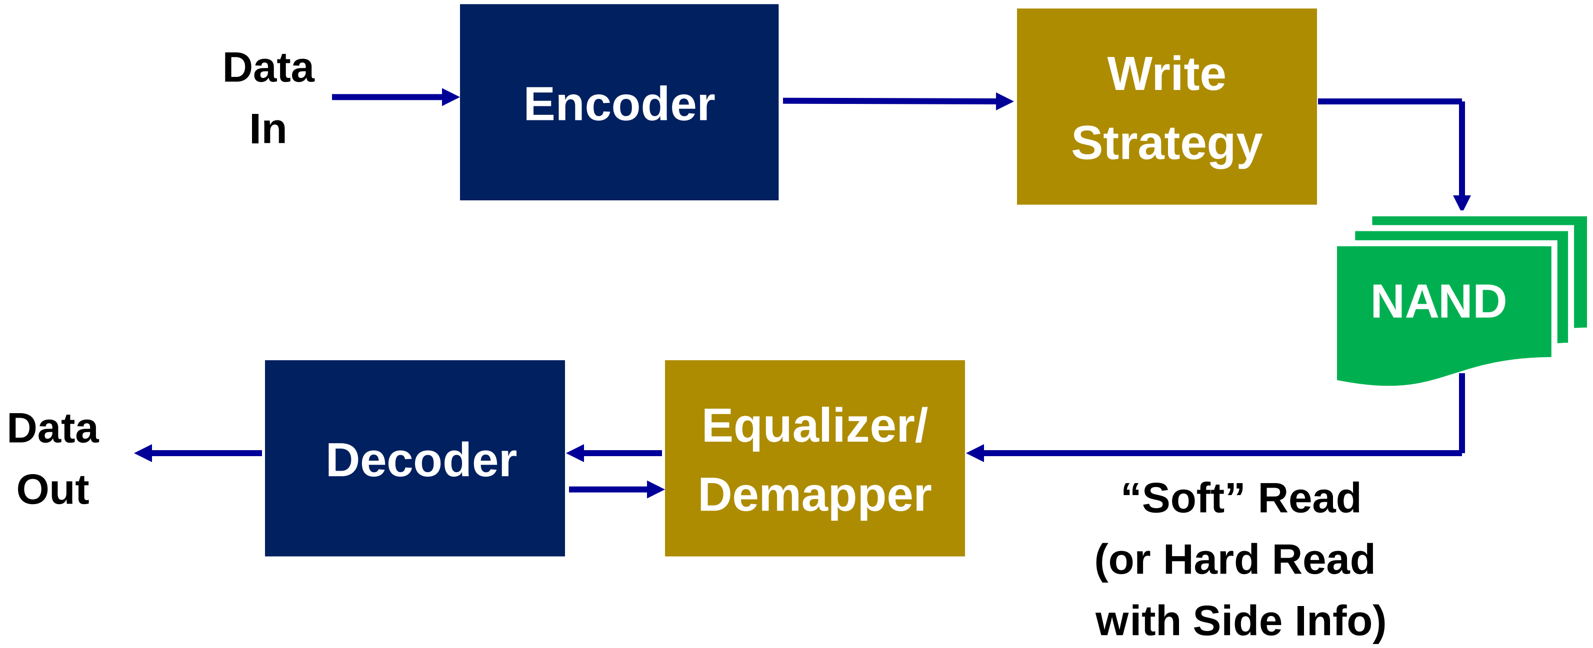

BCH Code Performance: Raw BER vs Corrected WER

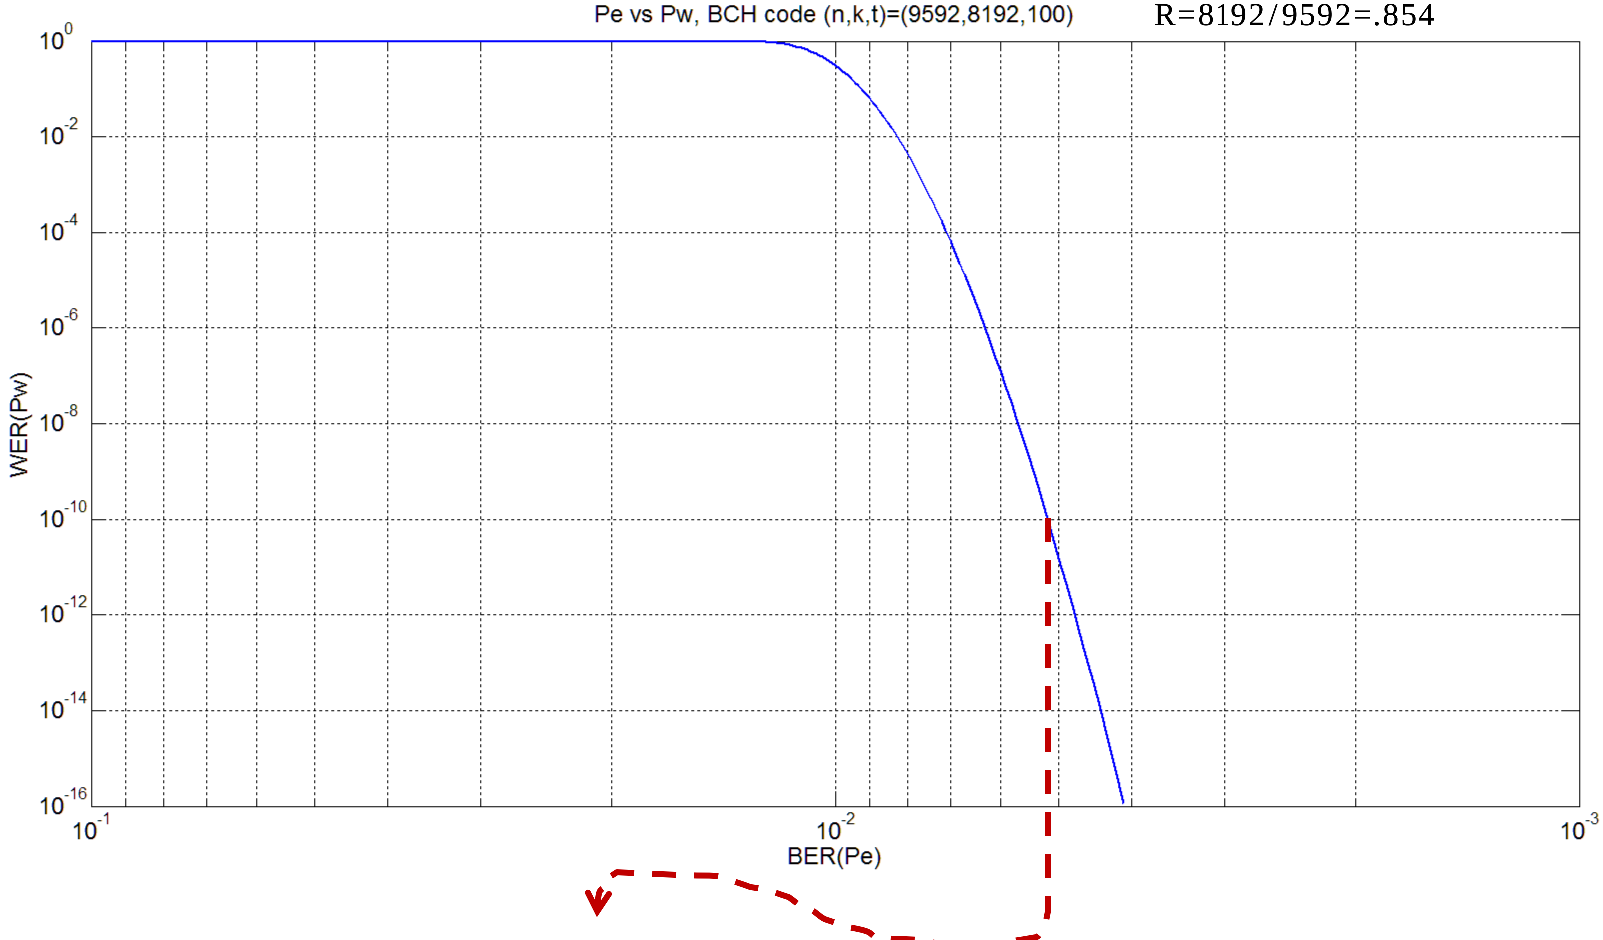

Need BER of ~5x10^(-3) or lower
to bring down WER to less than 10^(-10).
