Basic building blocks and architectures for realizable QCA devices

Michael Niemier

(with contributions from Amitabh Chaudhary, Danny Chen, Pranay Harsh, Sharon Hu, Peter Kogge, Craig Lent, Marya Lieberman, Wolfgang Porod, Ram Ravichandran, and Kevin Whitton)

Talk Outline

- Review basic constructs
  - Circuit constructs and clock
- Implementations
  - Molecular and Magnetic QCA
    - (systems with cells having only 1 orientation)
    - (systems with cells having 2 orientations)
- Basic building blocks for various implementations
  - …fundamental building blocks first…
- …and then architectures that use them…
  - …and also map well to QCA's device architecture
- Possible killer apps + what's next.

"Conceptual" QCA

A Device

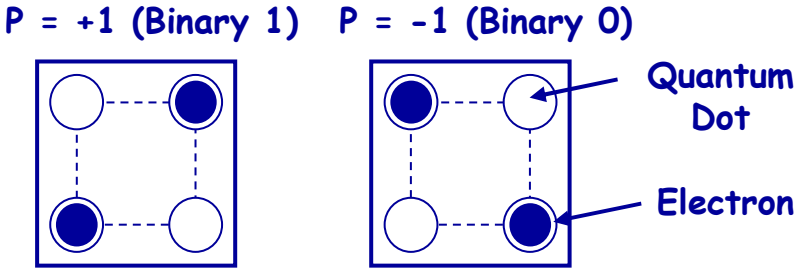

A Wire

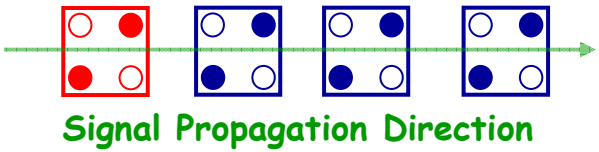

Majority Gate

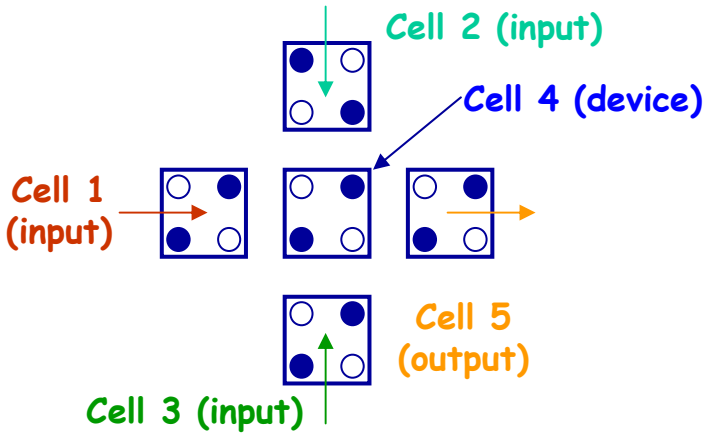

Wire Cross in the Plane

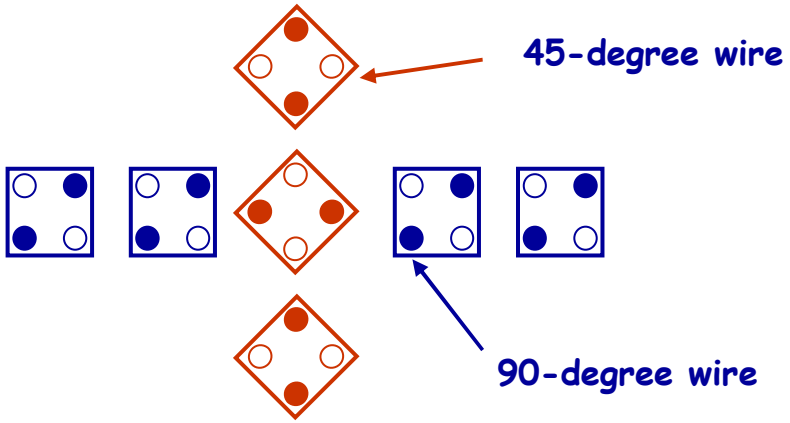

A 45-degree Wire

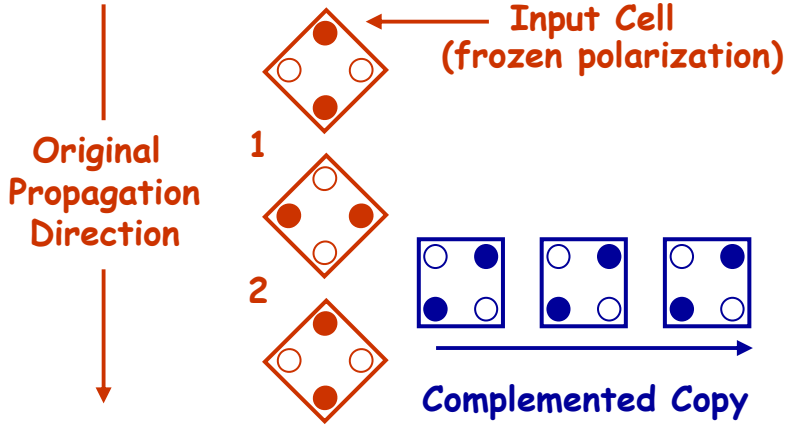

Clocked Molecular QCA

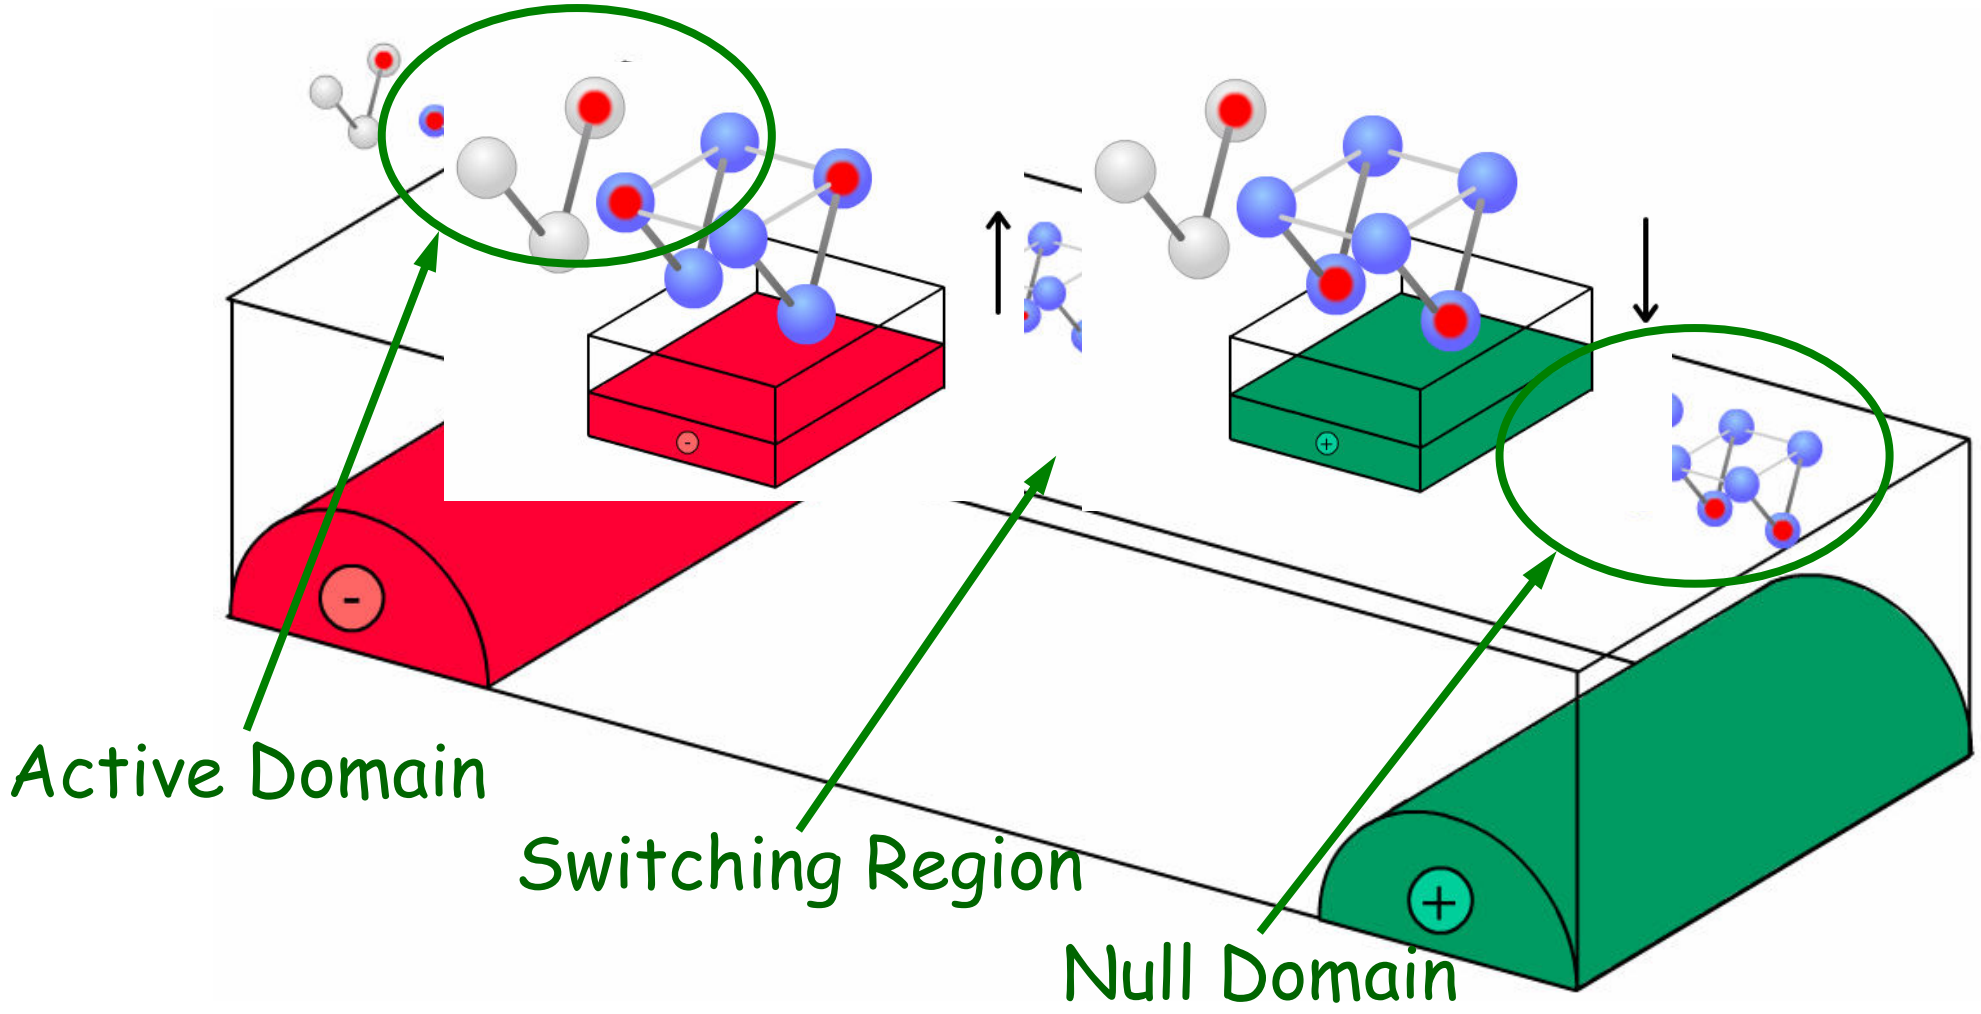

No current leads. No need to contact individual molecules.

Can use clock for I/O too…

How does a signal from "off chip" address an individual molecular QCA device that is approximately 1.2 *nm* x 1.2 *nm*?

Need a lithographic clock anyhow - use it to provide paths to permanent 0s and 1s.

T-Junction

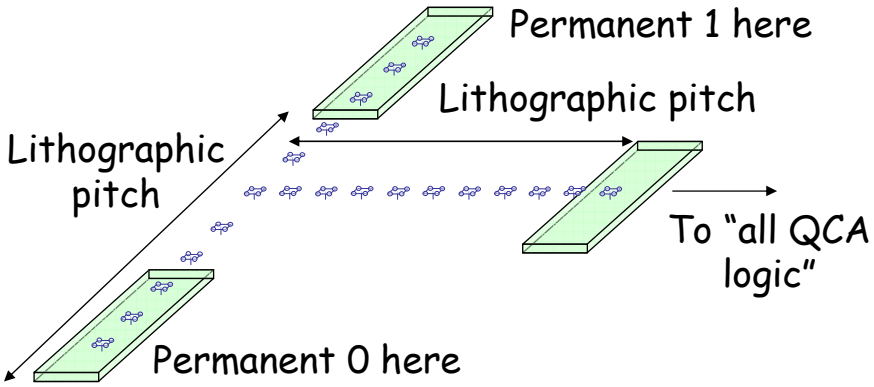

T-junction input mapped to 23 tile DNA raft

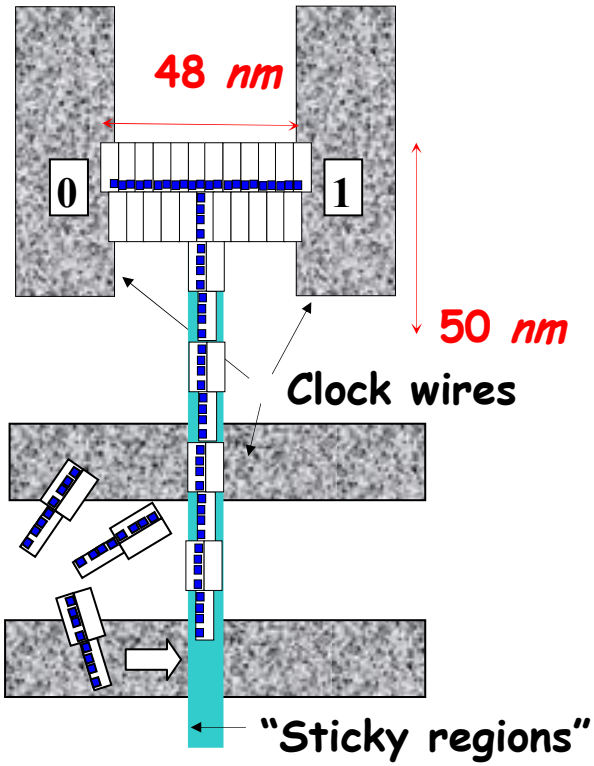

Implementations

- **Molecular**
  - See Craig Lent's talk…
- **Magnetic**
  - Bigger: 100s of nm [A,B]
  - Energy difference b/t 2 states ~ 100-200 k_bT [A,B]
    - (This is at room tempetature)
    - (Energy of $40k_bT$ needed to keep thermally induced errors < 1/year) [A]
  - Maximum dot dissipation ~ 10^{-17} J [A]
    - Microprocessor might dissipate ~ 1W [A]
  - Slower: ~100s of MHz for cross-chip frequency [A]
  - Could be integrated w/MRAM, insensitive to radiation [B]
  - Useful for space, military applications?

A: R.P. Cowburn and M.E. Welland, *Science*, Vol. 287, Issue 5457, 1466-1468, February 2000.

B: G.H. Bernstein et. al., *Microelectronics Journal*, 36 (2005) p. 619-624.

