AER Circuits, Systems, and Tools

Bernabe Linares-Barranco

Sevilla Microelectronics Institute (IMSE) - Spanish Research Council (CSIC)

Instituto de Microelectrónica de Sevilla
CICA, Av. Reina Mercedes s/n,
41012 Sevilla, SPAIN

Phone: 34-95-5056670
Fax: 34-95-5056686
E-mail: bernabe@imse.cnm.es

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems
- Some useful circuits:
  - calibration
  - LVDS interface
- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip
- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR
- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

Conventional Vision Sensing/Processing/Recognition

FRAMES

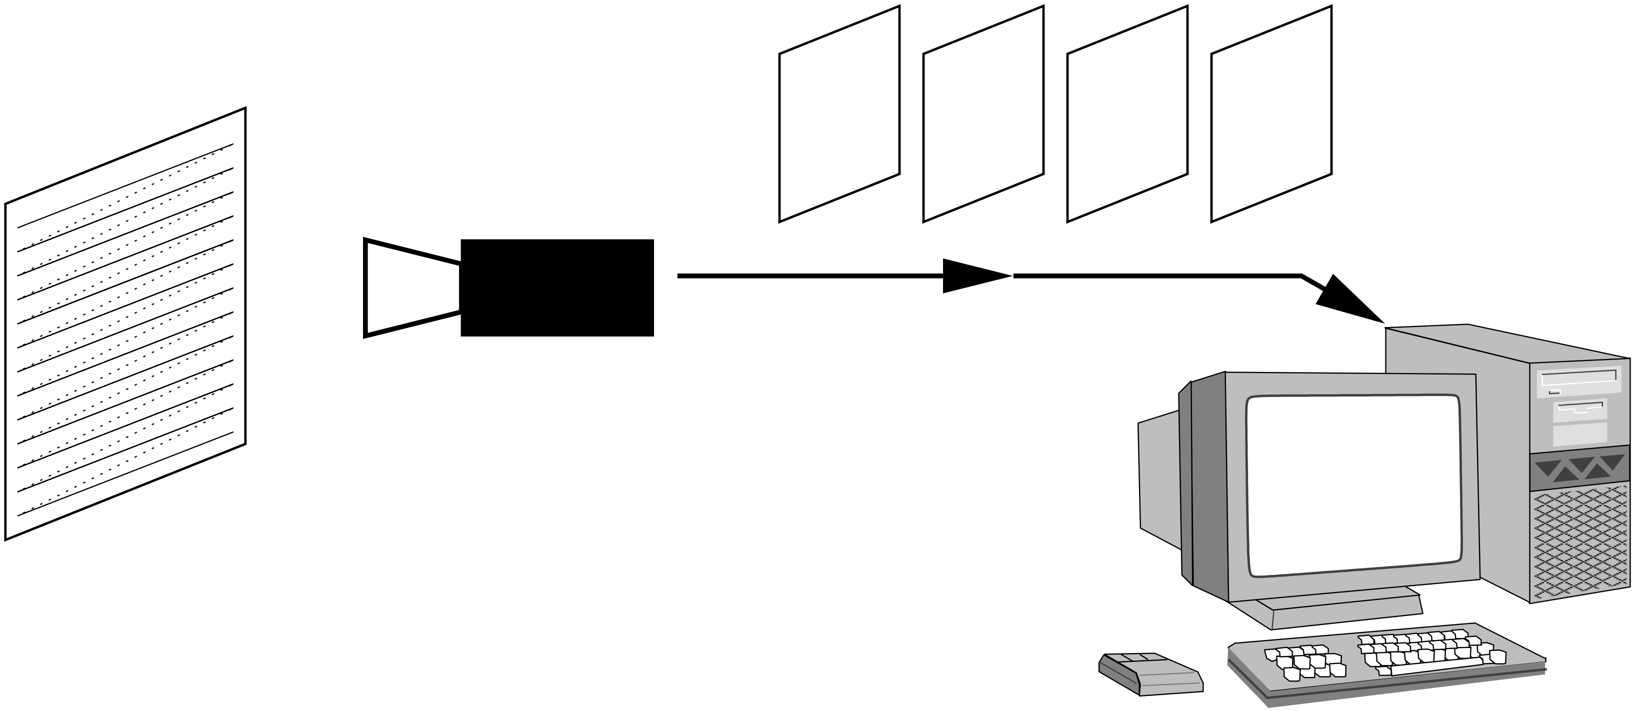

- Feature Extraction Stages
- Feature Combination Stages
- Classification/Decision Stages

Biology

Recognition
Delay
< 150*ms*

Simon Thorpe
Nature 1996

Biology

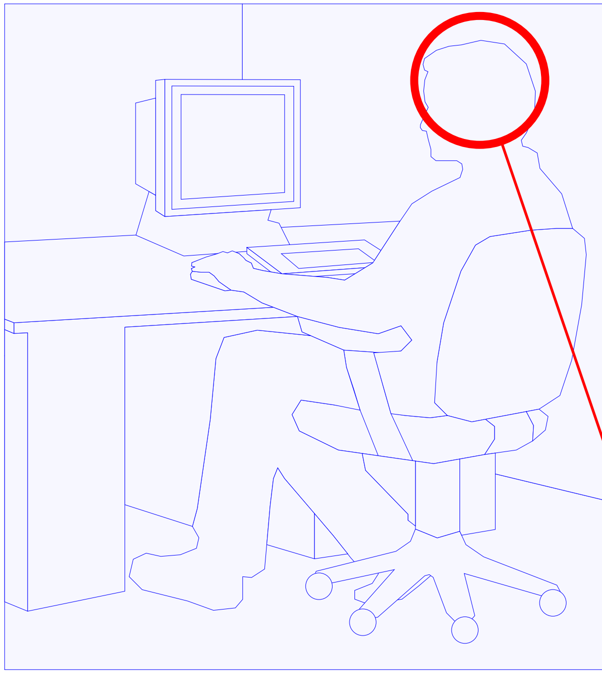

Simon Thorpe
Nature 1996

Recognition
Delay
< 150*ms*

- *feedforward*
- *1 spike/neuron*

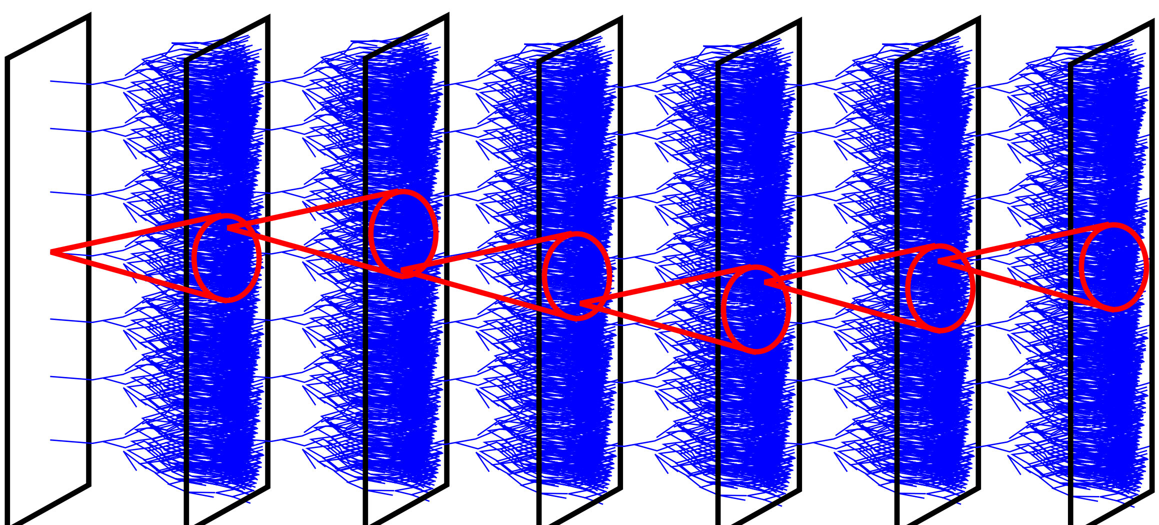

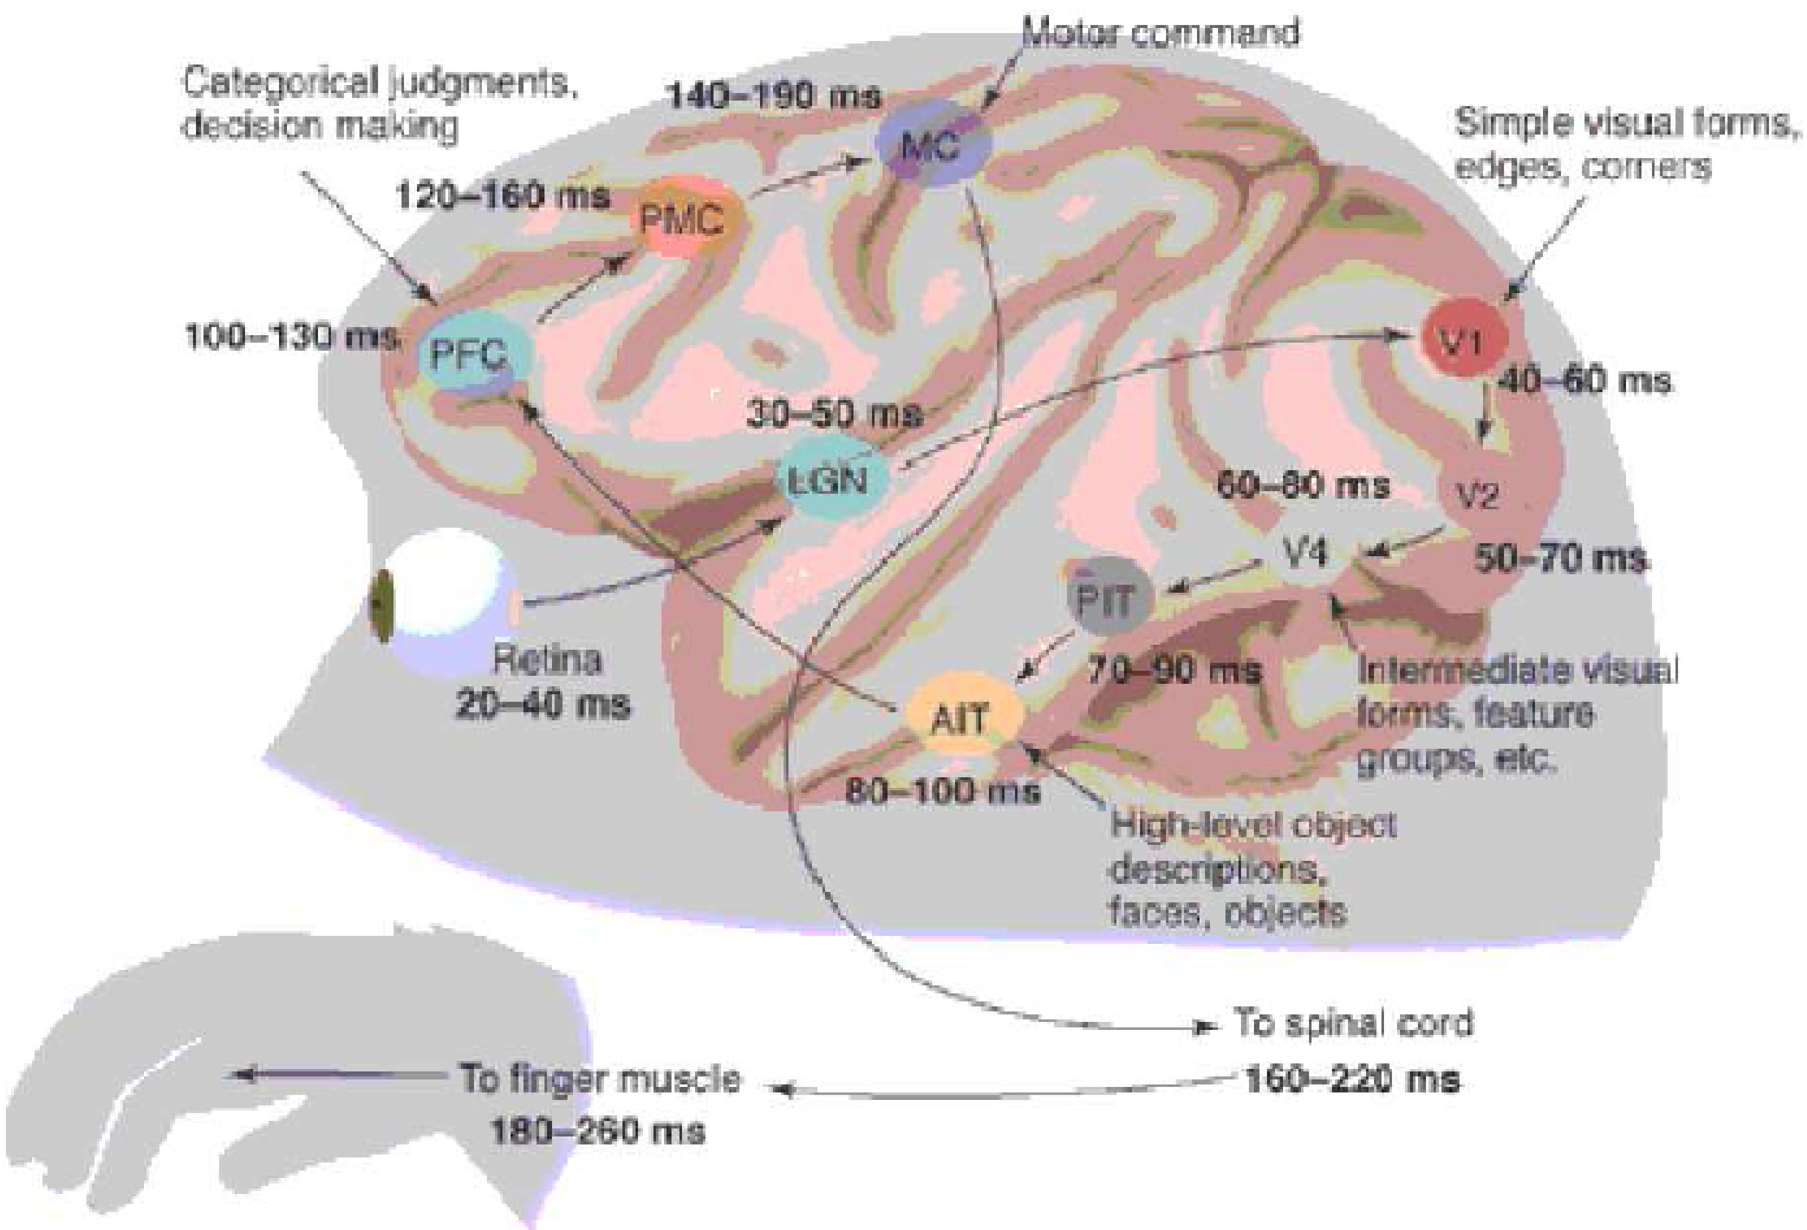
Motor command
Categorical judgments, decision making
140–190 ms
MC
Simple visual forms, edges, corners
120–160 ms
PMC
100–130 ms
PFC
V1
40–60 ms
30–50 ms
LGN
60–80 ms
V2
V4
50–70 ms
PIT
Retina
20–40 ms
70–90 ms
Intermediate visual forms, feature groups, etc.
AIT
80–100 ms
High-level object descriptions, faces, objects
To spinal cord
To finger muscle
160–220 ms
180–260 ms

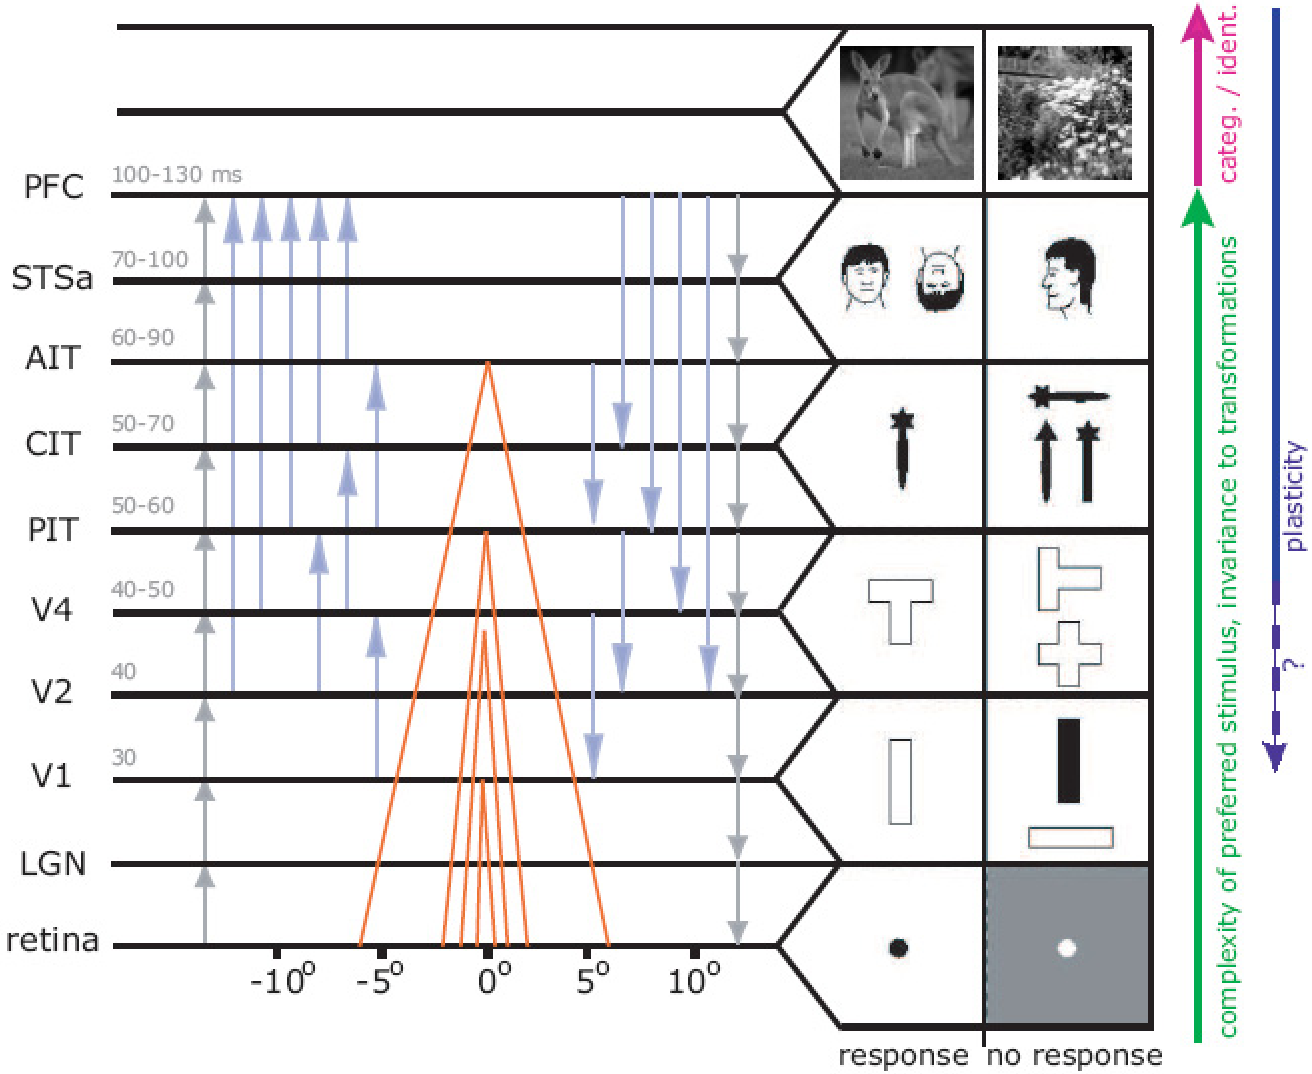
PFC
100-130 ms
STSa
70-100
AIT
60-90
CIT
50-70
PIT
50-60
V4
40-50
V2
40
V1
30
LGN
retina
-10°
-5°
0°
5°
10°
response
no response
categ. / ident.
complexity of preferred stimulus, invariance to transformations
plasticity
?

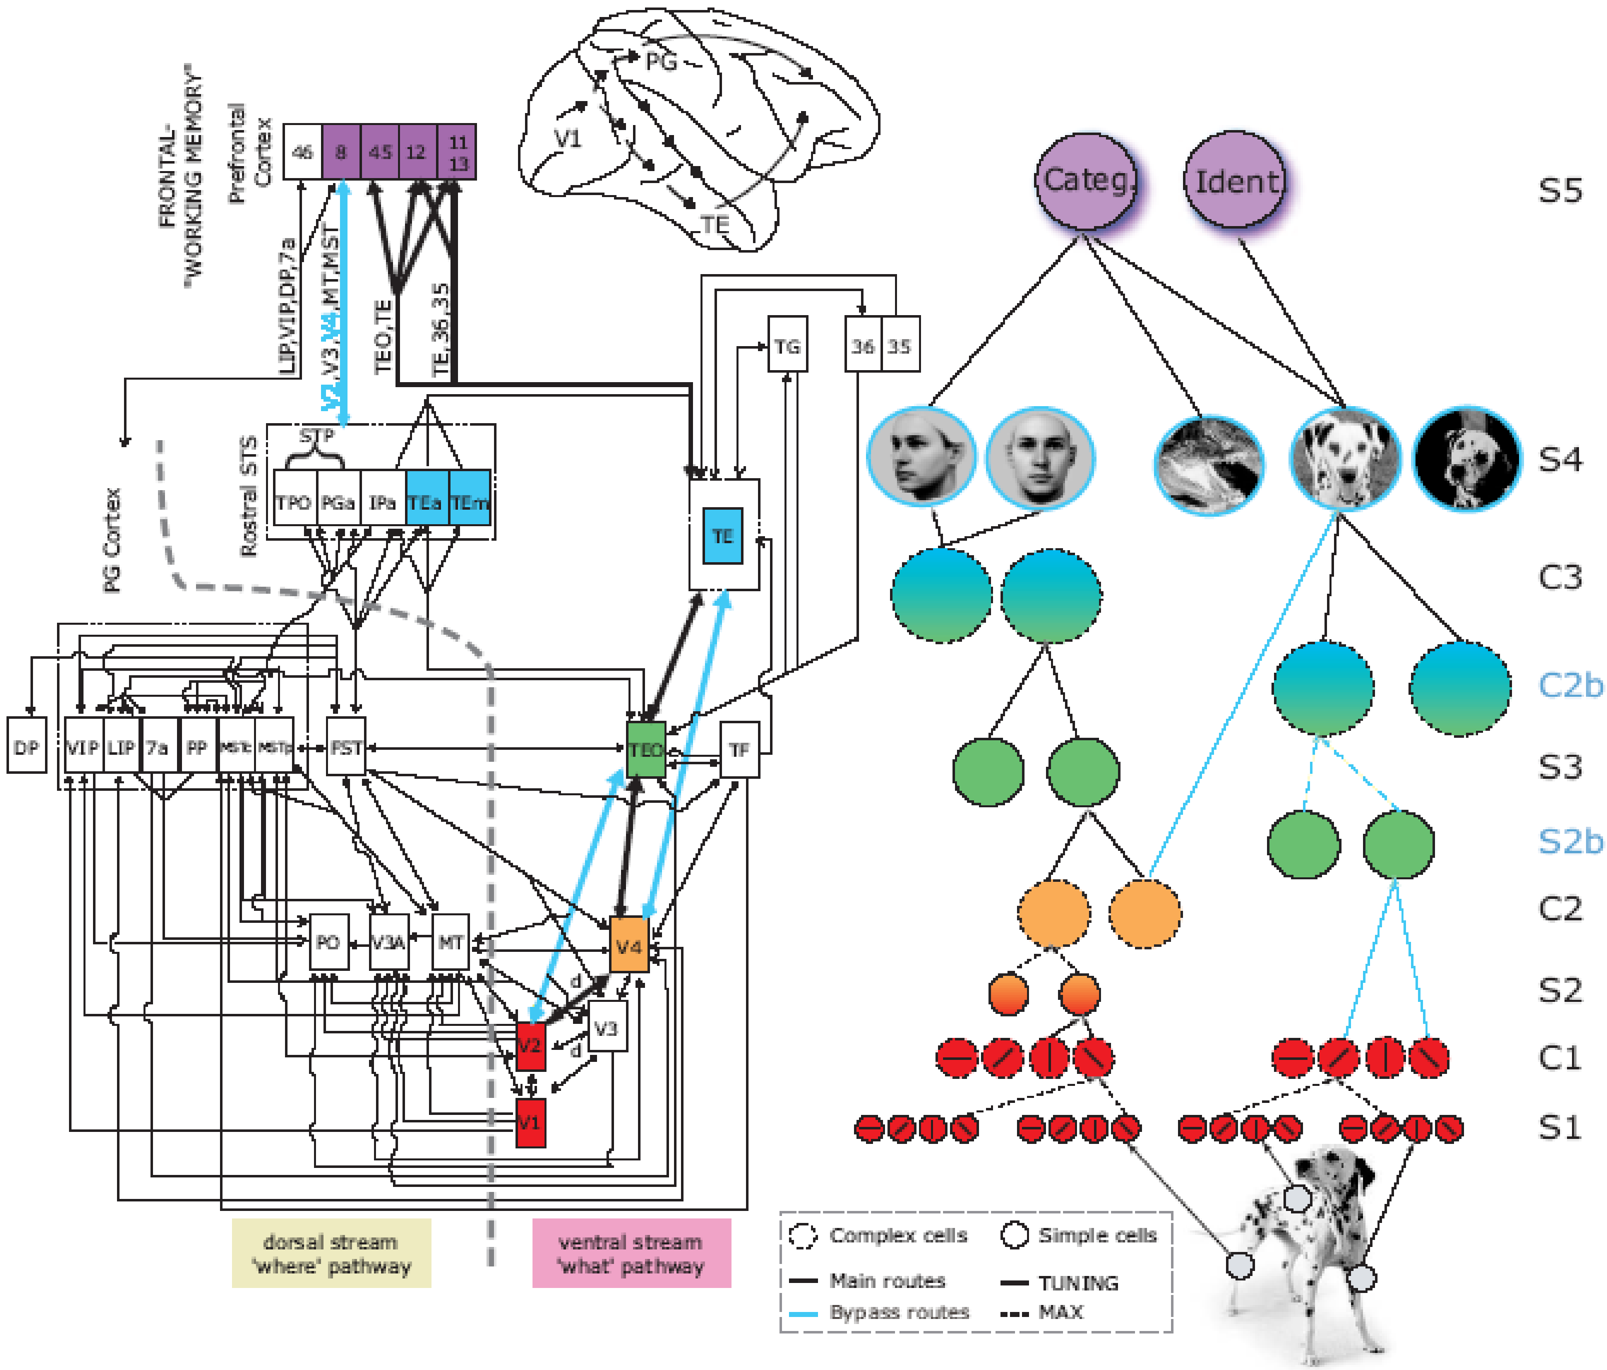

FRONTAL-
"WORKING MEMORY"
Prefrontal Cortex
46
8
45
12
11
13
V1
PG
TE
LIP,VIP,DP,7a
V2,V3,V4,MT,MST
TEO,TE
TE,36,35
PG Cortex
Rostral STS
STP
TPO
PGa
IPa
TEa
TEm
TG
36
35
TE
DP
VIP
LIP
7a
PP
MSTc
MSTp
FST
TEO
TF
PO
V3A
MT
V4
V3
V2
V1
d
dorsal stream
'where' pathway
ventral stream
'what' pathway
Categ.
Ident
S5
S4
C3
C2b
S3
S2b
C2
S2
C1
S1
Complex cells
Simple cells
Main routes
TUNING
Bypass routes
MAX

Serre & Poggio (MIT)

Ventral Stream Model for Immediate Recognition

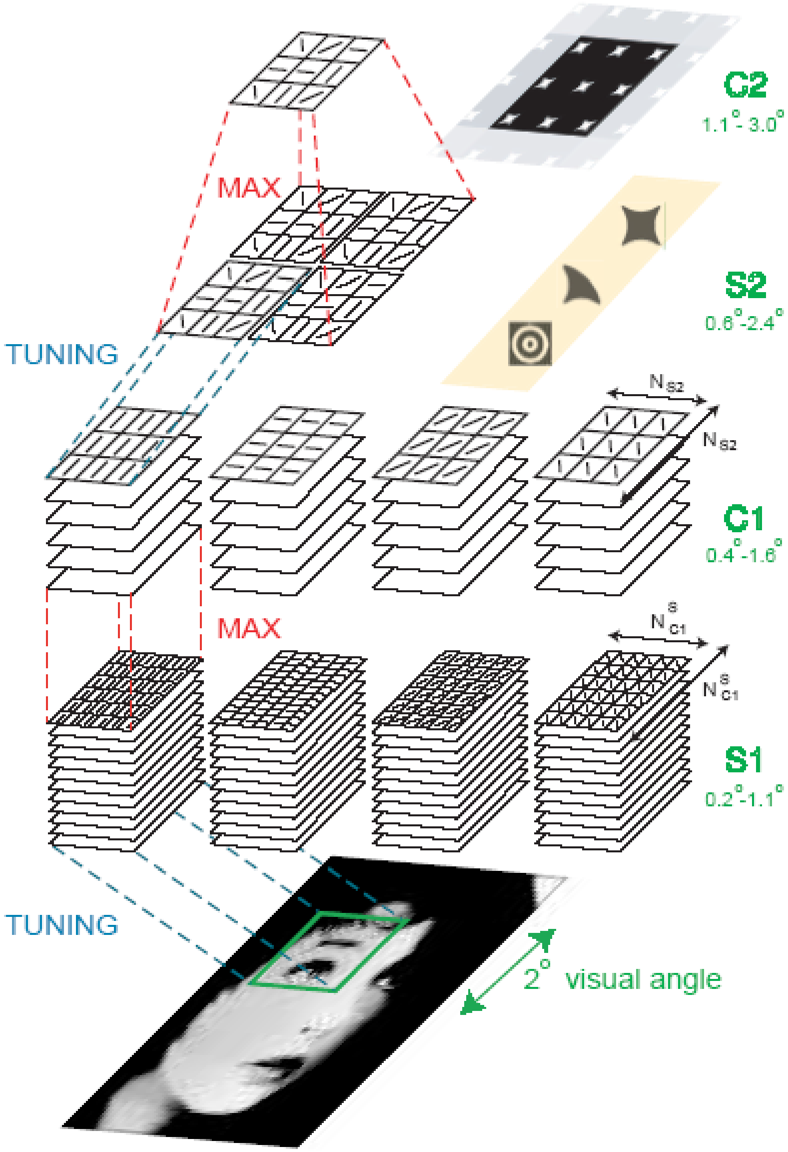

- projection field processing
- short-range & dense for first layers
- long-range & sparse for later layers
- hard-wired for first layers
- plastic for later layers
- first layers: massive 2D filtering for different angles and scales
- first layers: basic feature extraction
- later layers: grouping of features -> abstractions

AER (Address Event Representation)