Can handle less than 48 error bits on average per 9592 bits.

LDPC Code Performance

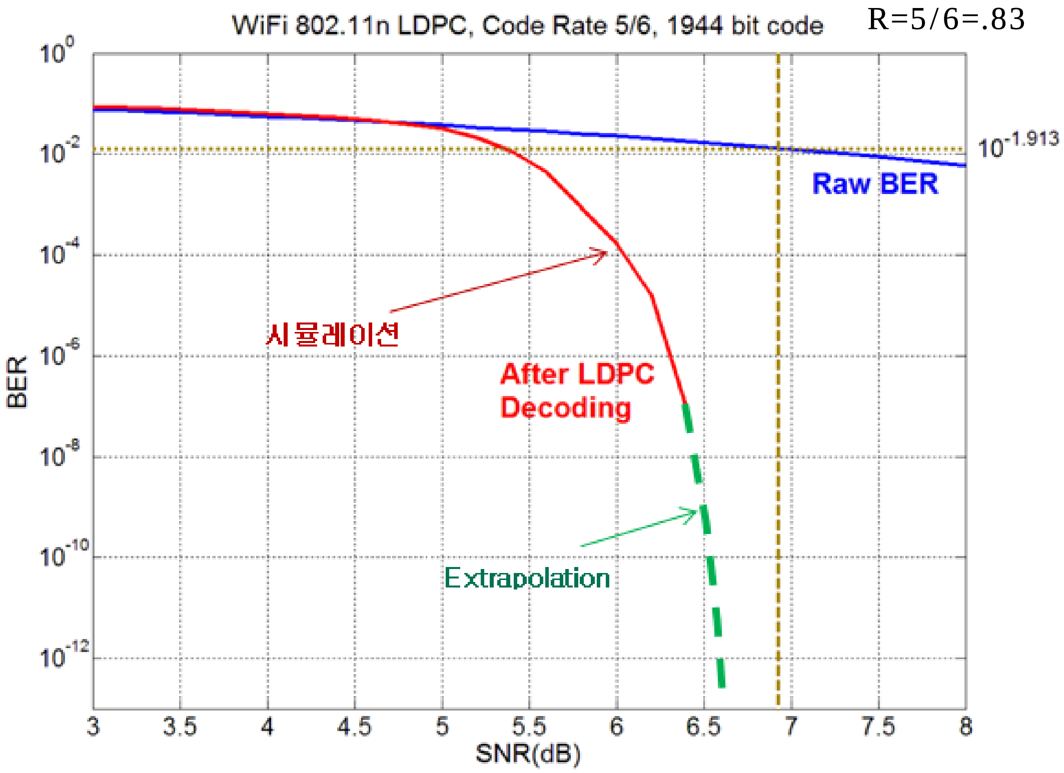

LDPC can handle 10^(-1.913)*9830=120 error bits on average per 9830 bits (Probably ideal and optimistic).