Magnetic QCA

Dipolar coupling in Co nanostructures

Single - domain

"Double" - domain

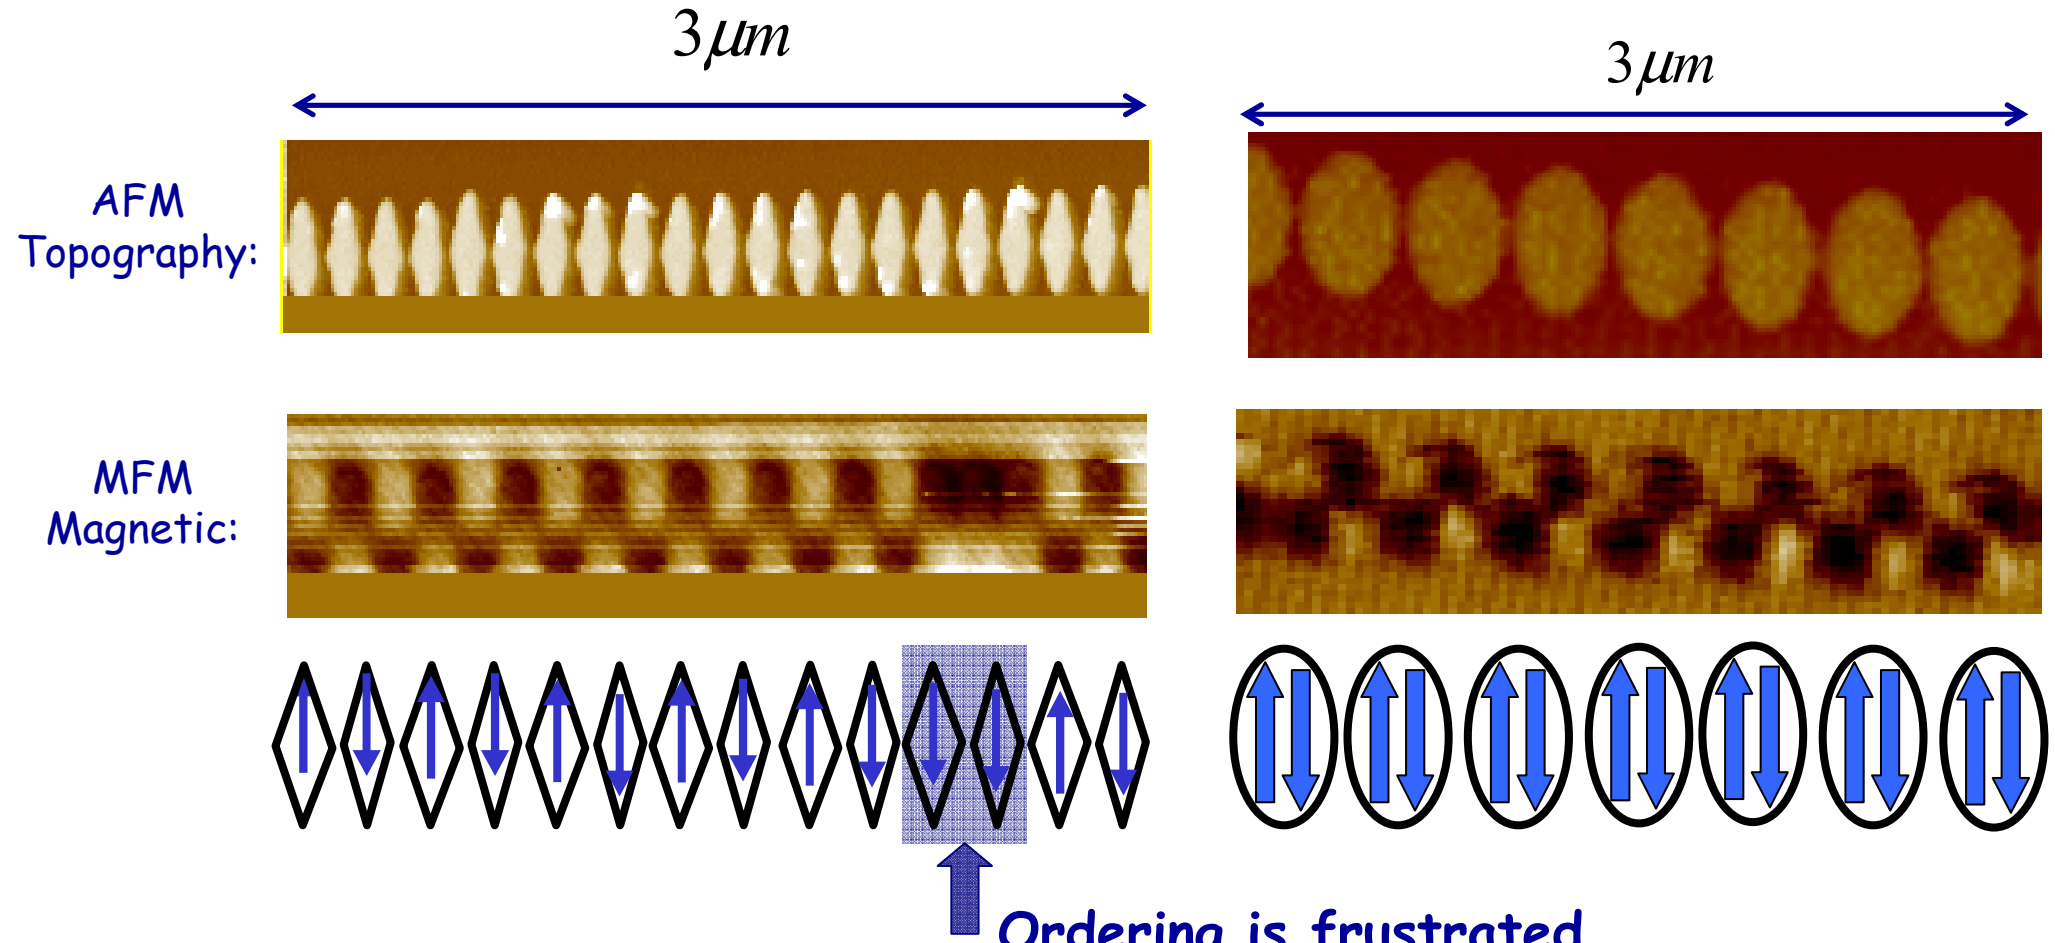

Coupling energy is potentially HUGE - I.e. 100kT (at 300K)

(note - our focus here mainly molecular, but basic building blocks + architectures should apply to both)

Molecular QCA - directed assembly

(not the only way -- but what I'll talk about…)

Idea: Integrate non-DNA components (devices + interconnect)

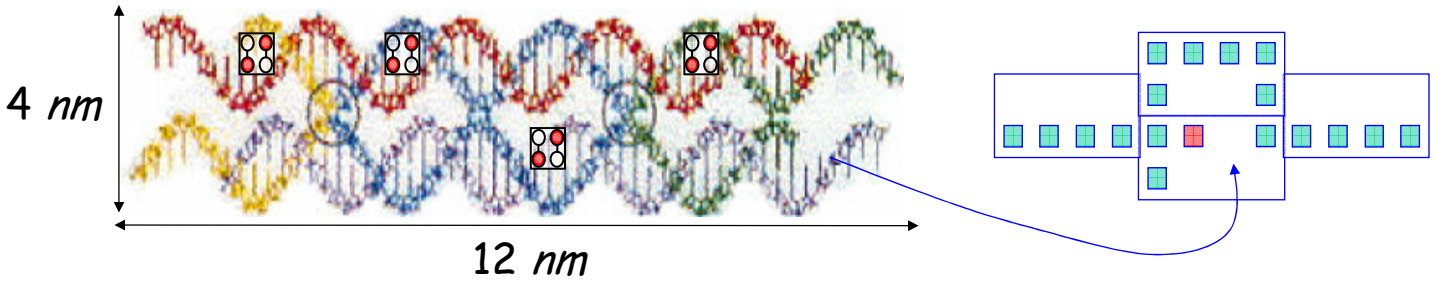

(Tiles have 8 accessible major groove sites)

Winfree and Seeman, 1998

(1) Assemble ~100 nm chunks of circuitry

No charge in wire means charge in cells not pushed up to active site; these cells are OFF