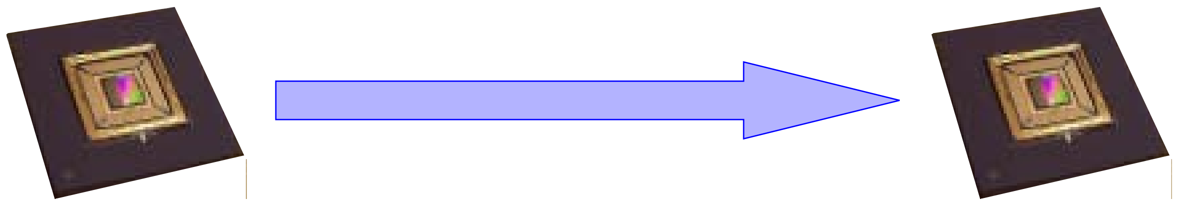

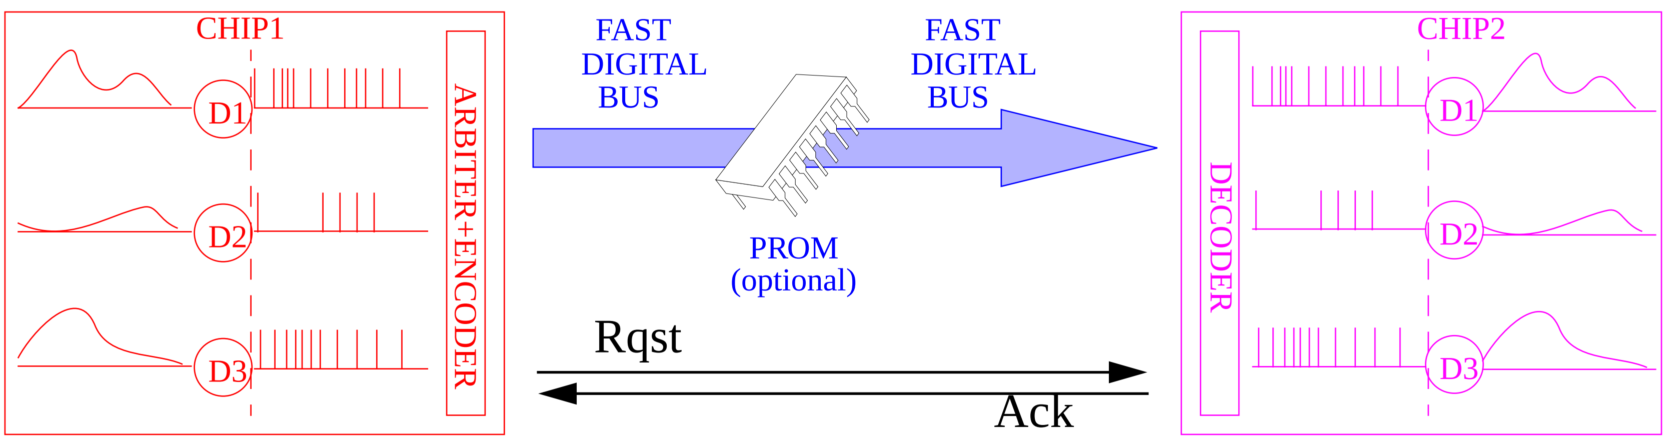

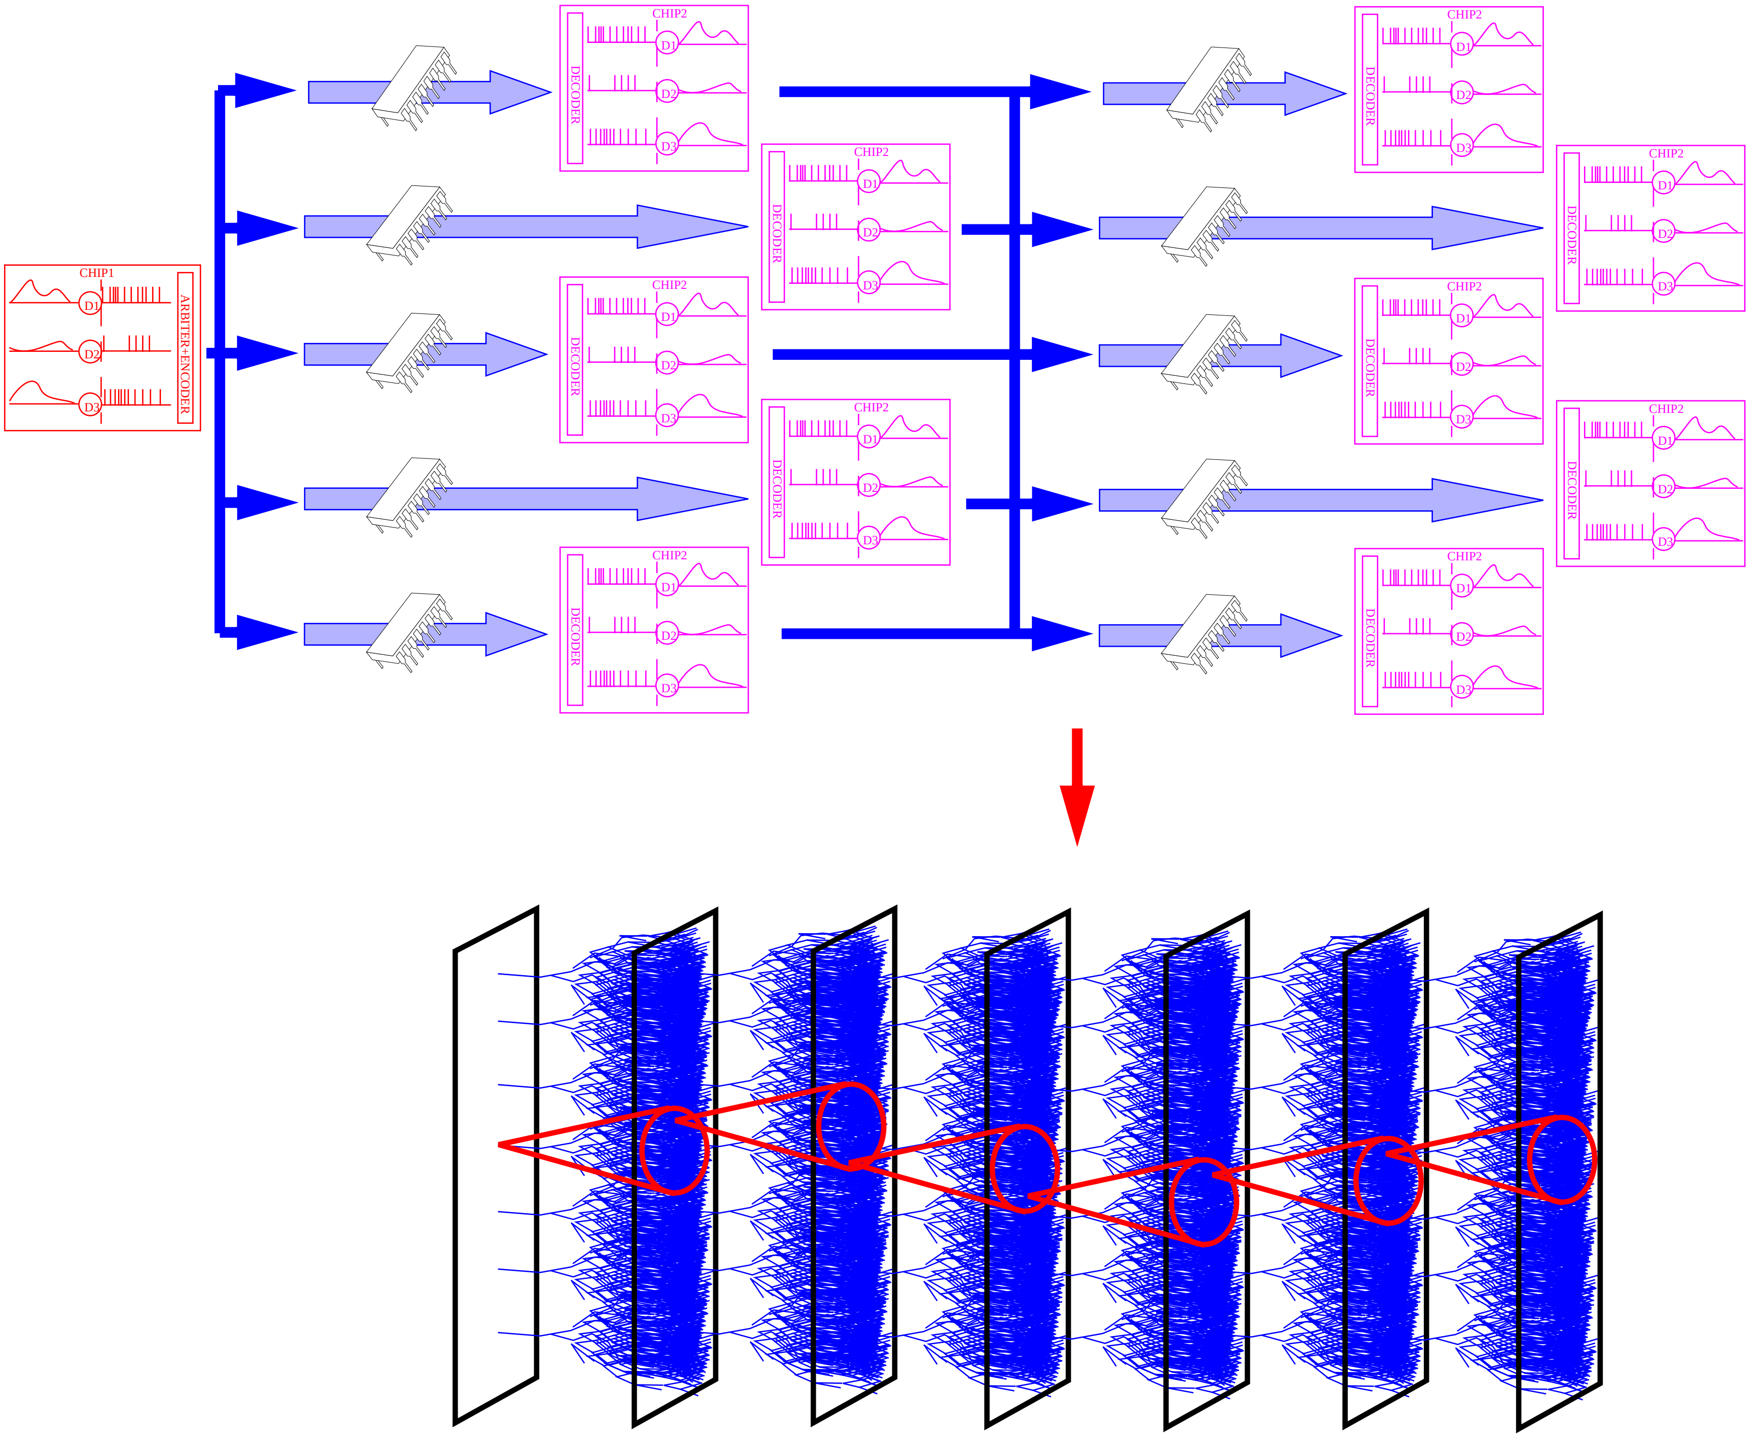
CHIP1
D1
D2
D3
ARBITER+ENCODER
CHIP2
DECODER
D1
D2
D3

Feature Extraction

(AER Convolution Chip)

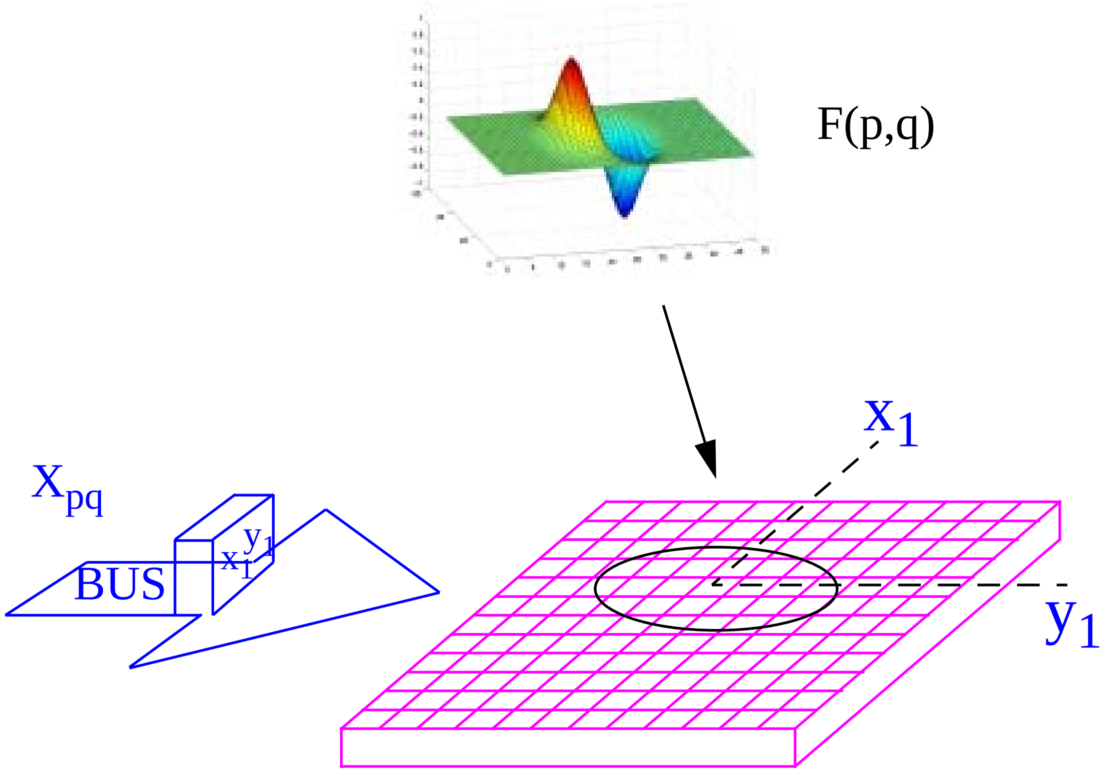

Matrix of integrators in
the receiver chip

CORTICAL TISSUE

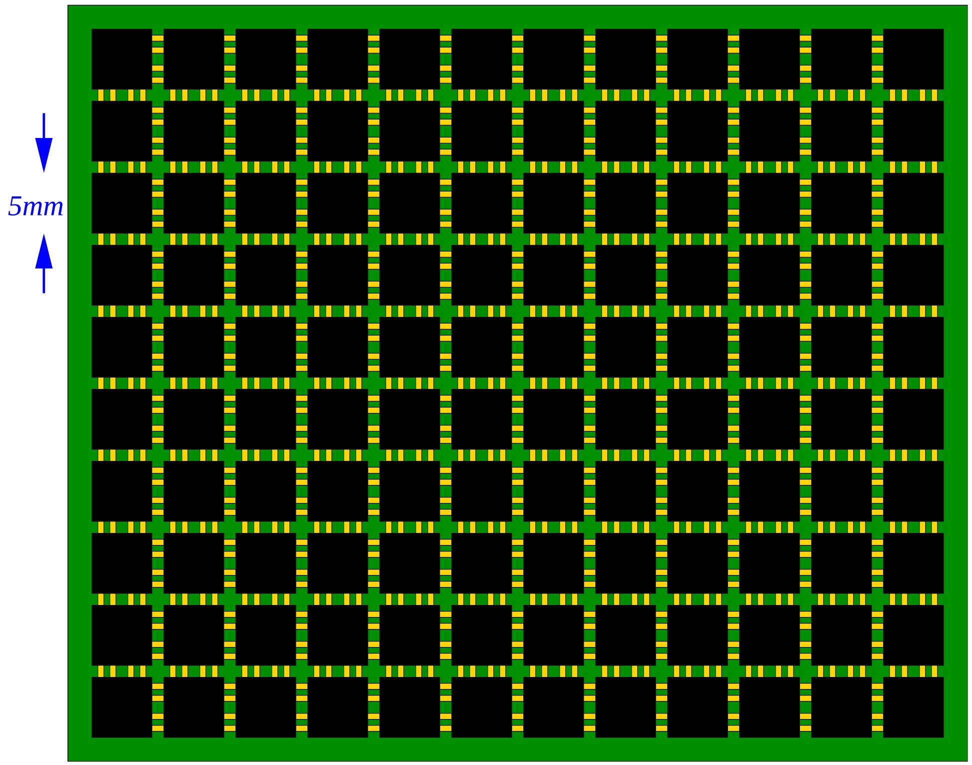

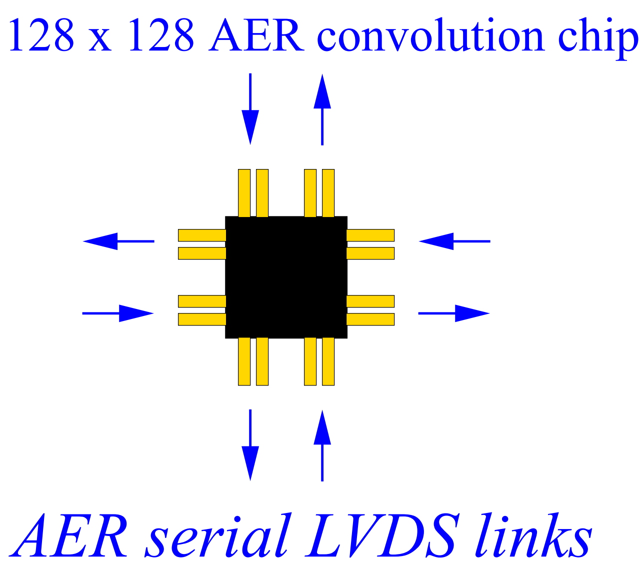

- Events are routed to neighbors through local on-chip routing tables
- Any arbitrary multi-layer feed-forward + feed-back hierarchy can be programmed
- LVDS links allow low-power high-speed event traffic
- Each tile could be a 128x128 programmable kernel convolution chip with local re-routing and remapping capability
- Hundreds of convolution chips can be fit in a ‘Cortical Tissue’ PCB

Computing Power of one such Cortical Tissue PCB

- 120 chips & 436 AER inter-chip links
- Each chip 128x128 neurons and kernel up to 128x128
- Total of 2M neurons
- Total of 32G synapses
- If each AER link requires 30ns per AE:
- 14Geps (interchip)
- 238 Tconnections/sec

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems
- Some useful circuits:
  - calibration
  - LVDS interface
- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip
- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR
- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

Calibration in Neuromorphic Cells

- large arrays
- small cell area
- very low currents (nano-pico amsp)
- high inter-pixel mismatch:

  $\sigma \approx 10\text{-}20\% \quad \Rightarrow \quad 6\sigma \approx 60\text{-}120\%$

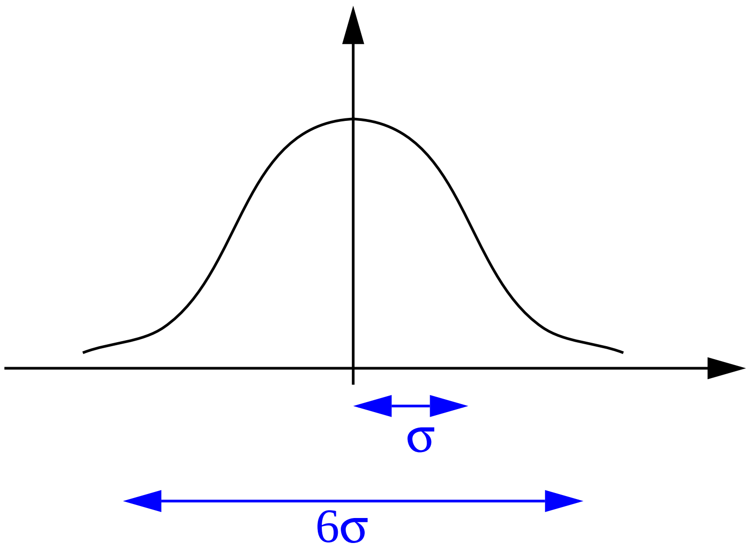

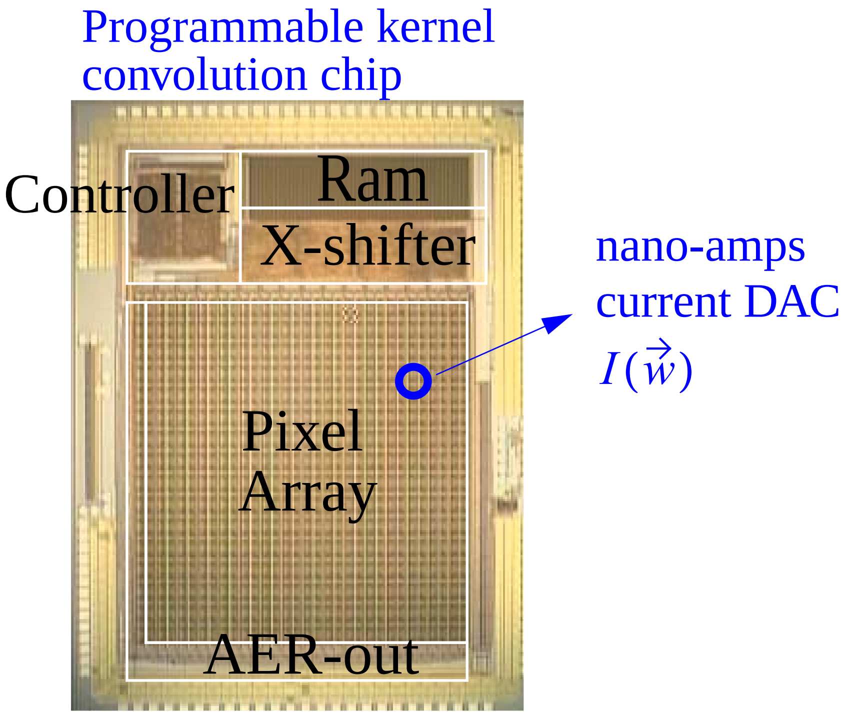

Compact Calibration Circuit

[*IEEE Trans. Neural Networks*, Sep. 2003]

- Ladder-based digi-MOS:

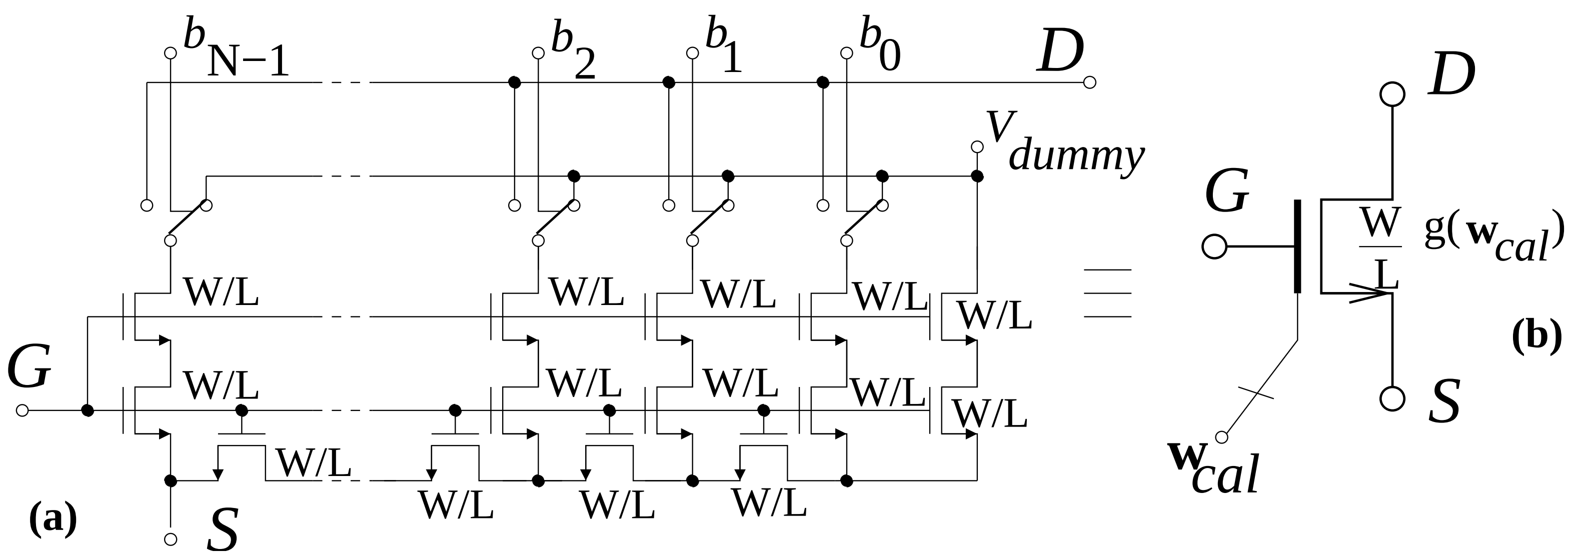

$3N + 1$ unit transistors

N = number of bits

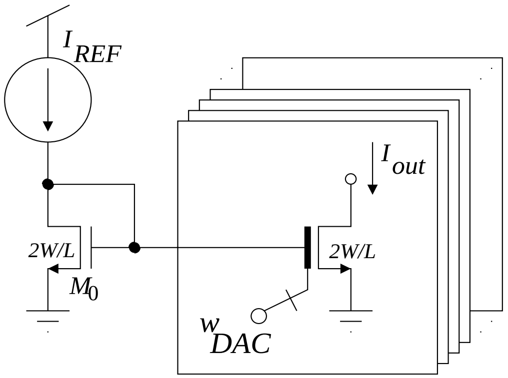

I_{REF}
I_{out}
2W/L
M_0
2W/L
w_{DAC}

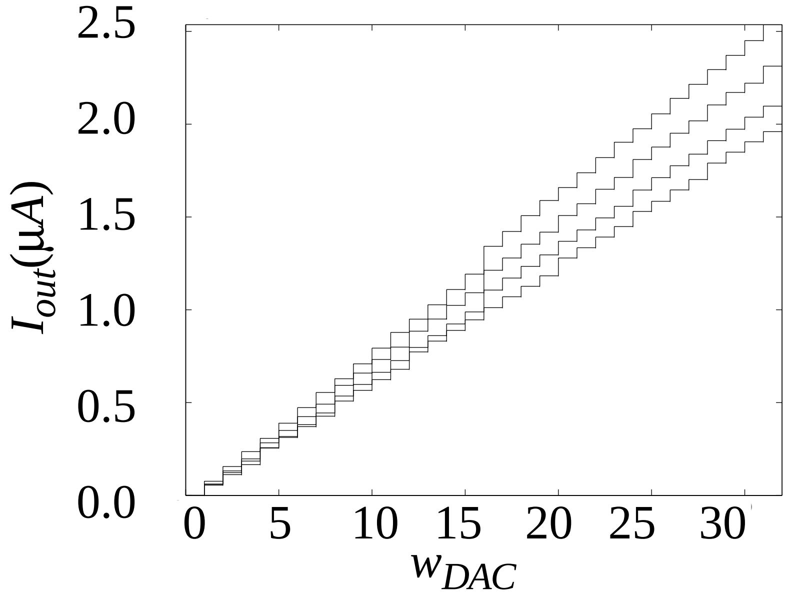

2.5
2.0
1.5
1.0
0.5
0.0
0
5
10
15
20
25
30
$I_{out}(\mu A)$
w_{DAC}

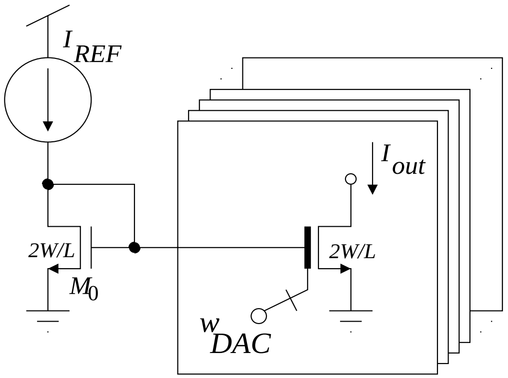

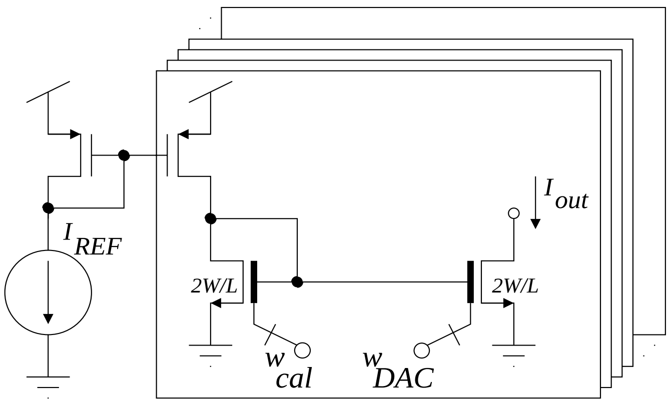

one point calibration

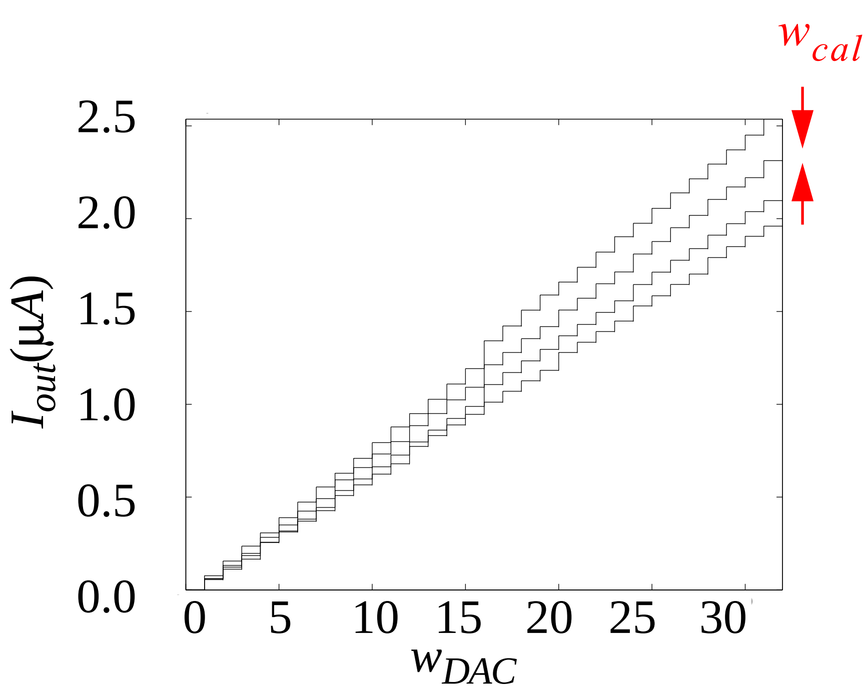

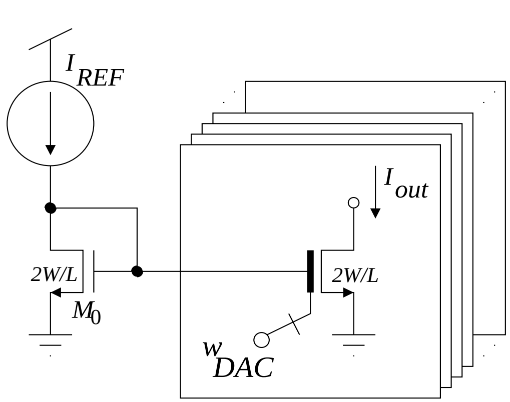

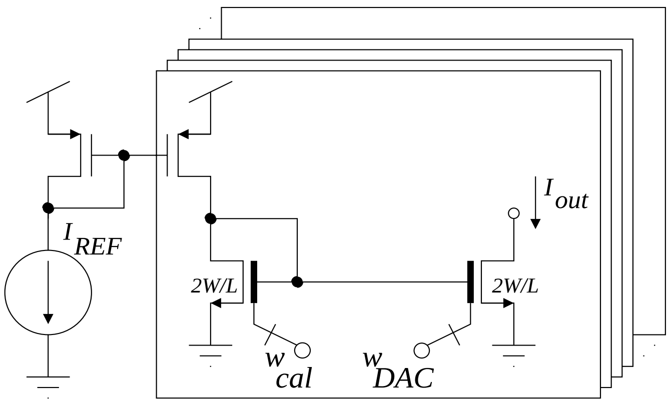

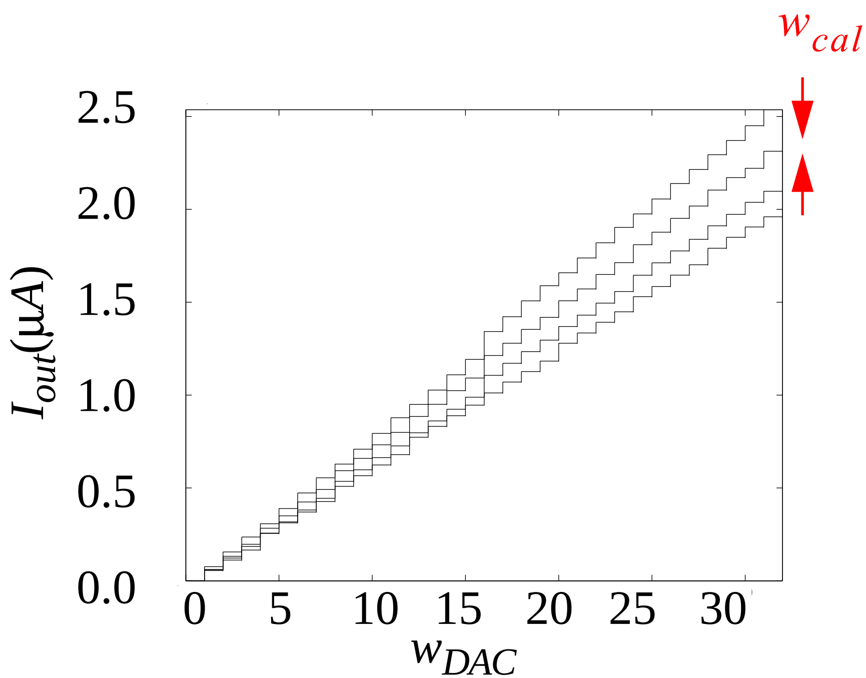

max error: 6σ~3-4%
(~ 4-bit precision)