Real LDPC Error performance (10G Ethernet) – no visible error floor (gives us hope!)

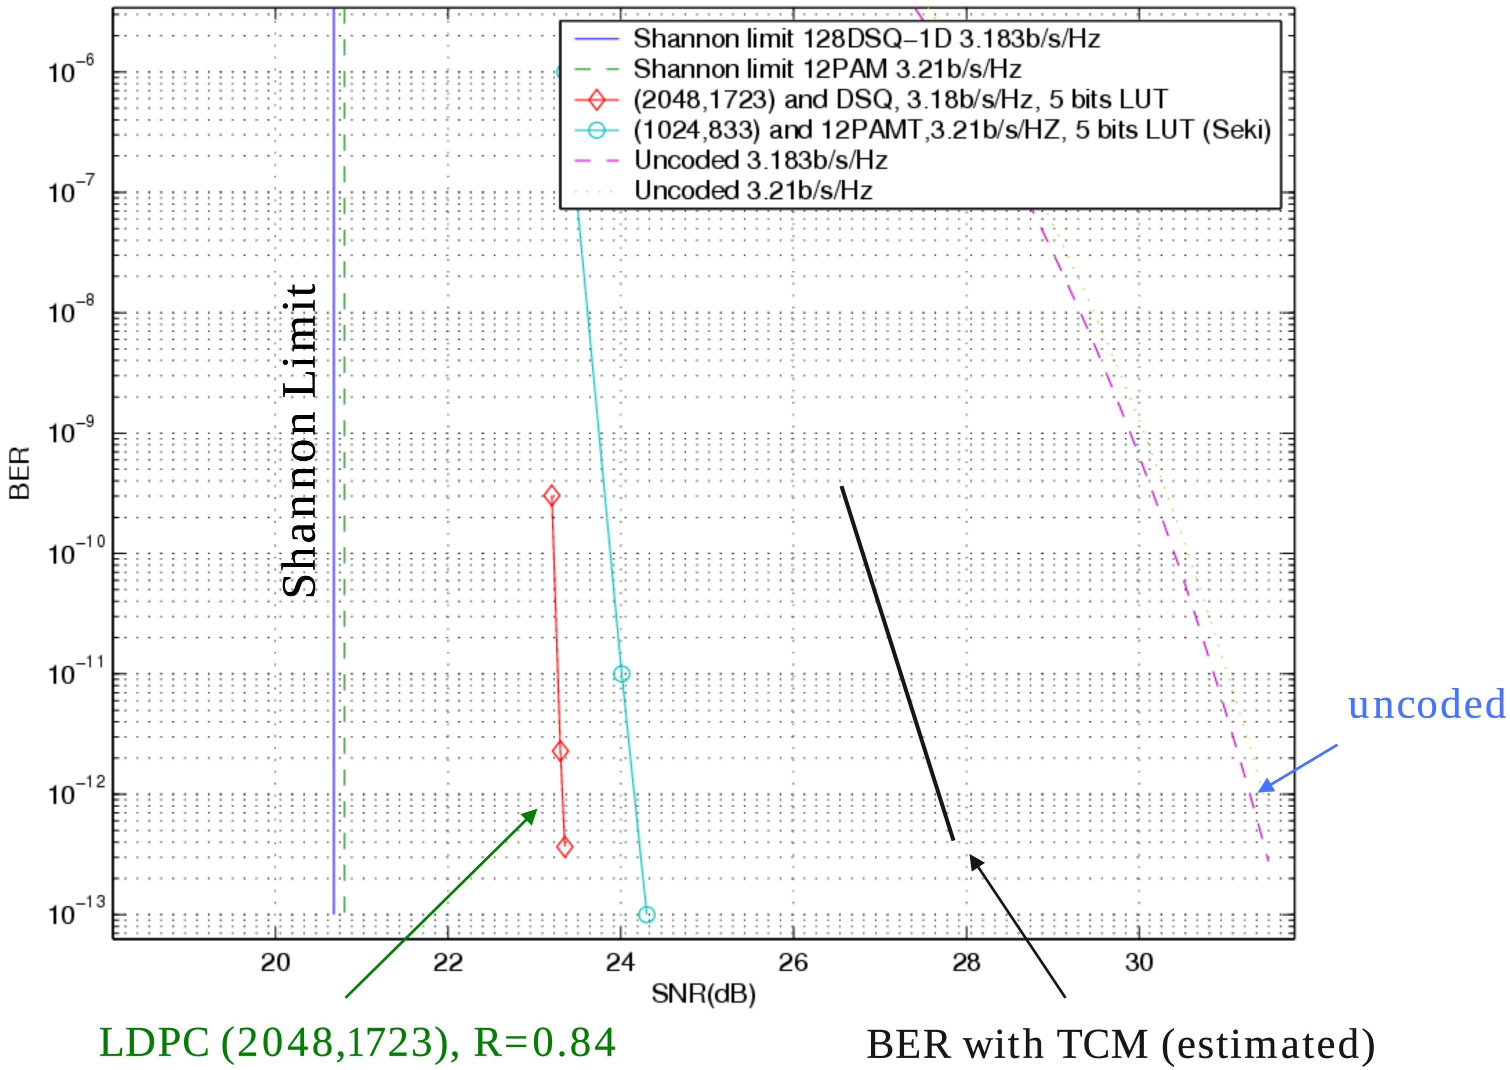

Source: G. Ungerboeck, 2010

Essential Ingredients in Advanced SSP

- Accurate channel characterization
- Soft information generation and processing

Channel Modeling: Signal-Level Characterization of Cell Correlation (Disturb)