Current in wire pushes cell charge to active region; turn cells ON

(2) Further assembly directed by lithography

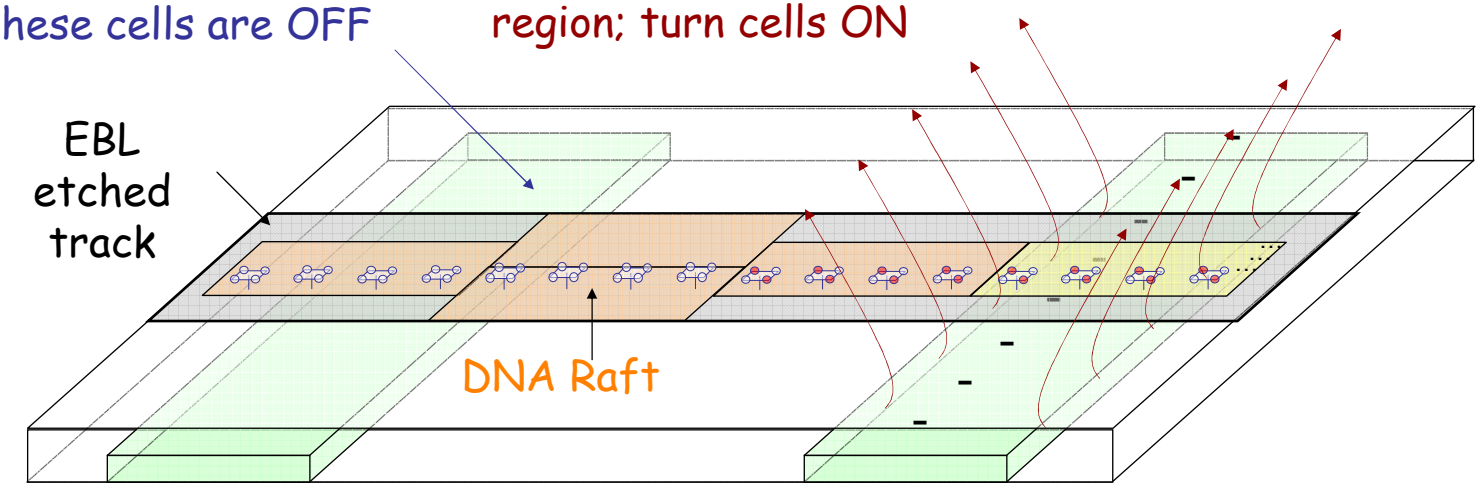

Experimental Liftoff of APTES/attachment of DNA rafts

From Marya Lieberman

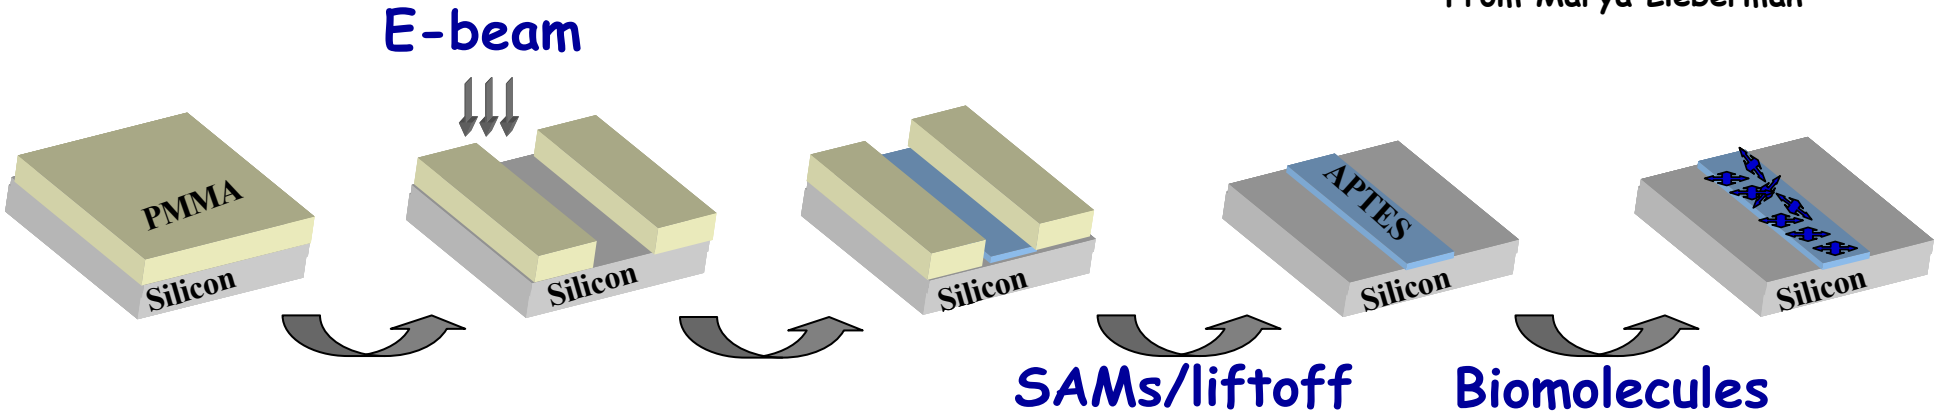

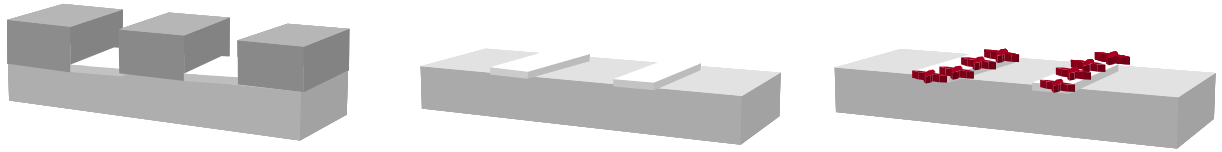

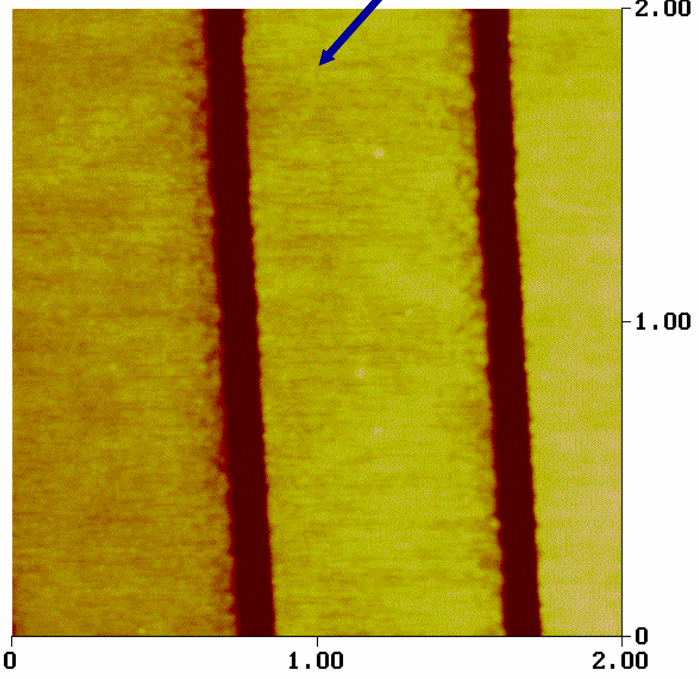

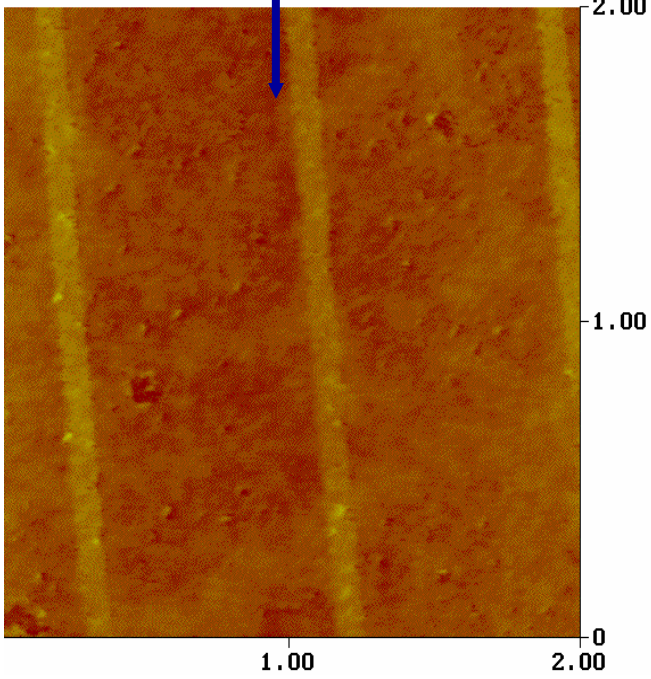

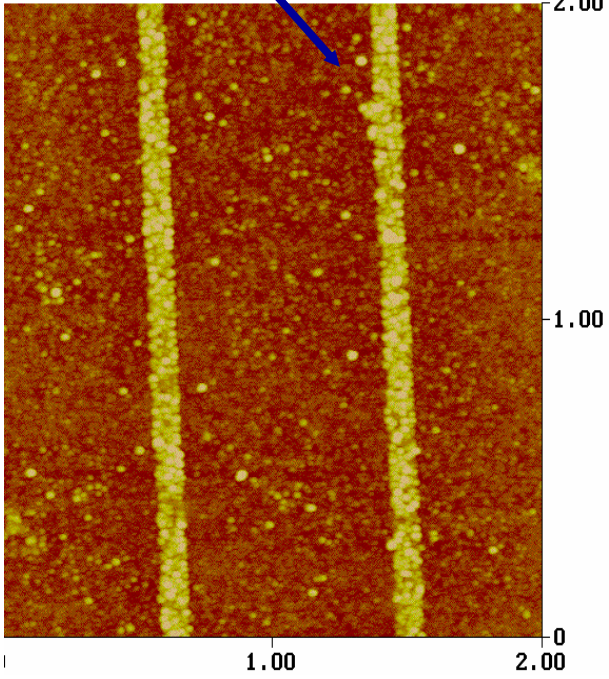

Cross-section views of rafts on EBL features

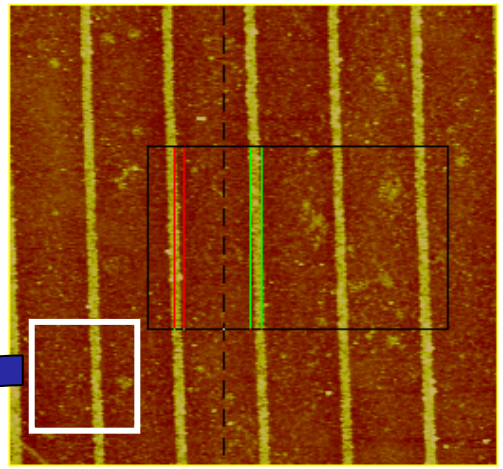

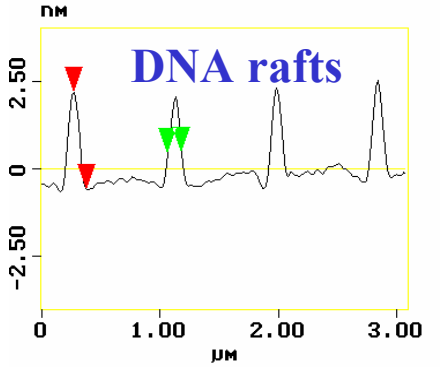

Average cross section

Line width 110 ± 10 nm

Height 2.7 nm

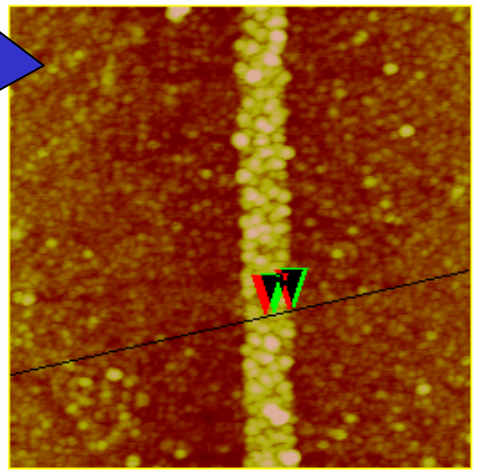

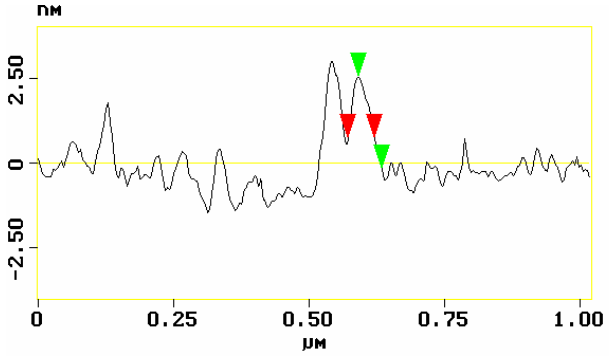

Line section

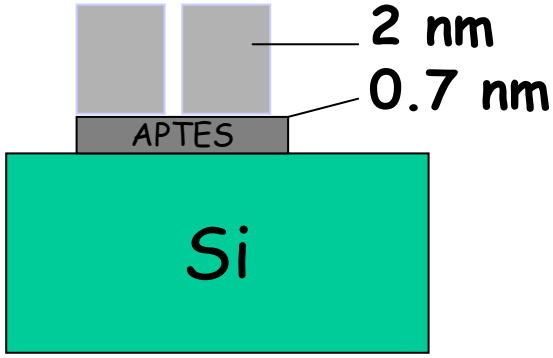

"Nanometer scale rafts bulit from DNA tiles," K. Sarveswaran, P. Huber, M. Lieberman, C. Russo, and C.S. Lent; *Proceedings of the 2003 3rd IEEE Conference on Nanotechnology,* 2003, p.417-20, vol. 2.

Jammin' on the surface

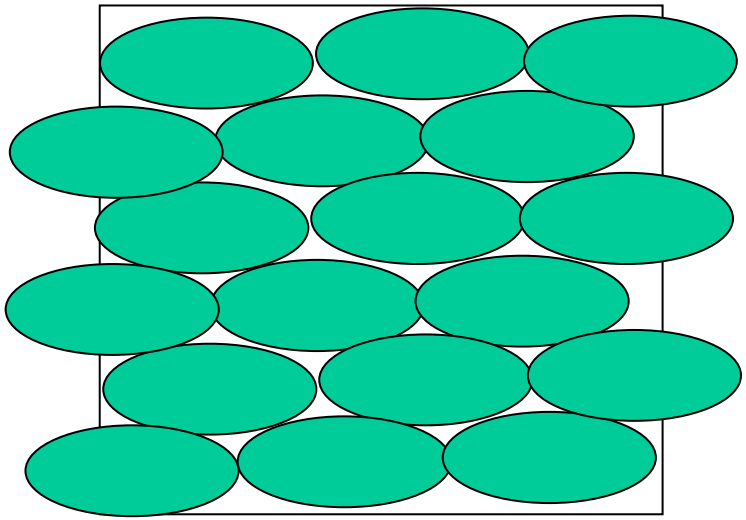

Θ =0.91
annealed

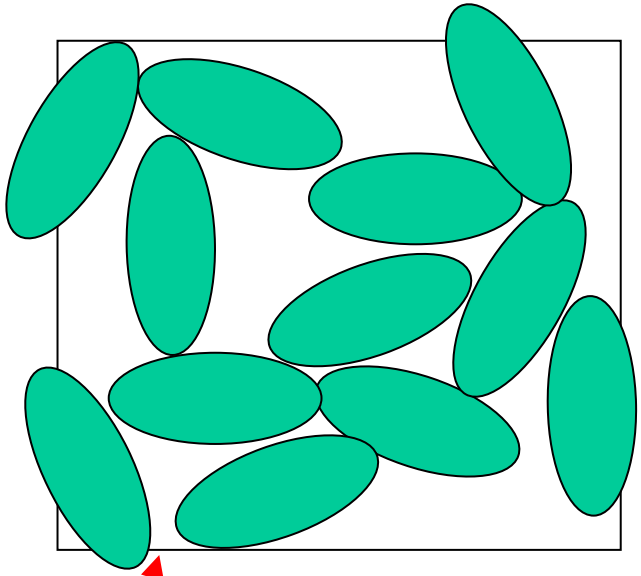

Θ =0.55 (α~4)
jammed

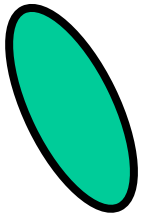

= "10 x 40 nm"
4-tile raft

Voit et al., JPC *97* 5212 1992

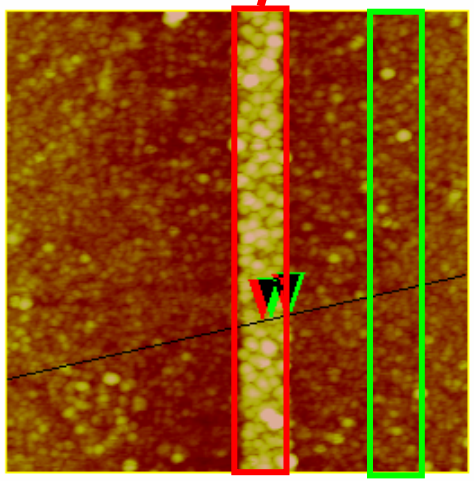

On APTES
Θ~0.40

On SiO_2:
Θ<0.02

Next: improve interactions between rafts:
(will actually involve CS theory…)

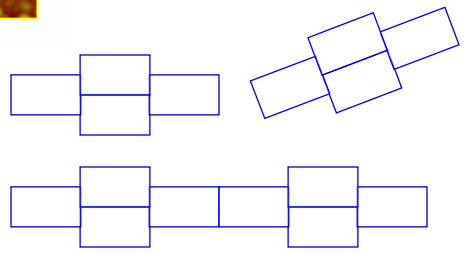

vs.

Molecular Systems - What's first?

Directed assembly…:

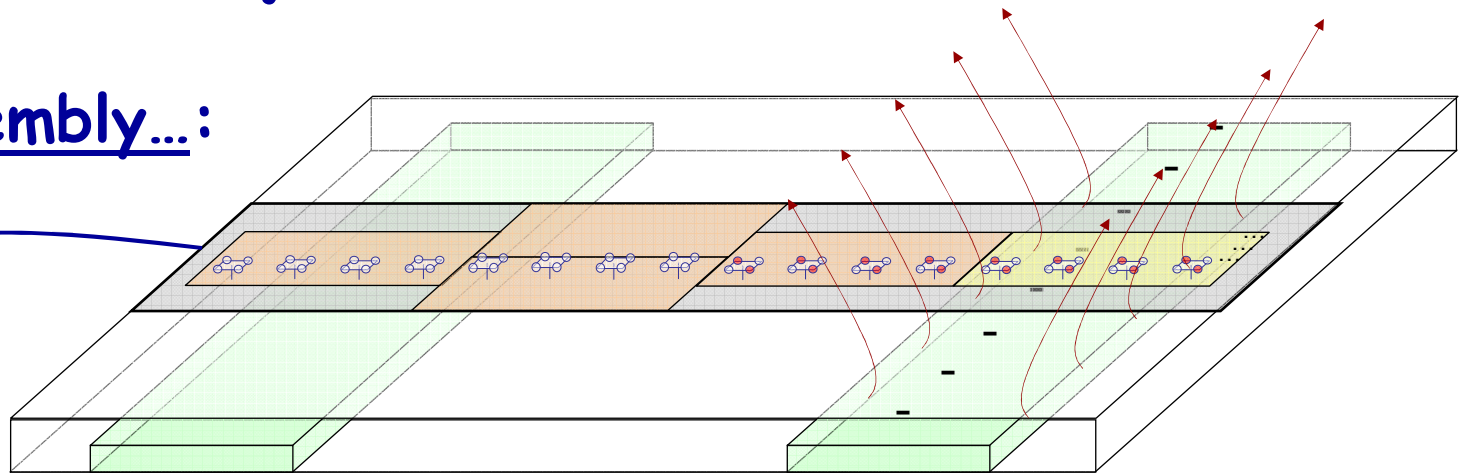

…but probably only one cell type on DNA raft…

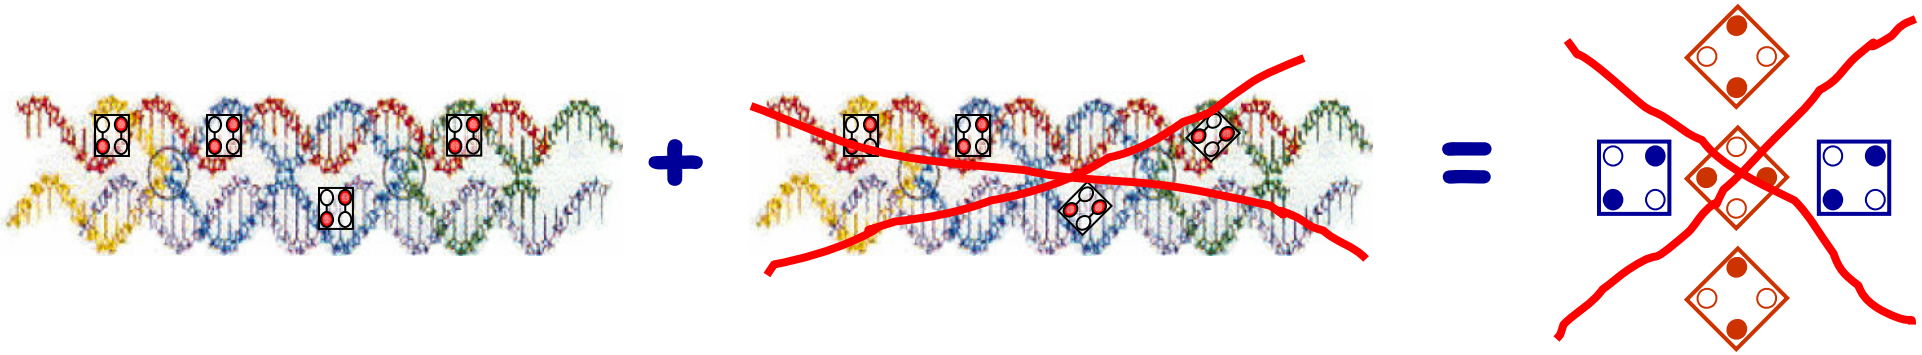

This first target is not even that restrictive…

Ways to cross wires…

Logical crossings

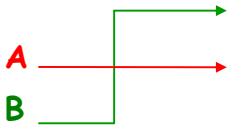

(a.)

A wire crossing

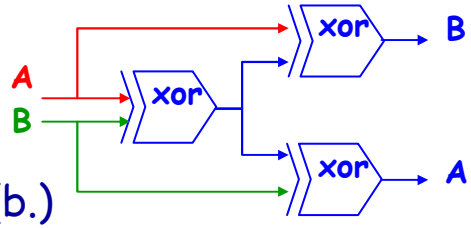

(b.)

A "logical" wire crossing

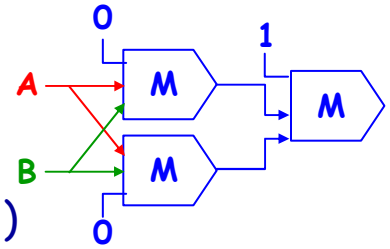

(c.)

XOR: (A and B') or (A' and B) (inherent crossing)

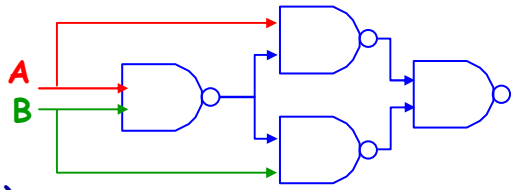

(d.)

Crossing can be made planar using NANDs

Duplication

- Make extra copies of logic to minimize crossings - especially if logic is so small…

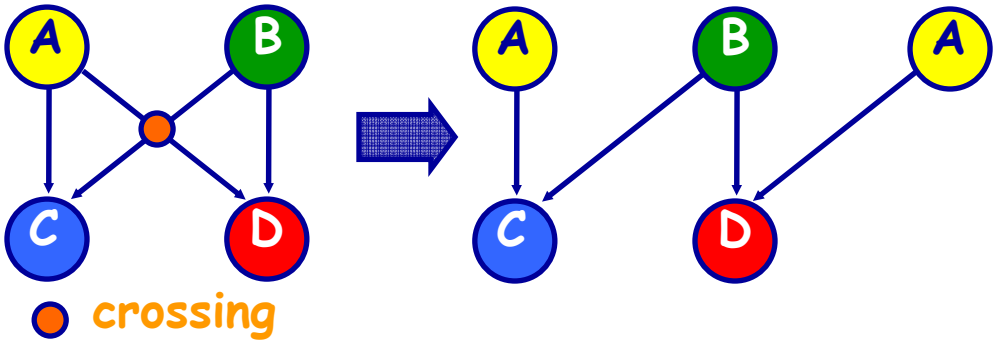

Time

- 2 signals share the same wire

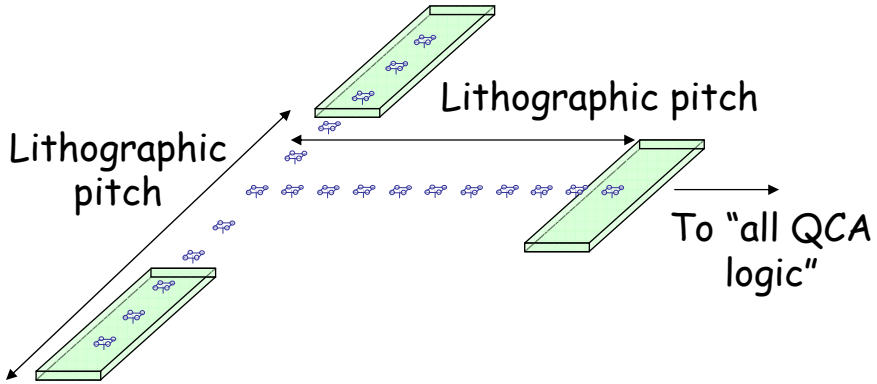

Logical Crossings

Statistical mechanics[A] tells us we need ~10-12 nm between parallel wires -- implies a 3 cell QCA pitch

NAND crossing can get (relatively) big…

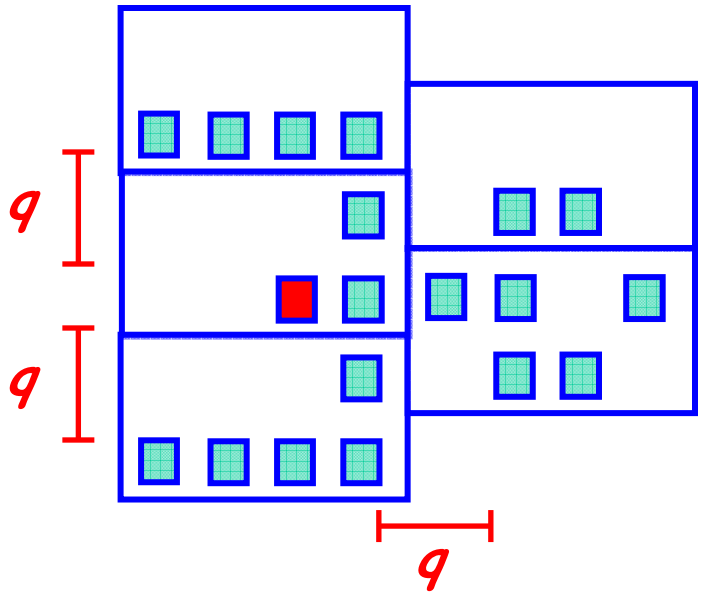

If *pitch* q increased to 3 (from 1), 2 more *tiles* required in *y* direction, 1 more in *x* direction - b/c of inverter

As 12 NANDs needed for logical X, this means at least 36 more tiles!