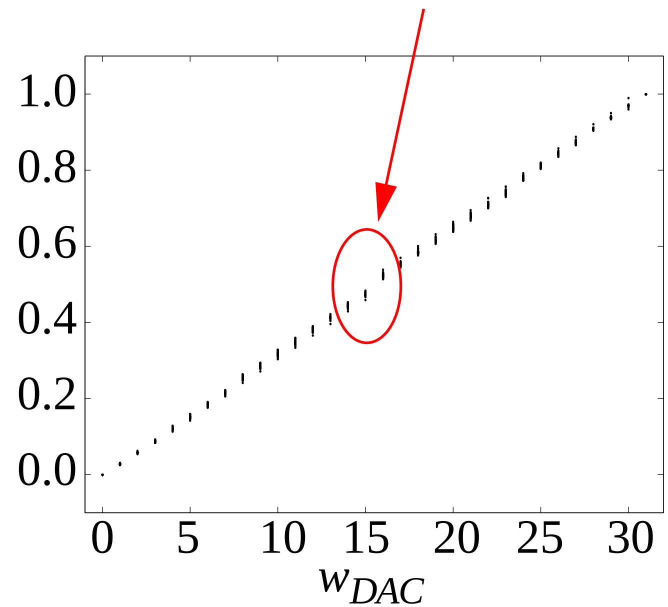

max current: ~2.5µA

unit transistor: 5 x 5µm^2

For Higher Precision

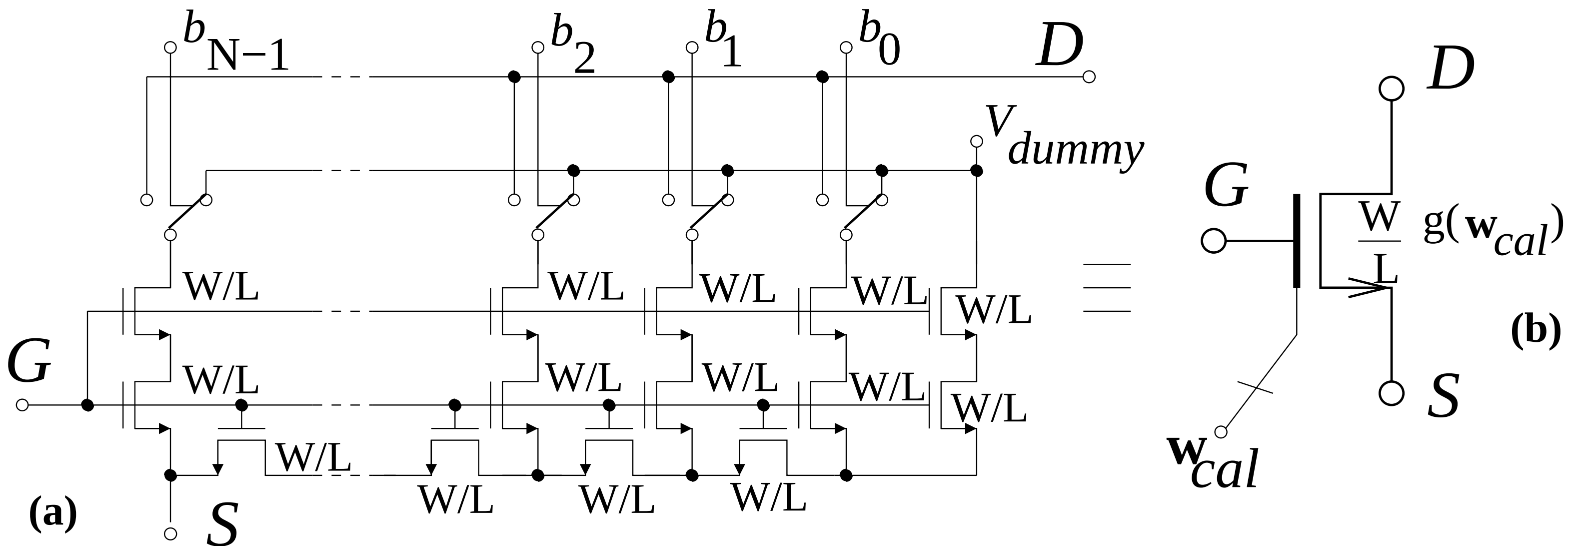

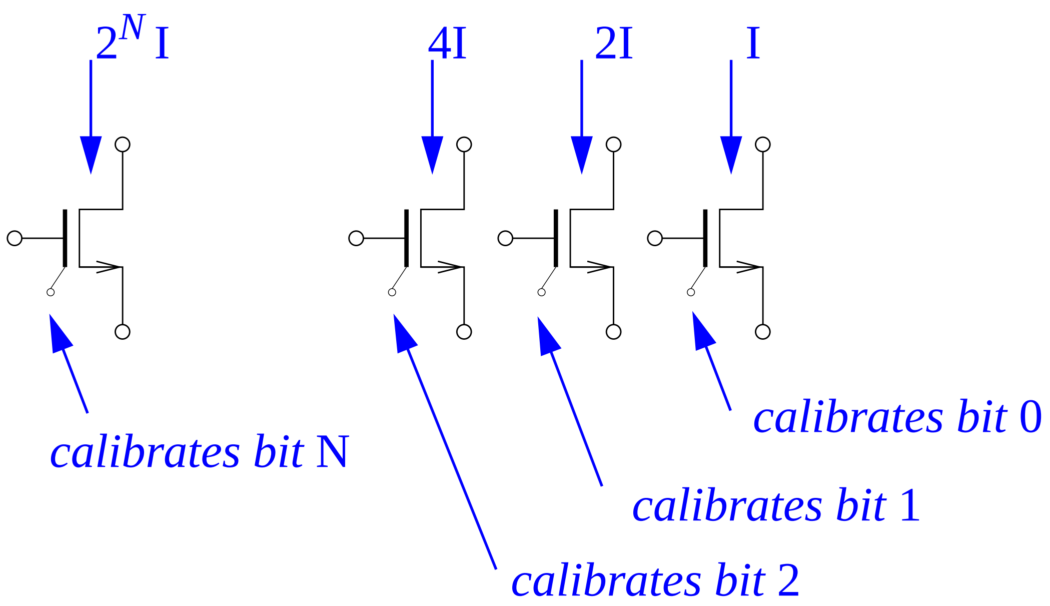

New Concept based on parallel/series MOS association

[Galup-Montoro et al., IEEE JSSC 1994]

From EKV/ACM models: $I_{DS} = \frac{W}{L}[f(V_G, V_S) - f(V_G, V_D)]$

- Generic: - parallel: $\left(\frac{W}{L}\right)_{eq} = \left(\frac{W}{L}\right)_A + \left(\frac{W}{L}\right)_B$

  - series: $\left(\frac{W}{L}\right)_{eq} = \frac{\left(\frac{W}{L}\right)_A \left(\frac{W}{L}\right)_B}{\left(\frac{W}{L}\right)_A + \left(\frac{W}{L}\right)_B}$

Consequently,

- Series association with equal *W*, $\rightarrow$ $L_{eq} = \sum L_i$

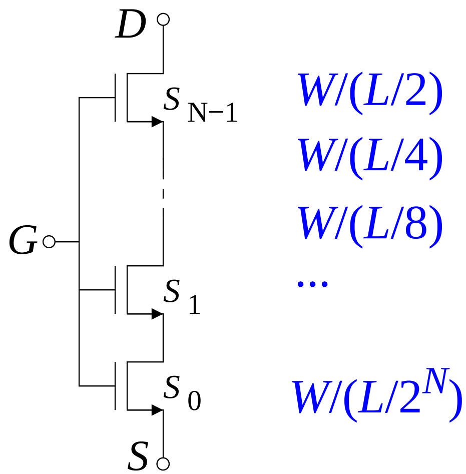

D
G
S
S N−1
S 1
S 0
W/(L/2)
W/(L/4)
W/(L/8)
...
W/(L/2^N)

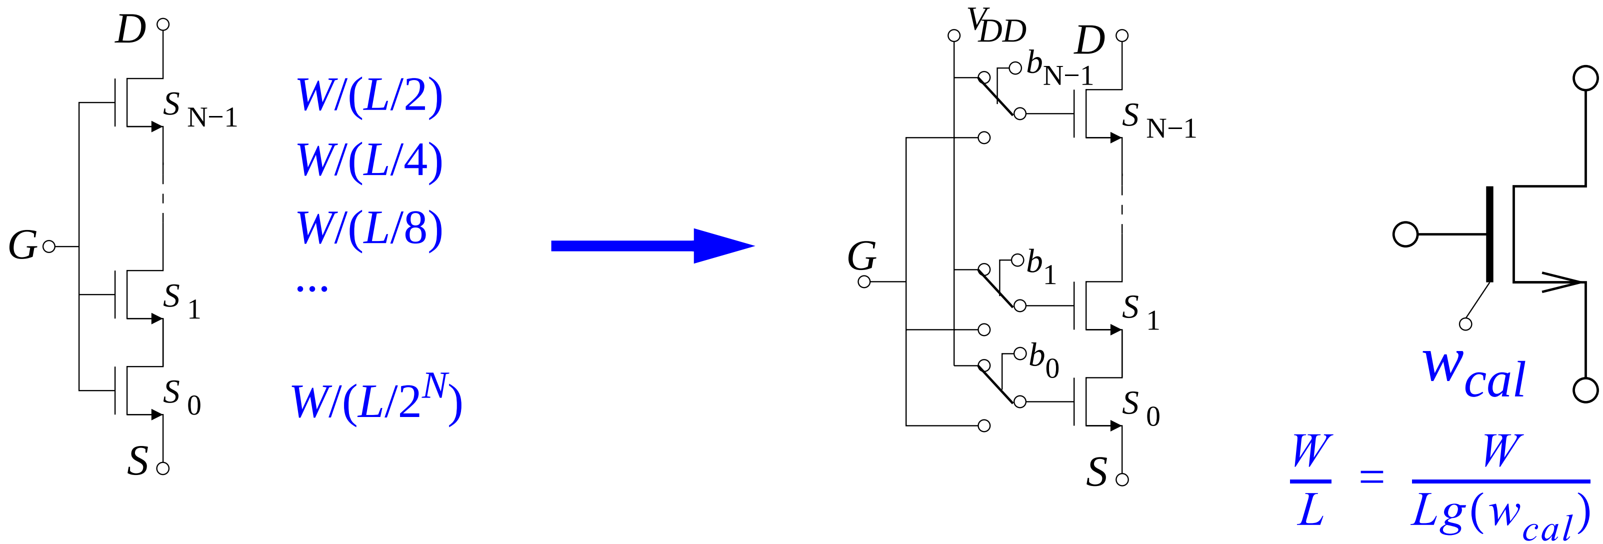
D
G
S
S N−1
S 1
S 0
W/(L/2)
W/(L/4)
W/(L/8)
...
W/(L/2^N)
V DD
b N−1
b 1
b 0
w cal
W/L = W/(Lg(w cal))

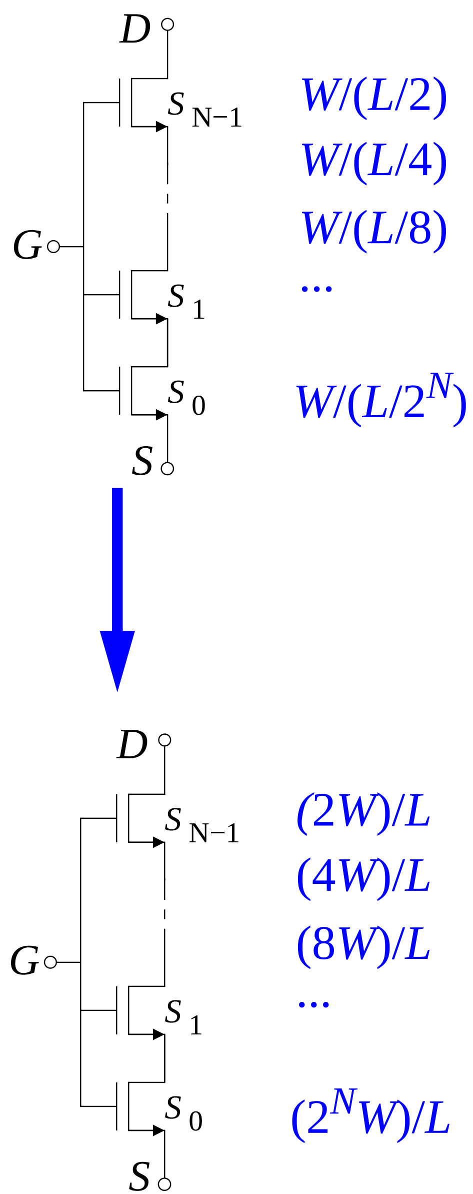

D
G
S
S N−1
S 1
S 0
W/(L/2)
W/(L/4)
W/(L/8)
...
W/(L/2^N)
D
G
S
S N−1
S 1
S 0
(2W)/L
(4W)/L
(8W)/L
...
(2^N W)/L

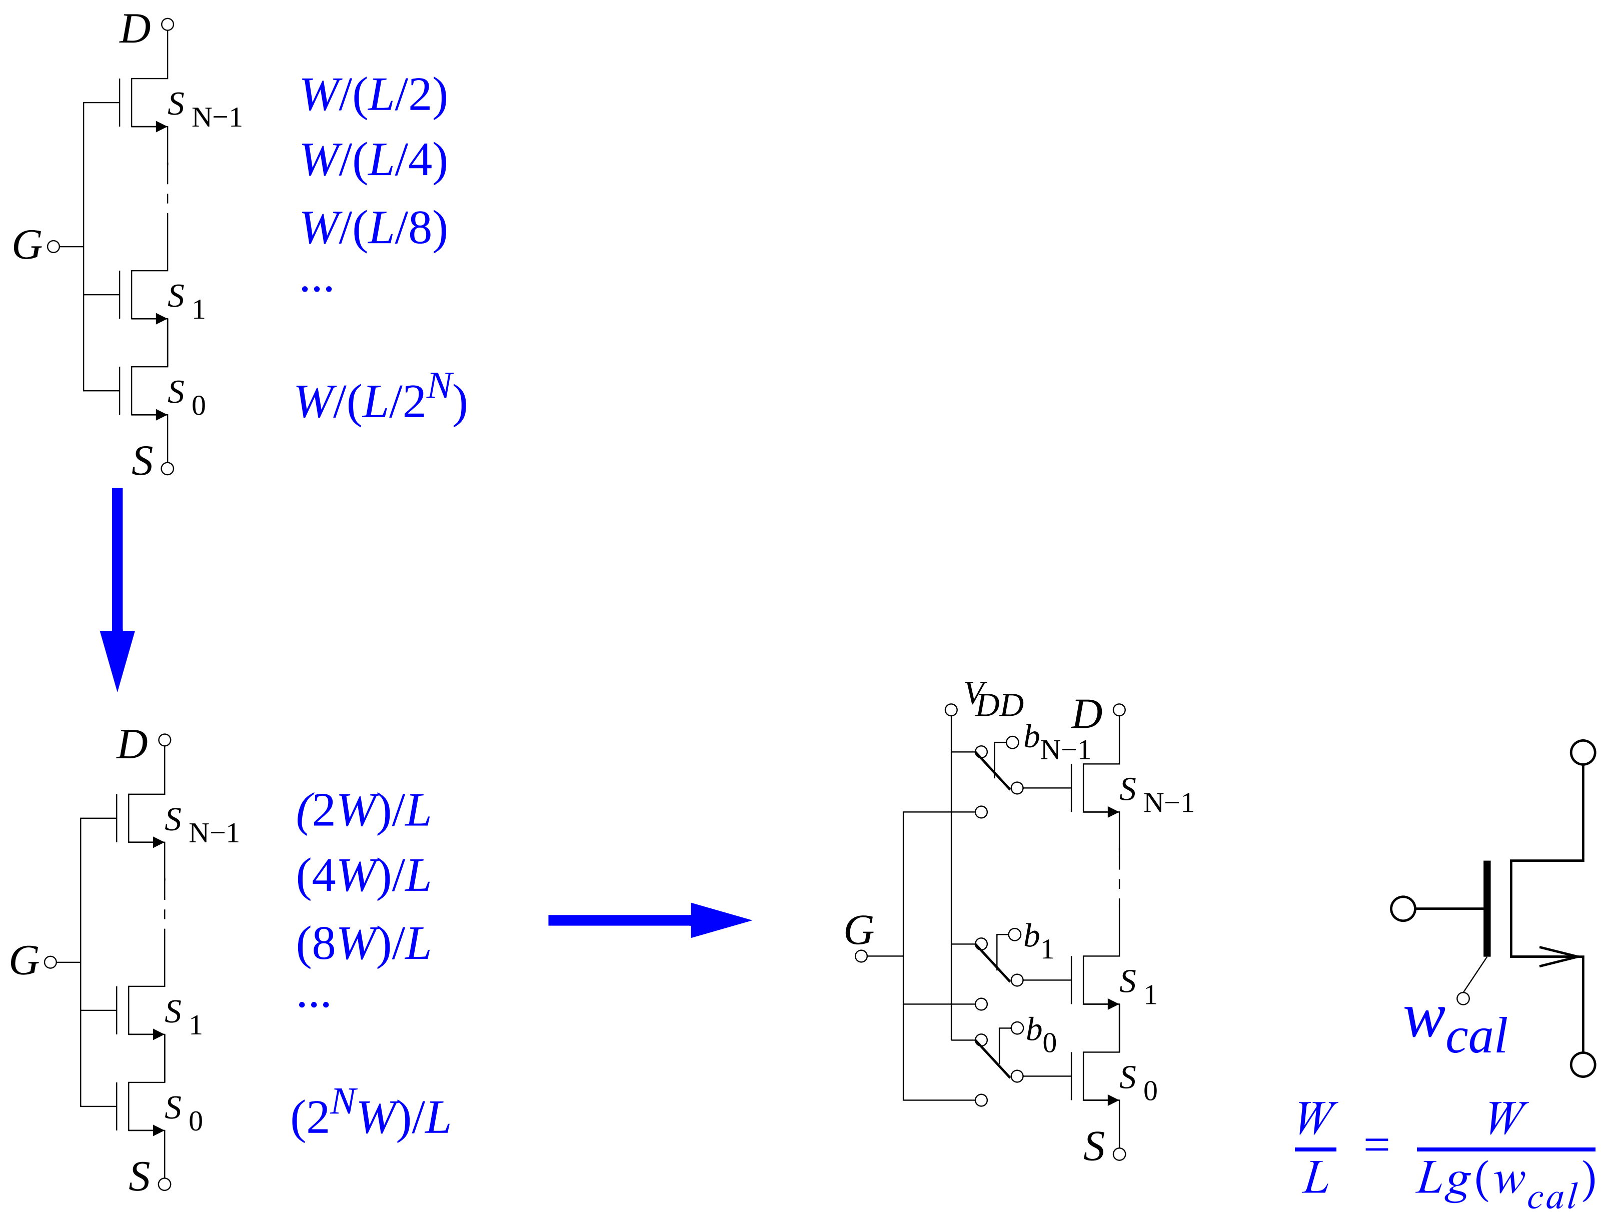
D
G
S N−1
S 1
S 0
S
W/(L/2)
W/(L/4)
W/(L/8)
...
W/(L/2^N)
(2W)/L
(4W)/L
(8W)/L
...
(2^N W)/L
VDD
b N−1
b 1
b 0
w cal
W/L = W/(Lg(w cal))

4-bit Monte Carlo Simulation

sub-pA current mirror

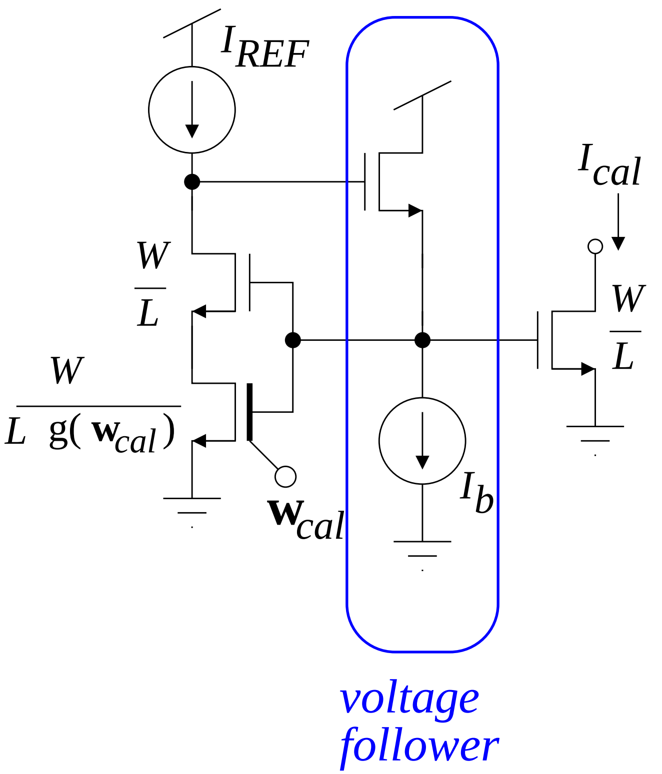

$$\frac{I_{cal}}{I_{REF}} = \frac{W/L}{W/(L + Lg(w_{cal}))}$$

$$I_{cal} = I_{REF} \times (1 + g(w_{cal}))$$

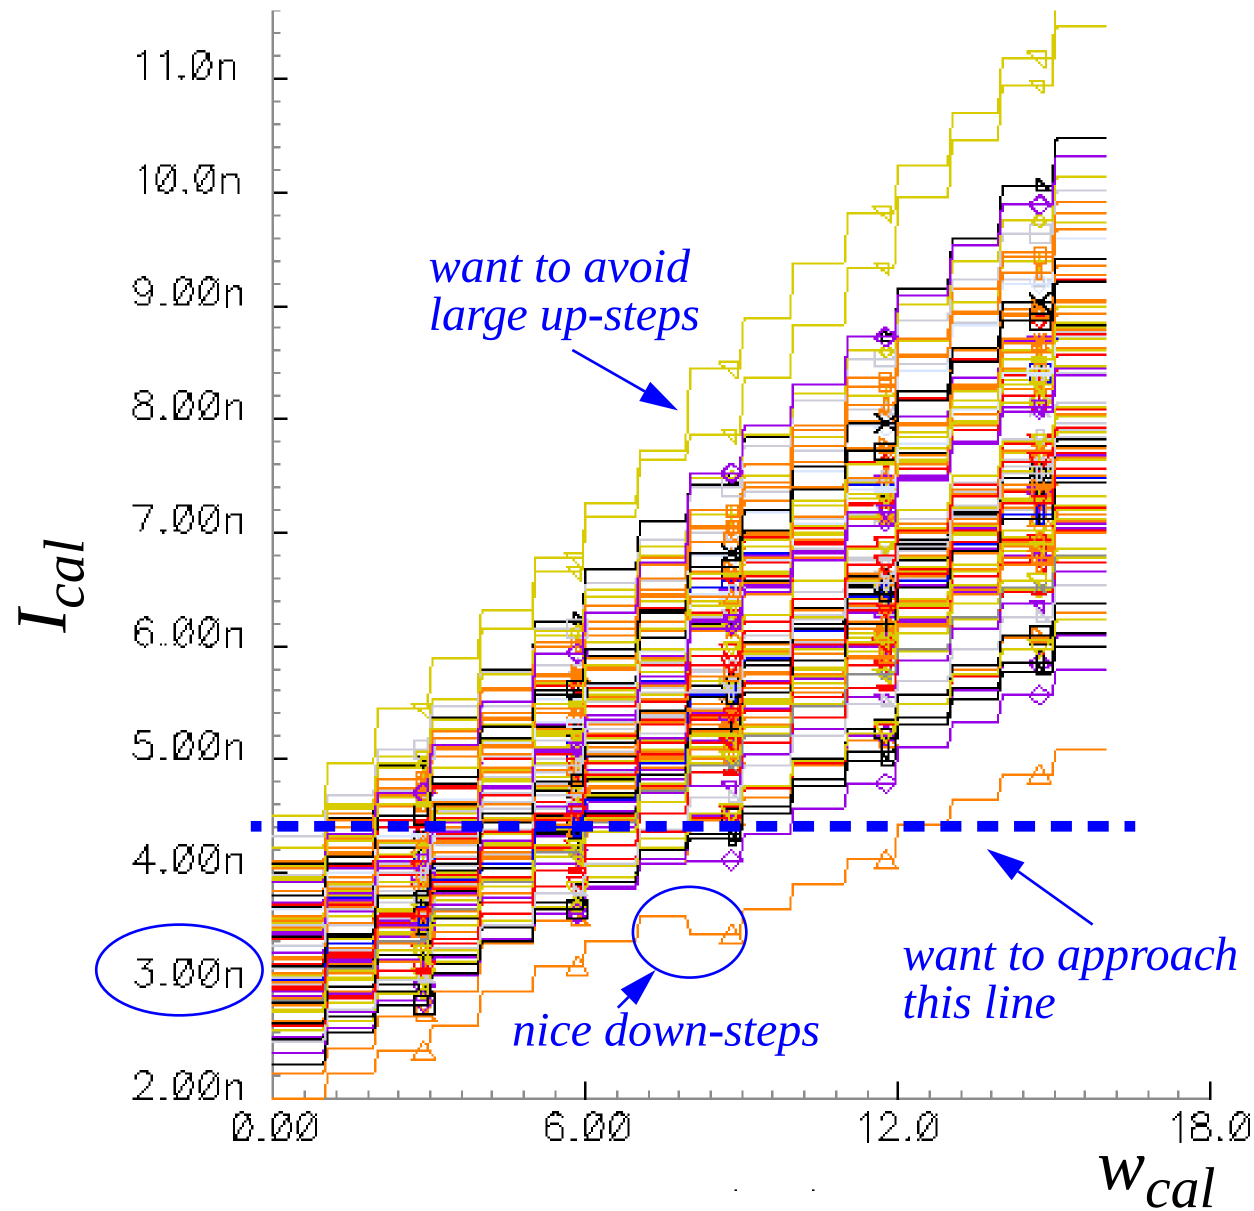

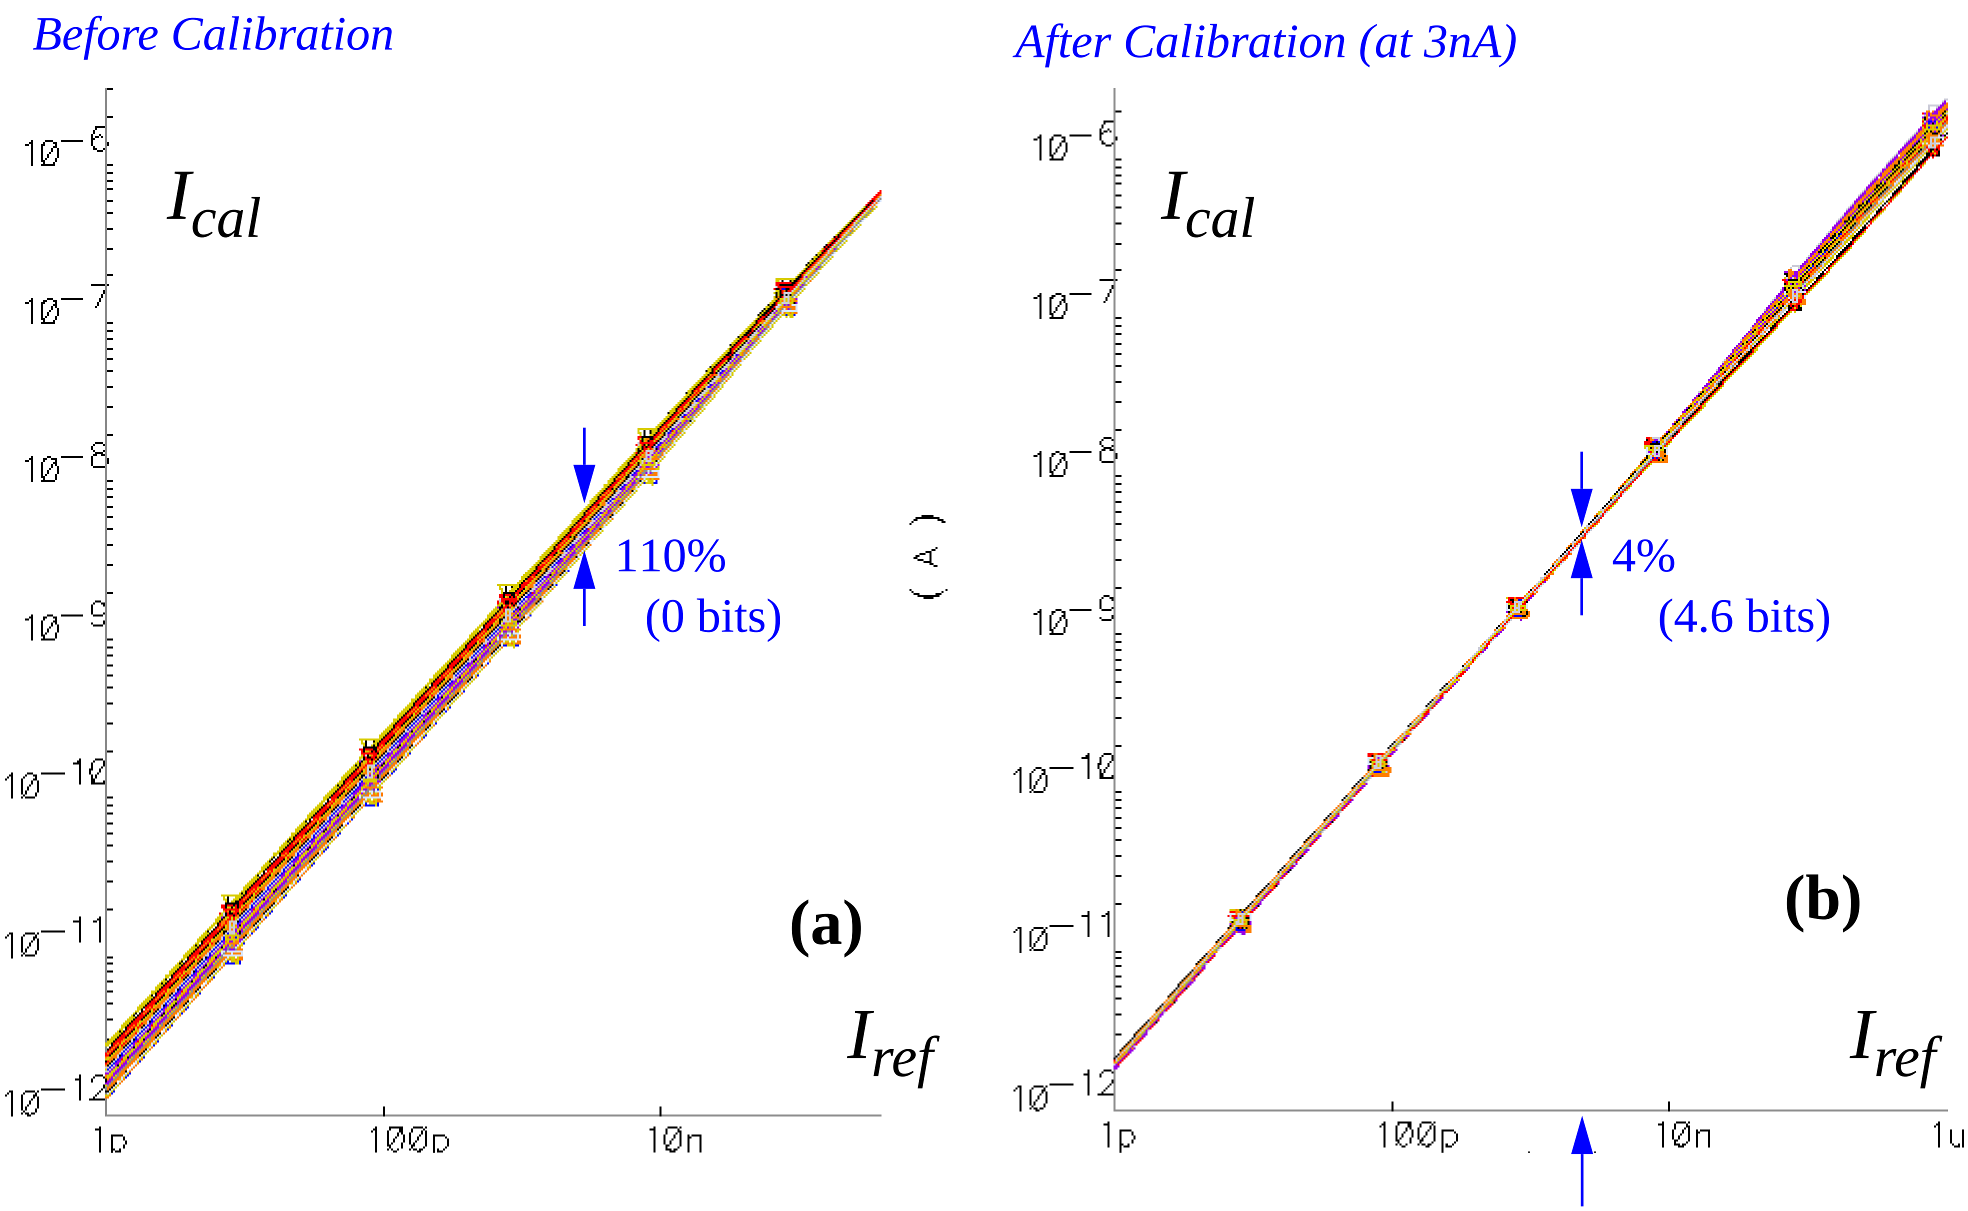

Before Calibration
After Calibration (at 3nA)
I_cal
110%
(0 bits)
(a)
I_ref
I_cal
4%
(4.6 bits)
(b)
I_ref
(A)
10^-6
10^-7
10^-8
10^-9
10^-10
10^-11
10^-12
1p
100p
10n
1u

- we don’t need nice precise stairs, but good coverage
- we like down-steps
- we like randomness
- we use same $W = 2\mu m$ and $L = \{3.0, 1.8, 1.8, 1.0, 0.7\}$ for a 5-bit digi-MOS

- we don't need nice precise stairs, but good coverage
- we like down-steps
- we like randomness
- we use same $W = 2\mu m$ and $L = \{3.0, 1.8, 1.8, 1.0, 0.7\}$ for a 5-bit digi-MOS