Channel Modeling: Cell Correlation

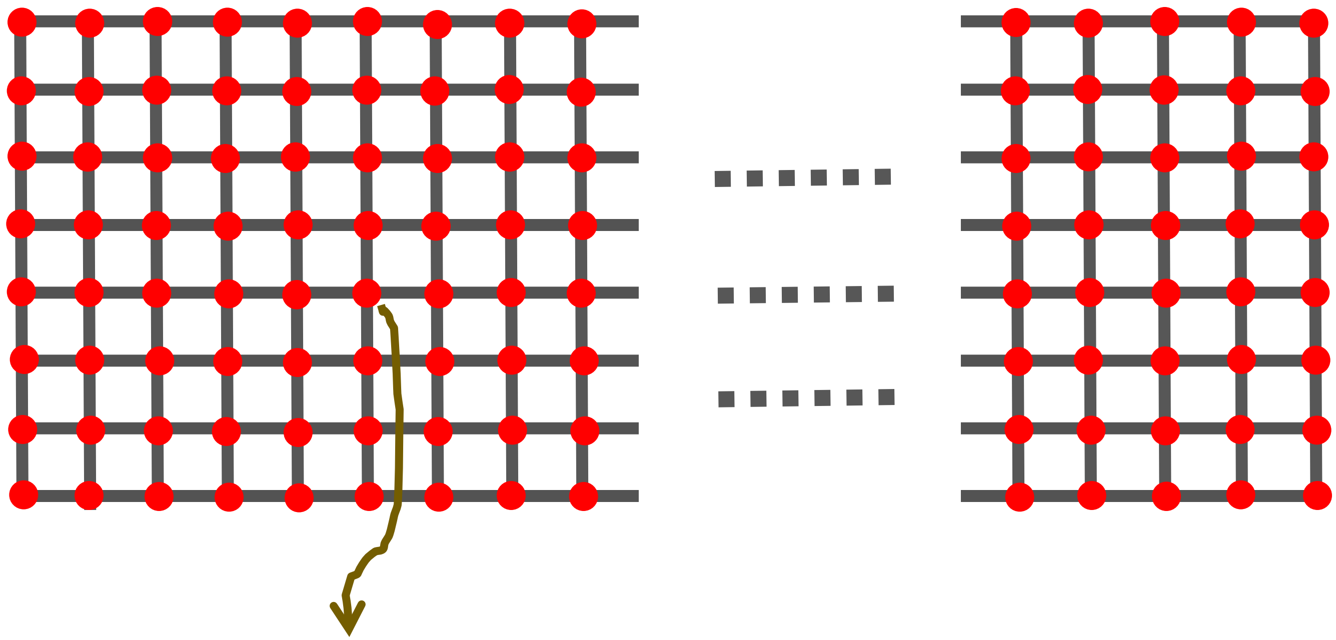

No cell correlation $r_{ij} = x_{ij} + n_{ij}$

Cell correlation through program/read disturb $r_{ij} = x_{ij} + \Delta(\underbrace{x_{ij}}_{\text{victim cell}}, \underbrace{x_{i_1 j_1}, x_{i_2 j_2}, x_{i_3 j_3}, x_{i_4 j_4}}_{\text{affecting cells}}) + n_k$

affecting cells

victim cell

Channel Identification Problem

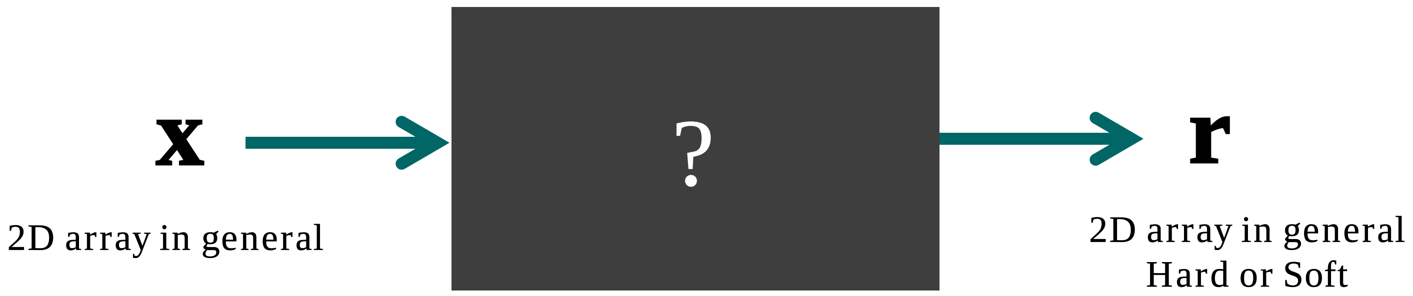

Feed the system with known data **x**.
Observe **r**.

Characterize the system enough, so for new data **x'** we would know what **r'** is.