…but, can remap this logic…

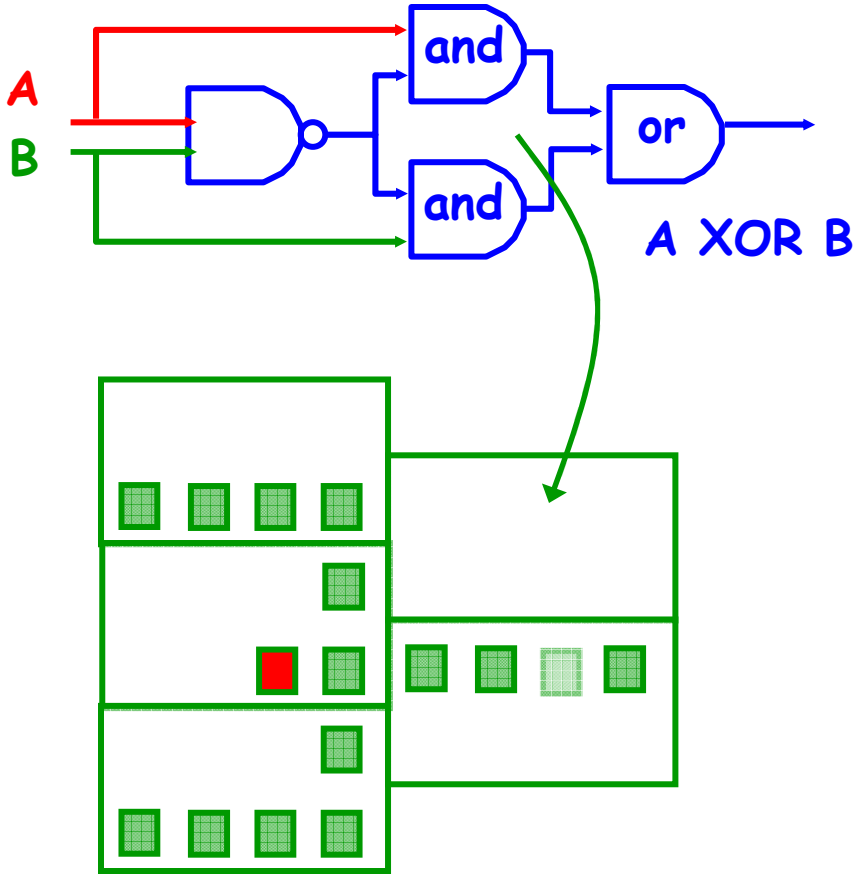

…to reduce area in *x* dimension

A: Based on : "Thermondynamic behavior of molecular QCA wires and logic devices", Lieberman and Wang, in IEEE T. Nano.

Logical Crossings

Design (all 3 cell pitch)	# of tiles in x	# of tiles in y	~ # of tiles	~ XOR area	~ area of crossing
NAND-based (1 cell thick)	8	15	120	5,760 nm^2	23,040 nm^2
Revised (1 cell thick)	4	9	36	1,728 nm^2	6,192 nm^2
Revised (2 cells thick)	7	13	91	4,368 nm^2	17,472 nm^2
Revised (3 cells thick)	8	17	136	6,528 nm^2	26,112 nm^2

What does this number *mean*?

[A] shows structures containing up to 200 correct tiles

What do *these* number mean?

[B] shows redundancy to defects

What about this number?

QuickTime™ and a
TIFF (LZW) decompressor
are needed to see this picture.

A: Rothemund PW, Papadakis N, Winfree E. , PLoS Biol. 2004 Dec;2(12):e424. Epub 2004 Dec 7

B: Enrique Blair, M.S. Thesis, 2003.

Logical Crossings

Some perspective on logical crossing area: Consider…

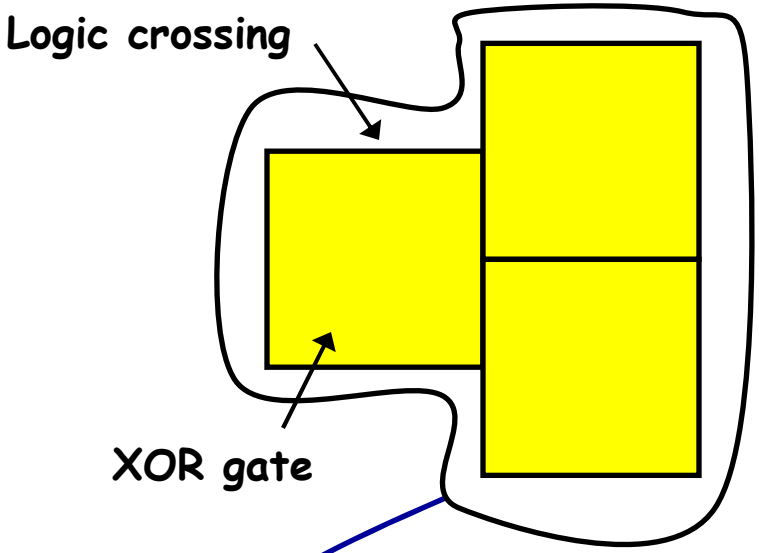

vs.

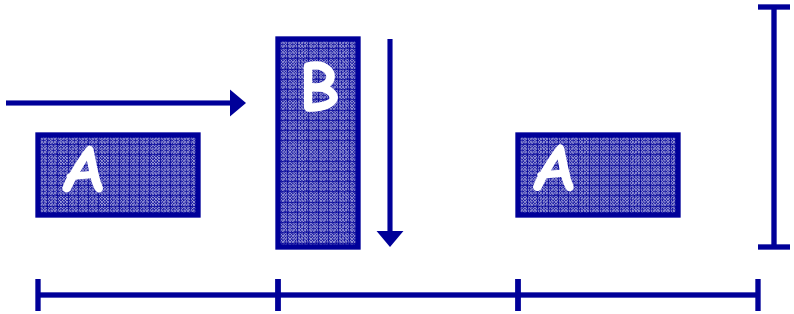

(Realistically, to cross 2 metal1 wires, must account for 3x1 metal1 "pitches")

This construct can be smaller than physical crossing @ 22 nm node.

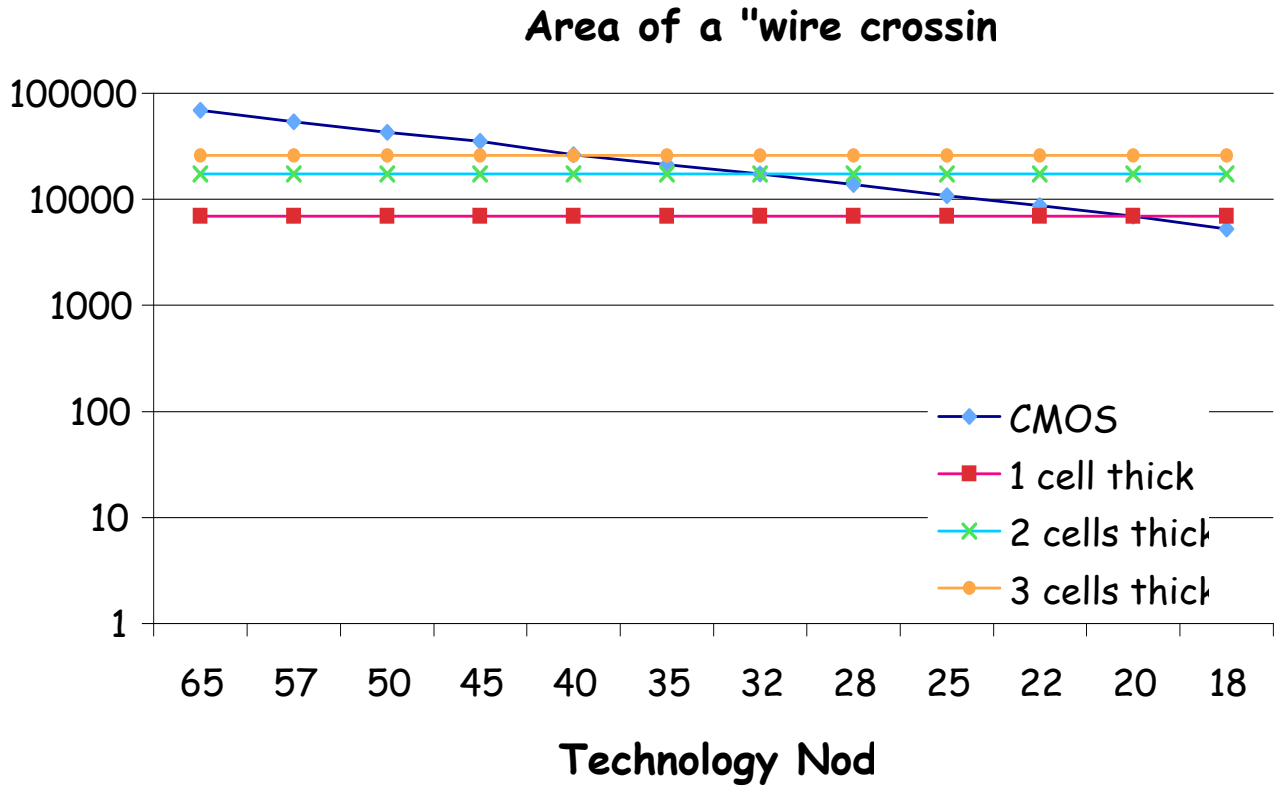

Calculated using metal 1 wiring pitch numbers from 2004 ITRS Interconnect Update.

An adder with fundamental blocks

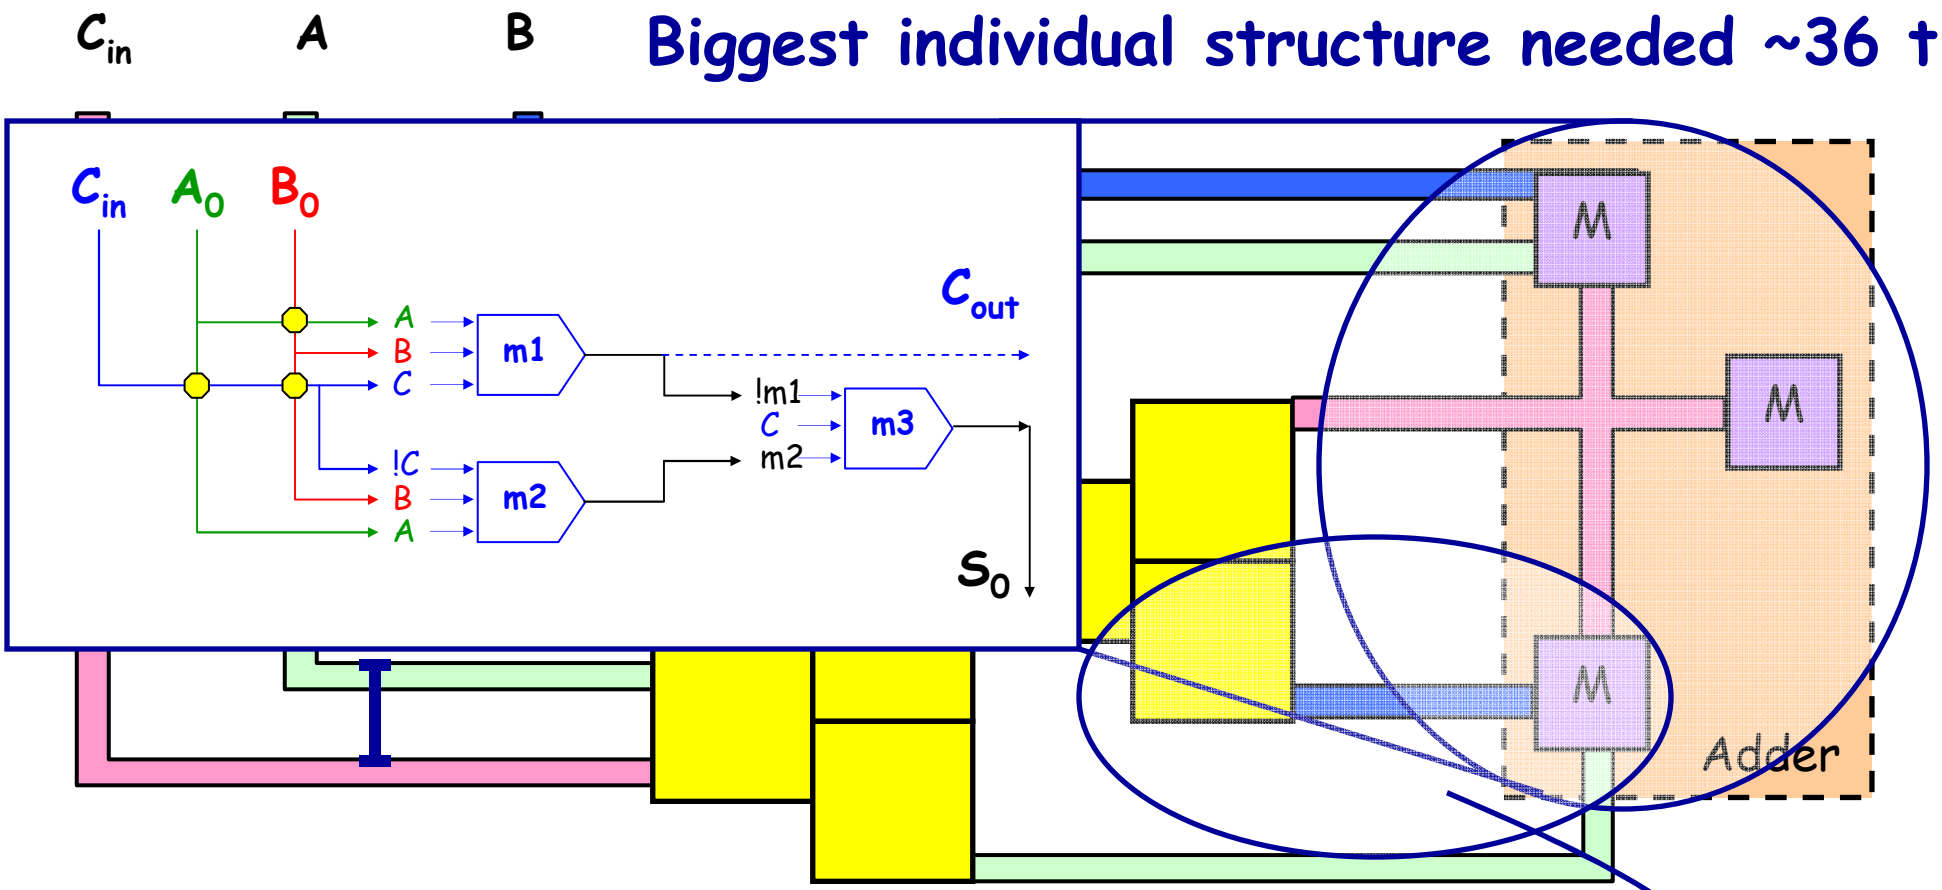

- Crossing area larger than EBL interconnect pitch matching
  - (I.e. these trenches depend on a pitch too…)
- Only adder itself might be a problem
  - Majority gates can be very small (5-7 DNA tiles)
- Solve by abutting rafts…