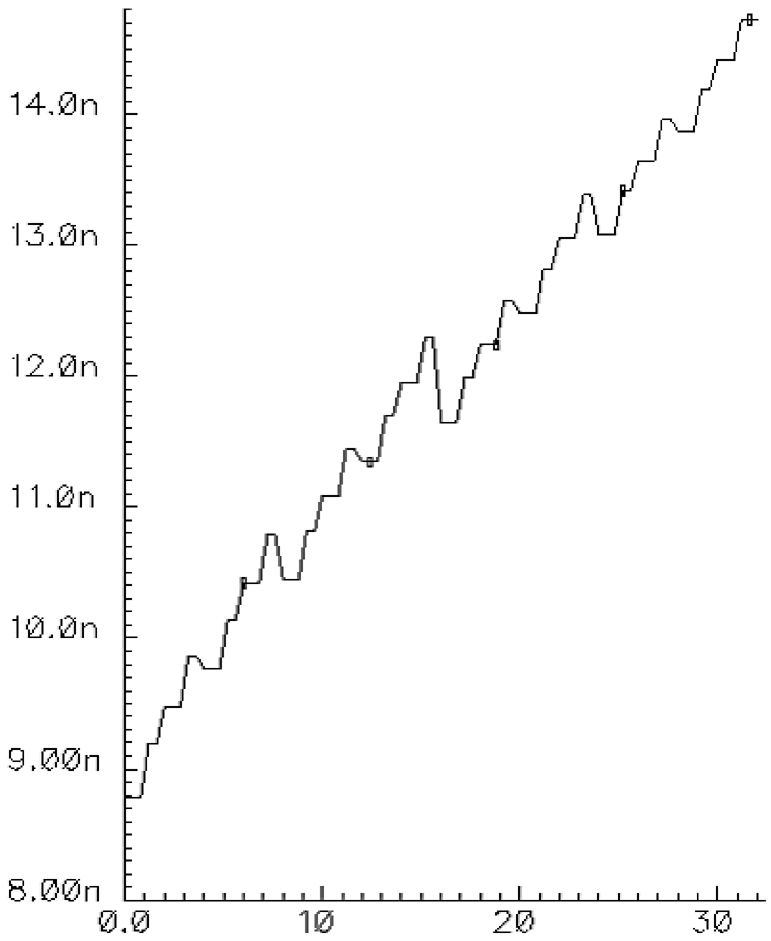

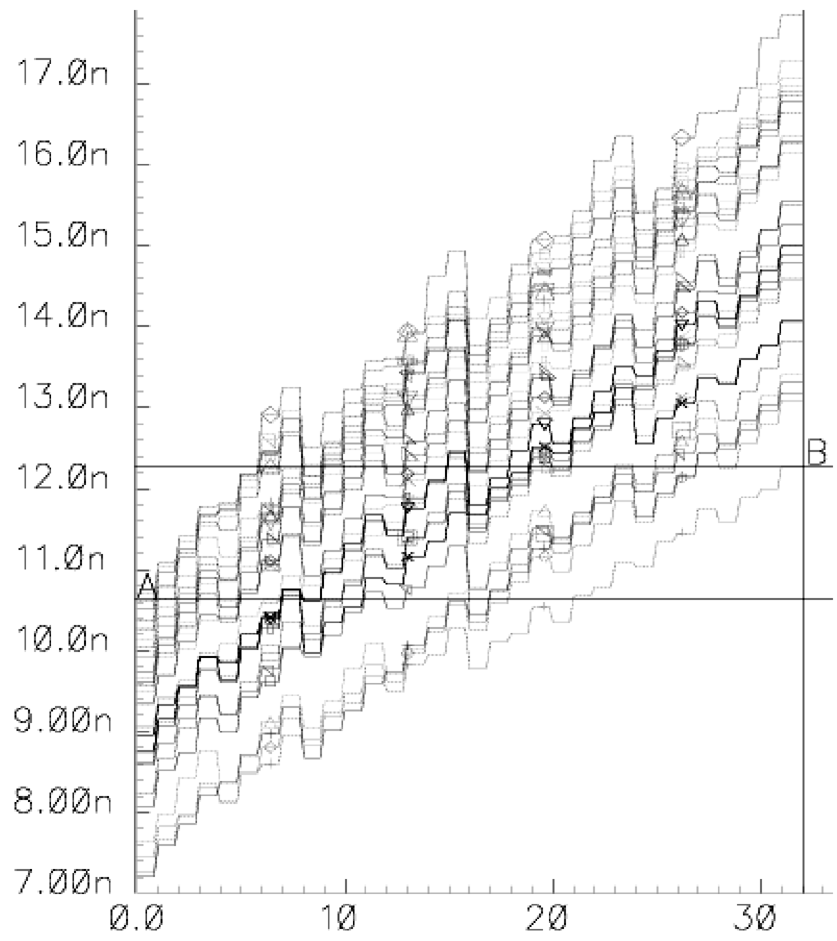

And we want two additional features:

- FEATURE-1: no recalibration when changing operating current
- FEATURE-2: take maximum advantage of calibration range: B>A but B~A

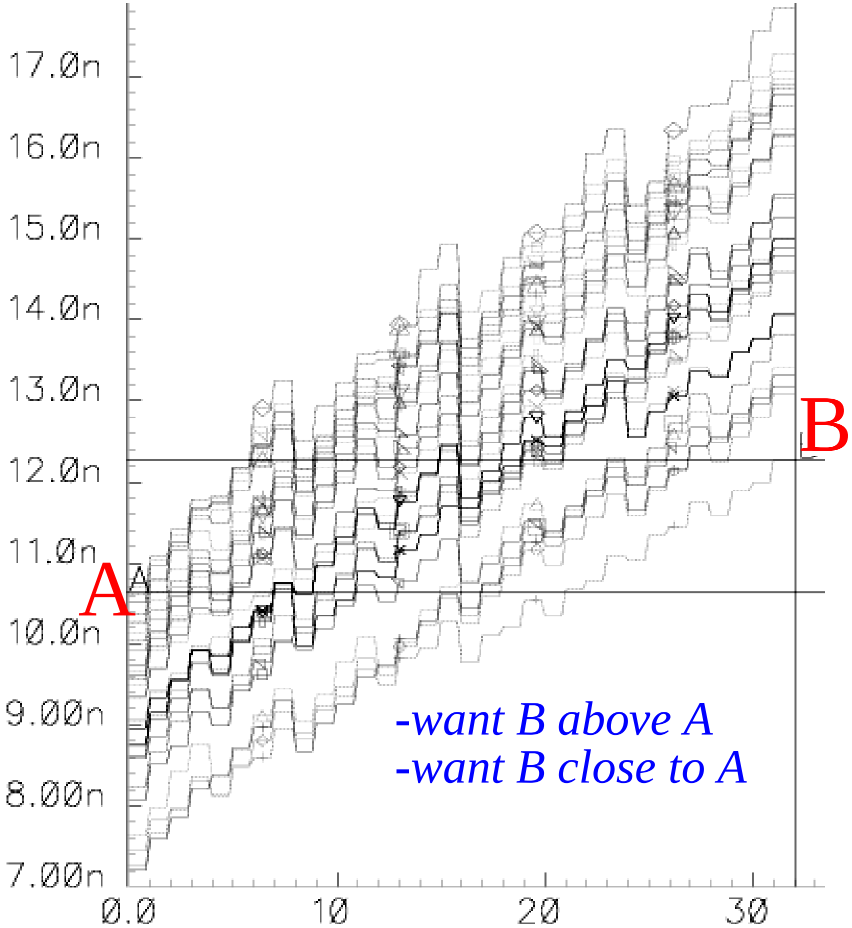

FEATURE-1: no recalibration

- for simple current mirror

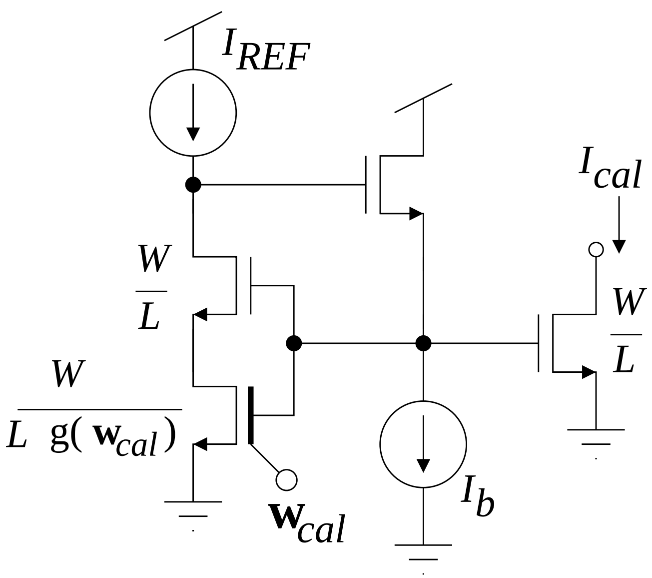

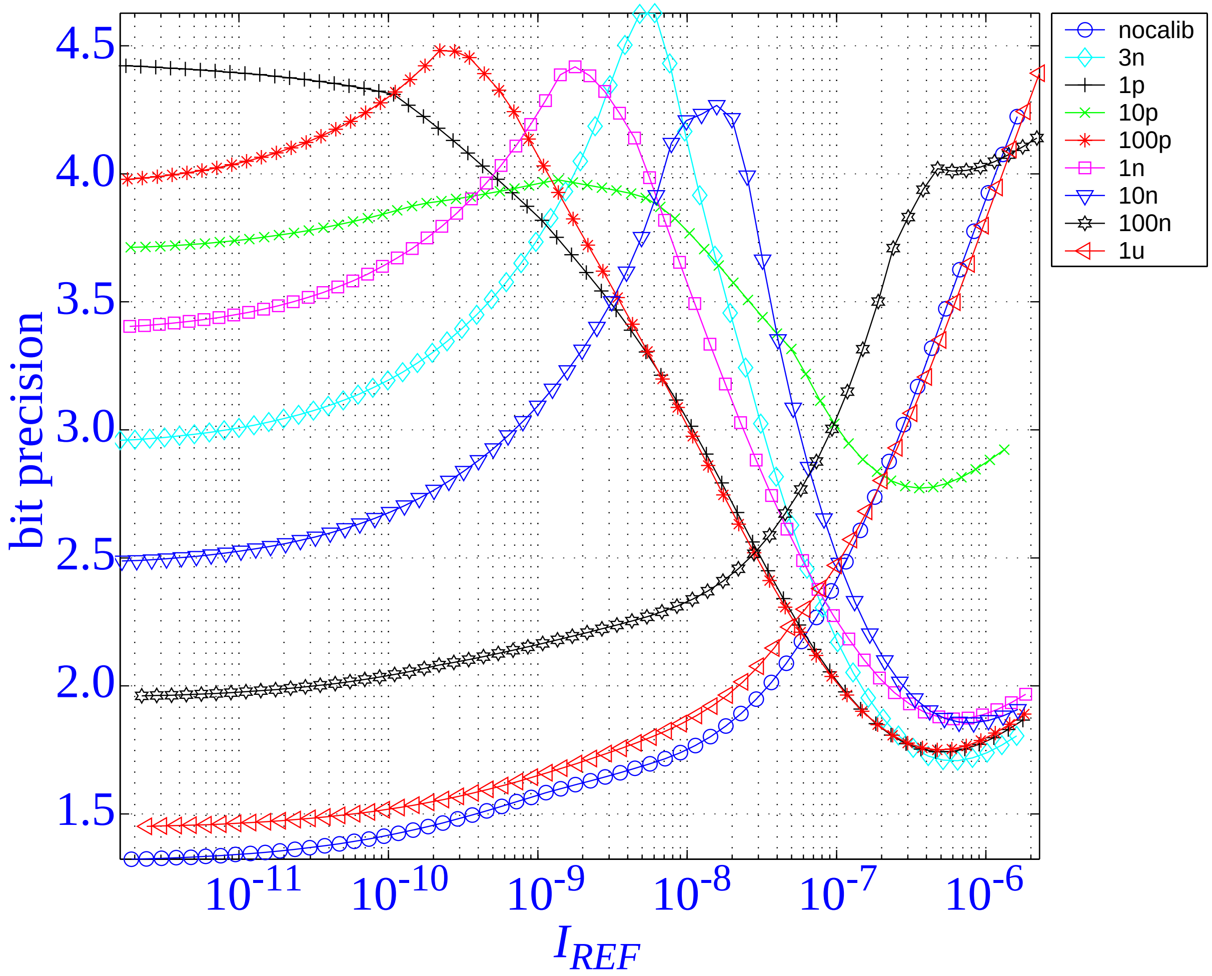

FEATURE-1: no recalibration

- by adding peripheral translinear tuning:

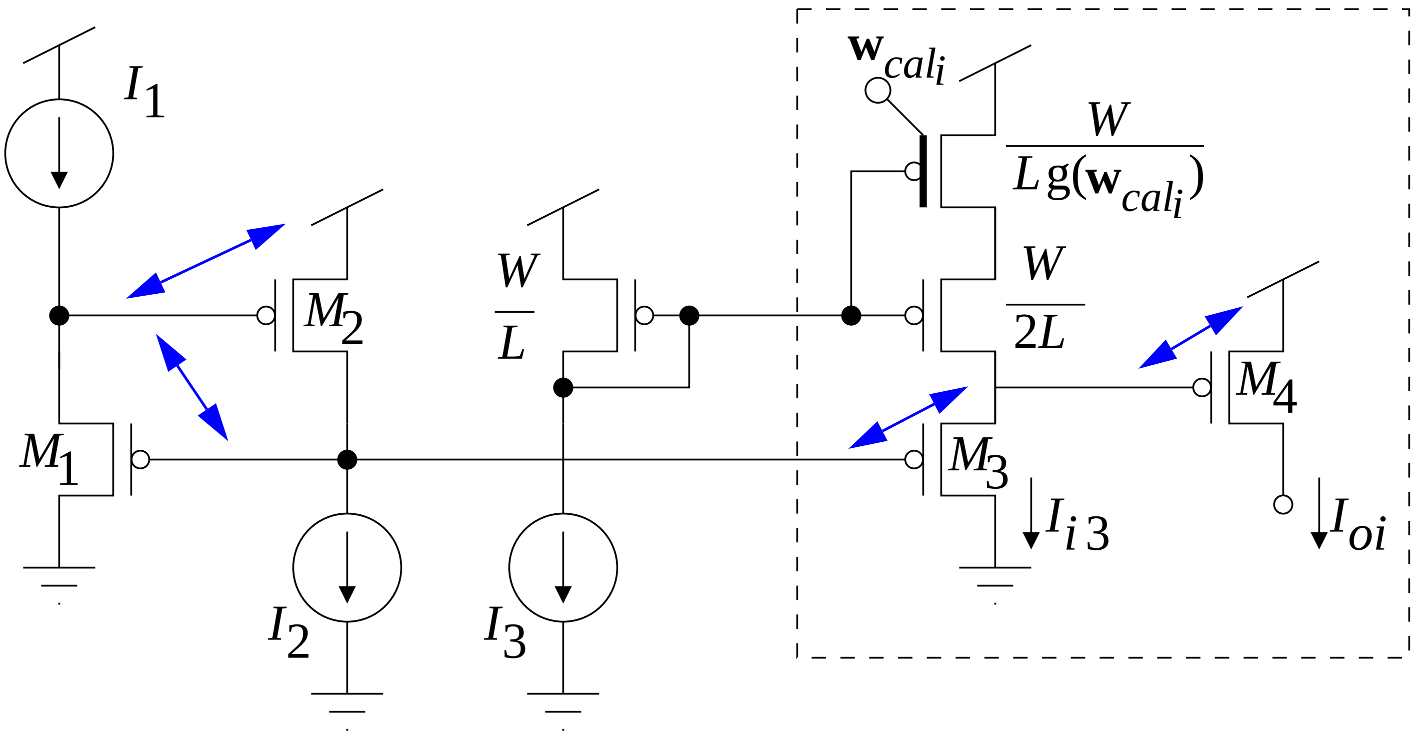

$$I_1 I_2 = I_{i3} I_4 \;;\; I_{i3} = I_3 / (2 + g(w_{cal}))$$

$$I_{oi} = \frac{I_1 I_2}{I_3}(2 + g(w_{cal}))$$

- I_3 is constant, so currents through branch I_{i3} is constant
- M_1 has similar bias condition than M_3, so I_1 is also kept constant
- M_2 has similar bias condition than M_4, so I_{oi} is scaled by changing only I_2

FEATURE-2: approach A and B

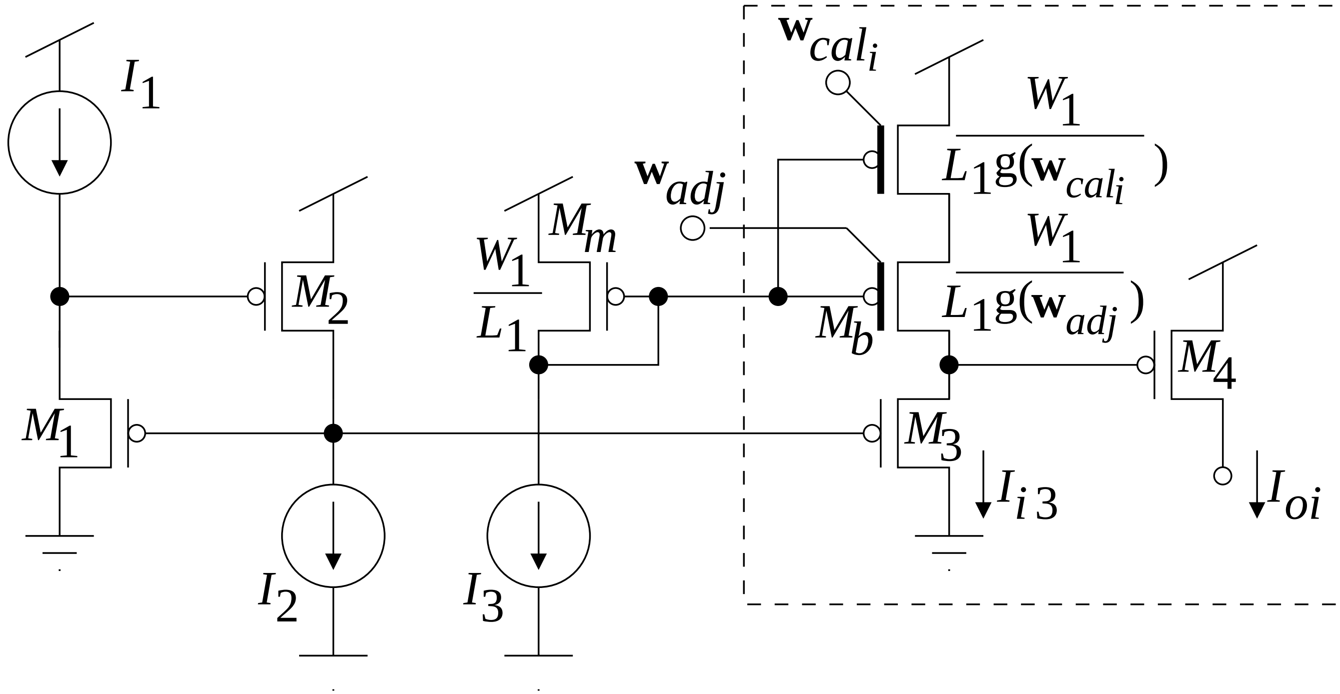

$$I_{oi} = \frac{I_1 I_2}{I_3}(g(w_{adj}) + g(w_{cal}))$$

global for array

local for each pixel

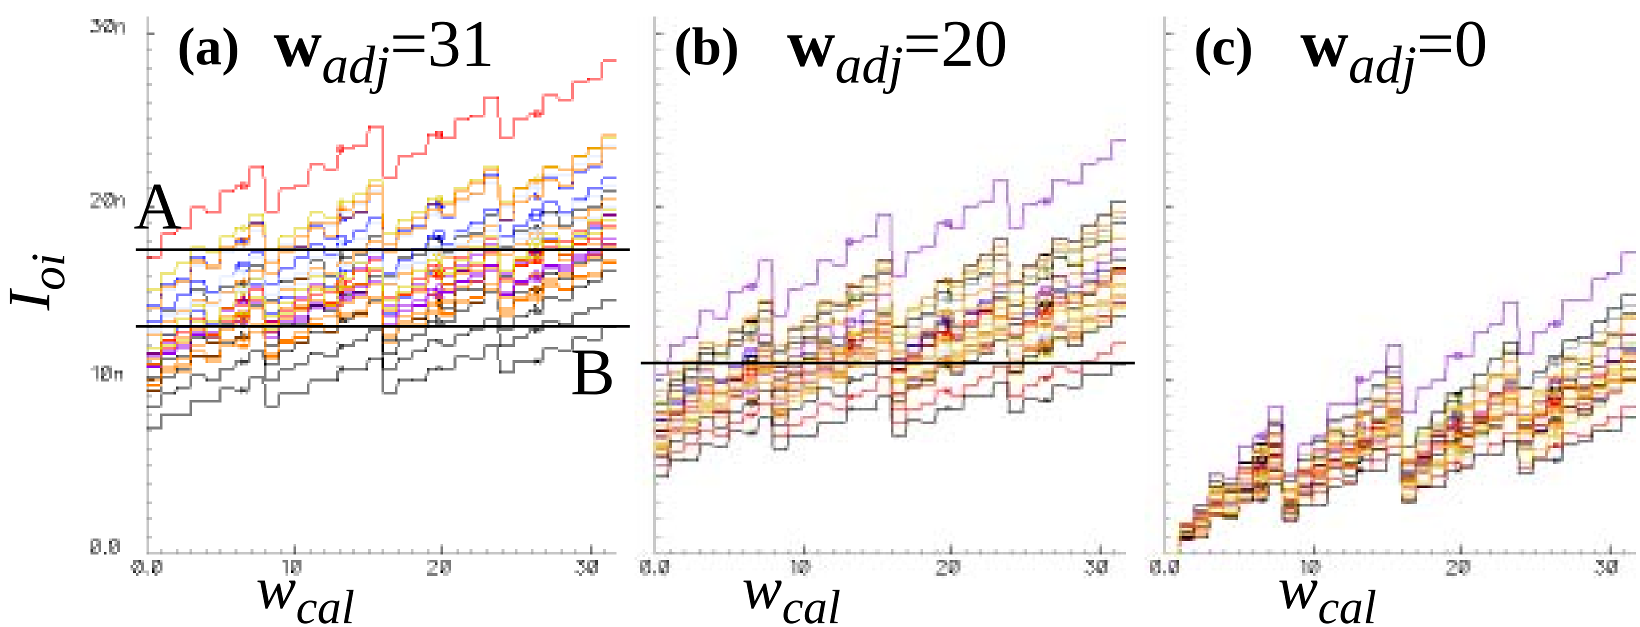

Experimental prototype CMOS 0.35um

- single current source calibrated at 10nA

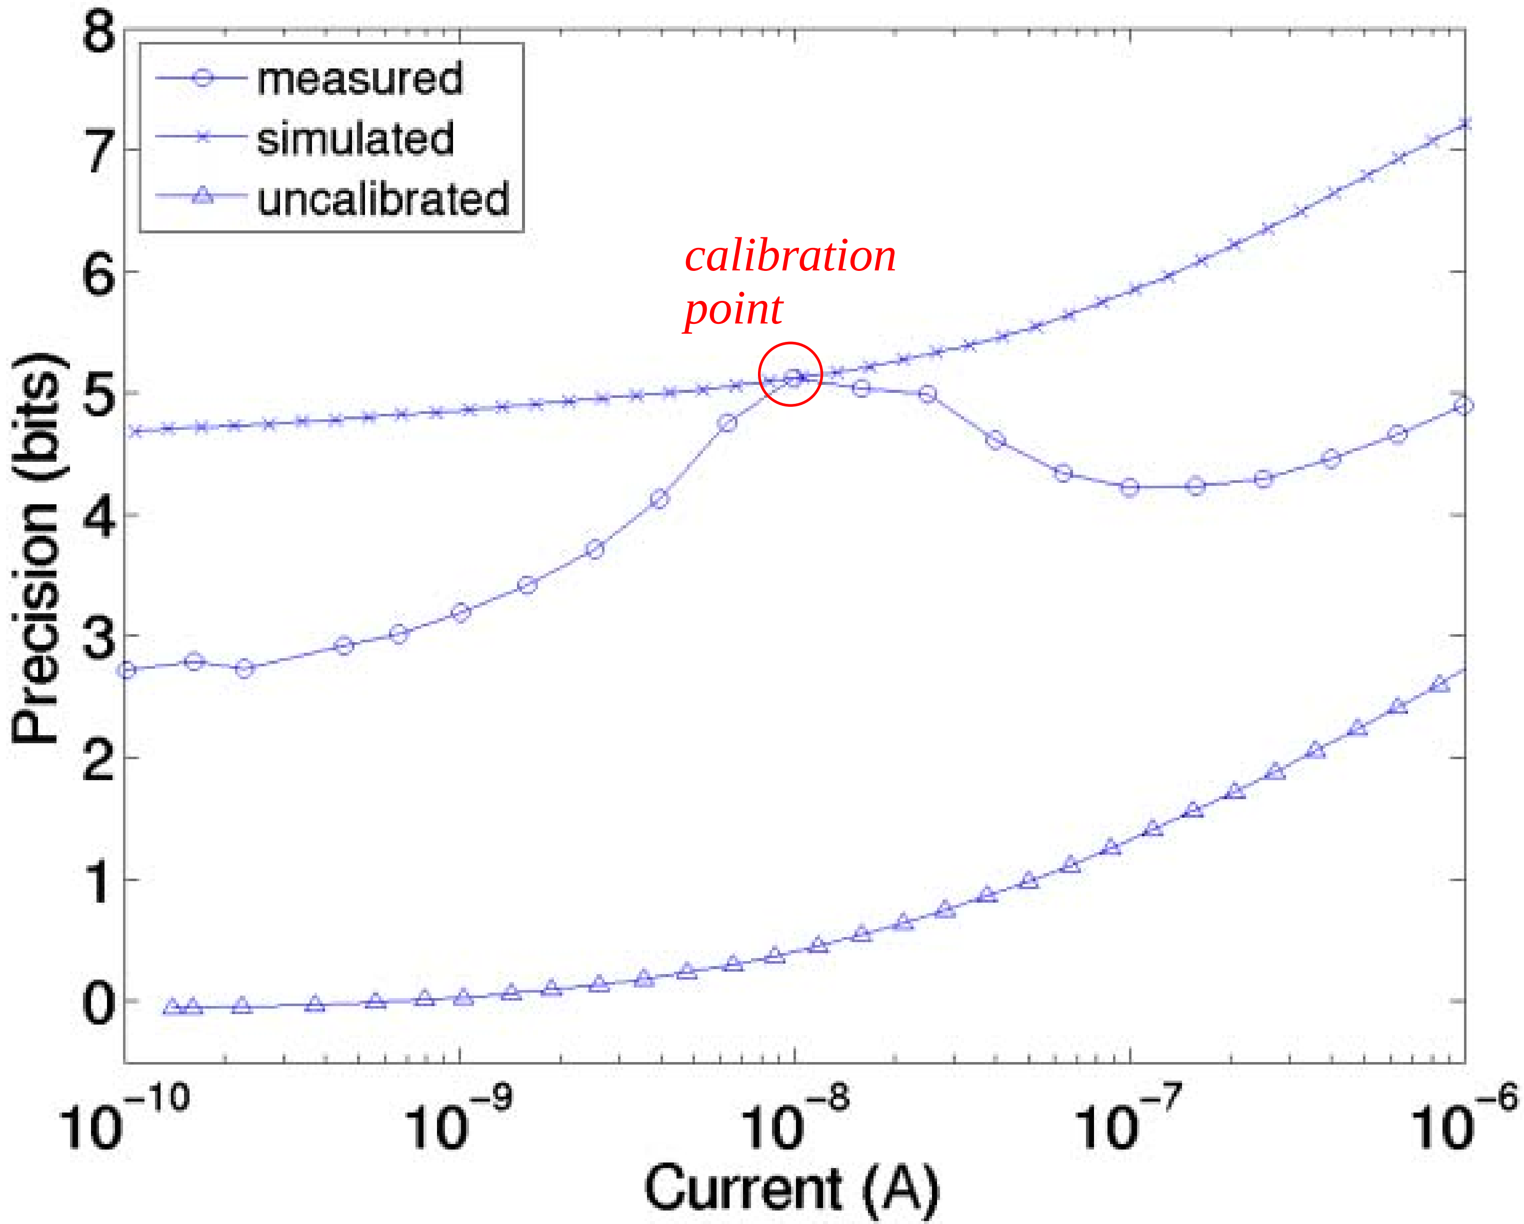

Experimental prototype CMOS 0.35um

- single current source calibrated at 10nA

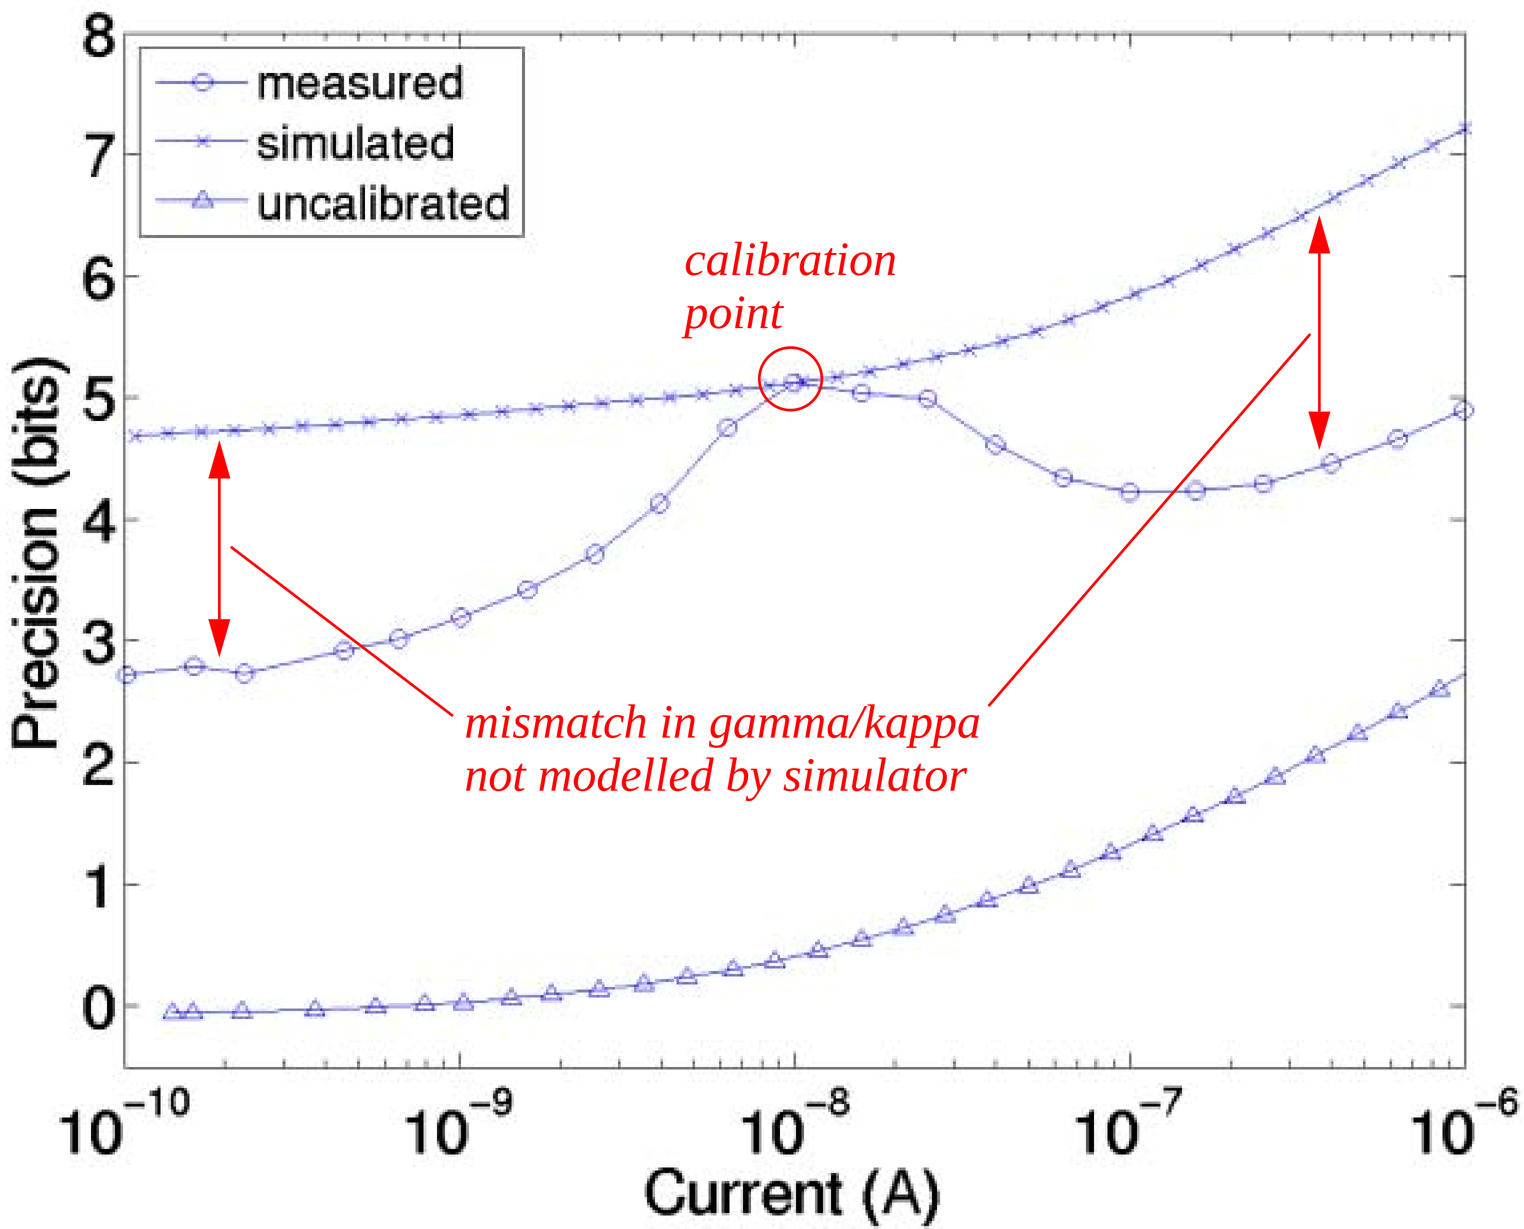

Experimental prototype CMOS 0.35um

- DAC: five current sources calibrated at 10nA, 5nA, 2.5nA, 1.25nA, 625pA and 16ºC