Channel Identification Problem

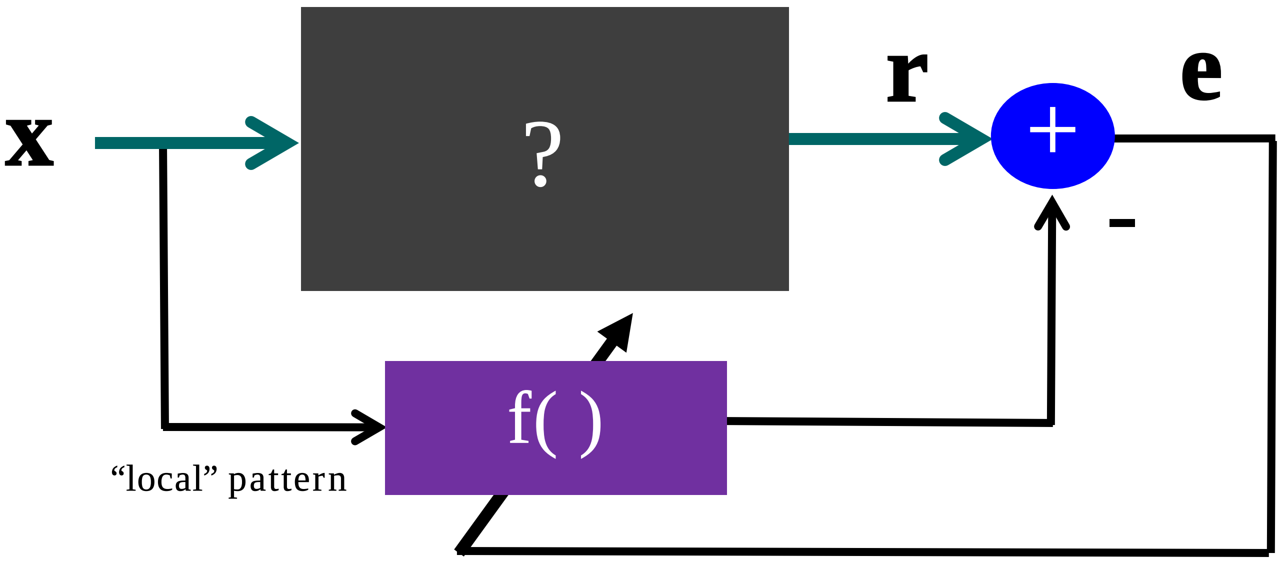

Feed the system with known data **x**.
Adjust f( ) until **e** is minimized (a sequential update algorithm is used).

Once **e** stabilizes, f( ) should resemble the system closely.

We in essence are fitting the unknown box with a partially unknown function (a certain structure is imposed)