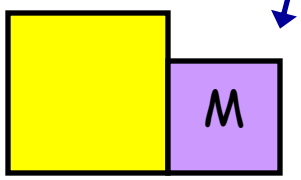

Logic Crossing - basic buildling blocks

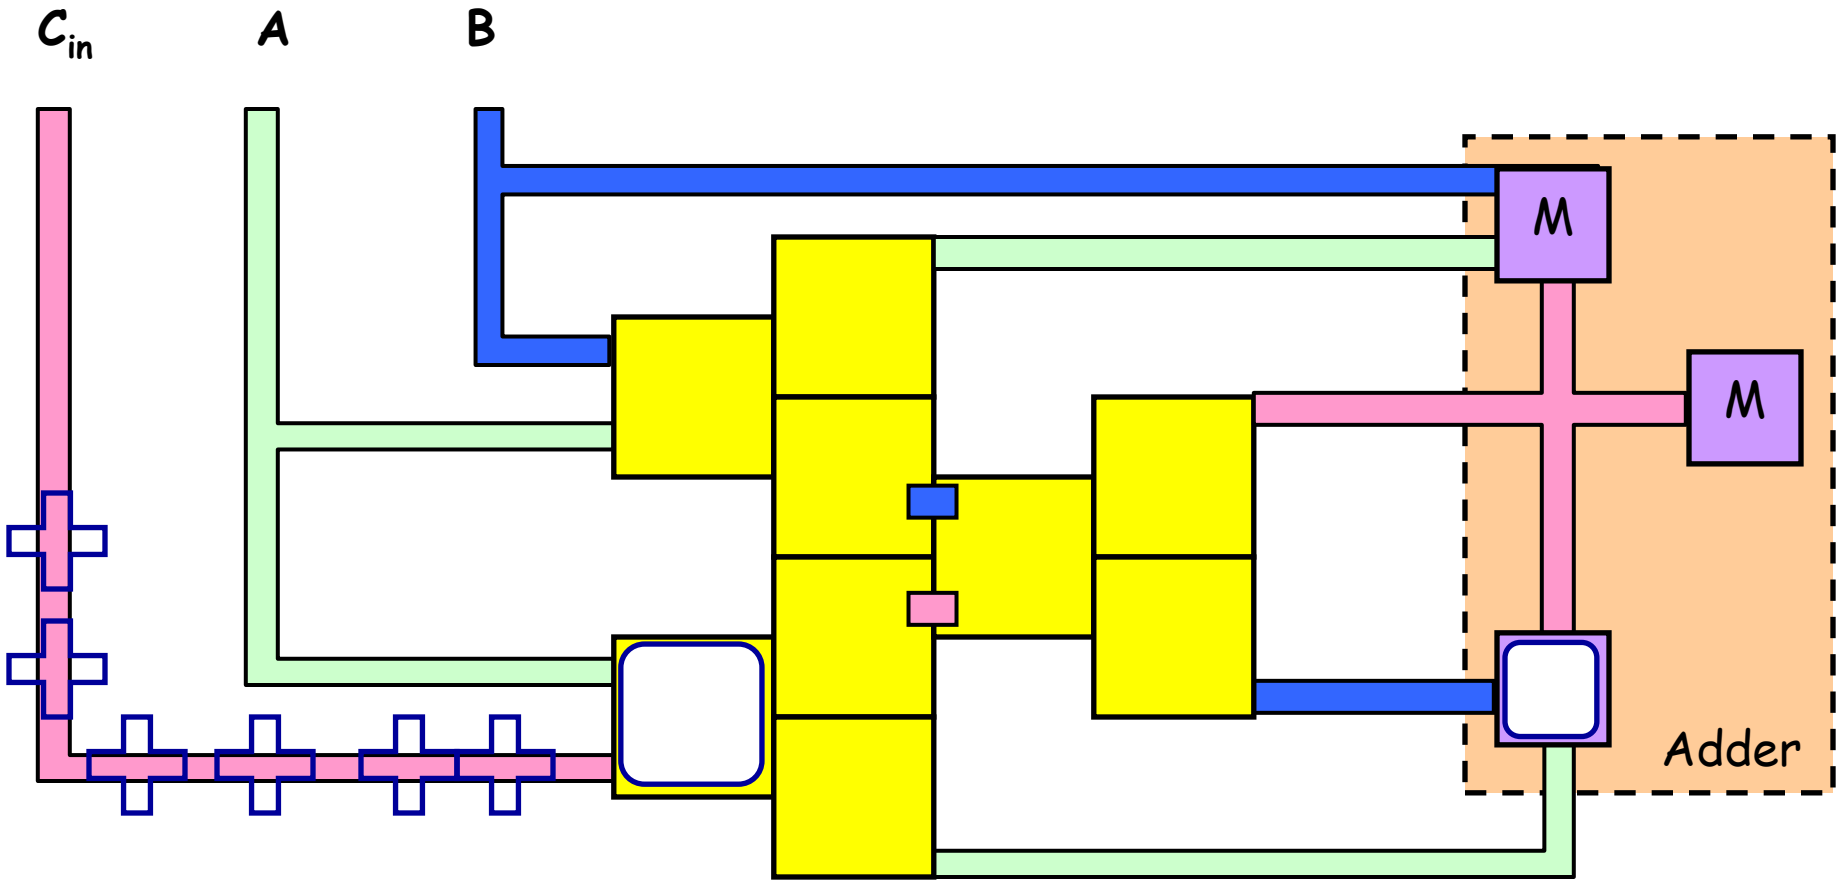

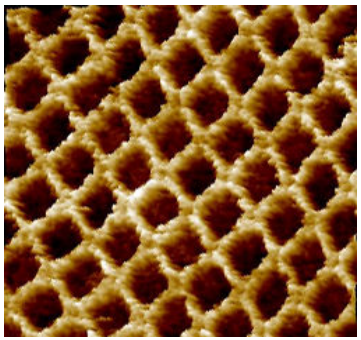

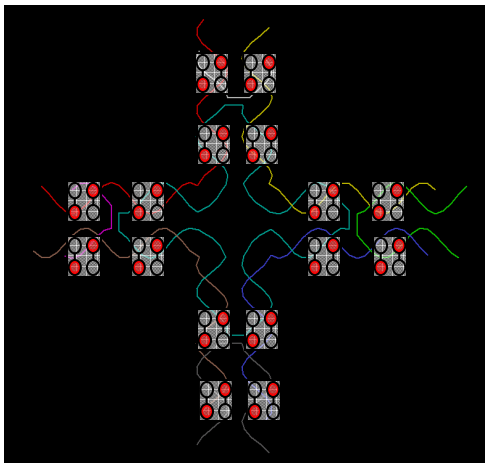

Idea: leverage Duke tile for wire…

- Can place QCA cells at all points…
- …have universal wiring tile…
- … simulated with stat. mech.

Goal: tiles self align in EBL trench

- "snap together" @ thermal equilibrium

John Reif, et. al -- Duke.

Duplication

- Idea: push (some) crosses to input[A]
  - Also let's us reuse some basic building/logic blocks
- Duplication…
  - …works in some cases (all IC local)…
    - I.e. in ISCAS benchmark, area *decreases* as all IC local
  - …not in others
    - Problem inherently can explode exponentially
- Doing this in select cases works…

Adder example revisited
(Logic small so multiple copies of some blocks OK…)

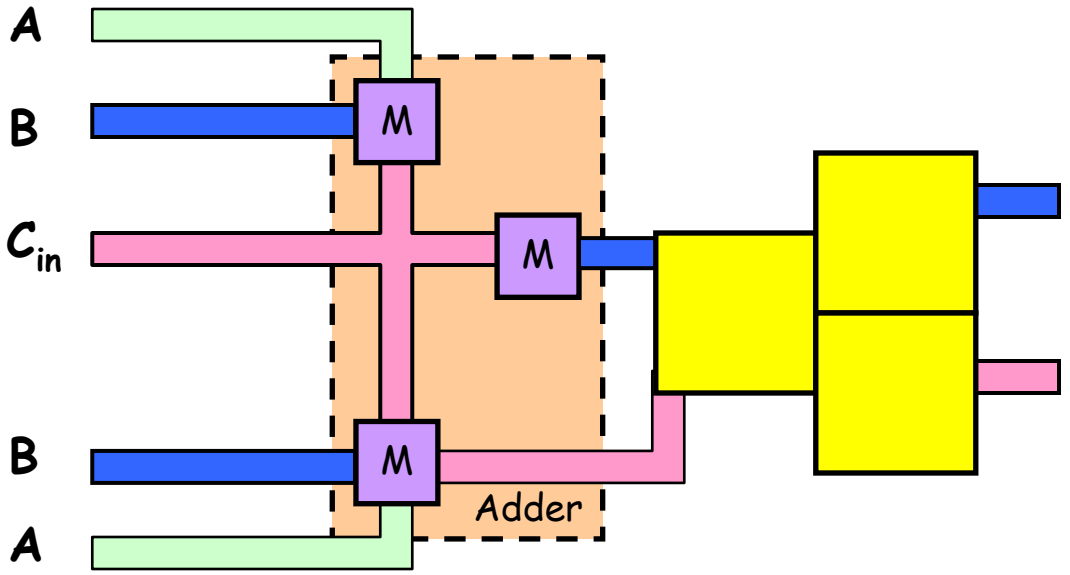

A: Chaudhary, Chen, Hu, Niemier, Ravichandran, Whitton -- to appear at ICCAD, 2005, San Jose, CA, Nov. 5-9.

Number of devices/cm^2 ?

Use previous info./designs for back of envelope calculation: how many QCA devices might be in 1 cm^2... Assume:

- Reasonable EBL pitch
- Parts mapped to DNA
- 3 cell QCA pitch
- Wiring tile (16 cells)
- Some redundancy

Note: Throughput dependent on granularity of clock

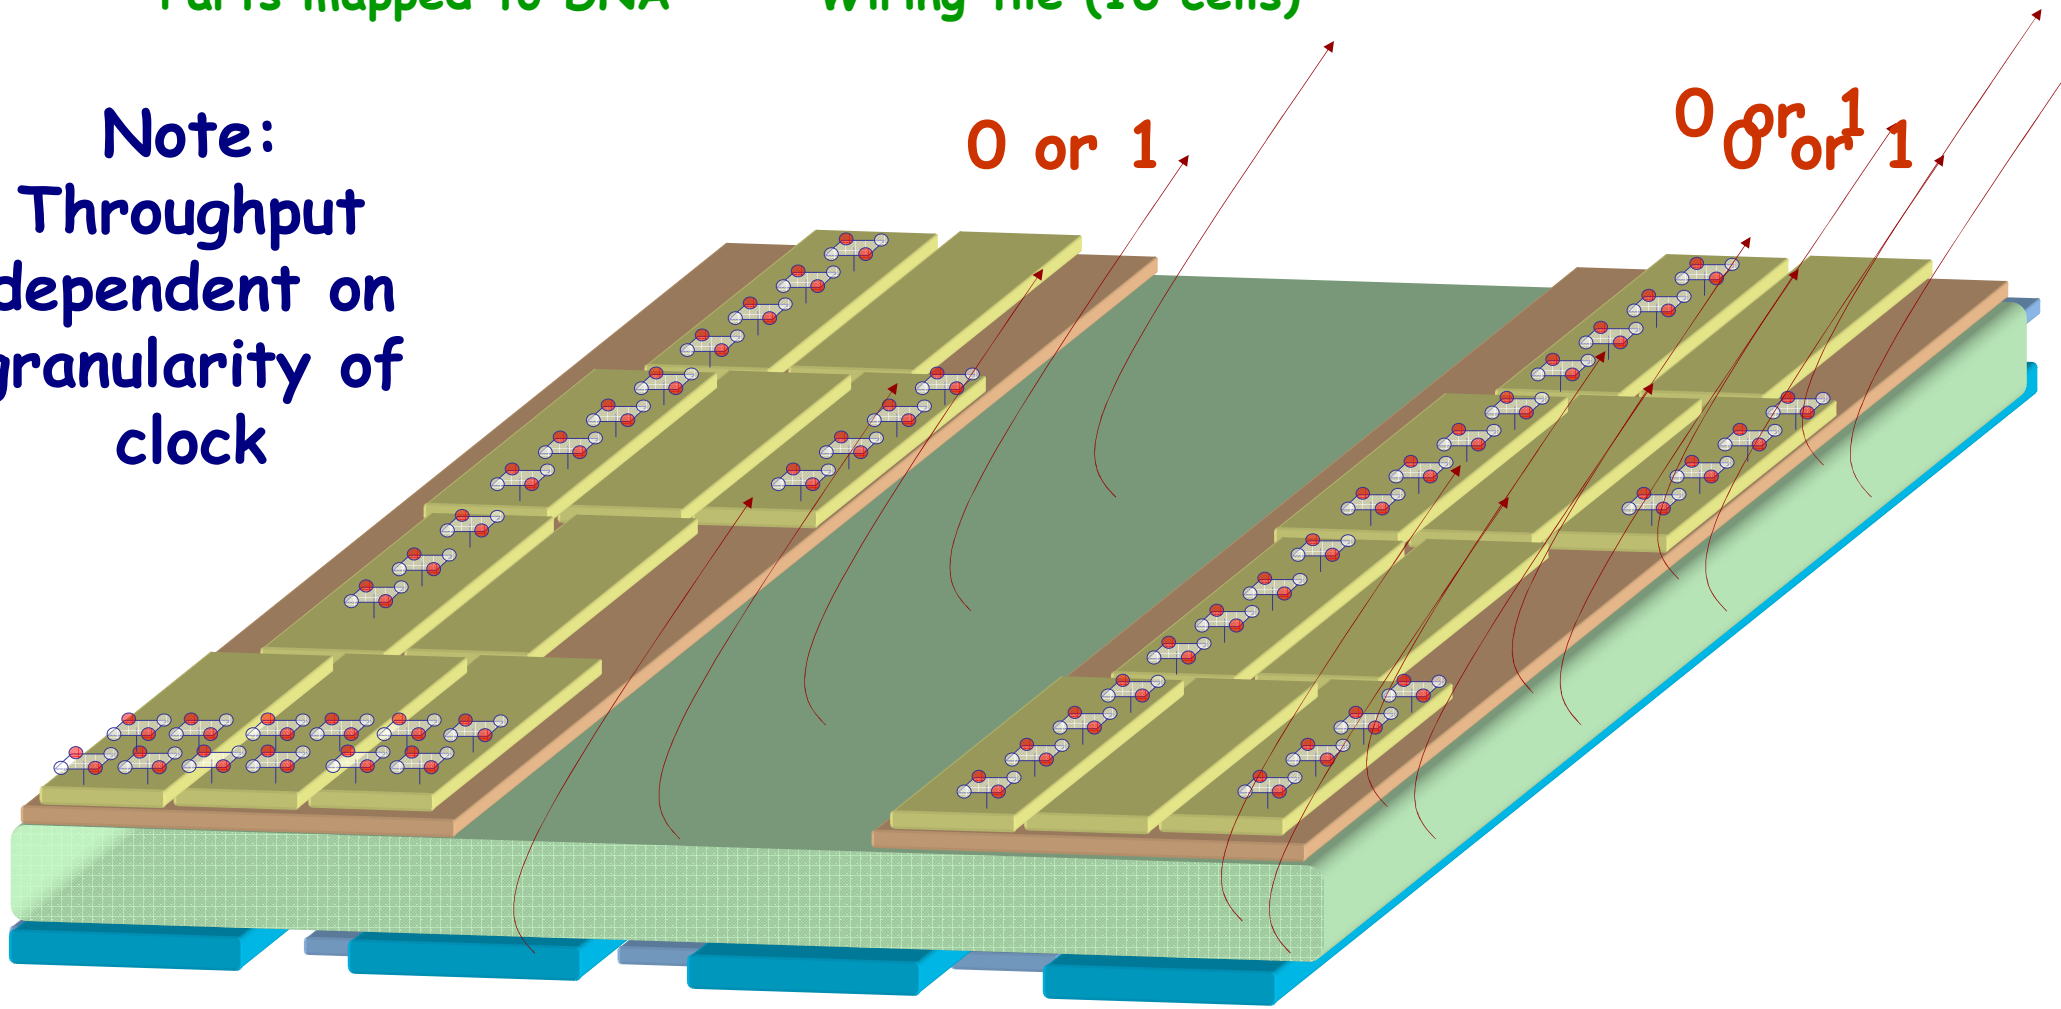

Number of devices/cm^2 ?

What do we get?

Design	devices/ bit	Area (cm^2)	~devices/ cm^2	% of *logical* devices
Adder with logical Xs	1750	8.5×10^{-10}	1.50×10^{12}	7%
Adder (duplication)	400	3.3×10^{-10}	1.20×10^{12}	22%
Adder - theoretical constructs	160	1.1×10^{-10}	1.47×10^{12}	35%

Huh?