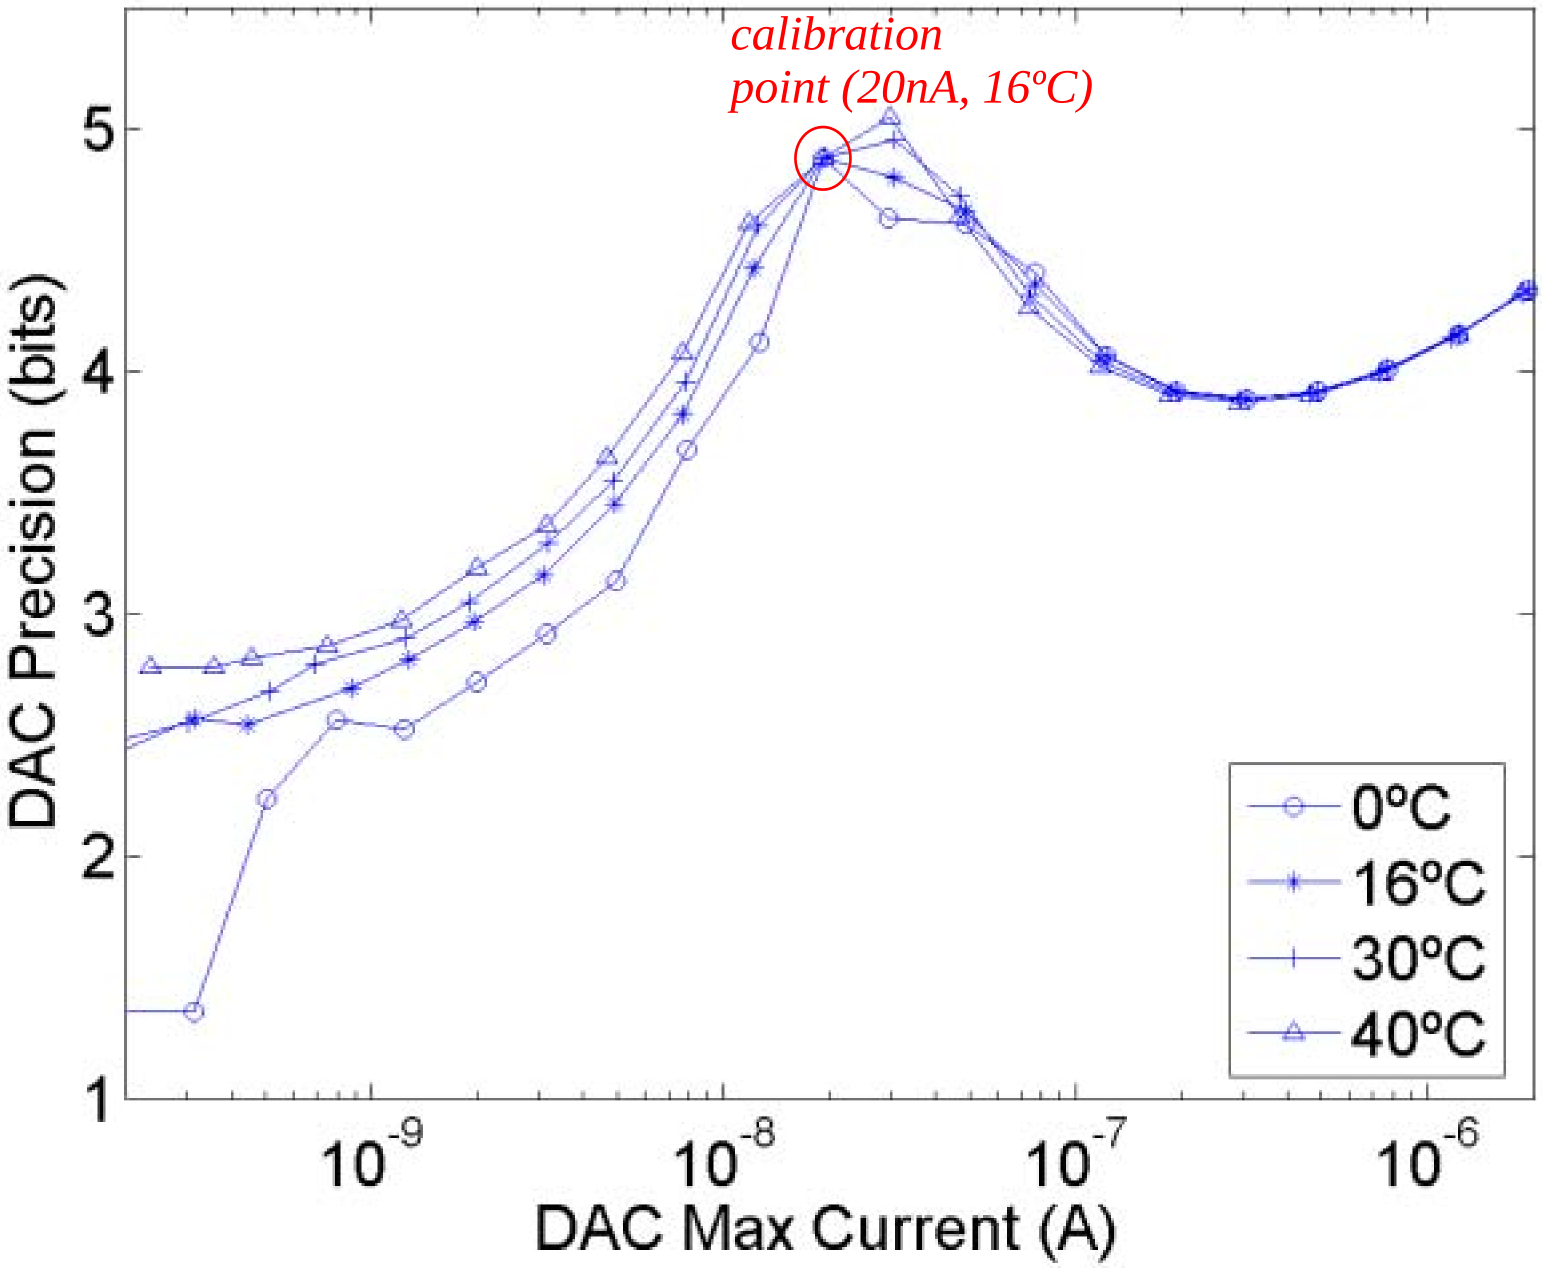

Another Translinear Tuning Circuit [TCAS-II, in Press]

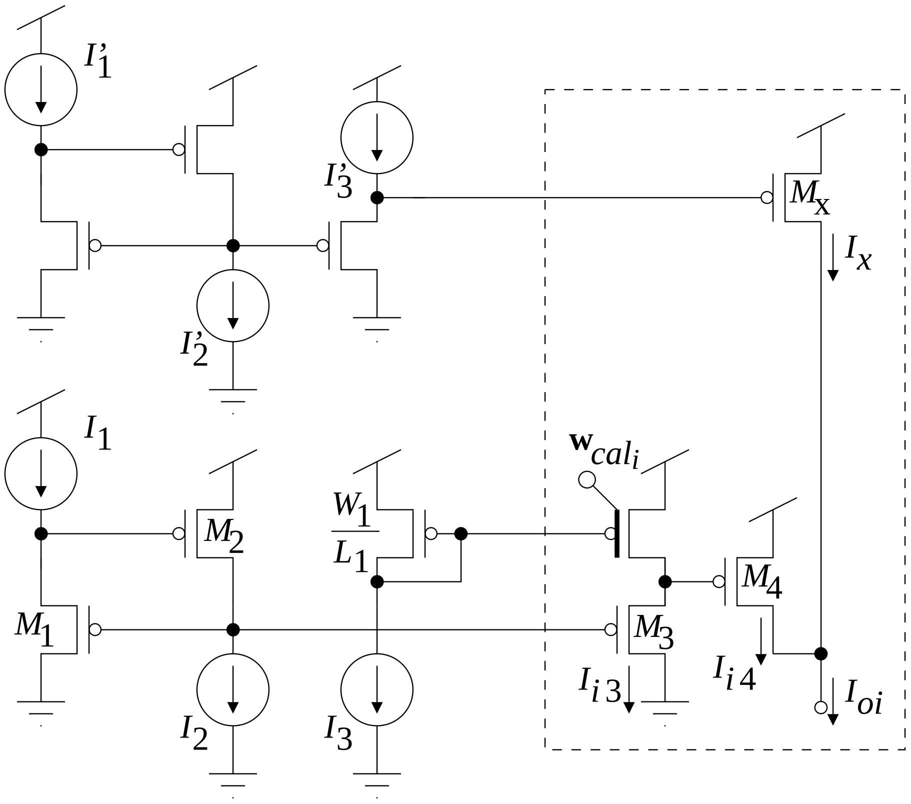

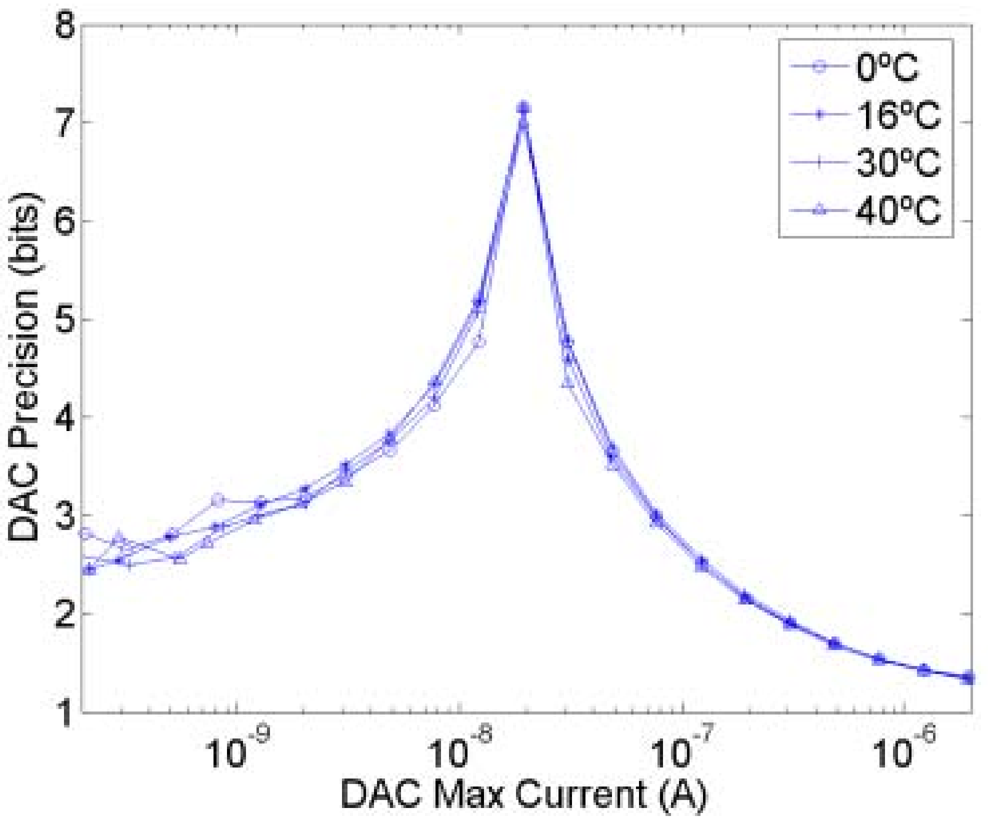

- achieves higher precision (7-bit) using a 5-bit circuit
- degrades more rapidly when changing bias conditions

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems

- Some useful circuits:
  - calibration
  - LVDS interface

- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip

- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR

- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

The bit serial LVDS AER interface

- Several options are possible:
  - Transmitting data and clock by different physical paths.
  - Recovering the clock using a PLL-based circuit.
  - Extracting the clock from the receiver data (e.g. using a Manchester coding).

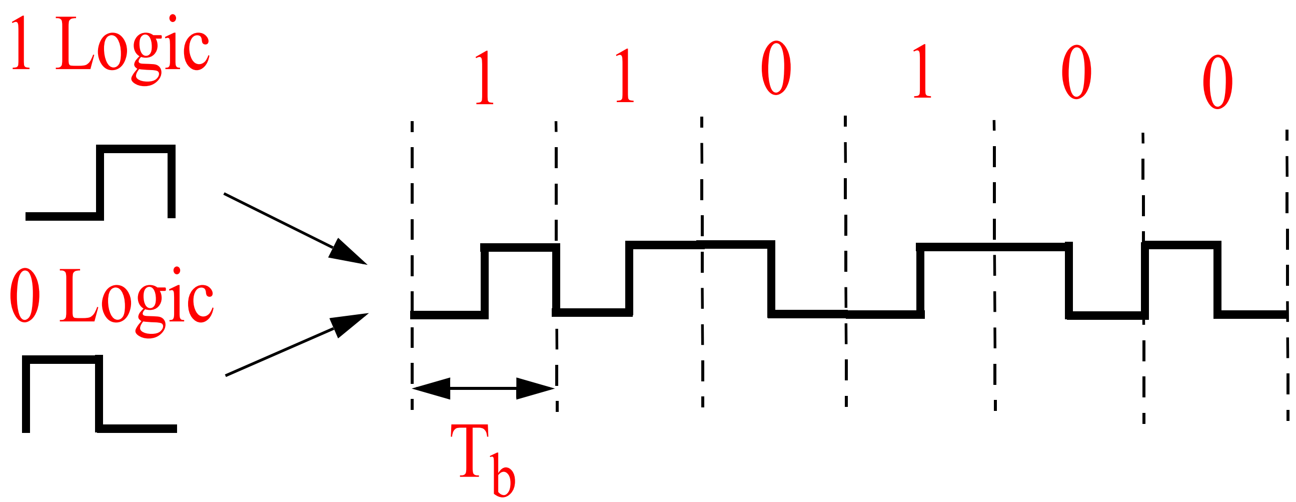

A transition takes place every T_b !

The bit serial LVDS AER interface

- In AER links we will need:
  - Keeping the receiver synchronized in the silent periods.
  - Detecting a new address start.
  - Implementing a fast and robust synchronization scheme.

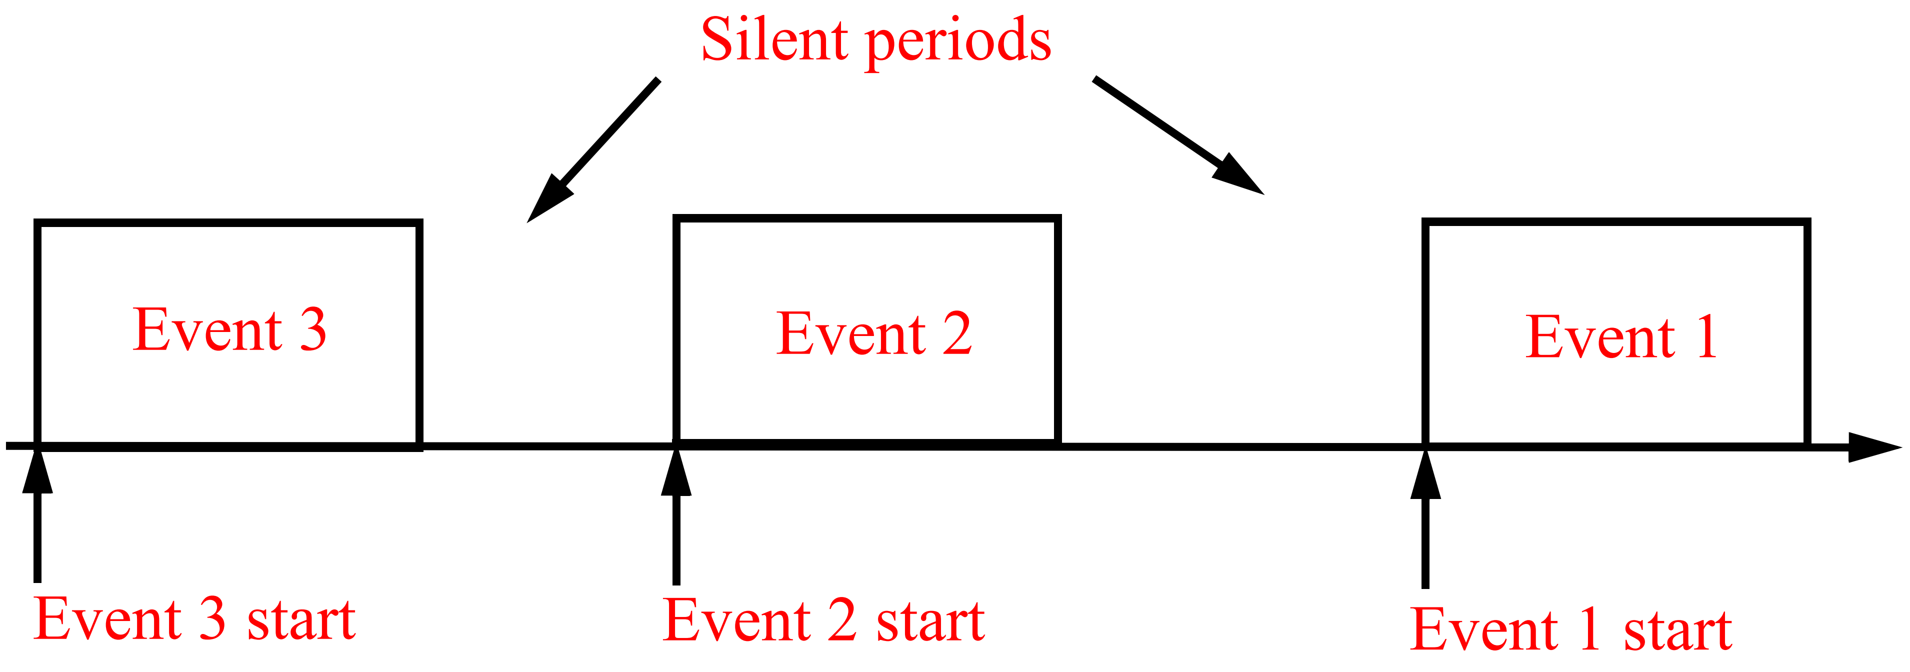

The bit serial LVDS AER interface (IV)

- Fast synchronization is a must in AER links because the events are generated in an asynchronous way.
- A Manchester coding scheme allows the receiver to recover the clock directly from the data flow.

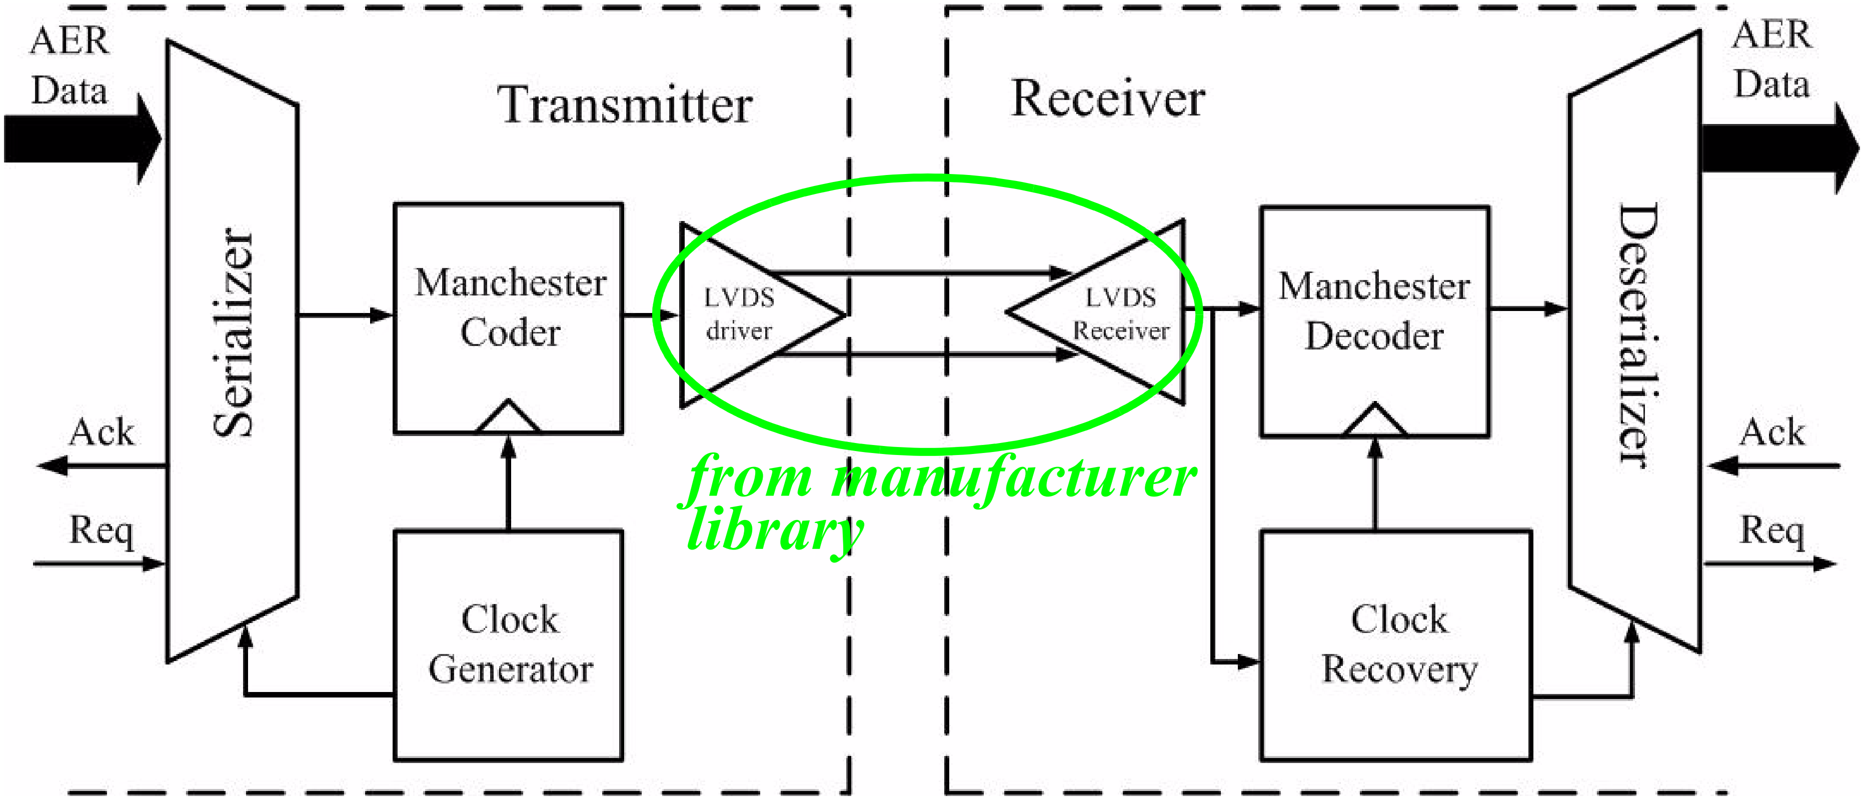

The bit serial LVDS AER interface (V).

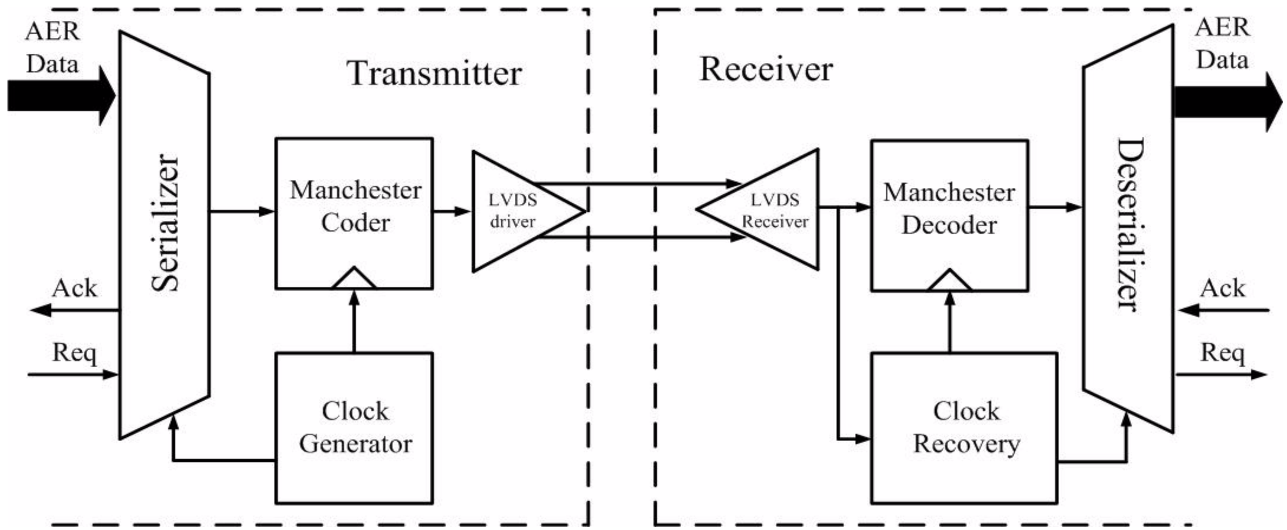

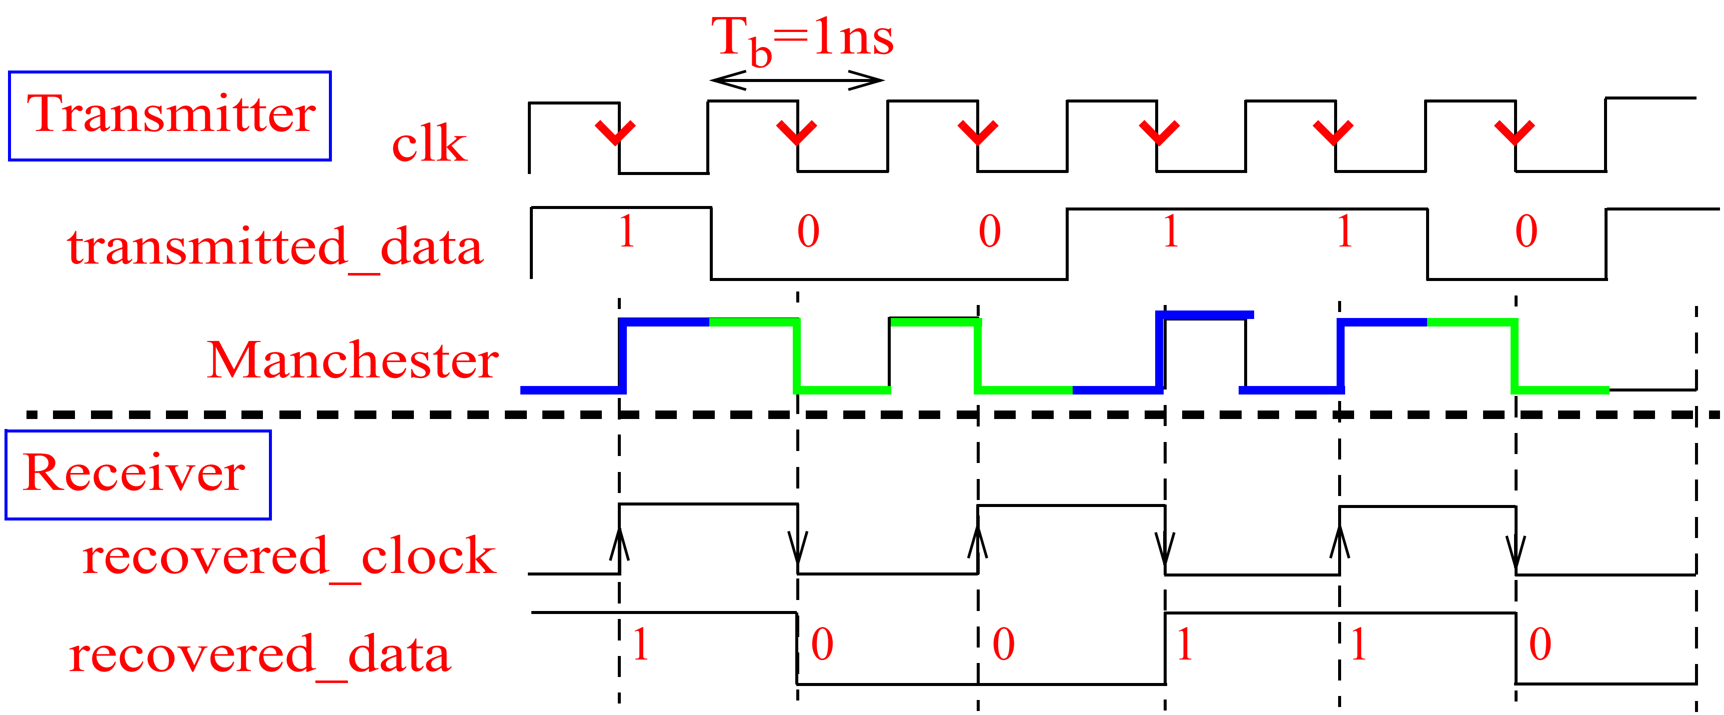

Transmitter circuit

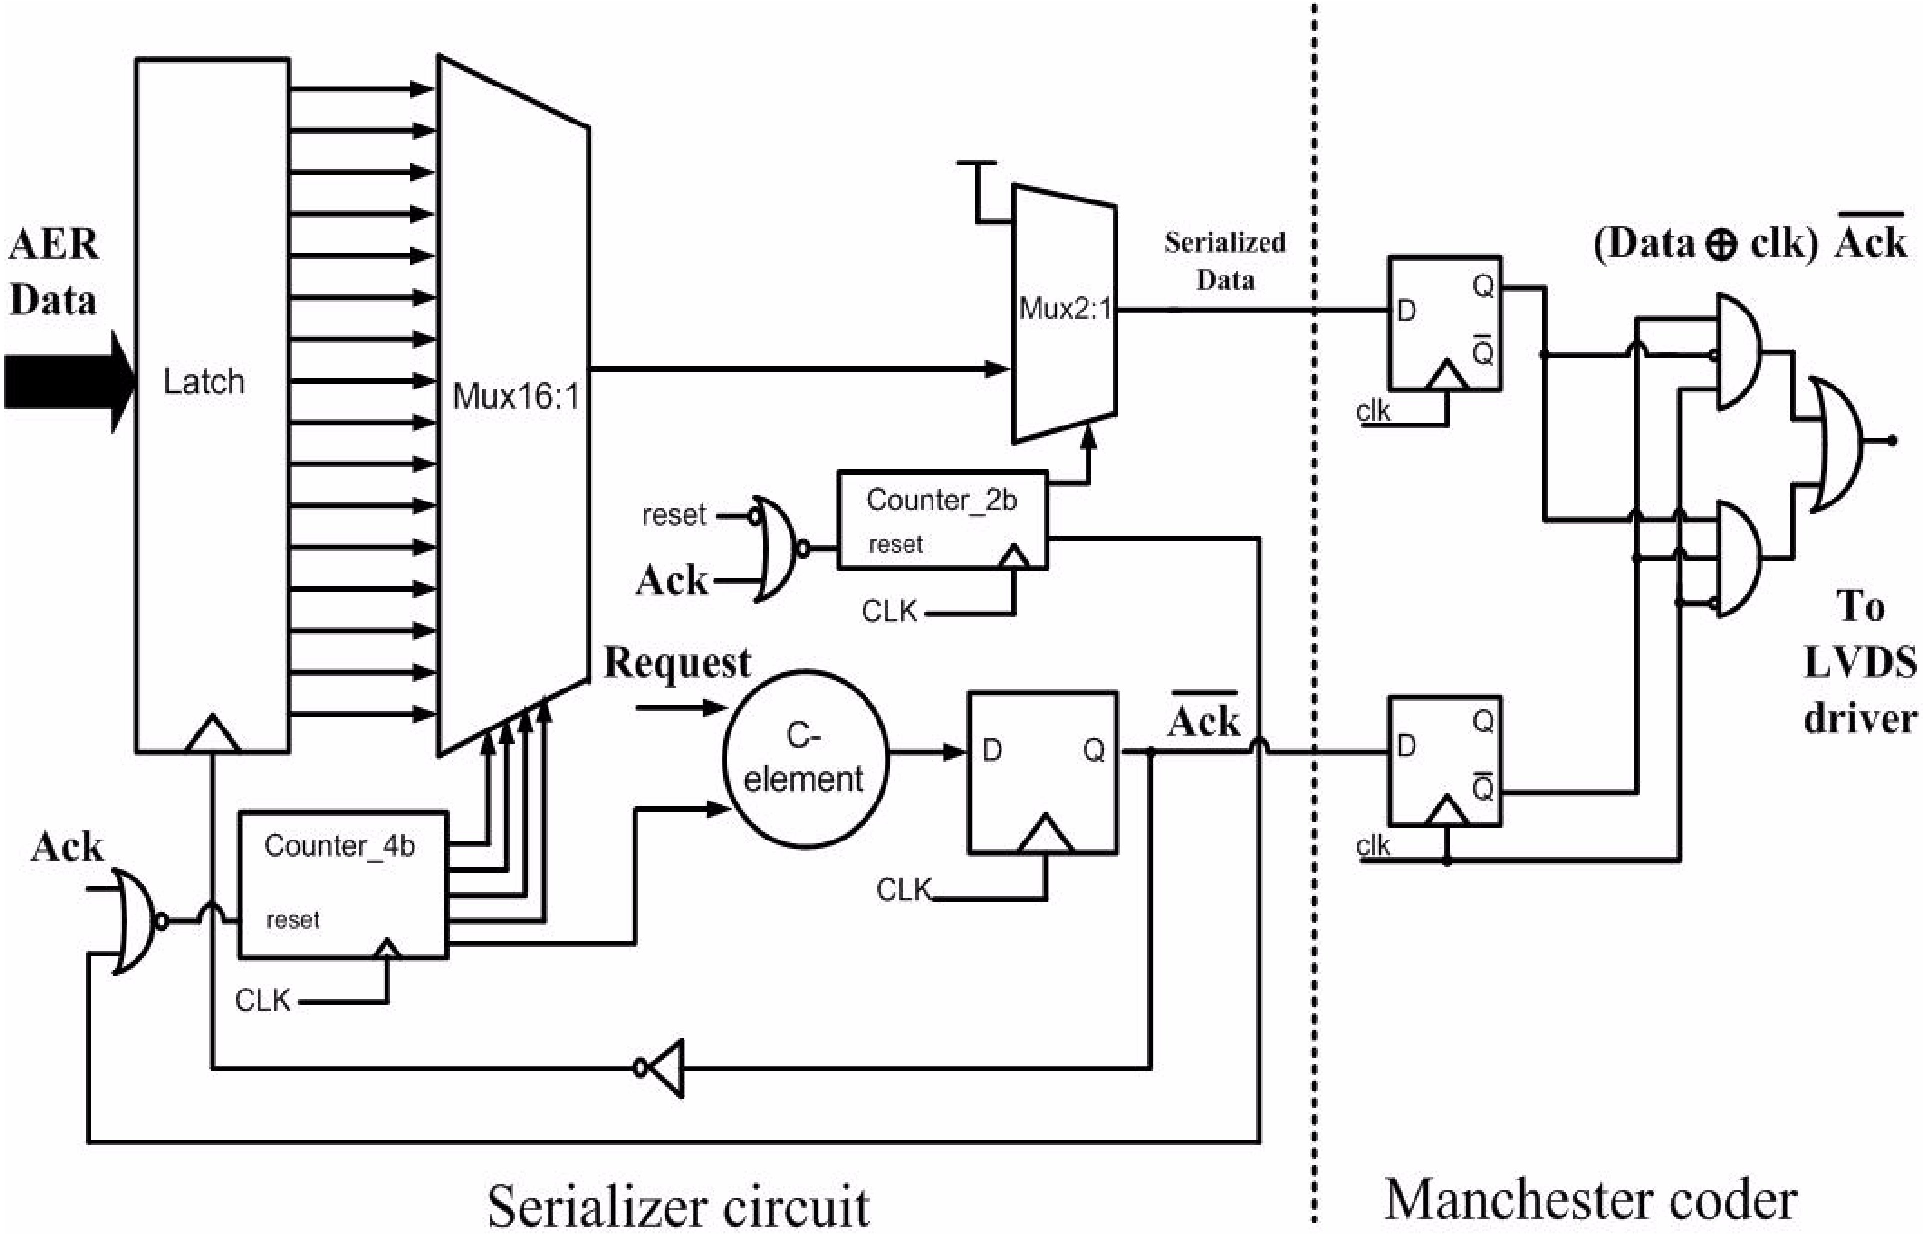

Receiver circuit

- The only requirement for the CDR design is that five delay elements must introduce a delay between $T_b/2$ and T_b.