Characterized Channel

Example of f( ): Linear FIR Filter

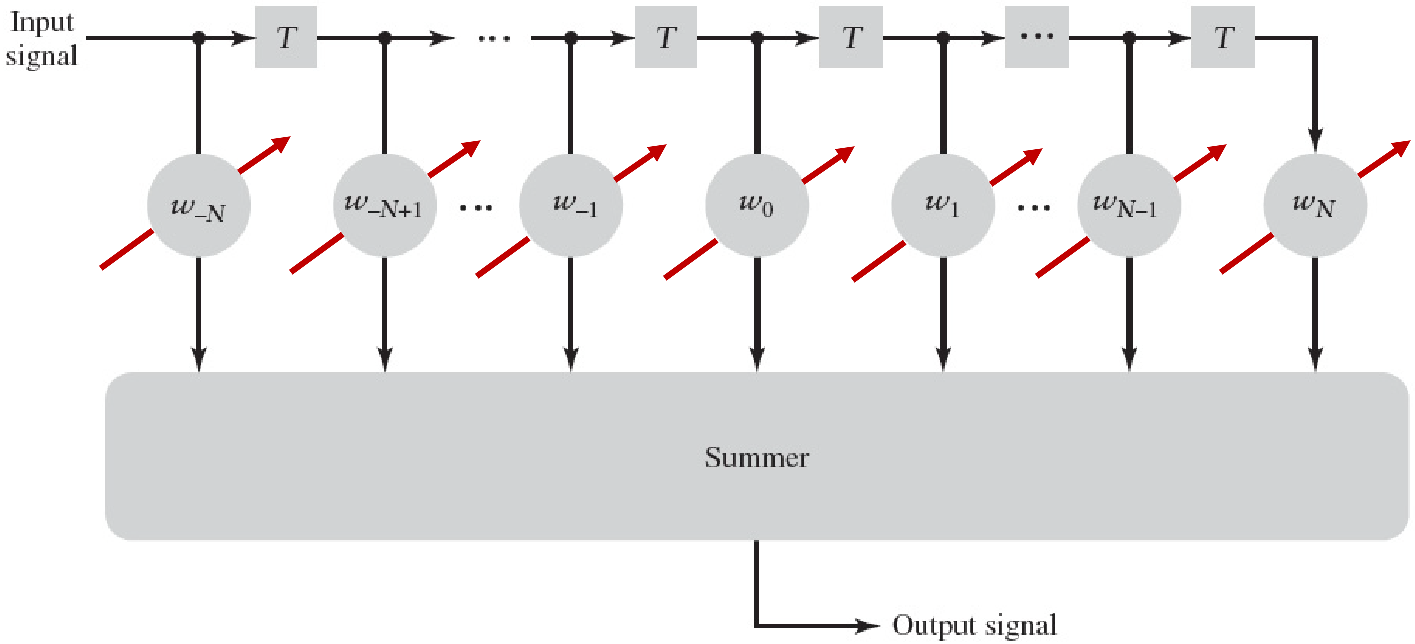

Popular model for one dimensional inter-symbol interference

A general f( ): RAM

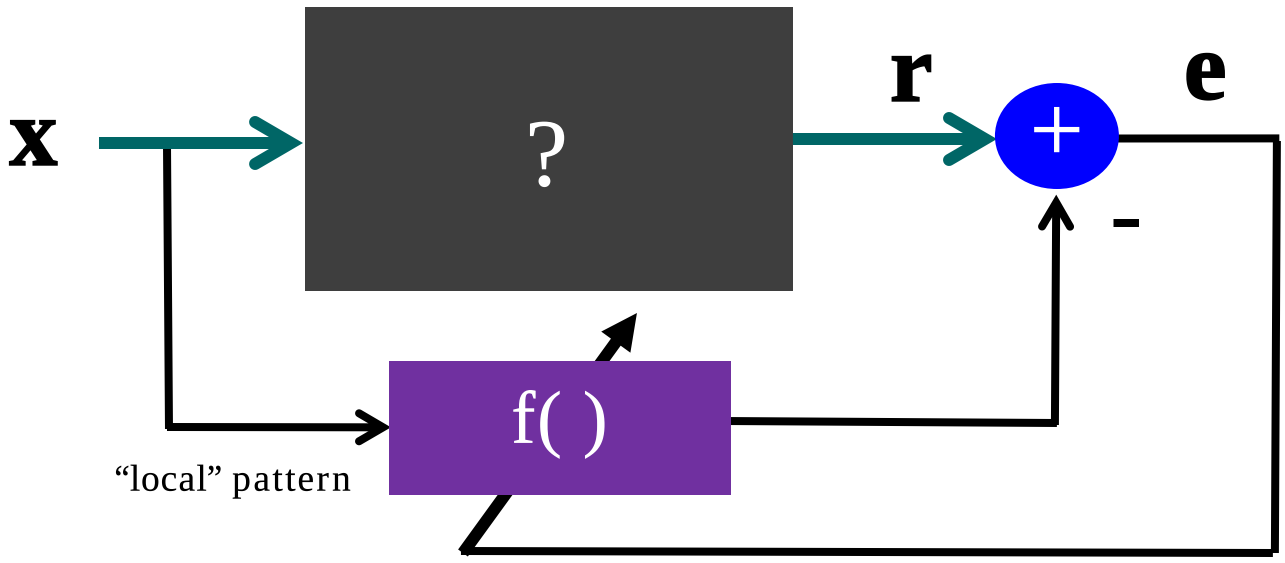

f( ) is a RAM and its contents get updated in the direction to minimize **e**. The write values of a cell and its affecting cells act as the pointer (address) to a particular location in the RAM. The RAM content is the read value of the victim cell.

After feeding the system and the model with a long data **x**, all locations of the RAM will have been updated and stabilized.