- Seemingly doesn't make sense…
- EBL for adder with logical crossings masks some wiring overhead
- + need to consider how many devices are logic vs. IC…

realistically higher - adder leverages majority gate function…

Also, must consider that this leverages traditional architecture/adder design + in QCA wires are made of devices

Architectures

- EVERYTHING is pipelined
  - In the past, instruction execution was pipelined

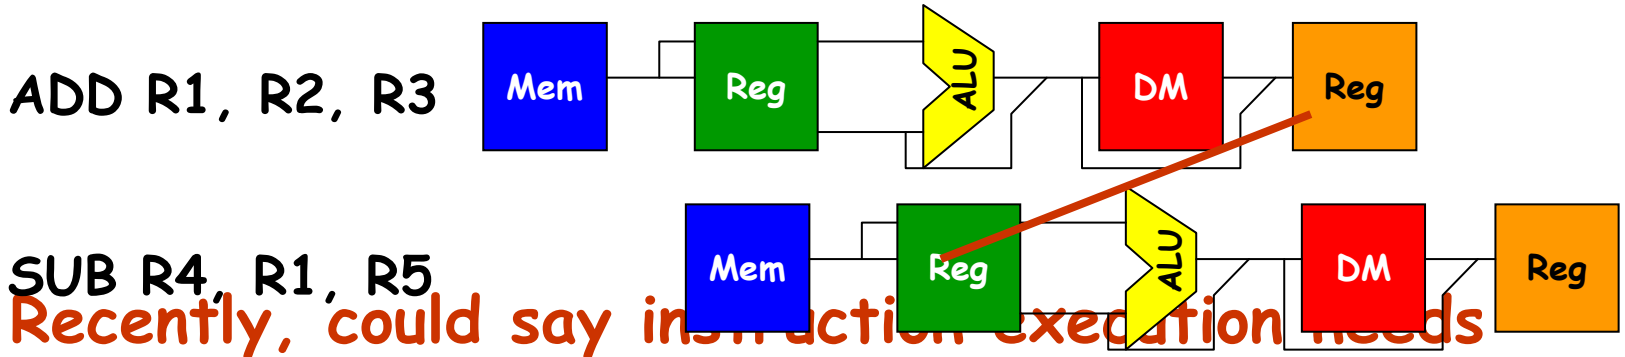

  - Recently, could say instruction execution needs pipelined wire[A]

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
TC	NXTIP	TC	Fetch	Drive	Alloc	Rename	Rename	Que	Sch	Sch	Sch	Disp	Disp	RF	RF	Ex	Figs	Br Ck	Drive

  - More recently, global interconnect pipelined[B]
  - Pipe depth depends in part on granularity of clock wires

A: Hinton, et. al. "The Microarchitecture of the Pentium 4 Processor," Intel Technology Journal, Q1, 2001, p. 1-12.

B: I.e J. Cong, Y. Fan, Z. Zhang, "Architectural-Level Synthesis for Automatic Interconnect Pipelining", DAC 2004, June 7-11

Architectures (cont.)

- Data can be/is latched on wires
  - Good and bad:
    - Lends itself to high throughput (example soon)…
    - …but medium + global IC can be difficult
  - Forwarding difficult at best…

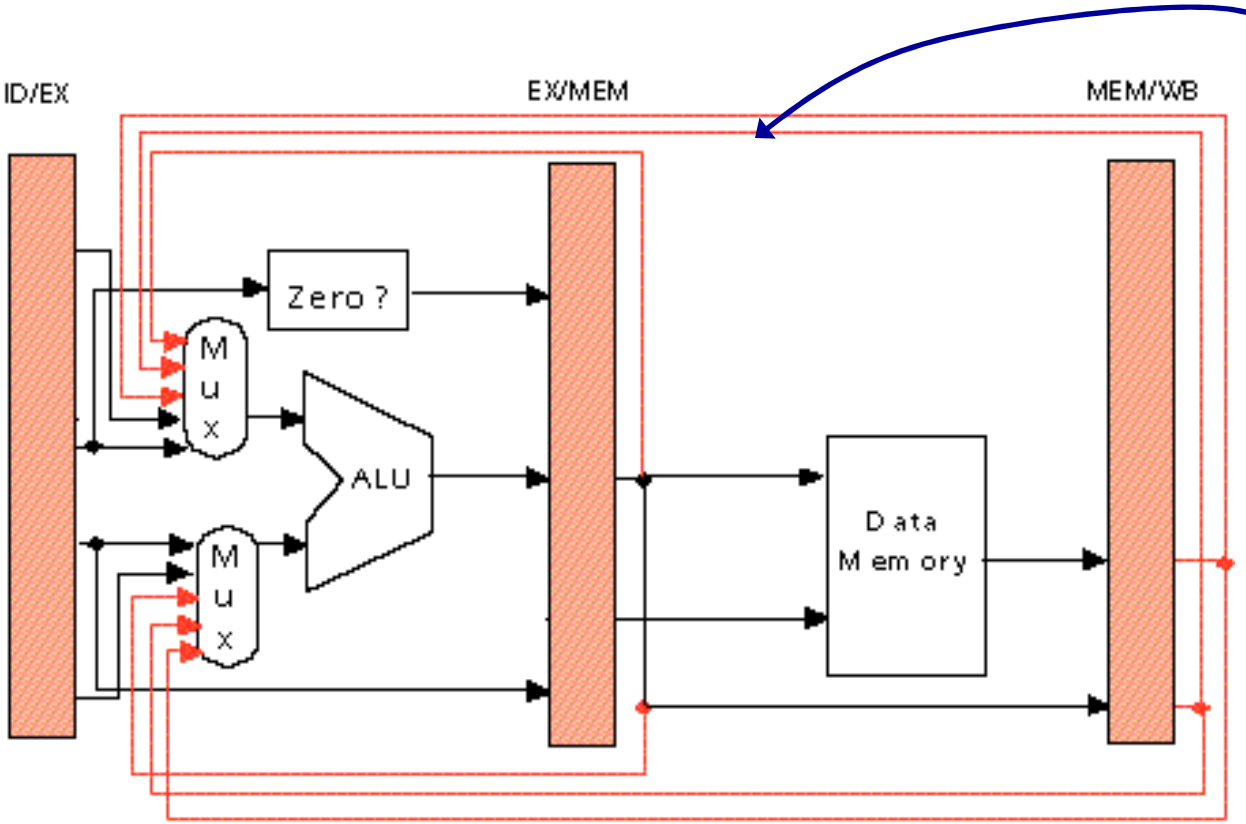

Wire would be driven by clock in opposite direction.

There is a latency in data flow.

Architectures (cont.)

- Defects
  - Must consider when computing at the nano-scale…
    - …especially anything that is self-assembled
  - Simple, regular, and replicable offers some protection
    - I.e. broken wire or missing tile or defective tile
- We'll discuss:
  - PLAs, reconfigurable, systolic, and counterflow
    - PLA
      - again, seemingly "simple" clock + some inherent redundancy
      - *NOT* best architecture for QCA - but illustrates what might work quite well…
    - Systolic and counterflow seem to map well…
      - …no global IC + potential for simpler clock structures…
    - What else?

Example PLA design (AND plane)

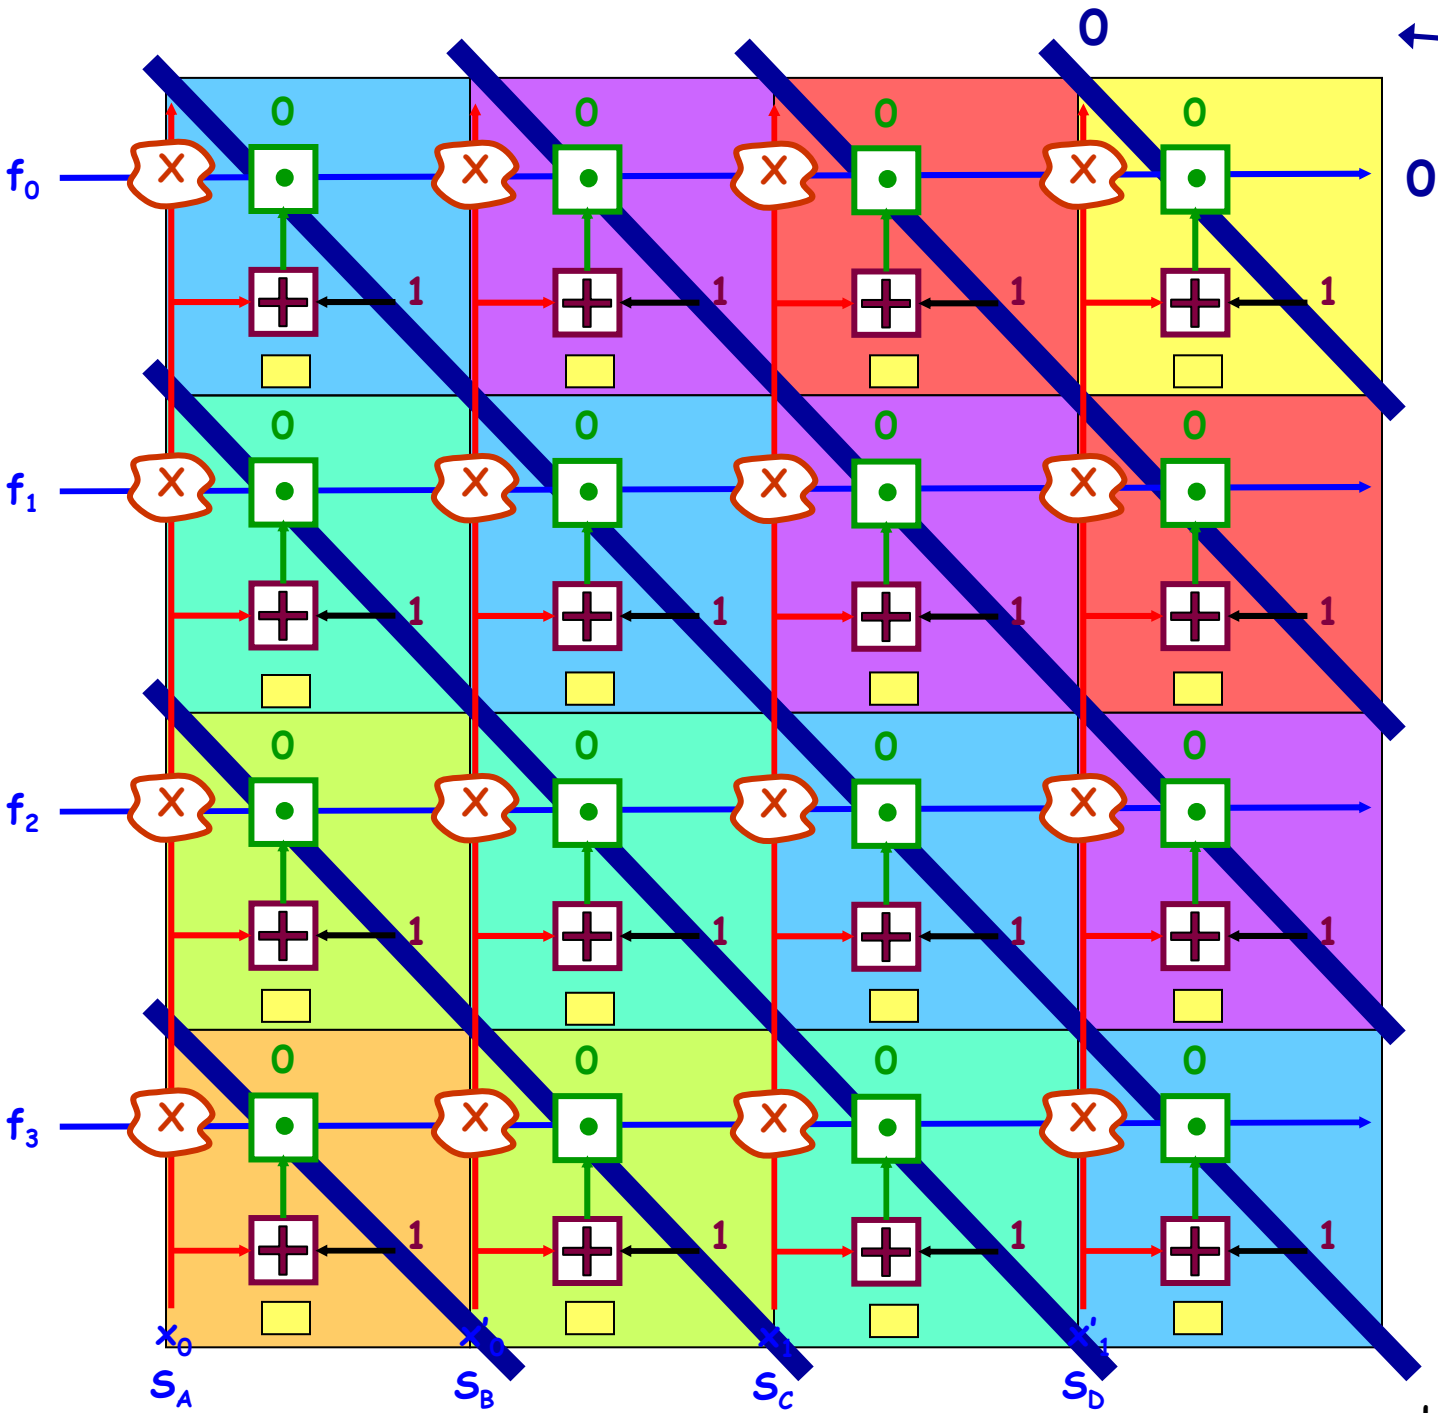

Program by clocking 2 0's, clocking top/right to bottom/left… (reconfigurable)

Execute by clocking from bottom/left to top/right

Work with Sharon Hu, University of Notre Dame

Example PLA design ("throughput")