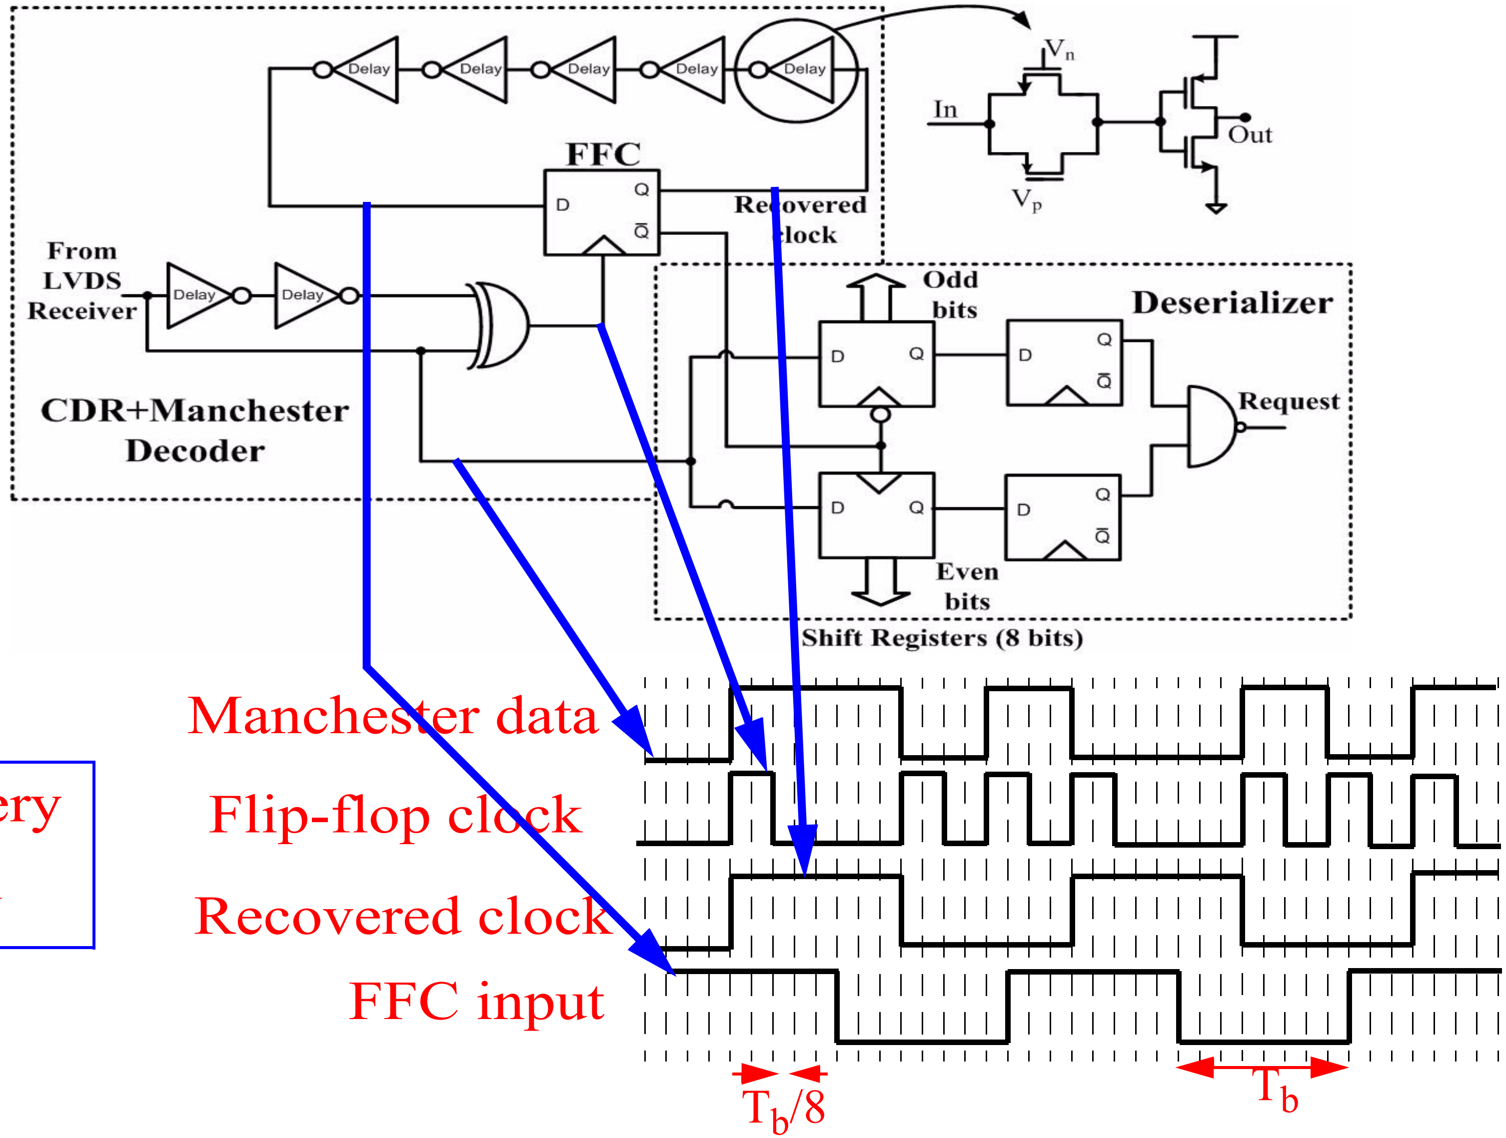

Clock recovery cronogram

Receiver circuitry

- A Delay Locked Loop is used to fix the delay introduced by the inverters. The phase difference between the reference clock and a 360º-delayed version of it is compared and the delay elements control voltage is changed depending on the phase error.

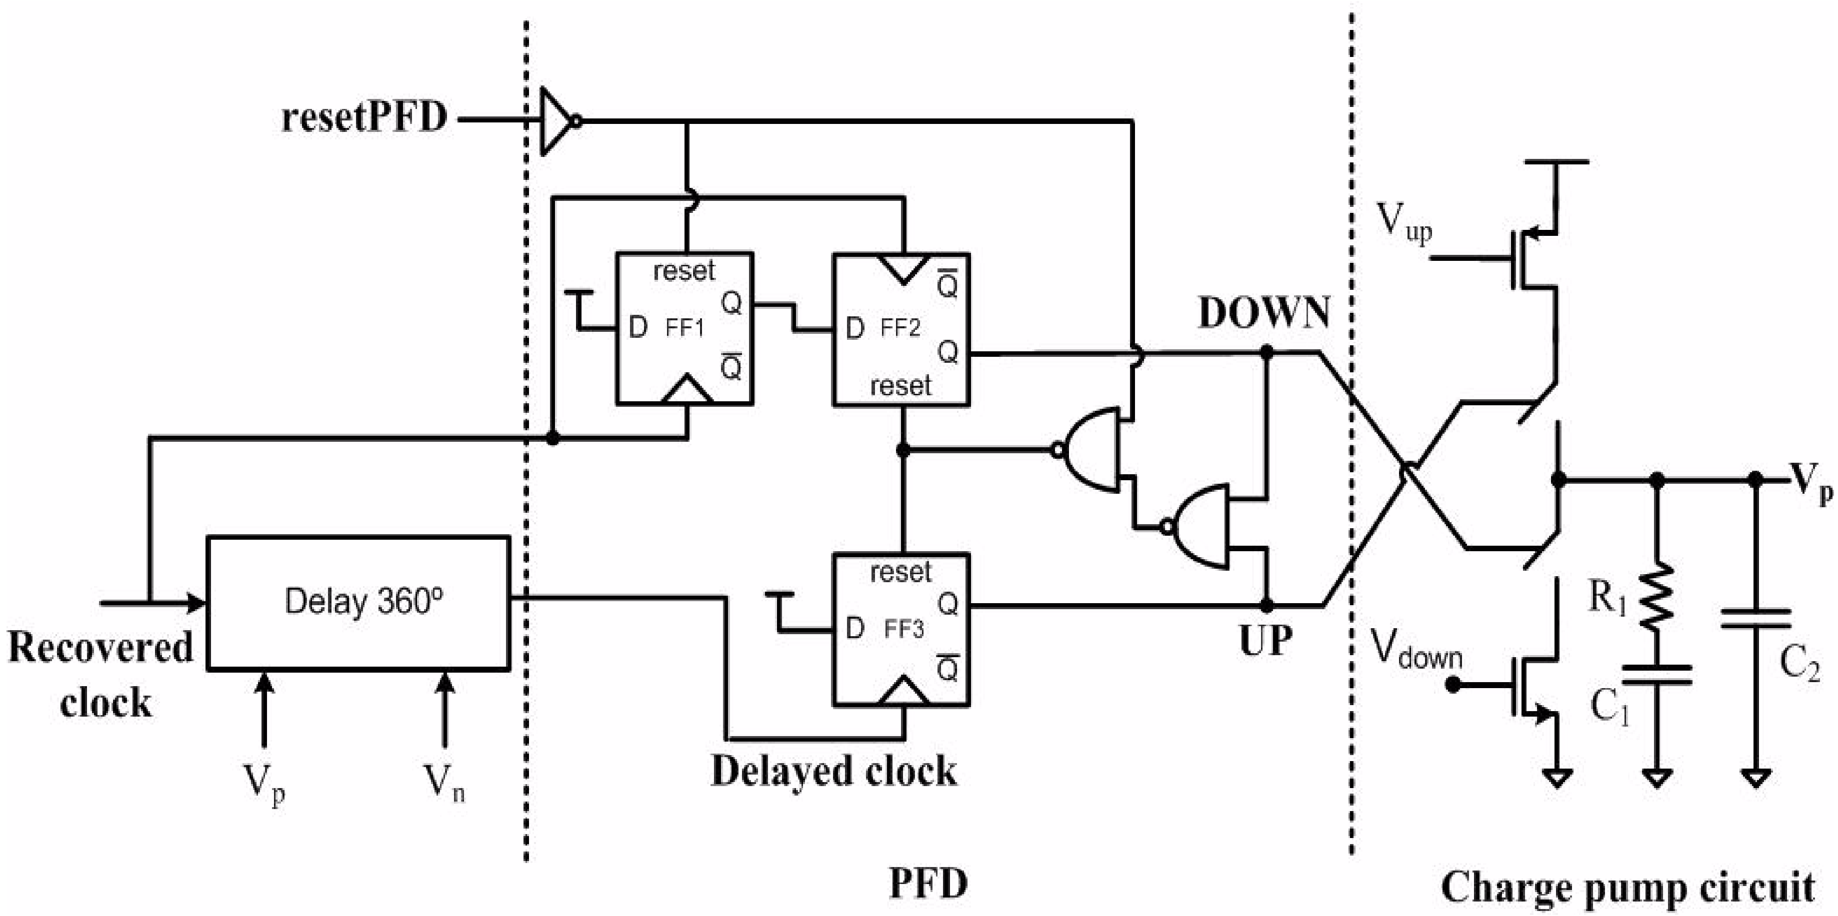

Burst mode operation (II)

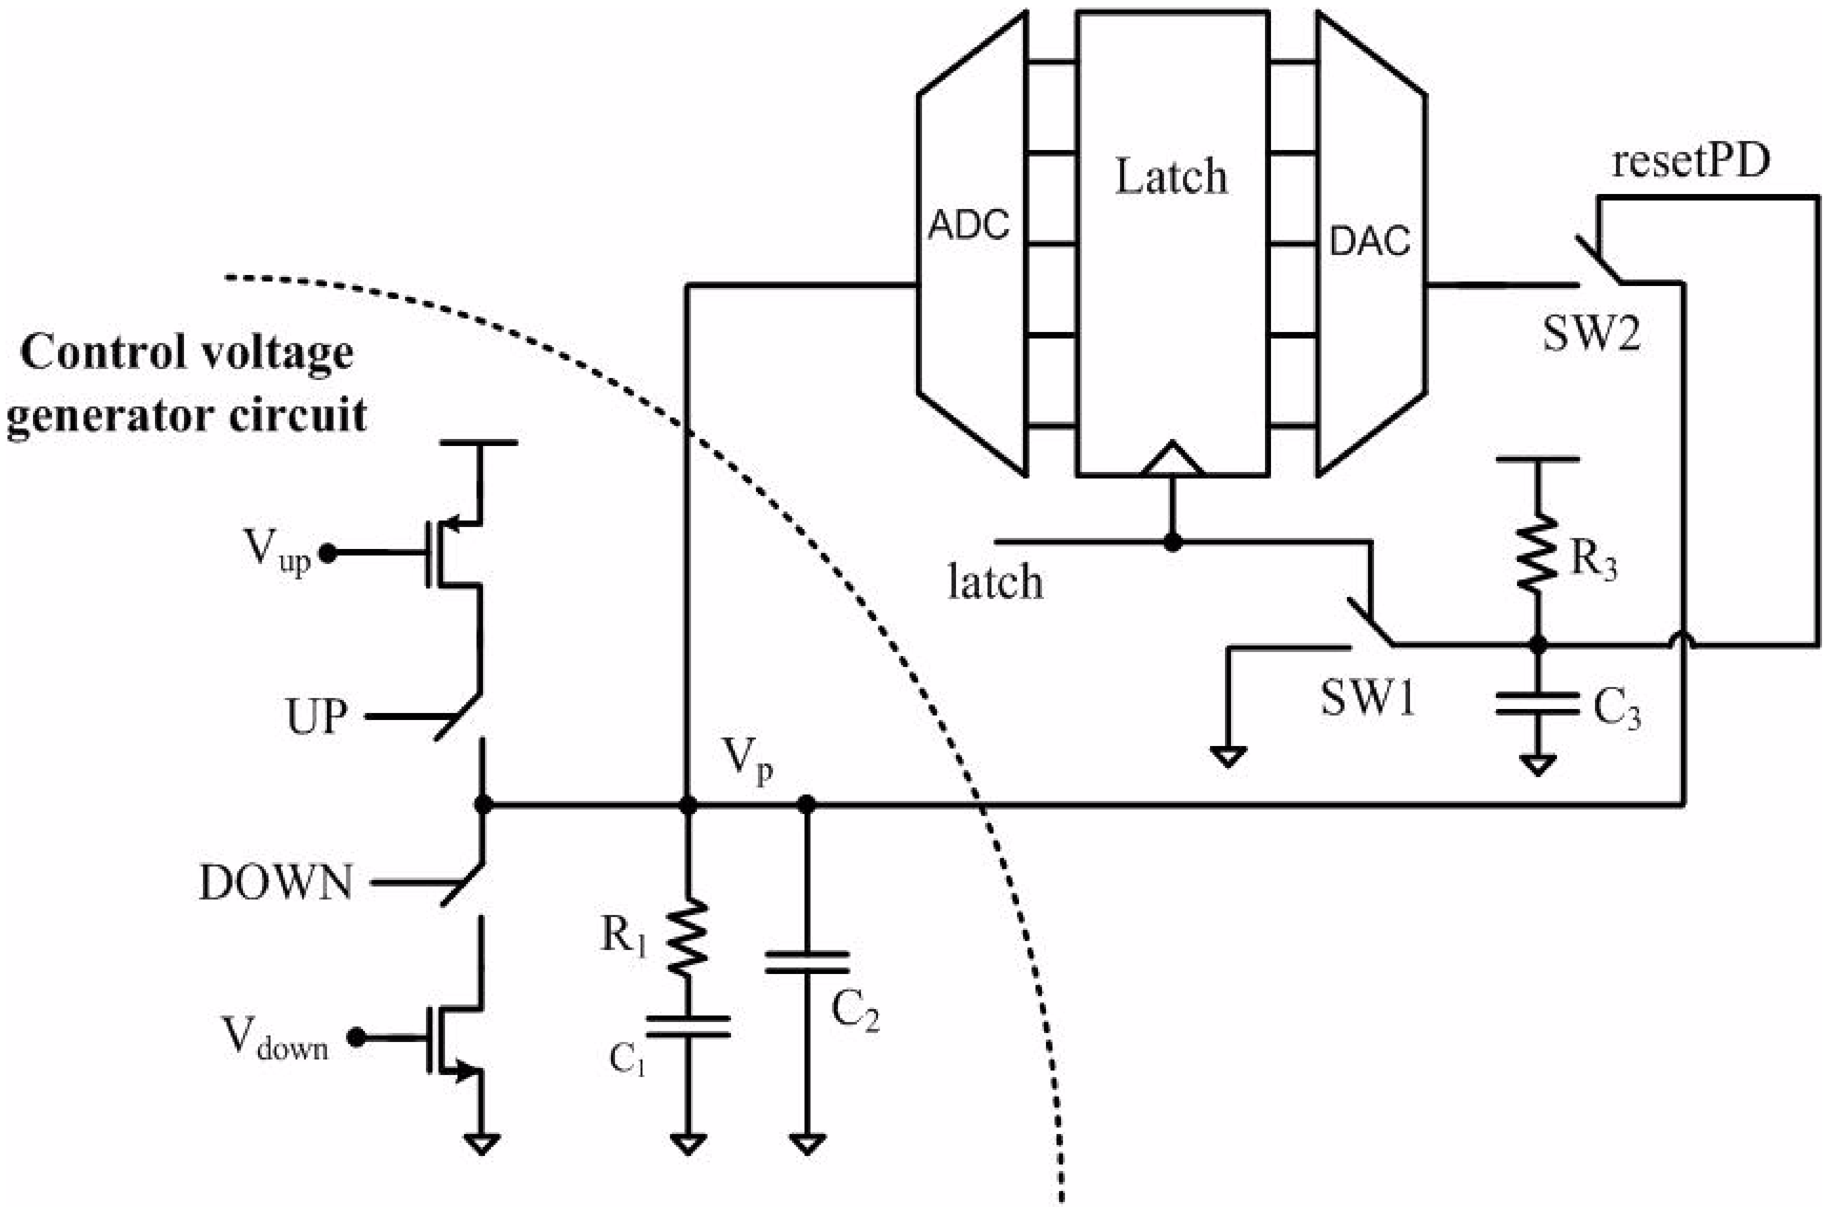

Simulation

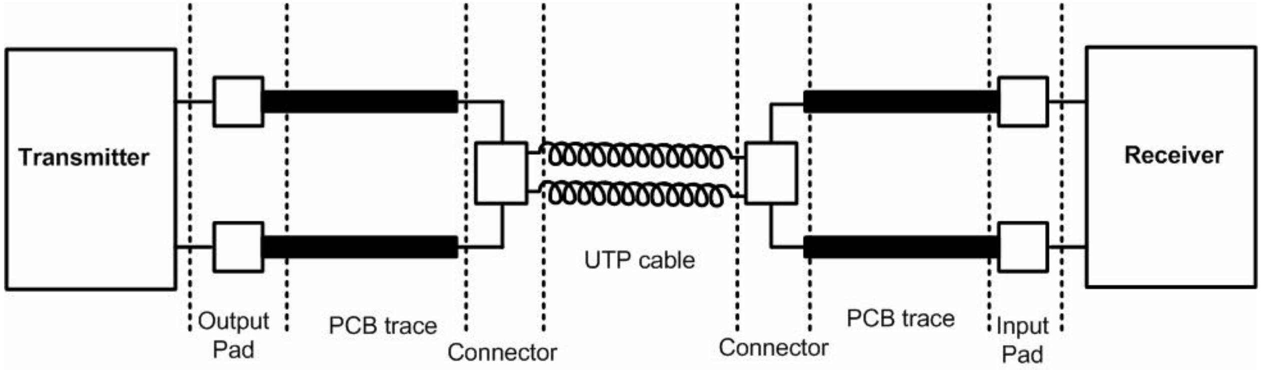

- ST 90nm CMOS
- 50cm cat5E UTP cable
- 5cm microstrip traces
- LVDS pads
- ESD protection circuits
- LVDS drivers available from ST 90nm library
- connectors
- simulated for all technology process corners
- temperature range 0-80ºC
- 5% variation in Supply Voltage

Simulation results

Signals involved in the clock recovery when the loop is locked

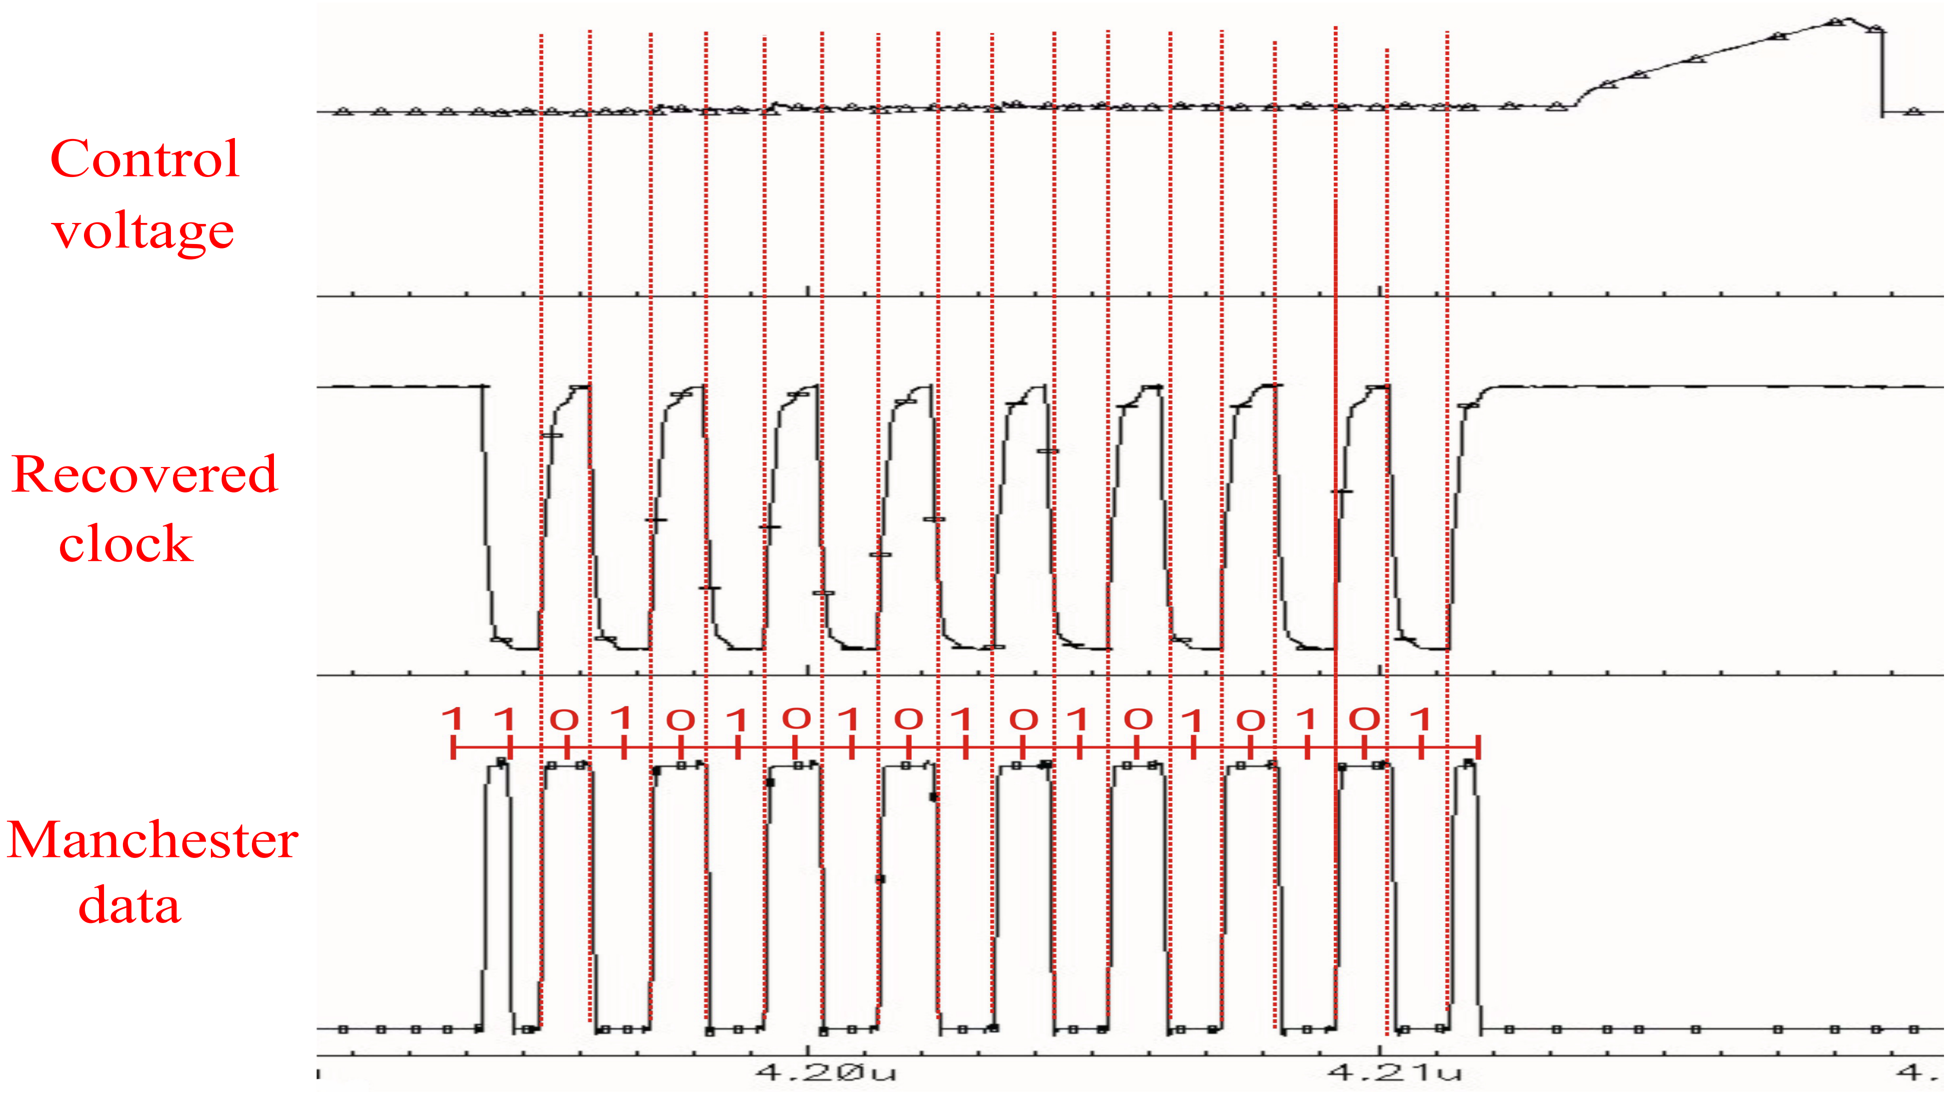

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems
- Some useful circuits:
  - calibration
  - LVDS interface
- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip
- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR
- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

What we want:

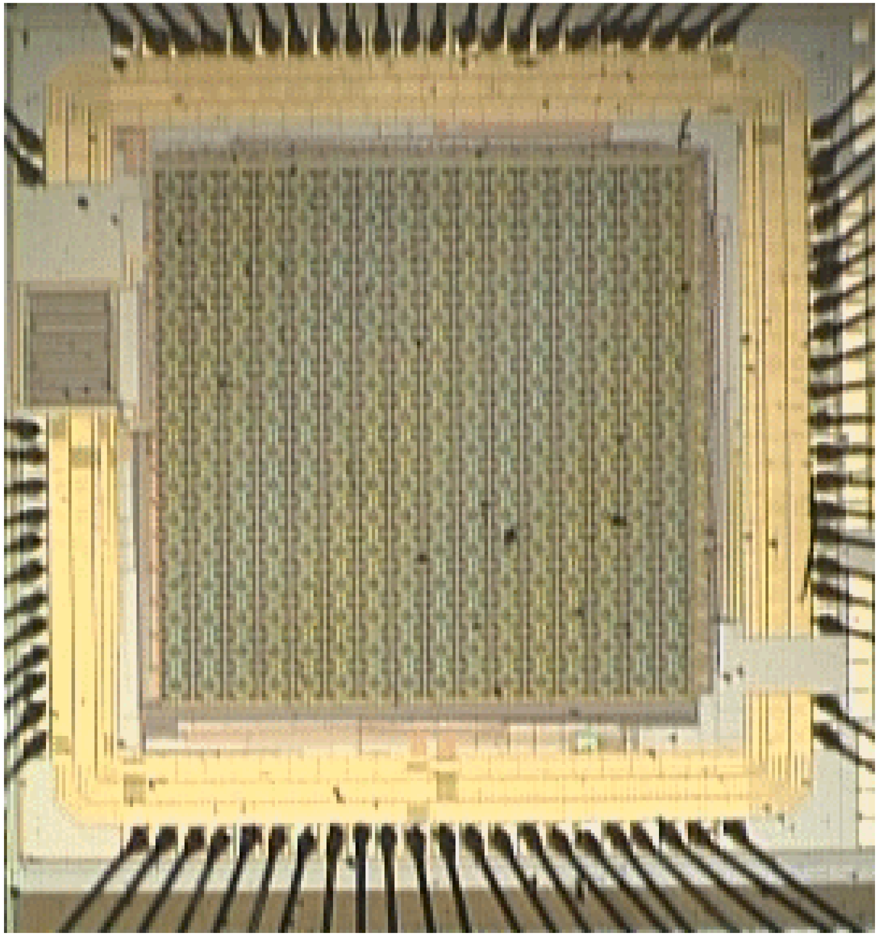

- Retina with AER output
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

What we want:

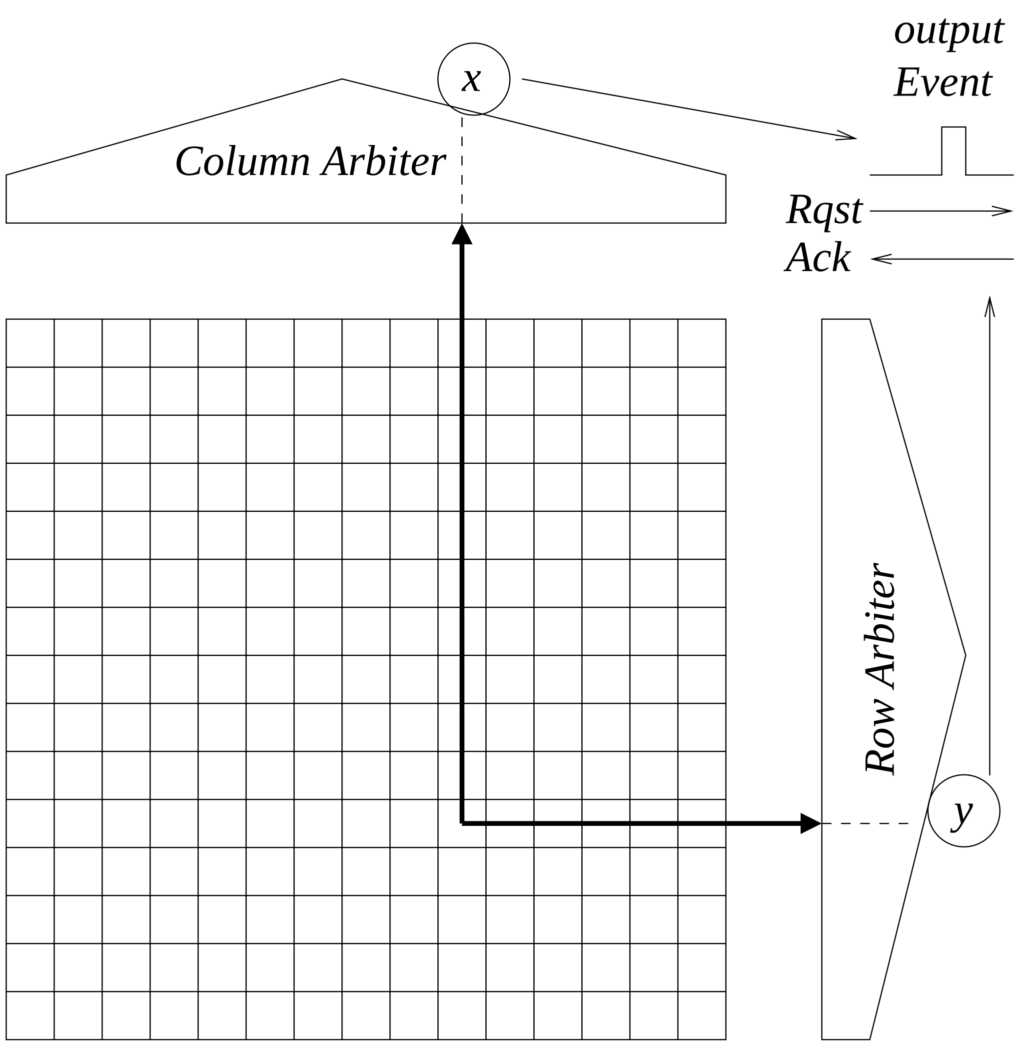

- **Retina with AER output**
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