RAM Update Process: An Example of "Local" Pattern

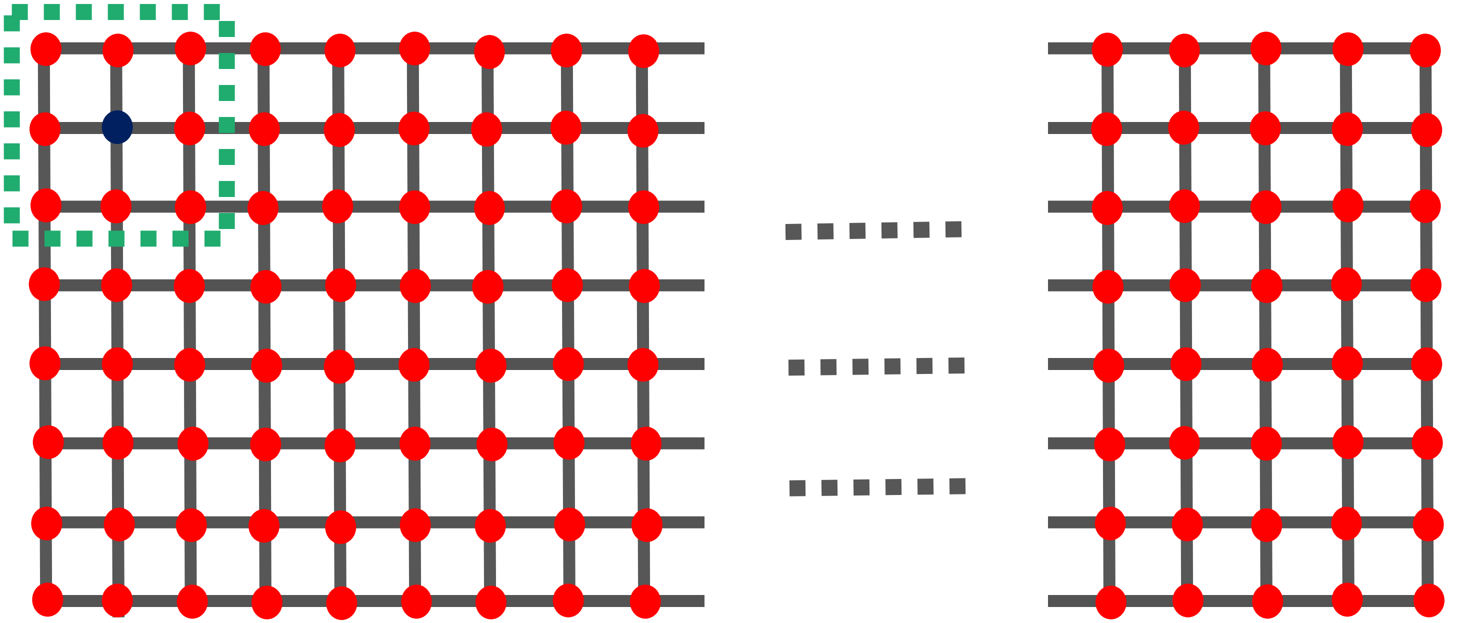

RAM Update Process: An Example of "Local" Pattern

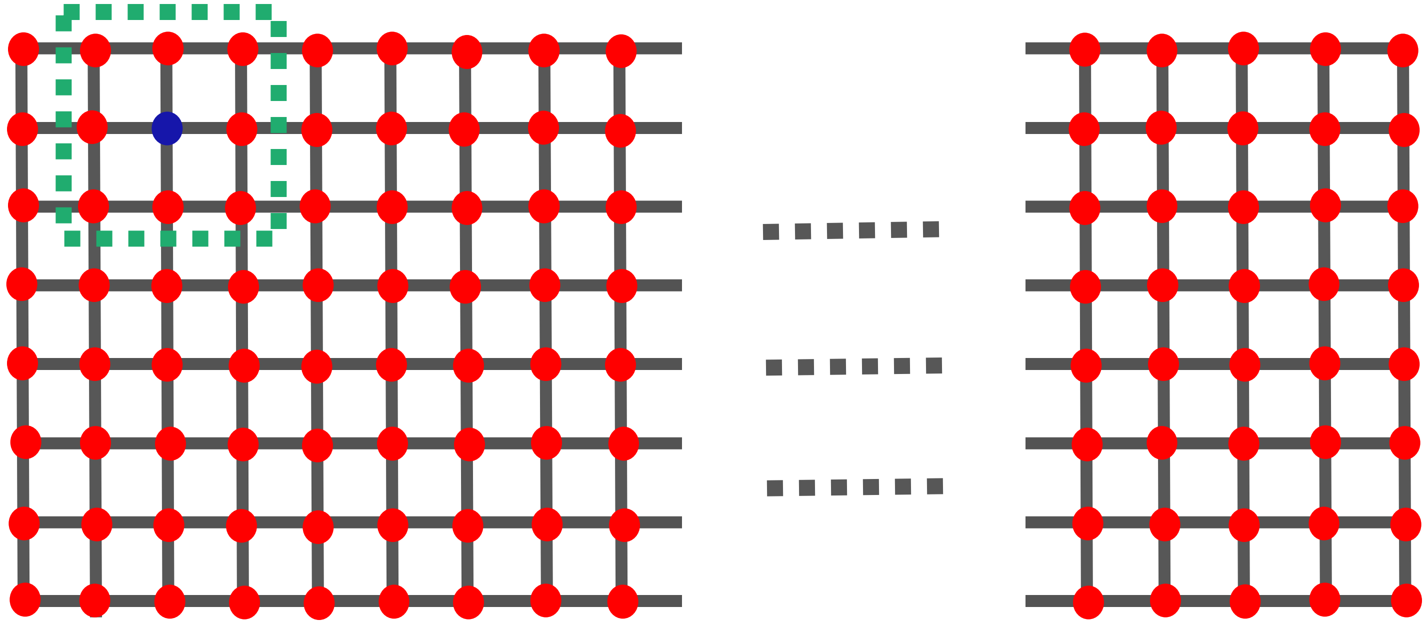

RAM Update Process: An Example of "Local" Pattern

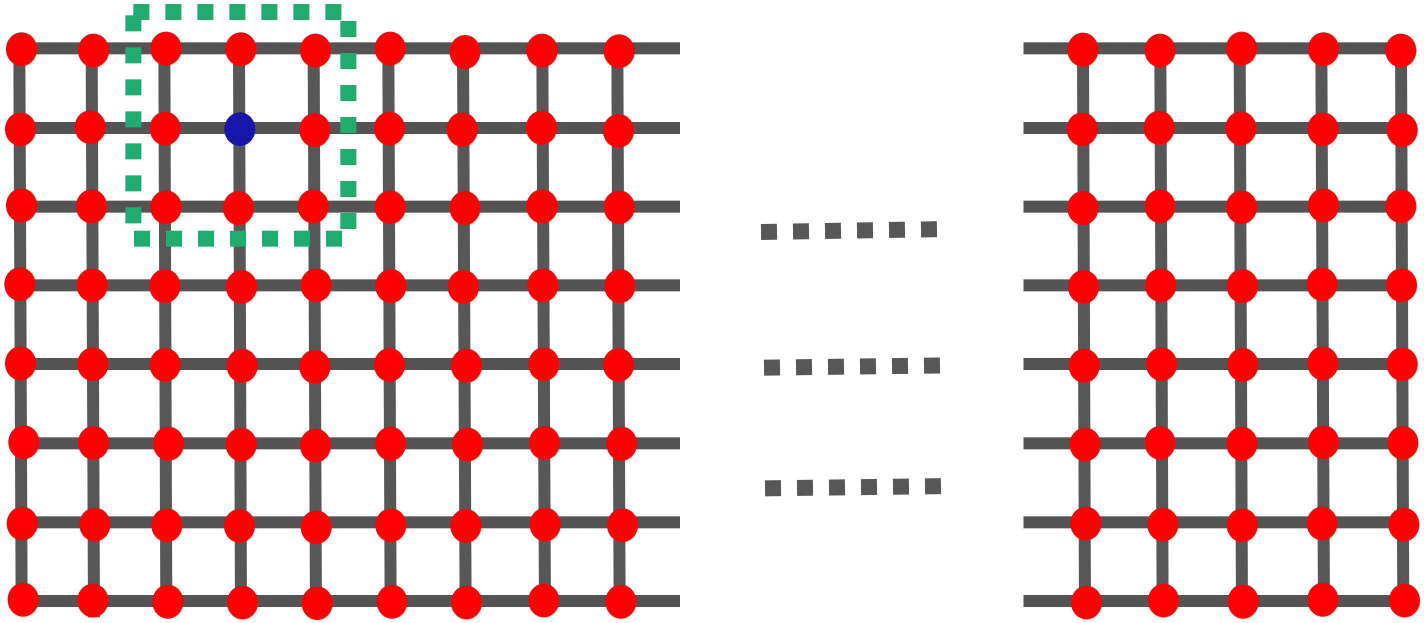

RAM Update Process

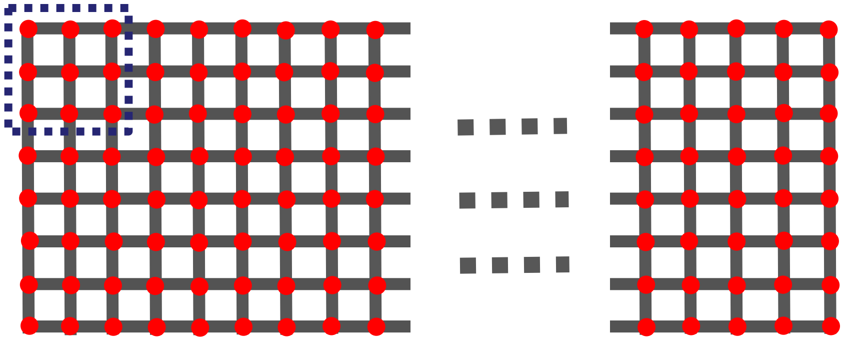

$$f\Big(\overbrace{\underbrace{x_{ij}}_{\text{victim cell}},\ \underbrace{x_{i,j-1}, x_{i-1,j}, x_{i,j+1}, x_{i+1,j}}_{\text{affecting cells}}}^{\textit{address}\ \text{(surrounding pattern)}}\Big) \approx r_{ij}$$

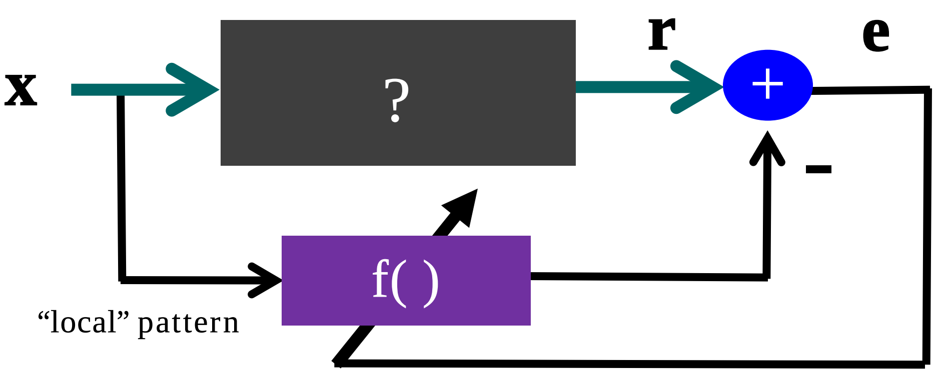

add1	add2	add3	add4
add5	add6		