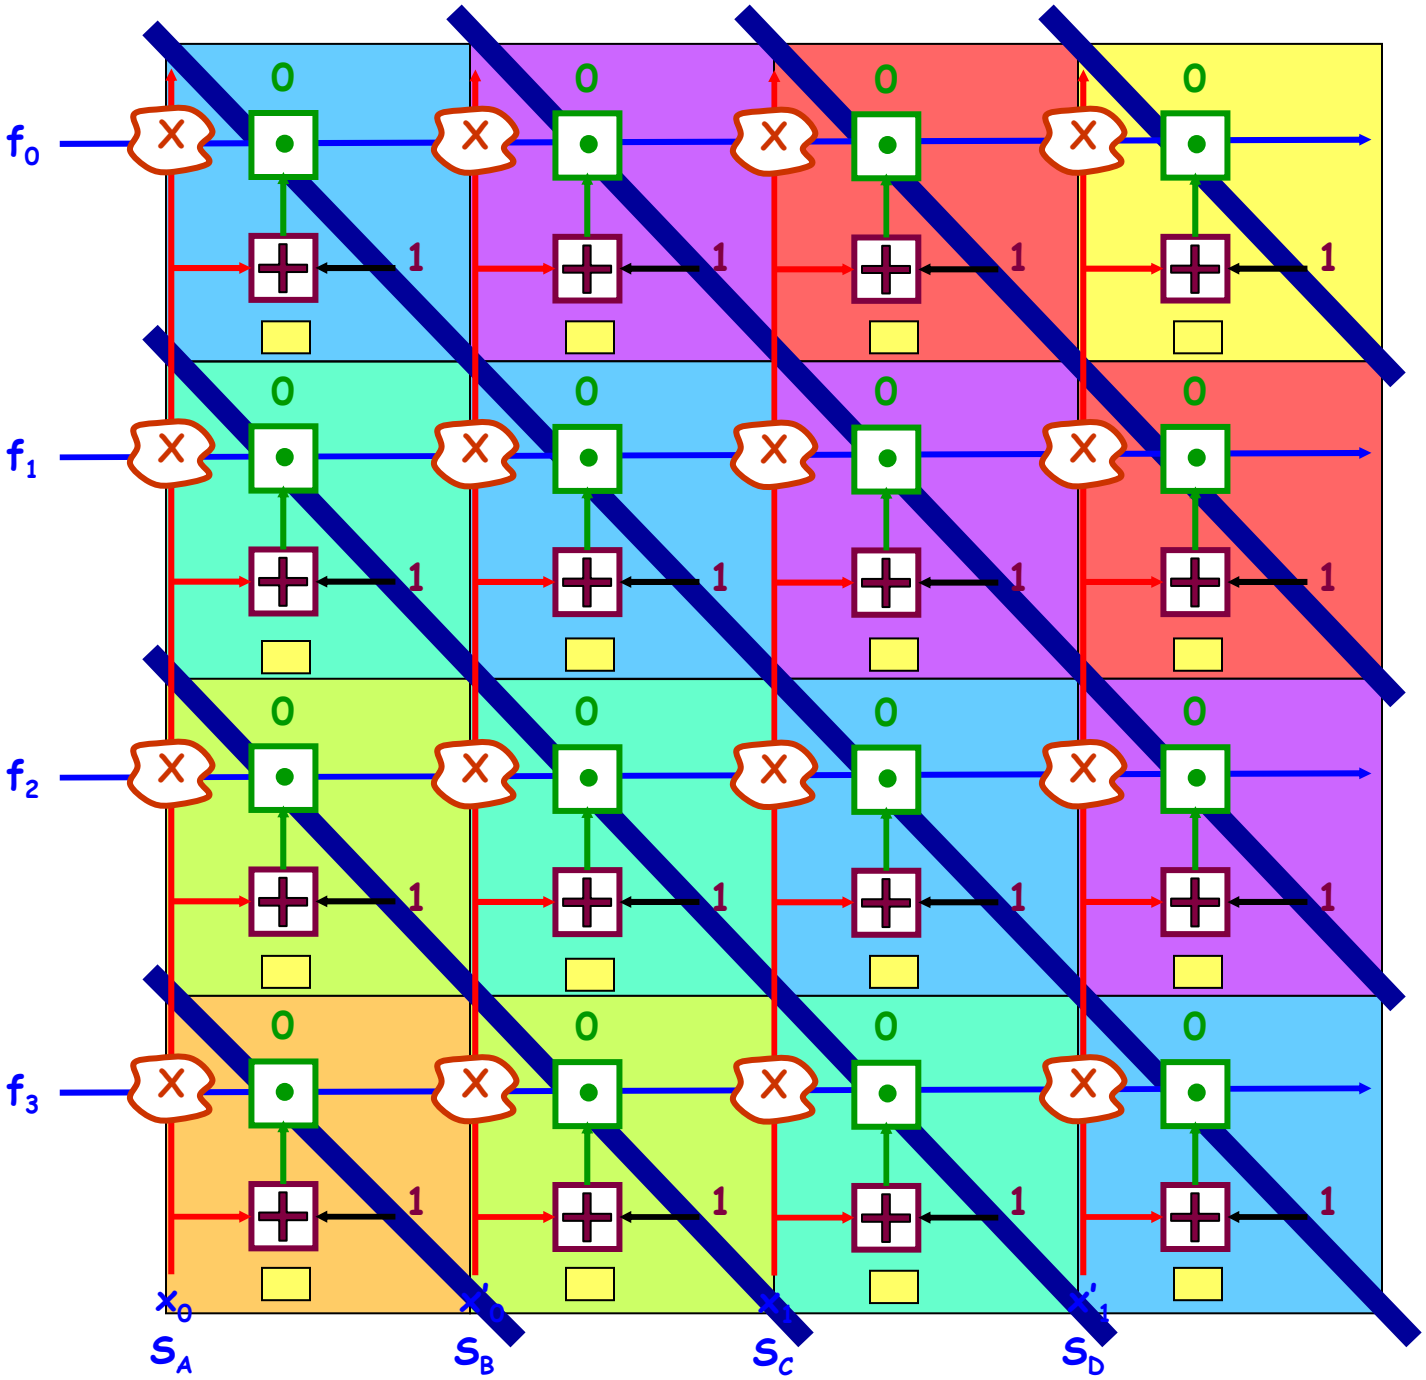

Dependent on granularity of CMOS clock.

If each colored region clocked separately, get two *f* values per "clock cycle"

If not, depends on granularity of clock…

Work with Sharon Hu, University of Notre Dame

Example PLA design ("parts")

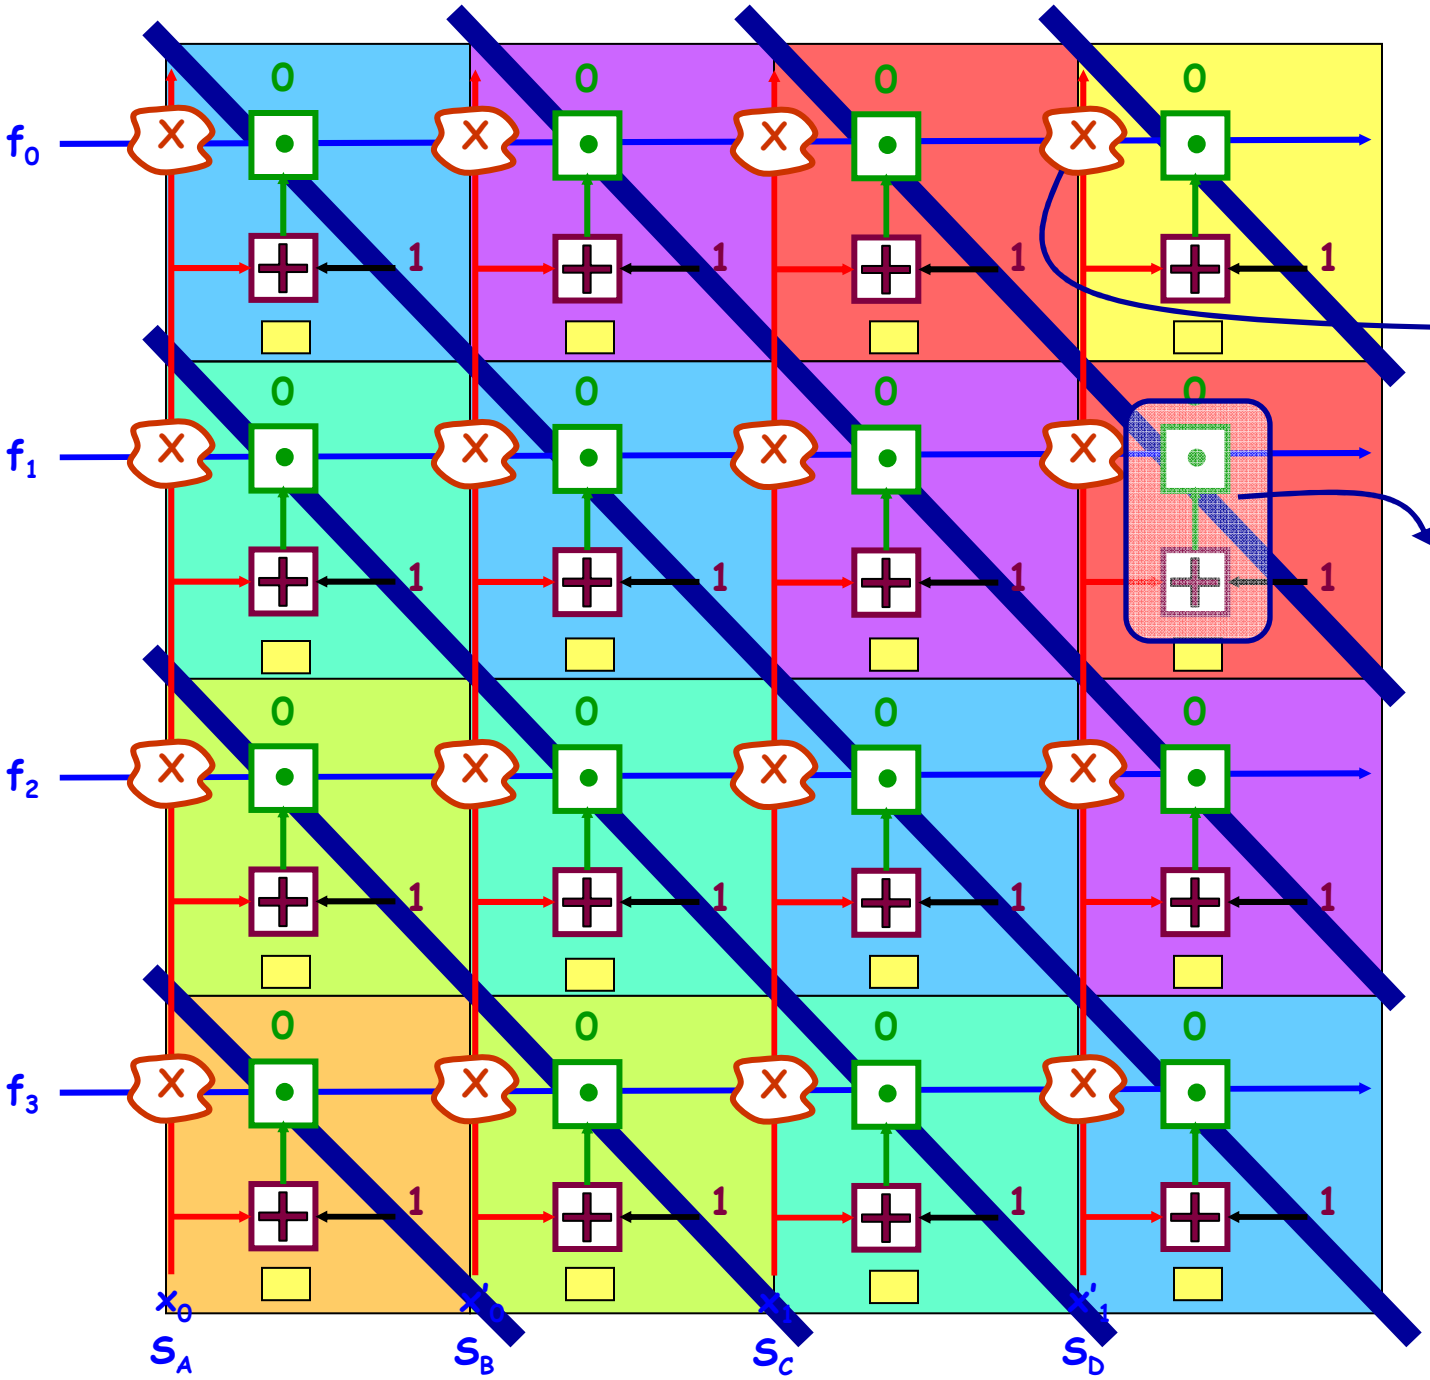

If one cell type, need 1 logical X per node…

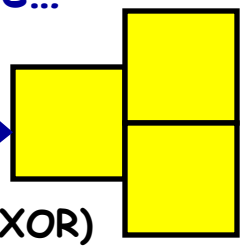

(36 tiles/XOR)

AND and OR gates are minimal - 4-6 DNA tiles…

Everything else just wiring tiles…

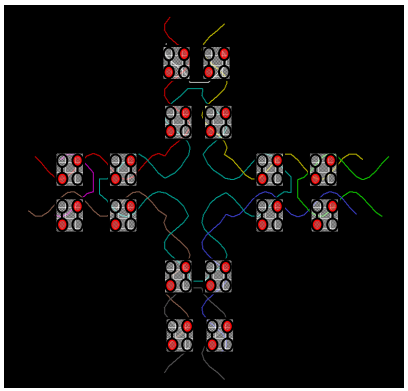

Work with Sharon Hu, University of Notre Dame

Example PLA design (area)

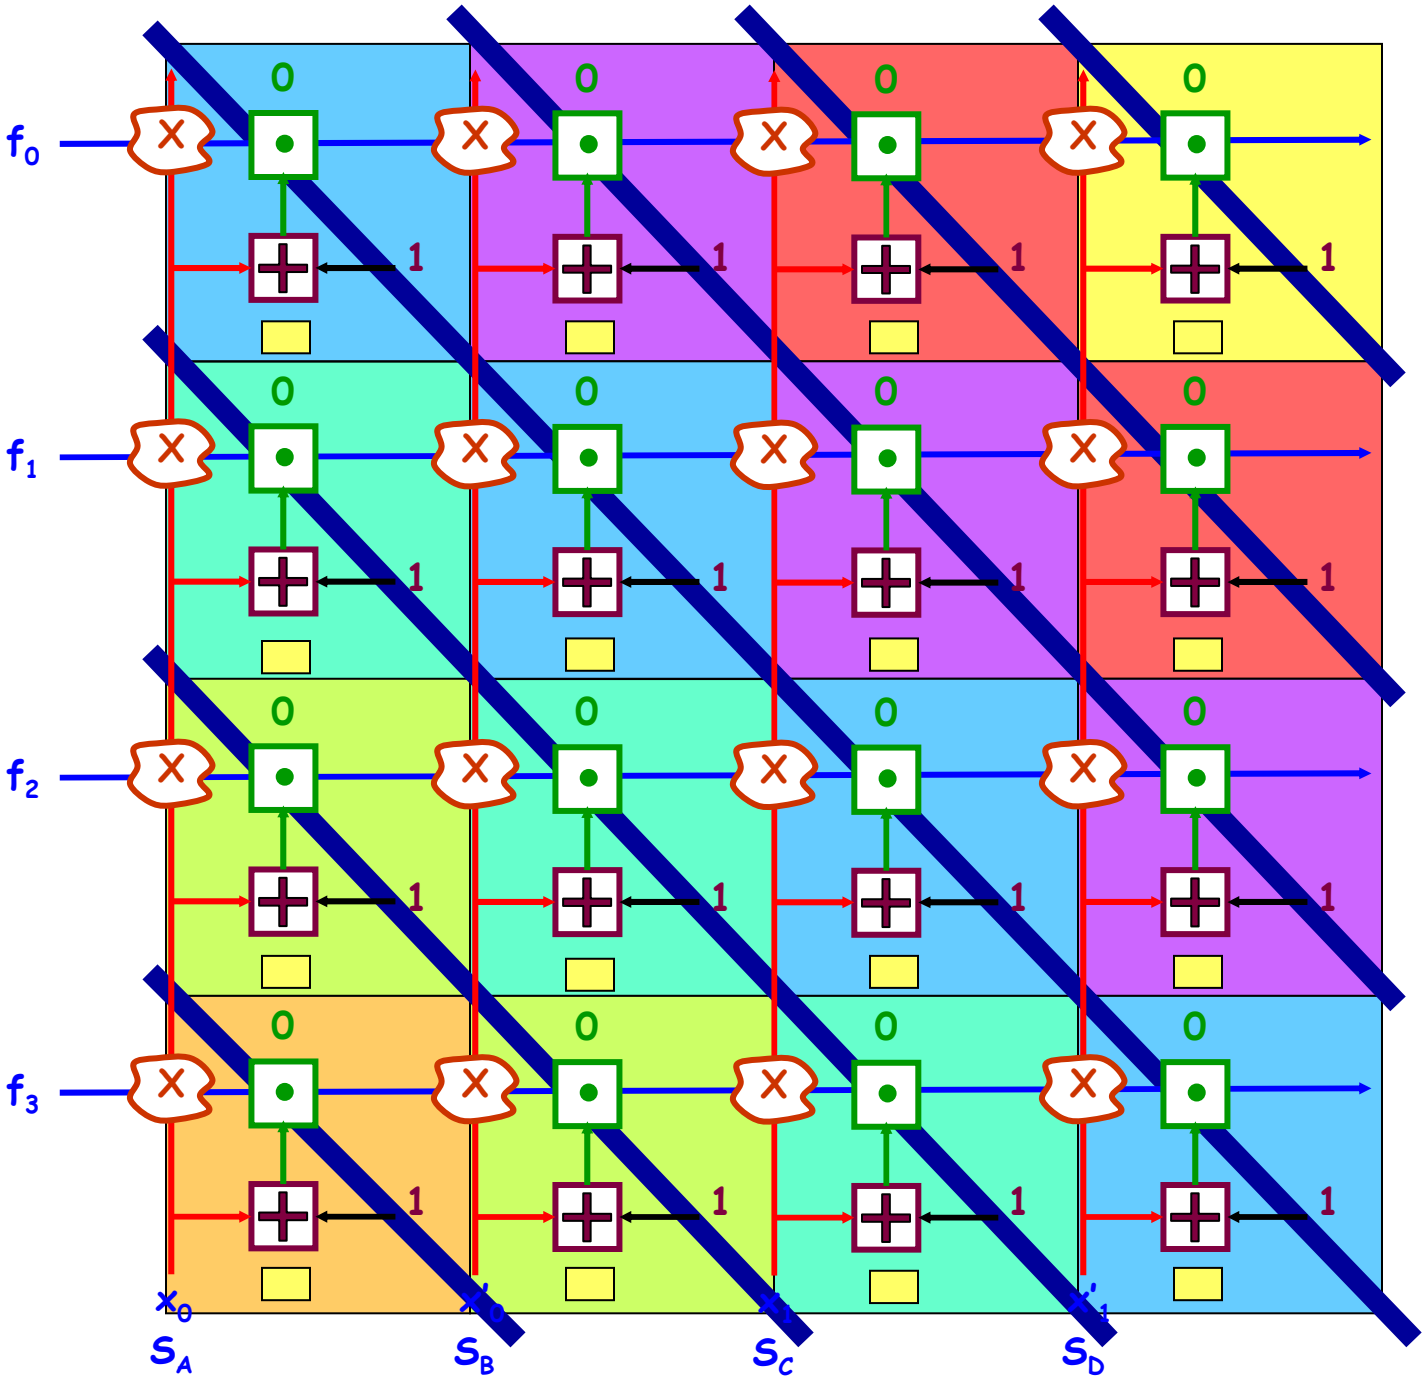

Back of the envelope calculations…

Considering EBL, each node would be about 150 nm x 150 nm

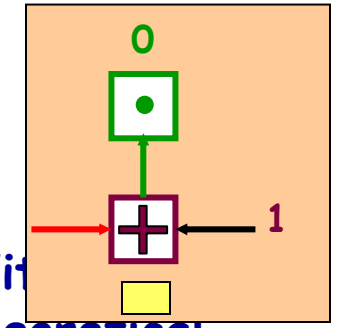

Wit… theoretical constructs, each node *possibly* 110 nm x 60 nm…

Work with Sharon Hu, University of Notre Dame

Example PLA design (programming)