What we want:

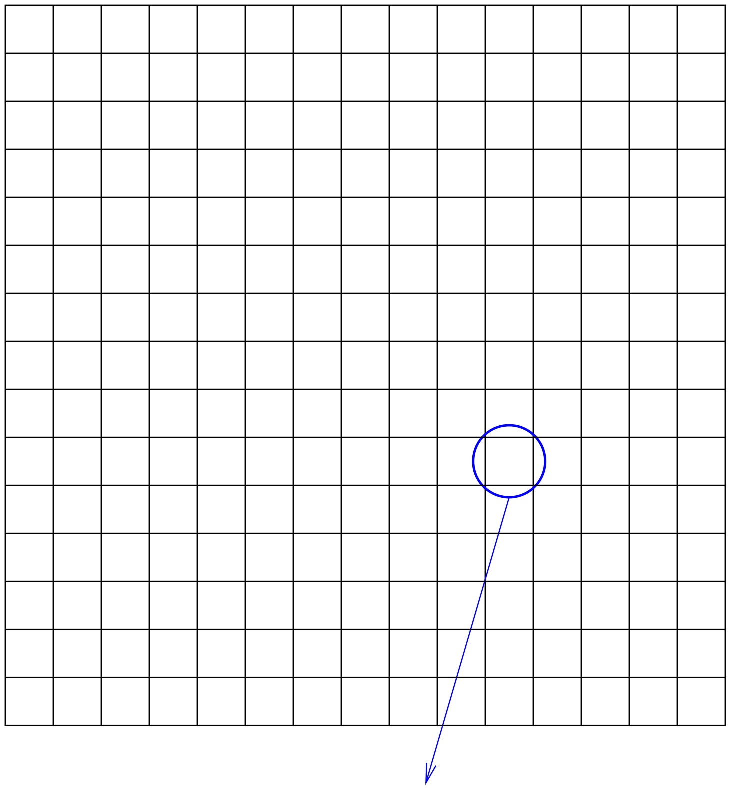

$$OutputFreq = f(I_{photo}(x, y), I_{neighbours})$$

- Retina with AER output
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

What we want:

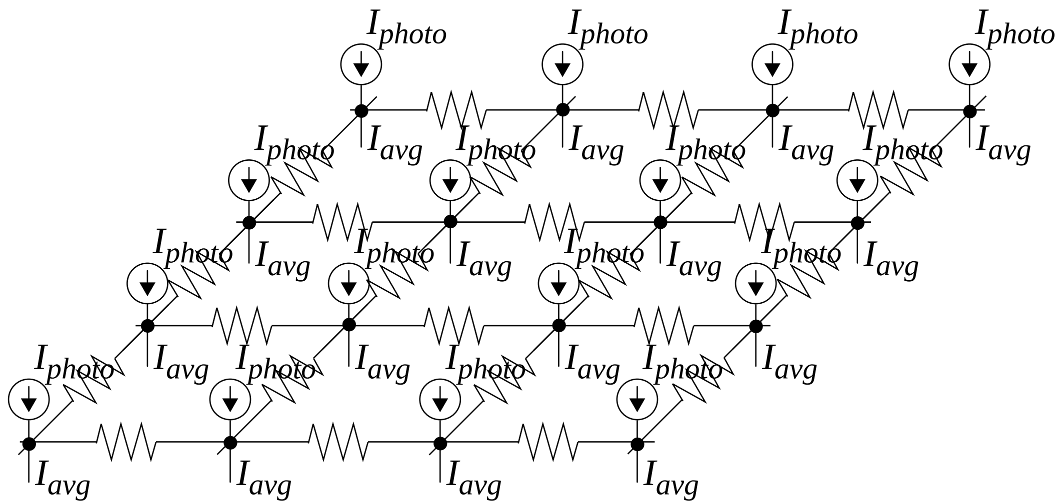

use of diffusers

- Retina with AER output
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

What we want:

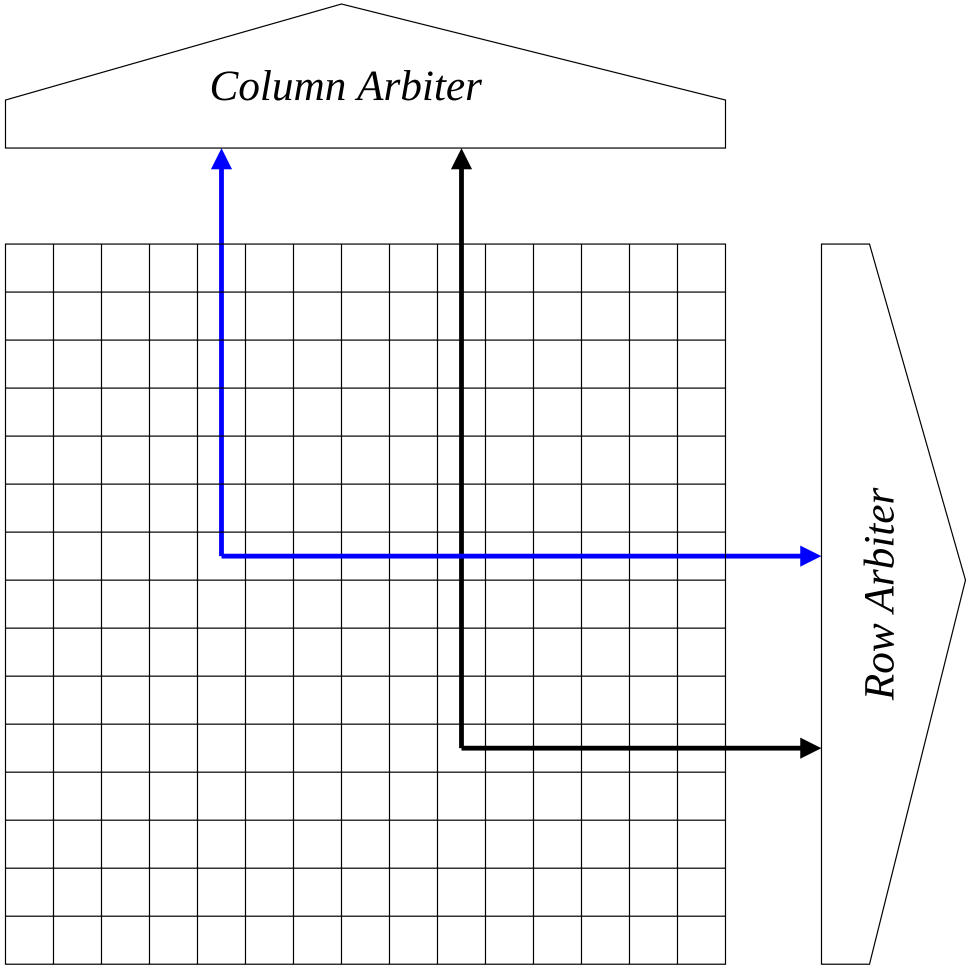

- each pixel decides when to generate an event
- there is no global periodic reset (no frames)

- Retina with AER output
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

What we want:

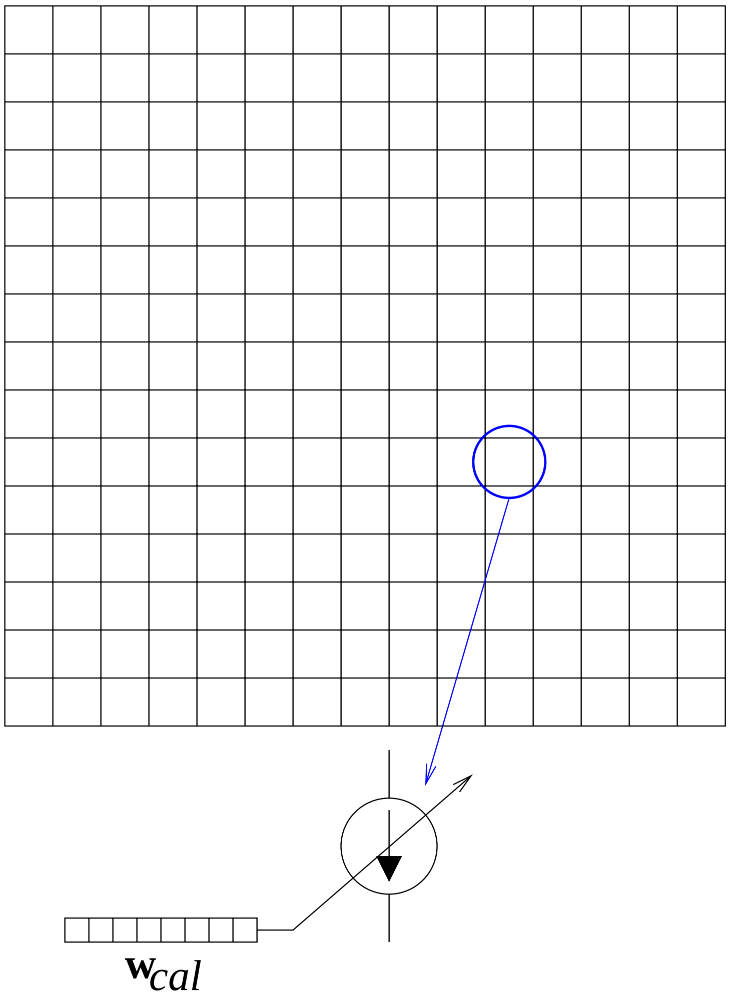

in-pixel calibration

- Retina with AER output
- Output frequency proportional to instantaneous *Spatial Contrast*
- *Spatial Contrast* computation not limited to nearest neighbors
- Fully Asynchronous output (no frames)
- low mismatch (FPN)

Calibration Technique

Active Current Generation

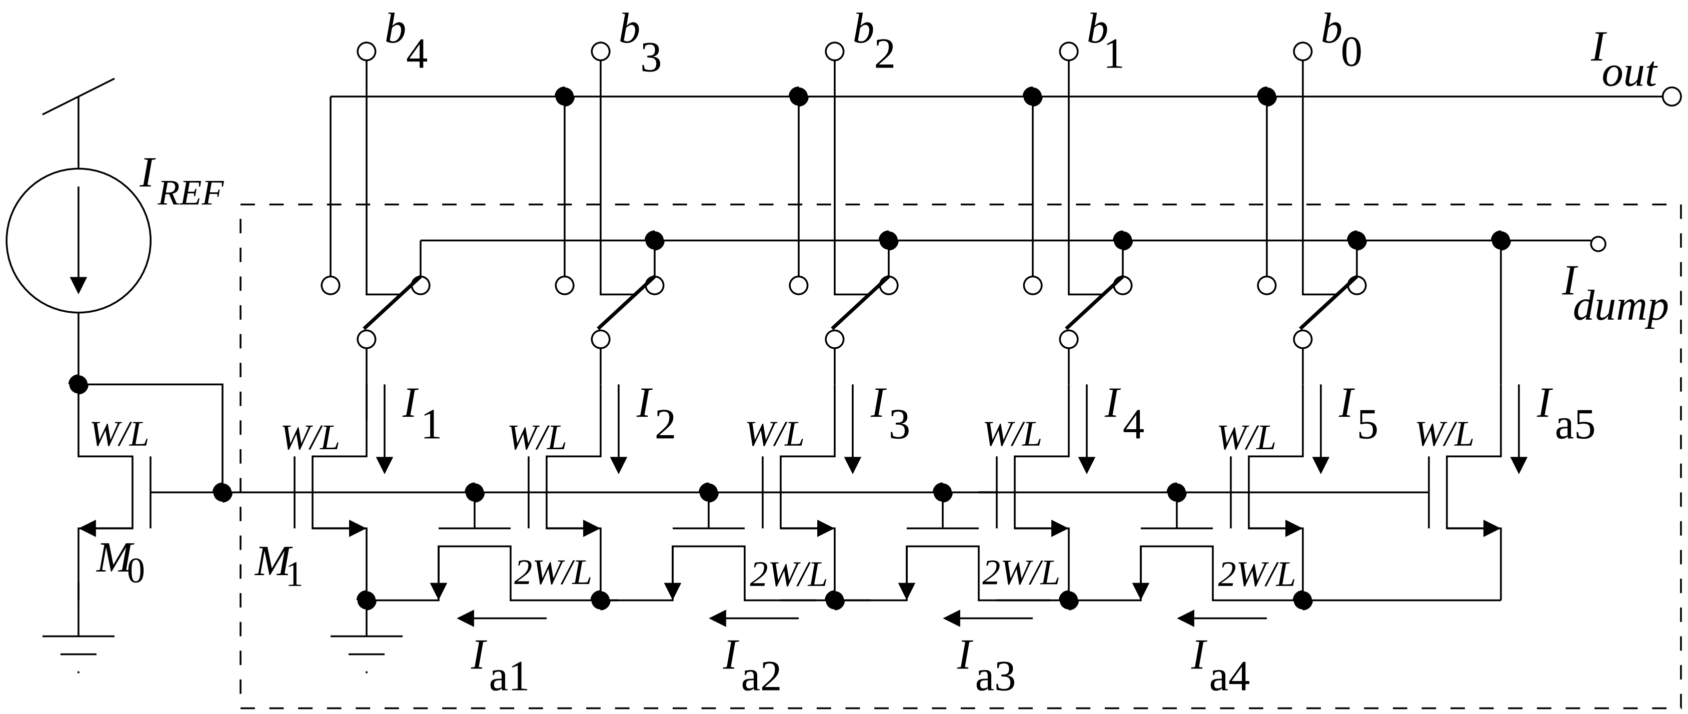

- Active current sources, controlled digitally
- Can be used as a Current DAC

Digi-MOS

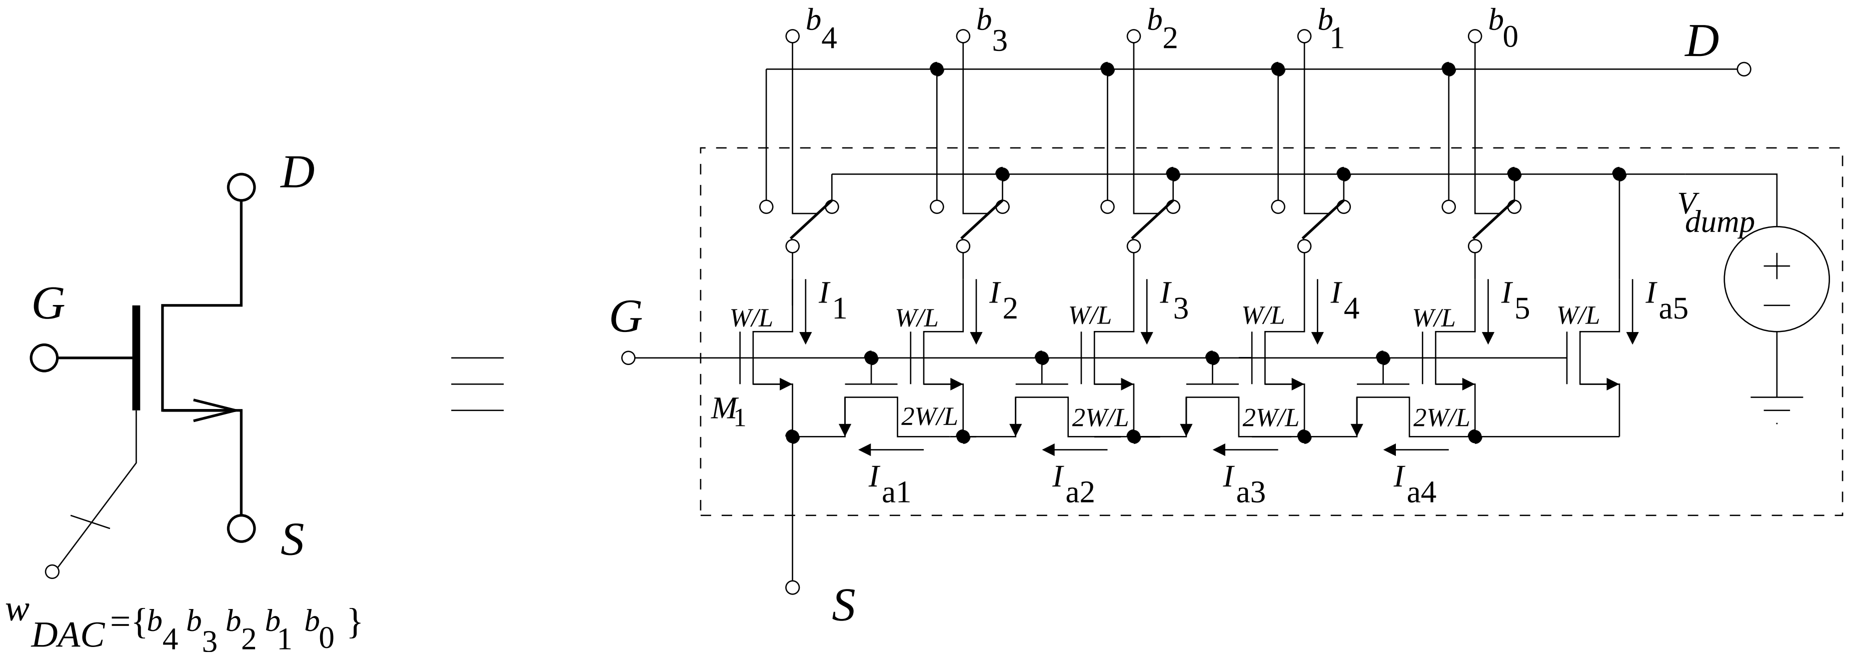

D
G
S
w DAC ={b 4 b 3 b 2 b 1 b 0 }
b 4
b 3
b 2
b 1
b 0
I 1
I 2
I 3
I 4
I 5
I a5
W/L
2W/L
M1
V dump
I a1
I a2
I a3
I a4

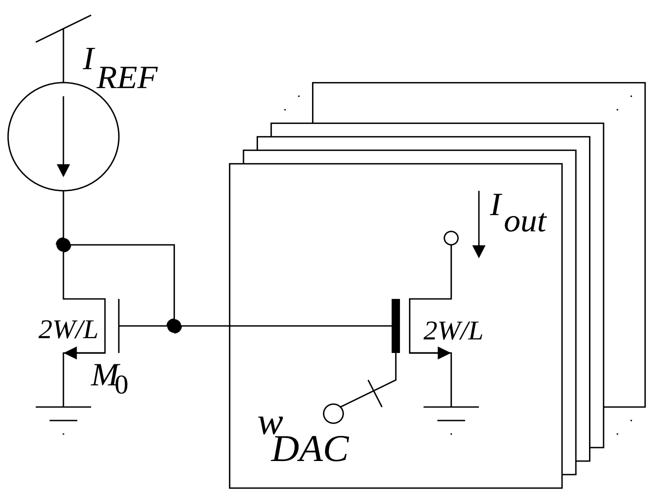

I REF
2W/L
M0
I out
w DAC

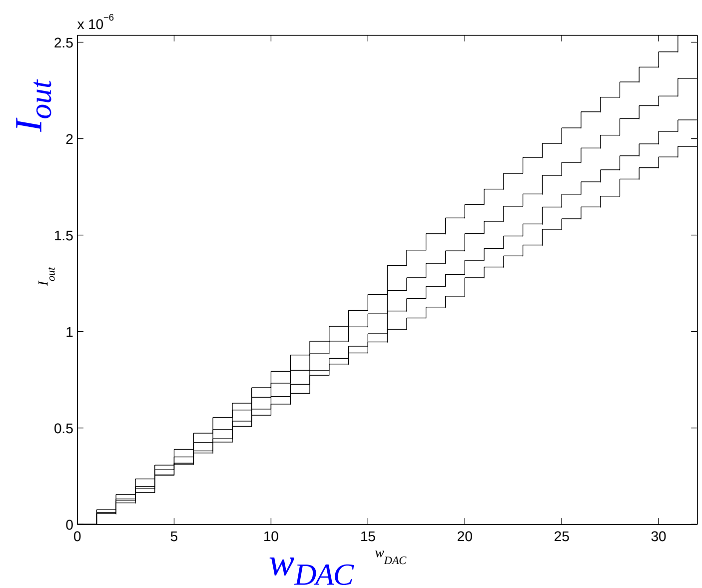

x 10^−6
I out
w DAC

Example Layout for 0.35μ*m* CMOS 5-bit digi-MOS

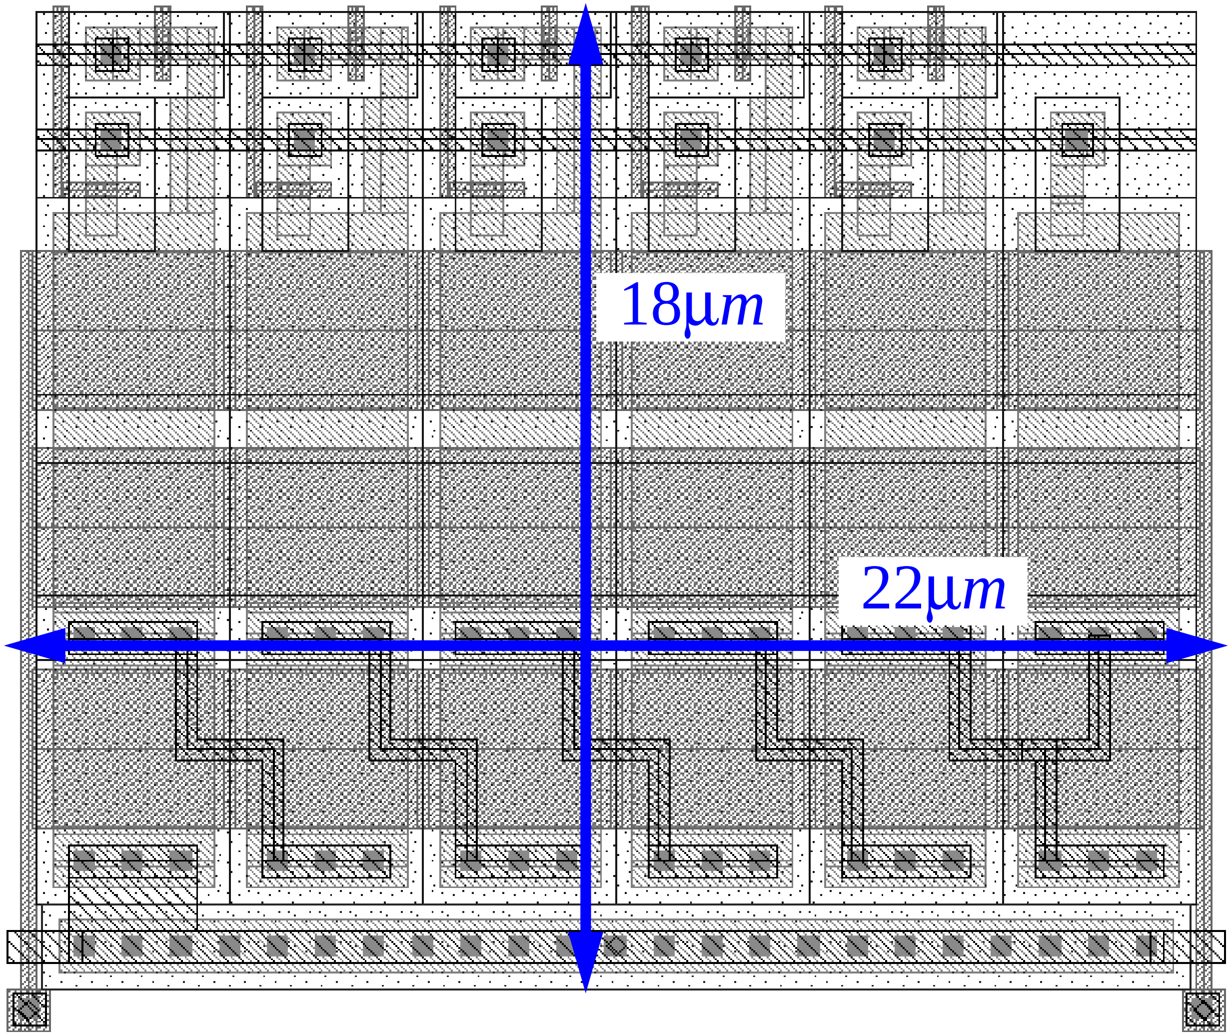

- unit transistor $W = L = 3\mu m$

Spatial Contrast Computation

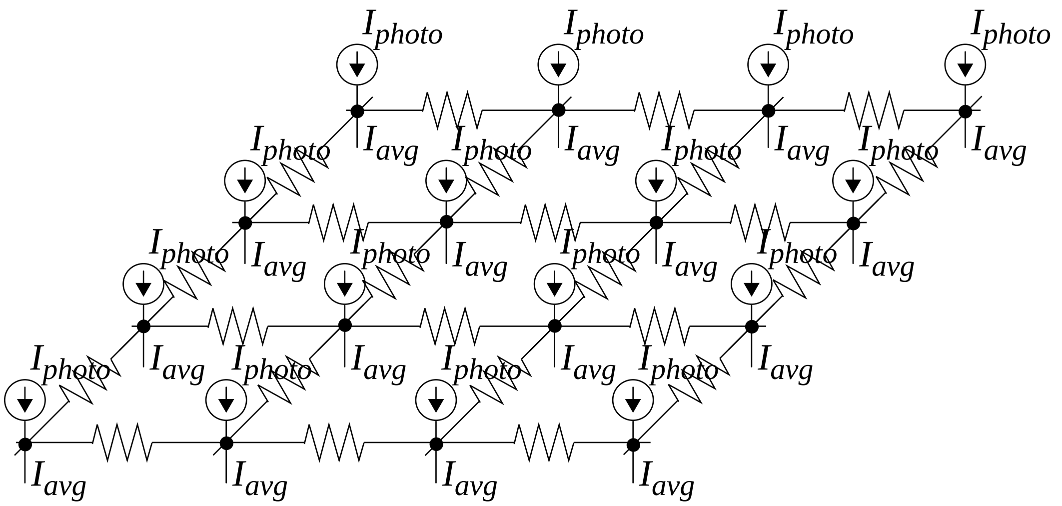

Michelson Contrast: $I_{cont}(x, y) = I_{ref}\dfrac{I_{photo}(x, y) - I_{avg}(x, y)}{I_{photo}(x, y) + I_{avg}(x, y)}$

Weber Contrast: $I_{cont}(x, y) = I_{ref}\dfrac{I_{photo}(x, y) - I_{avg}(x, y)}{I_{avg}(x, y)}$

Simple Ratio Contrast: $I_{cont}(x, y) = I_{ref}\dfrac{I_{avg}(x, y)}{I_{photo}(x, y)}$

Calibrating for Mismatch

Sums/Subtractions & Multiplications/Division:

$$I_o = I_1 \frac{I_2 - I_3}{I_4} \quad \rightarrow \quad I_o + \Delta_o = (I_1 + \Delta_1)\frac{(I_2 + \Delta_2) - (I_3 + \Delta_3)}{(I_4 + \Delta_4)}$$

$$I_o + \Delta_o \approx \frac{I_1 I_2}{I_4}(1 + \Delta_1 + \Delta_2 - \Delta_4) - \frac{I_1 I_3}{I_4}(1 + \Delta_1 + \Delta_3 - \Delta_4)$$

Calibrating for Mismatch

Sums/Subtractions & Multiplications/Divisions:

$$I_o = I_1 \frac{I_2 - I_3}{I_4} \quad \rightarrow \quad I_o + \Delta_o = (I_1 + \Delta_1)\frac{(I_2 + \Delta_2) - (I_3 + \Delta_3)}{(I_4 + \Delta_4)}$$

$$I_o + \Delta_o \approx \frac{I_1 I_2}{I_4}(1 + \Delta_1 + \Delta_2 - \Delta_4) - \frac{I_1 I_3}{I_4}(1 + \Delta_1 + \Delta_3 - \Delta_4)$$

Only Multiplications/Divisions:

$$I_o = I_1 \frac{I_2}{I_4} \quad \rightarrow \quad I_o + \Delta_o = (I_1 + \Delta_1)\frac{(I_2 + \Delta_2)}{(I_4 + \Delta_4)}$$

$$I_o + \Delta_o \approx \frac{I_1 I_2}{I_4}(1 + \Delta_1 + \Delta_2 - \Delta_4)$$

only one calibration current per pixel

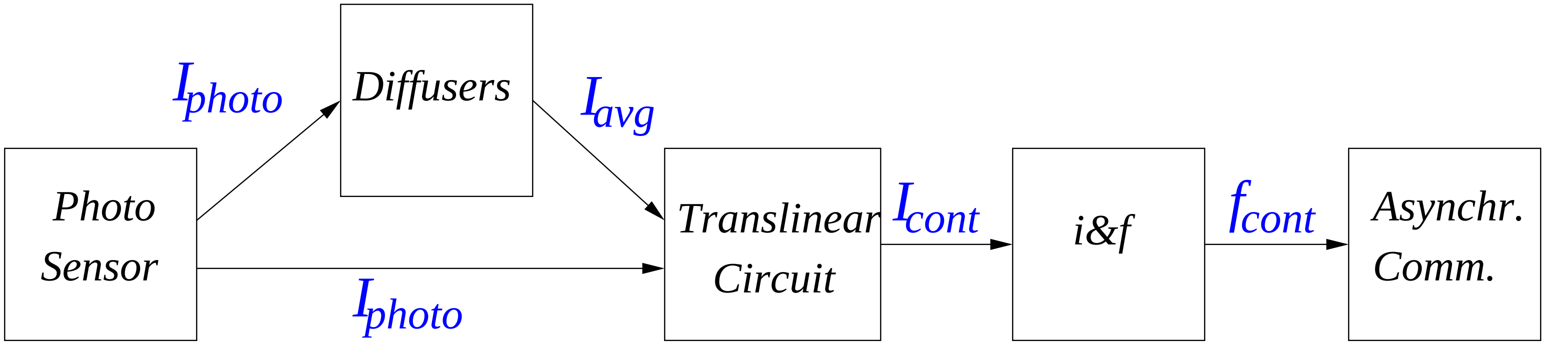

Photo
Sensor
I_{photo}
Diffusers
I_{avg}
I_{photo}
Translinear
Circuit
I_{cont}
i&f
f_{cont}
Asynchr.
Comm.