$$f(adress1) \leftarrow f(adress1) + \mu \cdot r(address1) \cdot e(address1)$$

Do this for all addresses.

A general f( ): RAM

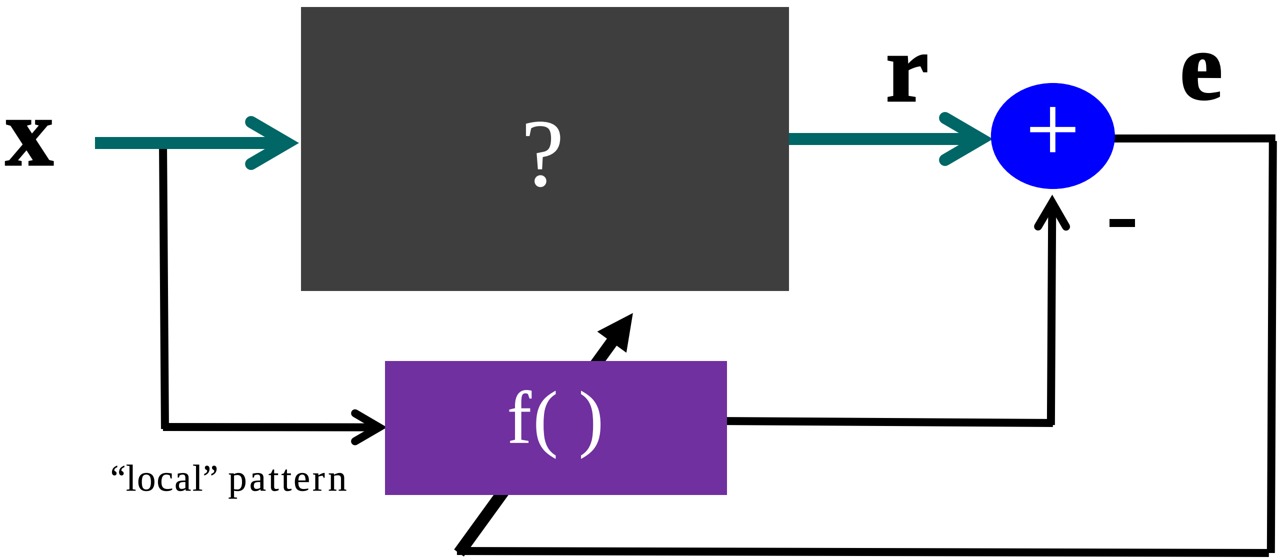

Assumption: correlation among a cell and its affecting cells is position-shift-invariant.

If x_k is affected by x_{k-1} and x_{k+1} in a certain way, then x_{k+1} will be affected by x_k and x_{k+2} in the same way.

Some of the discussion points …

- Are there such thing as the rough “local” pattern and can device people figure this out, if so?
  - Exact function can be obtained via statistical identification as described
- Feasibility of extracting soft information