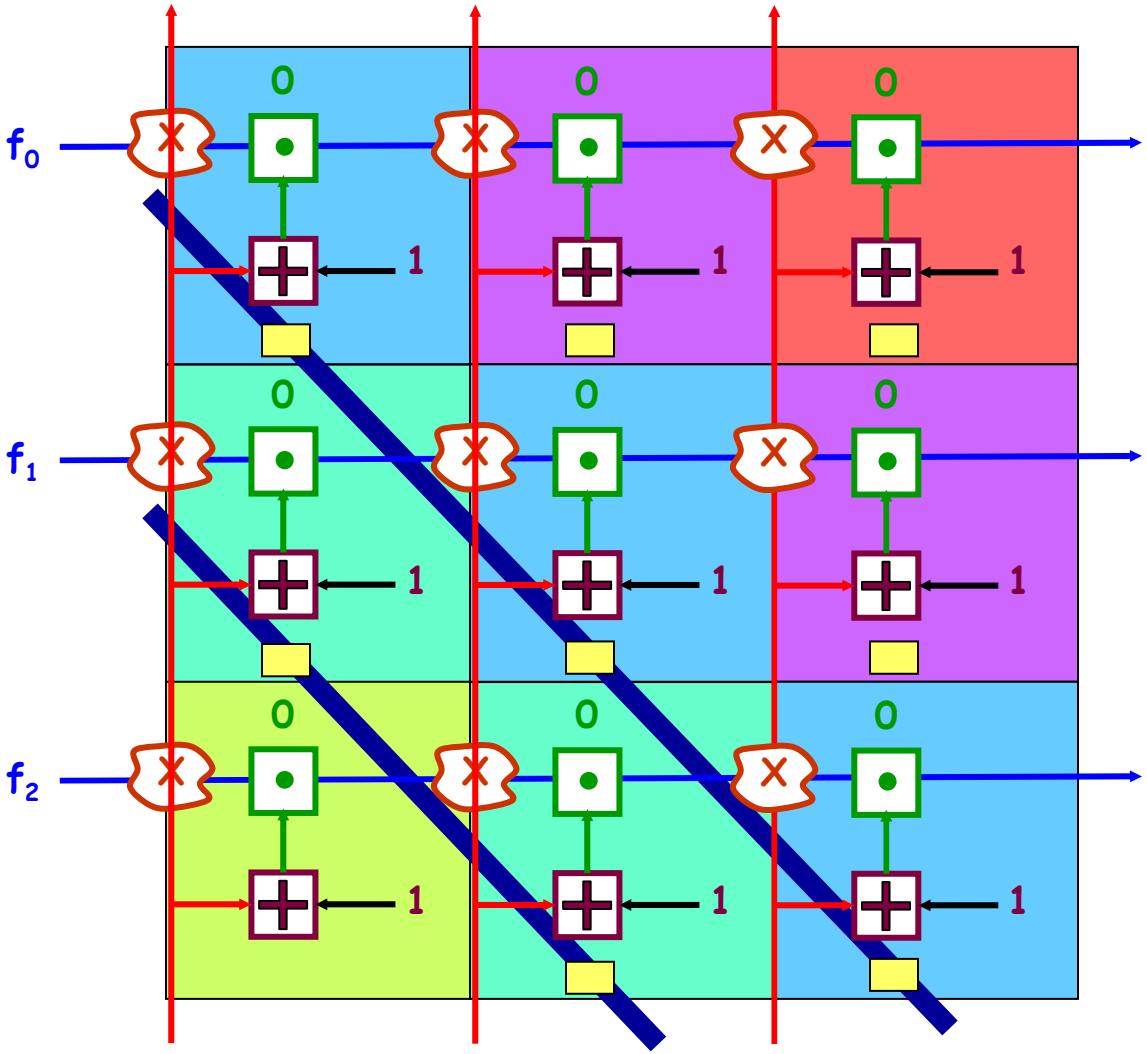

Idea applied in reverse - wires kept low to always keep part of a circuit off…[A]

Can *physically* program by keeping certain clock wires permanently high…

QuickTime™ and a TIFF (LZW) decompressor are needed to see this picture.

Work with Sharon Hu, University of Notre Dame

B: Enrique Blair, M.S. Thesis, 2003.

Example PLA Design (AND and OR)

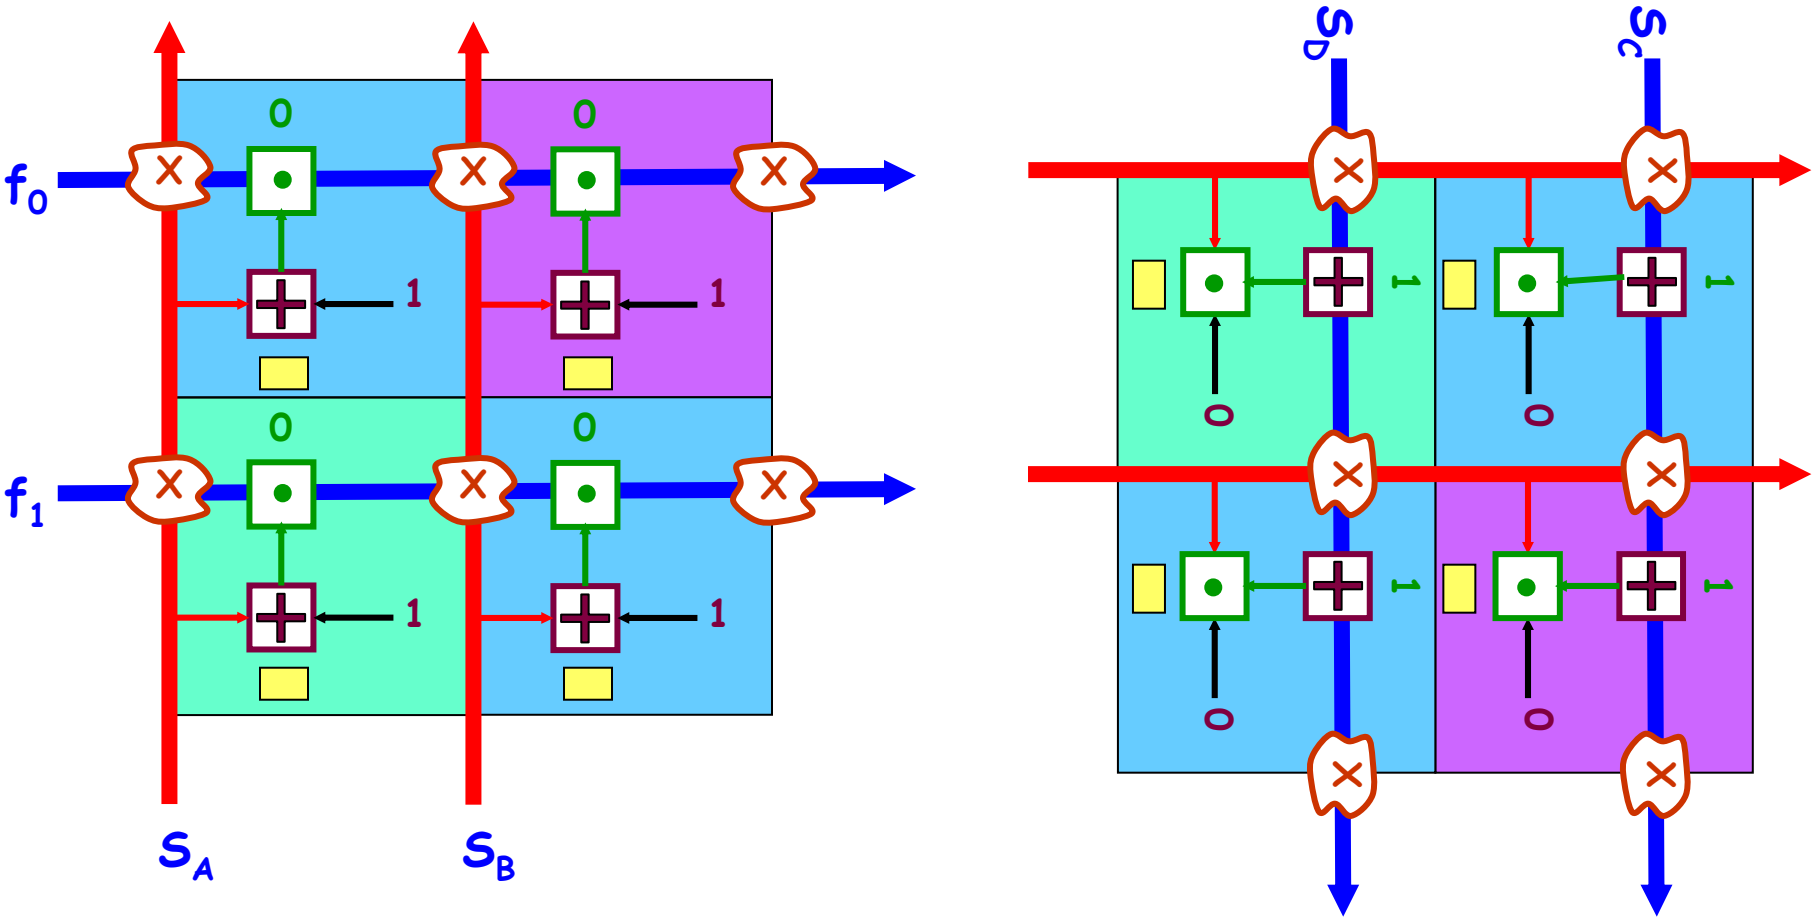

OR plane is almost exactly the same structure, just reversed…

PLA - counterflow

No control signal is stored

Can reprogram during computation

"Throughput" remains unchanged…

In 4x4, get two *f* values every "clock cycle"

AND terms "rotate"

Clock more complex

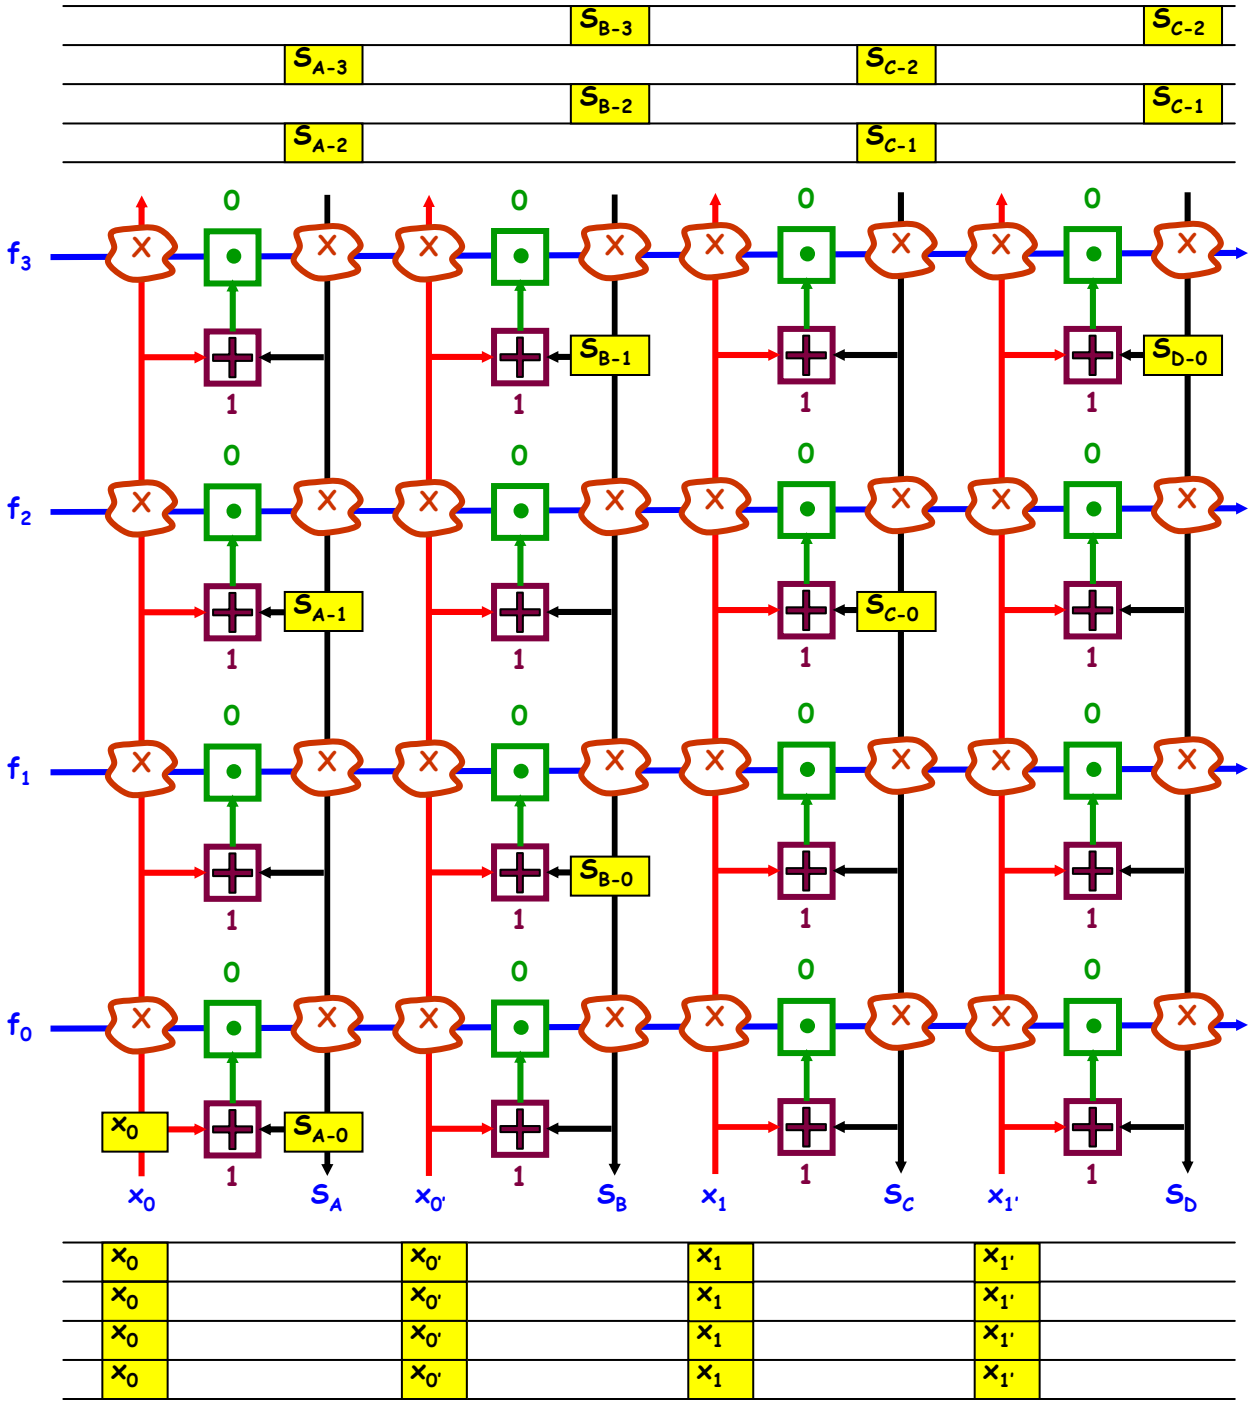

Max throughput when pipe 1/2 full (Sutherland)

Conclusions

- Most architectural work should apply to all implementations
  - Even with first target, can do interesting things at reasonable scales…
- Can design a processor + memory to…
  - Conventional von Neumann architecture probably not most efficient[A]…
- CS work should guide PS as to what parts to build 1st…
- Density numbers good for (probably) bad architectures…
  - …and a gate is only 6 cells and all IC is cells…

A: Niemier, Kogge, ISCA 2001

Conclusions

- Systolic[A], wave-like, counterflow architectures all insinuated by PLA slides… (Doug Berger's work too…)
  - (Some) applications that might map well to QCA
    - Signal processing - FIR, IIR
    - Matrix arithmetic, Eigenvalue calculations
    - Non-numeric applications: graph algorithms, language recognition, polynomial division, etc.
- …interesting designs look possible with even the simplest of constructs…
  - In working group yesterday…
    - Intel successful in part b/c they found a way to build lots of the same basic part with high yield…
    - …apply this lesson here…

A: "Why Systolic Architectures," H.T. Kung, IEEE Comuter, January, 1982, 37-46.