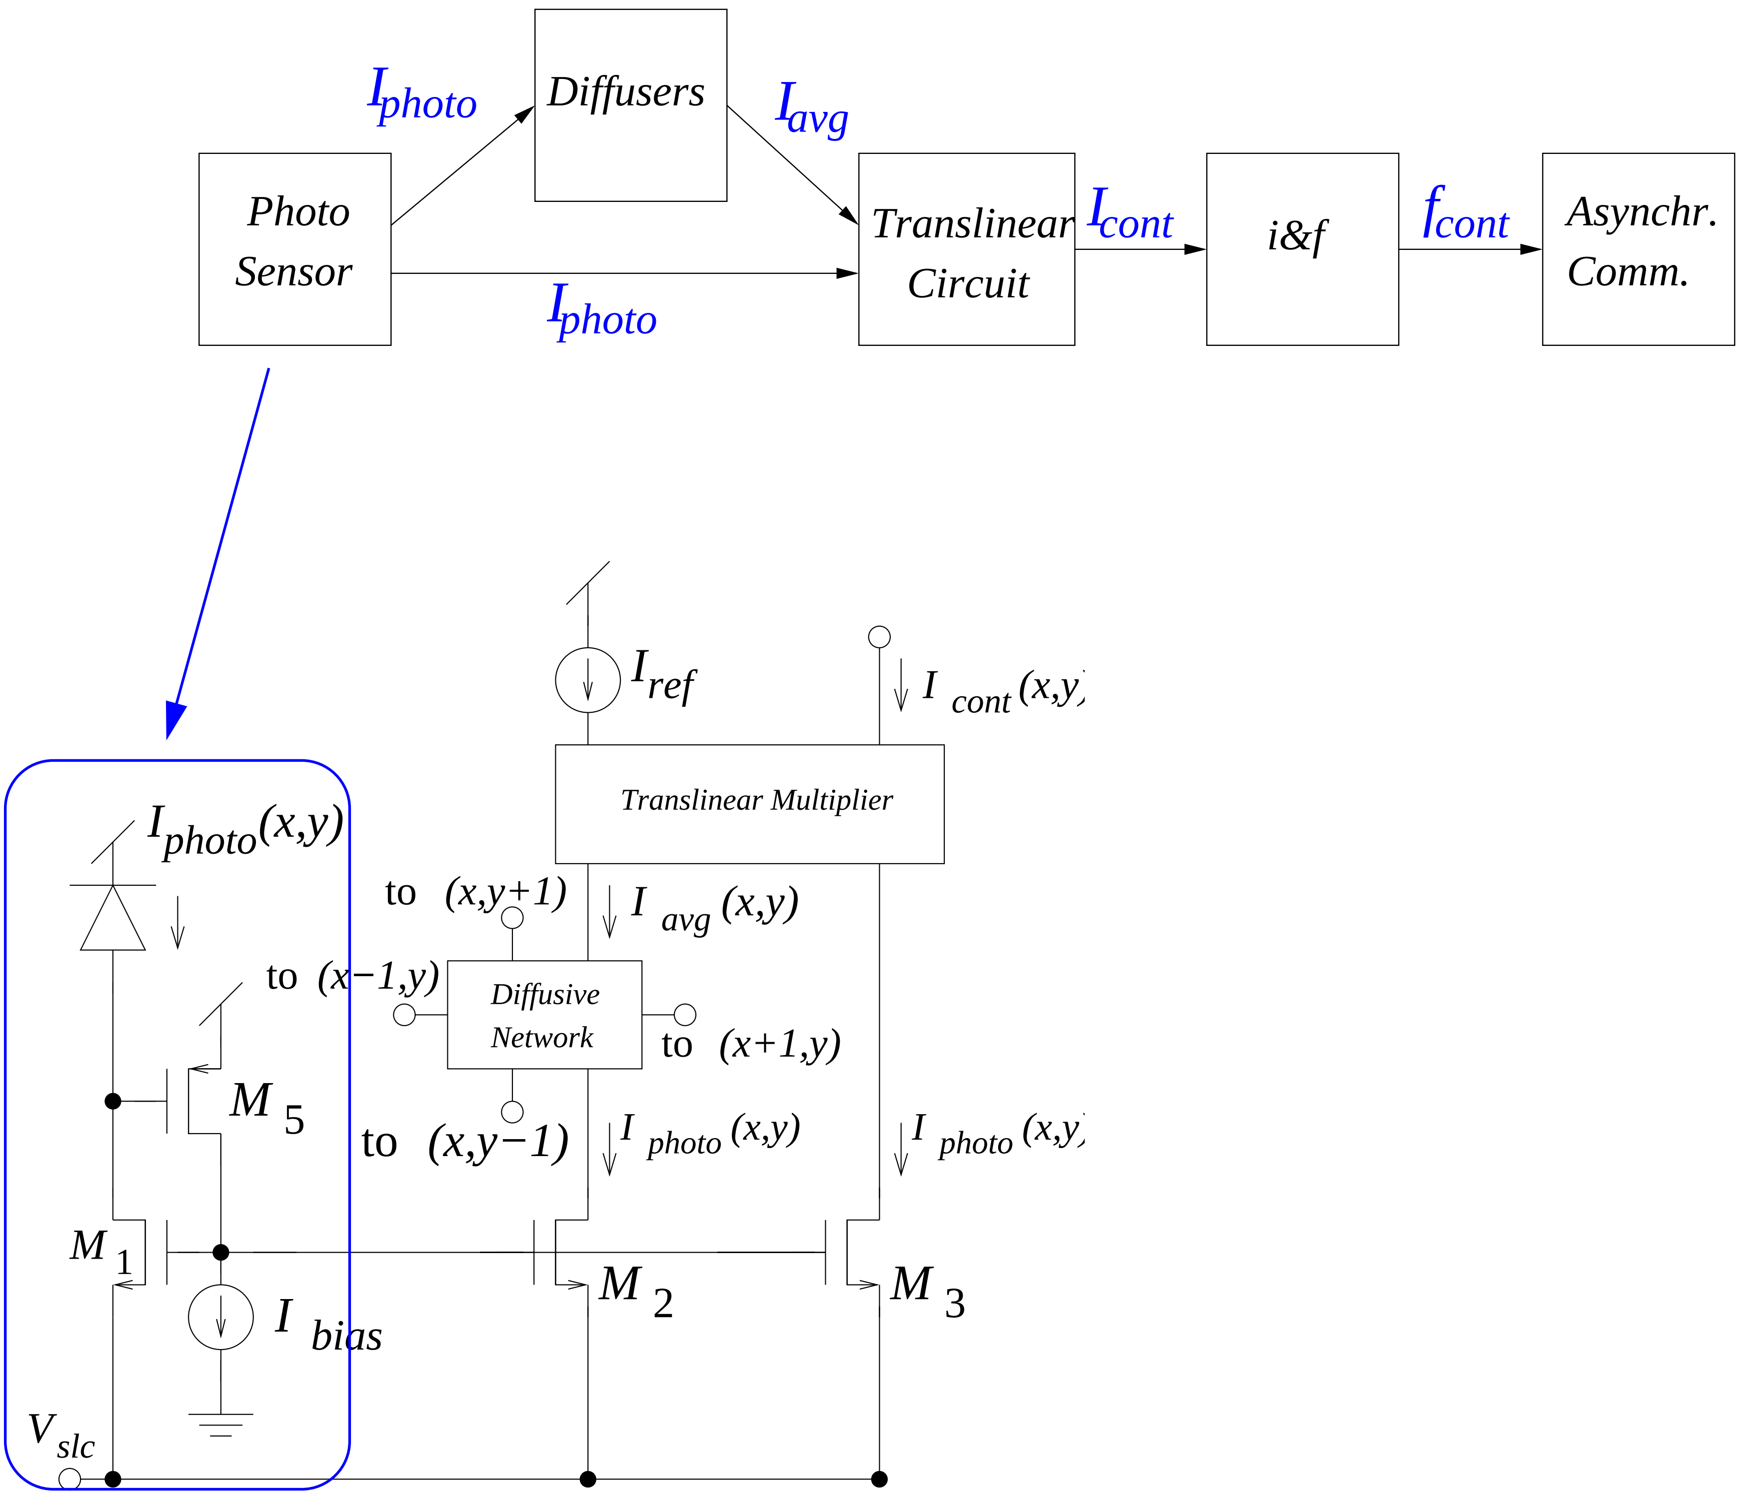

Photo Sensor
Iphoto
Diffusers
Iavg
Iphoto
Translinear Circuit
Icont
i&f
fcont
Asynchr. Comm.
Iref
I cont (x,y)
Translinear Multiplier
Iphoto(x,y)
to (x,y+1)
I avg (x,y)
to (x−1,y)
Diffusive Network
to (x+1,y)
M 5
to (x,y−1)
I photo (x,y)
I photo (x,y)
M 1
M 2
M 3
I bias
V slc

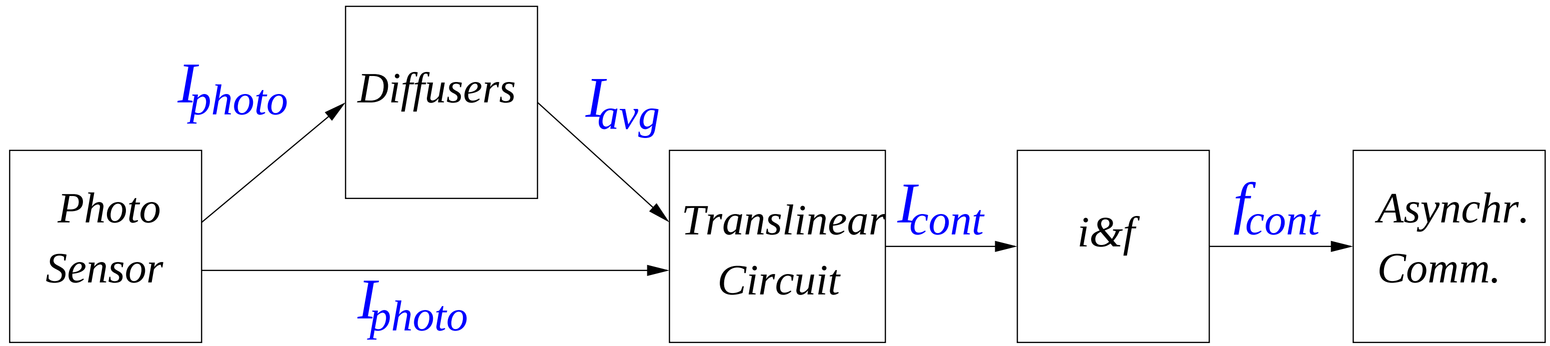
Photo Sensor
I_{photo}
Diffusers
I_{avg}
I_{photo}
Translinear Circuit
I_{cont}
i&f
f_{cont}
Asynchr. Comm.

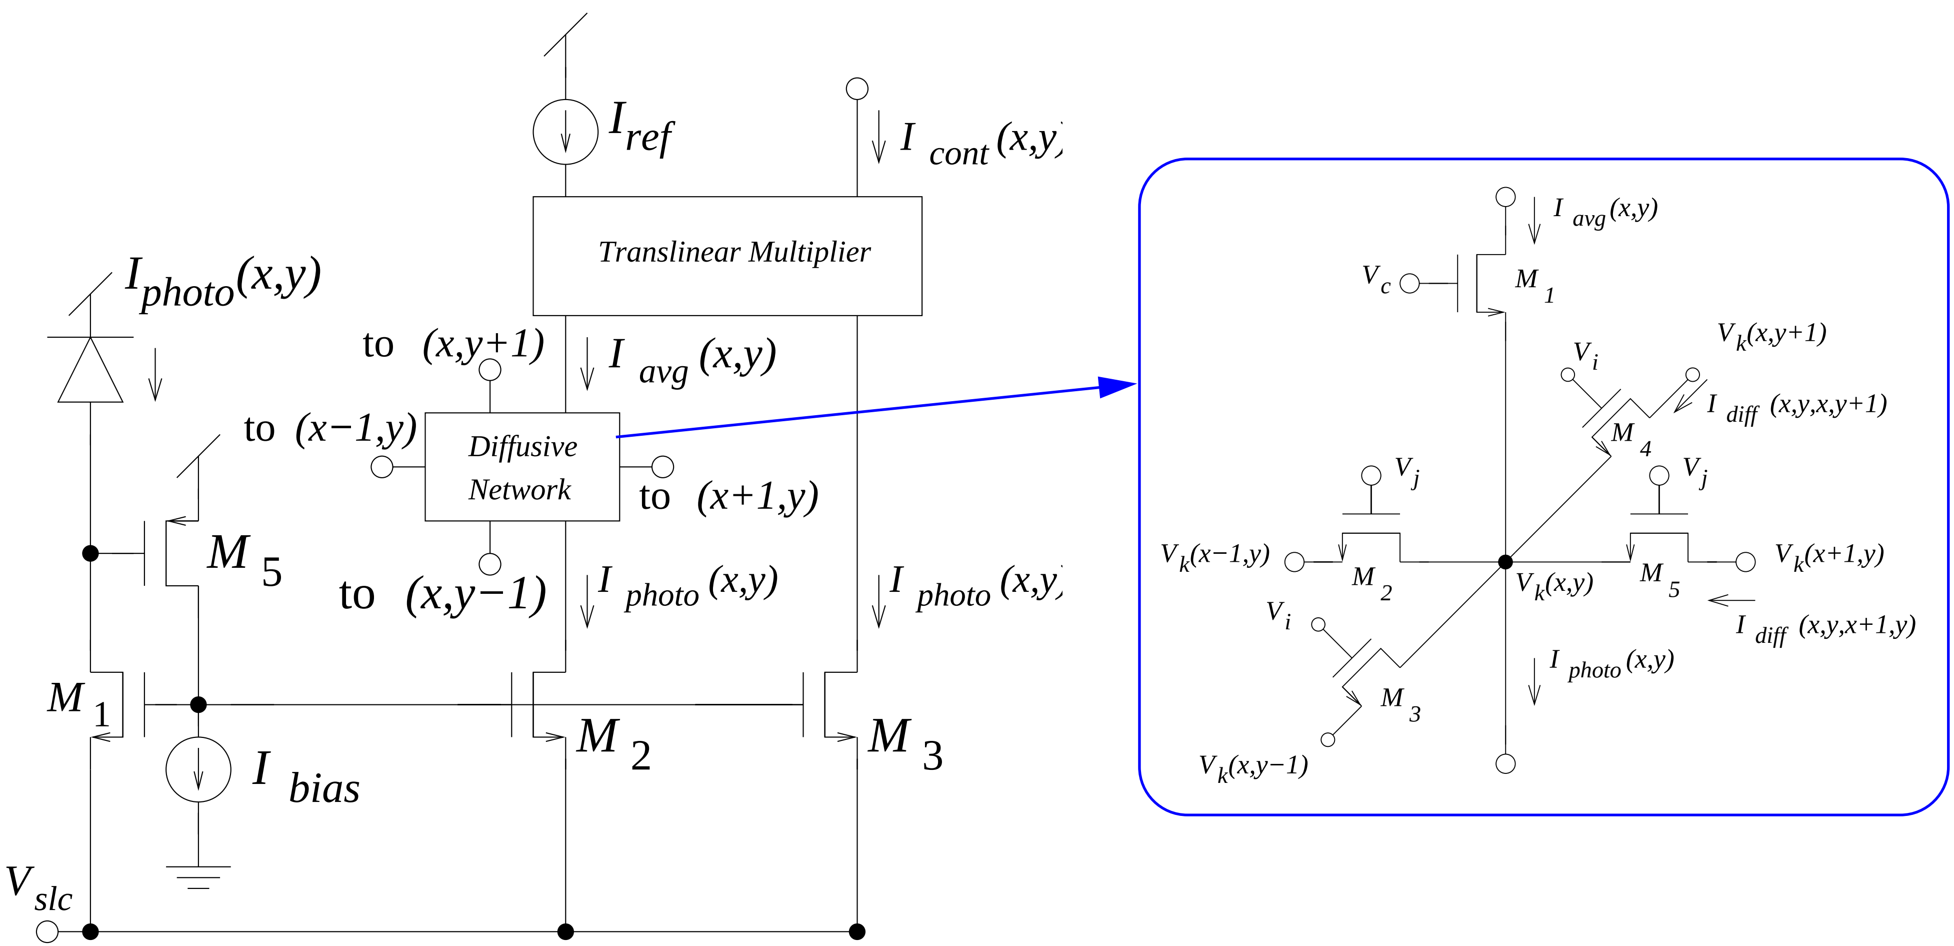
I_{ref}
$I_{cont}(x,y)$
Translinear Multiplier
$I_{photo}(x,y)$
to (x,y+1)
$I_{avg}(x,y)$
to (x−1,y)
Diffusive Network
to (x+1,y)
M_5
to (x,y−1)
$I_{photo}(x,y)$
$I_{photo}(x,y)$
M_1
M_2
M_3
I_{bias}
V_{slc}
$I_{avg}(x,y)$
V_c
M_1
$V_k(x,y+1)$
V_i
$I_{diff}(x,y,x,y+1)$
M_4
V_j
V_j
$V_k(x-1,y)$
M_2
$V_k(x,y)$
M_5
$V_k(x+1,y)$
V_i
$I_{diff}(x,y,x+1,y)$
$I_{photo}(x,y)$
M_3
$V_k(x,y-1)$

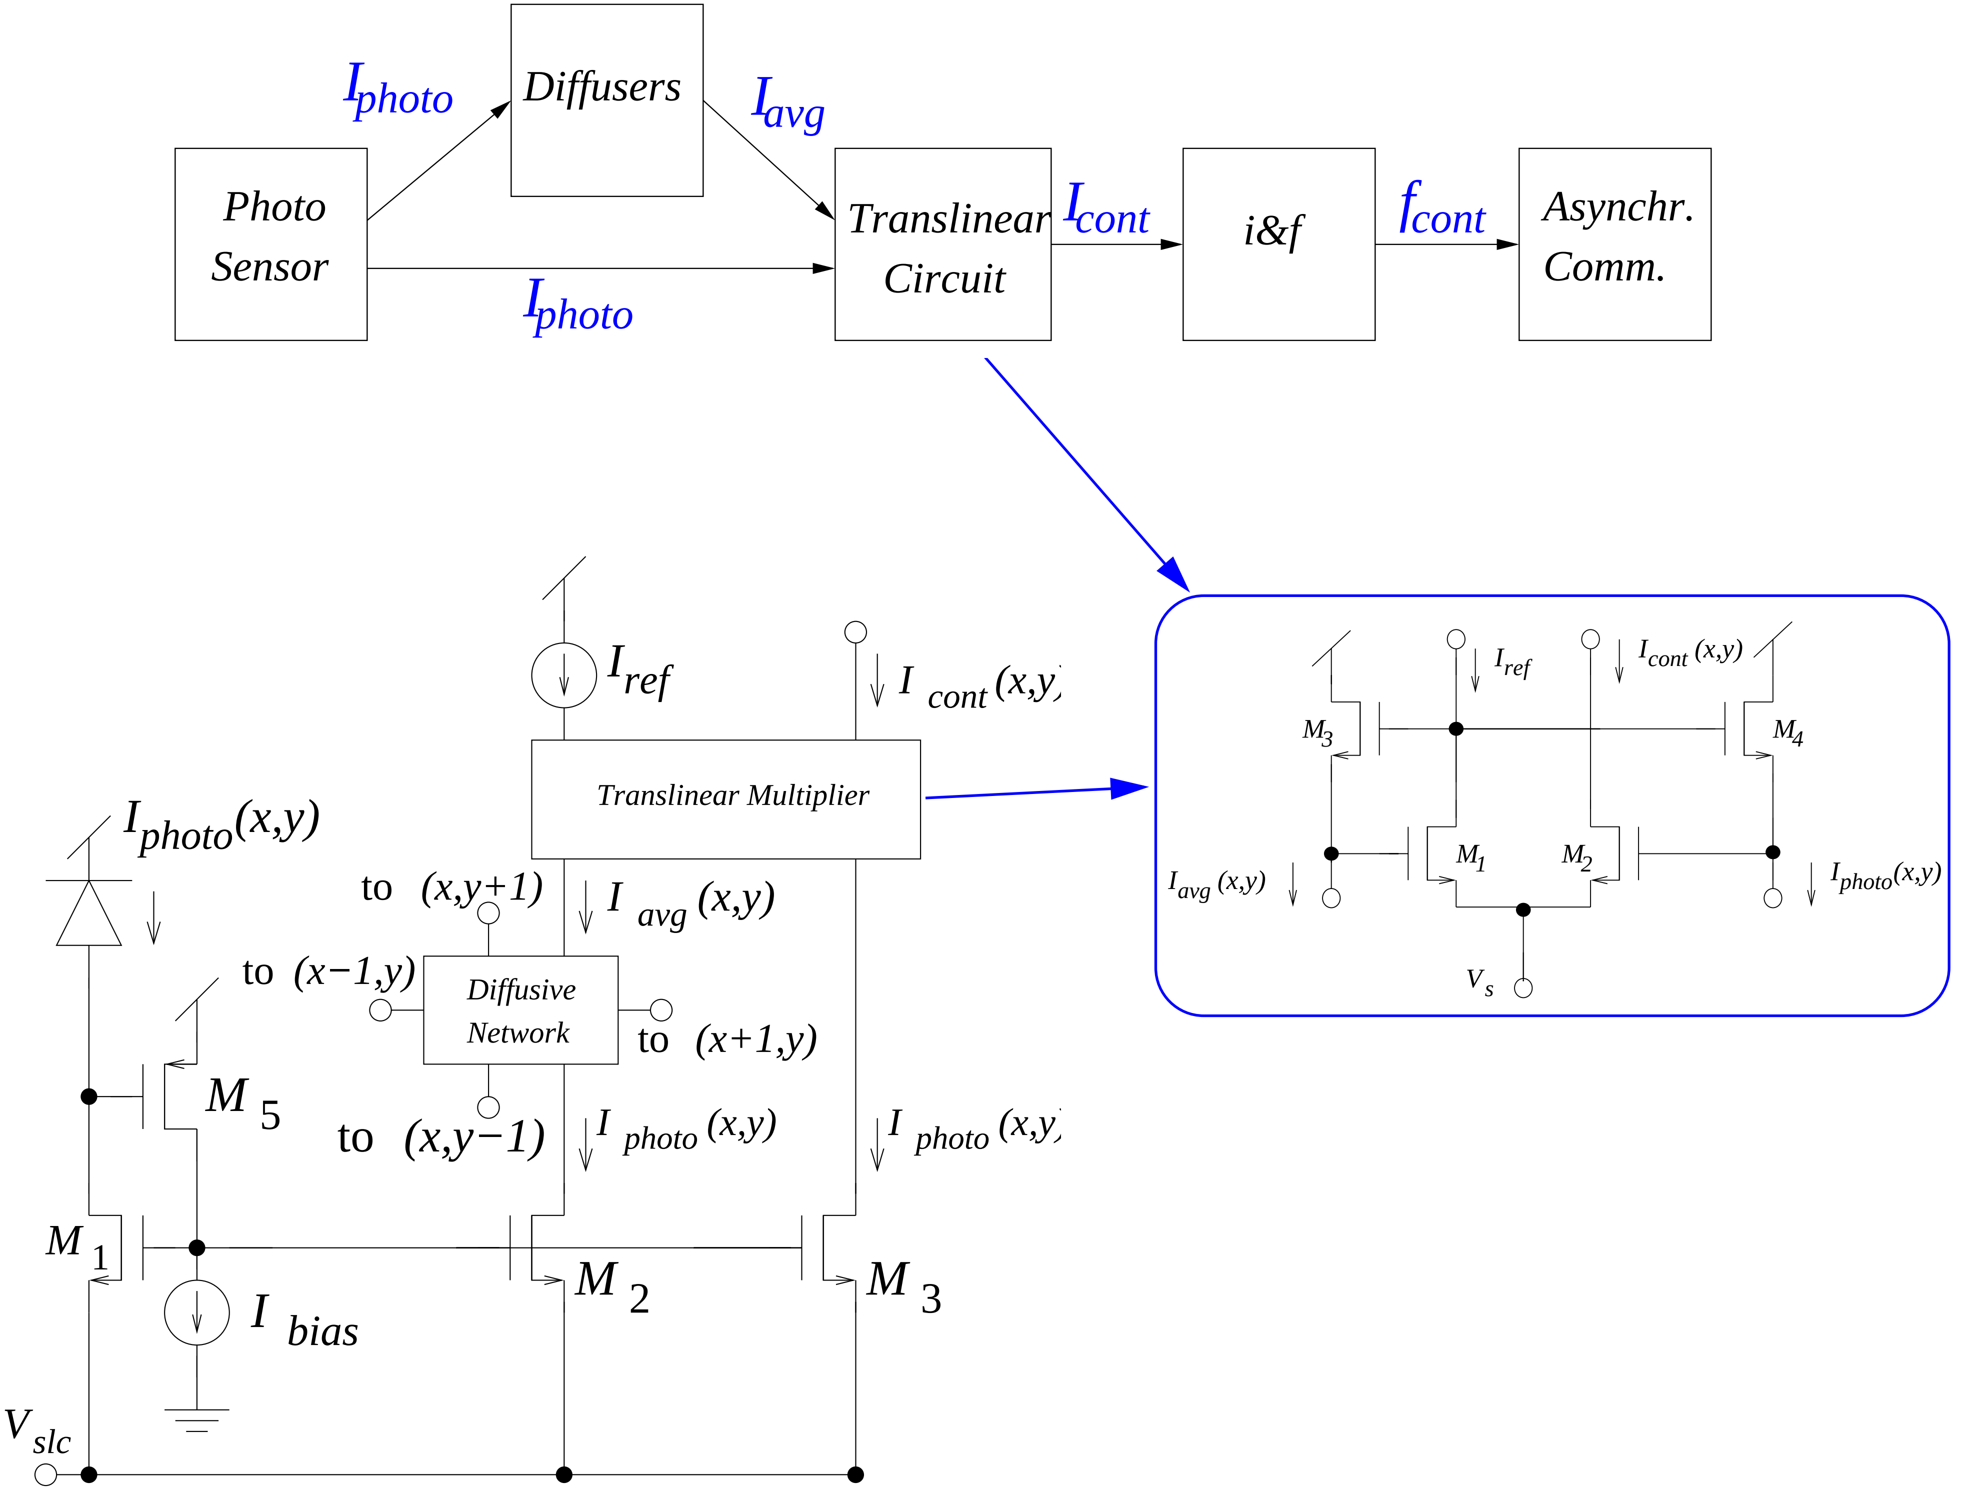

Photo Sensor
Iphoto
Diffusers
Iavg
Iphoto
Translinear Circuit
Icont
i&f
fcont
Asynchr. Comm.
Iref
I cont (x,y)
Translinear Multiplier
Iphoto(x,y)
to (x,y+1)
I avg (x,y)
to (x−1,y)
Diffusive Network
to (x+1,y)
M 5
to (x,y−1)
I photo (x,y)
I photo (x,y)
M 1
I bias
M 2
M 3
V slc
M3
M4
M1
M2
Iavg (x,y)
Iphoto(x,y)
Vs

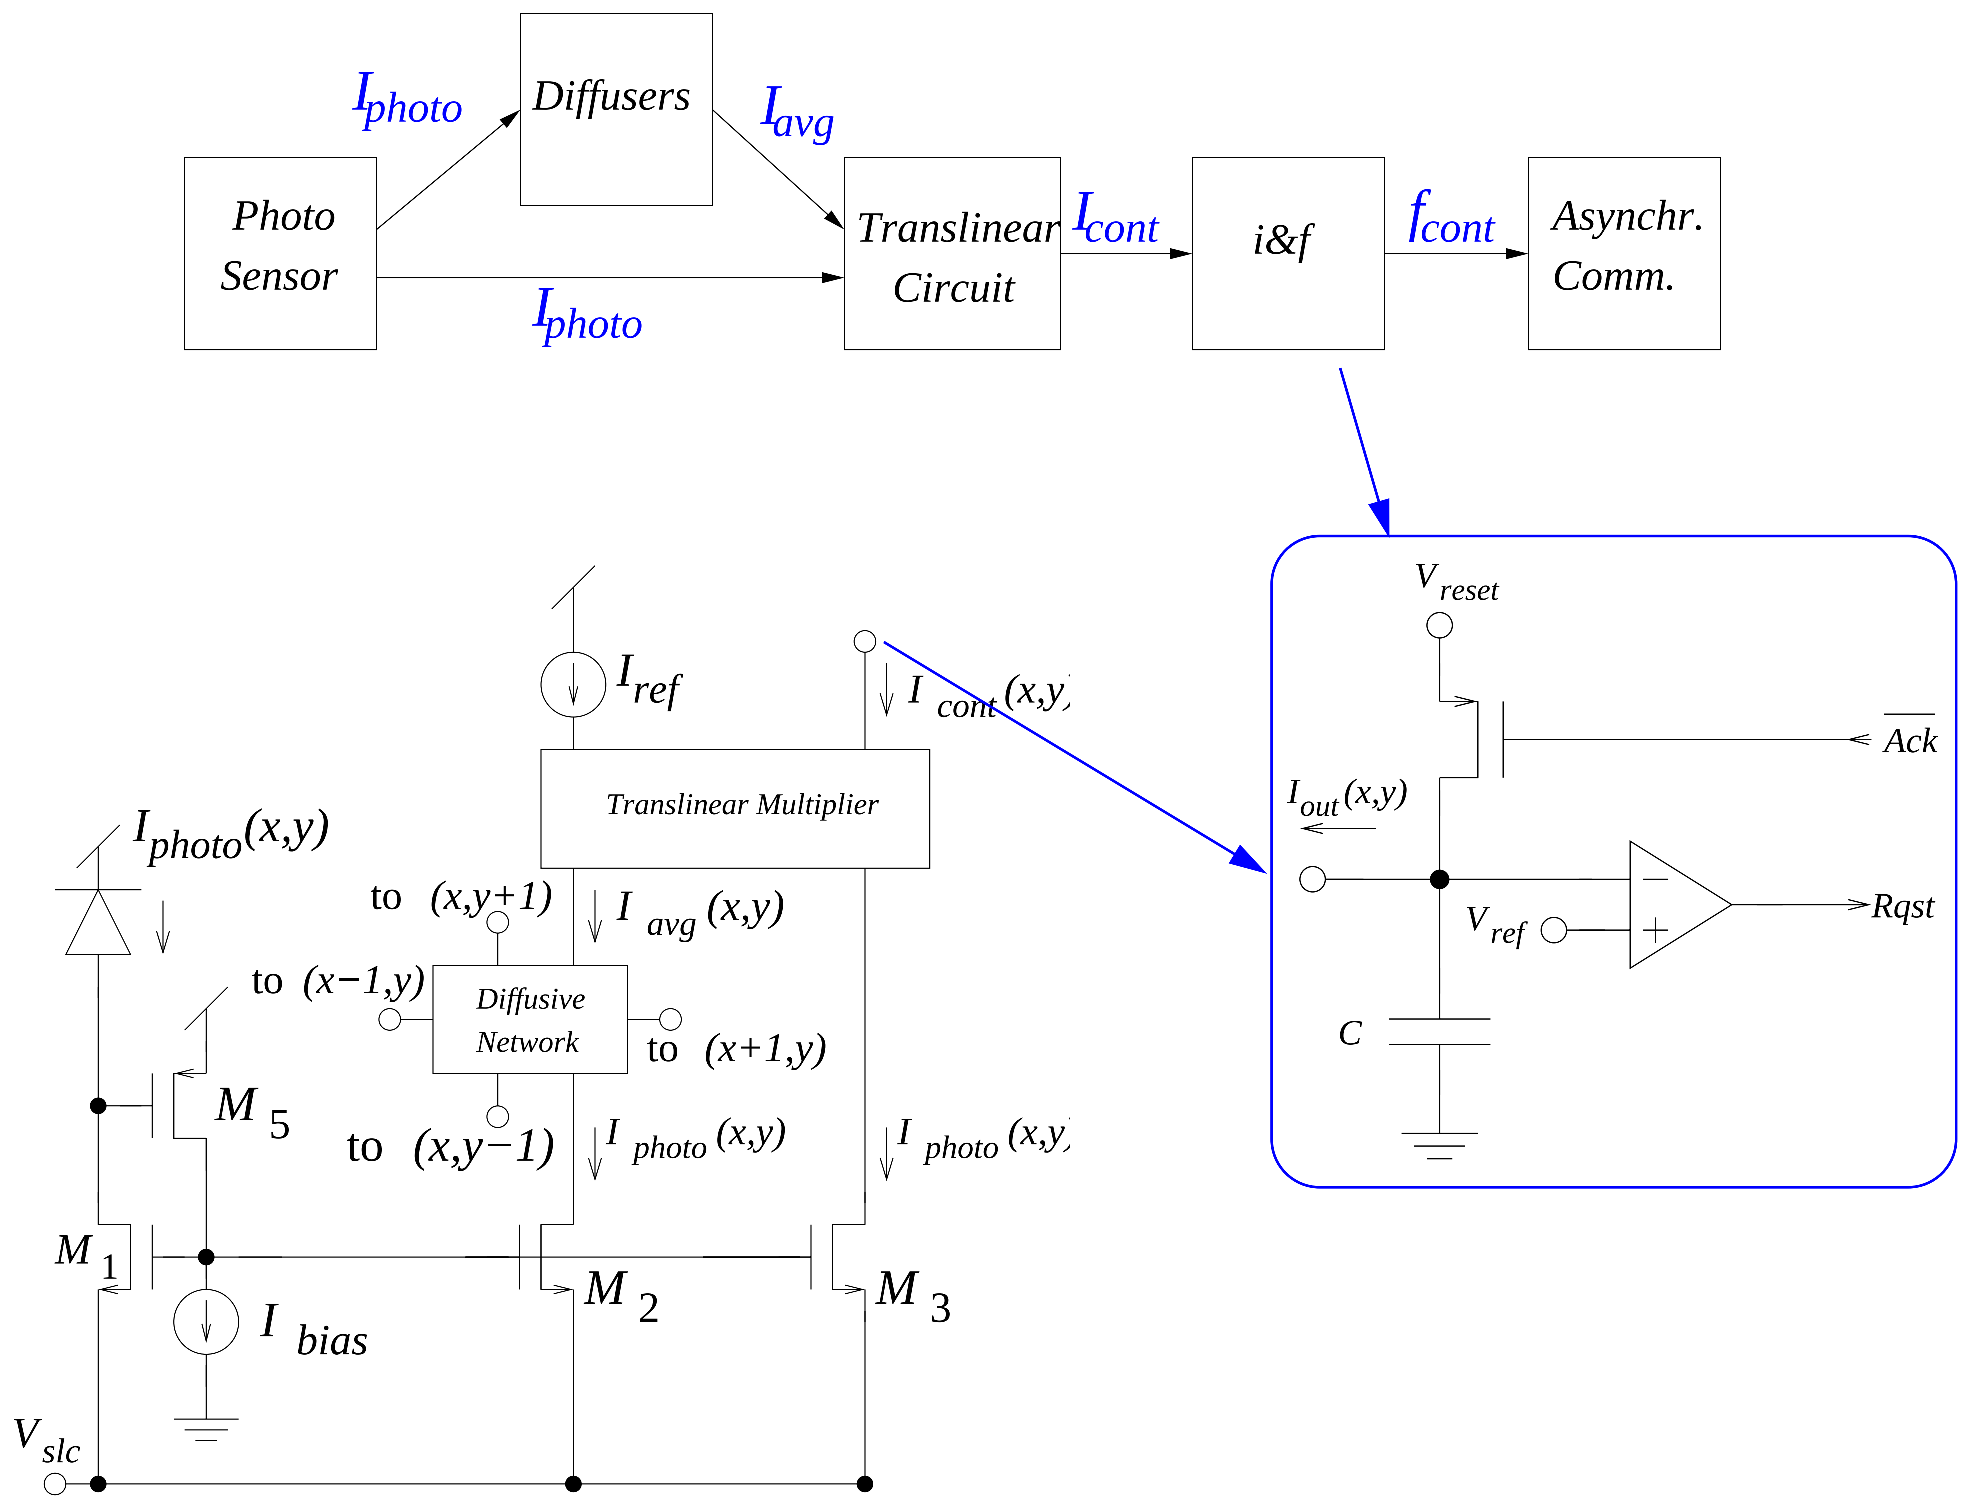

Photo Sensor
Iphoto
Diffusers
Iavg
Iphoto
Translinear Circuit
Icont
i&f
fcont
Asynchr. Comm.
Iref
Icont(x,y)
Translinear Multiplier
Iphoto(x,y)
to (x,y+1)
Iavg(x,y)
to (x−1,y)
Diffusive Network
to (x+1,y)
M5
to (x,y−1)
Iphoto(x,y)
Iphoto(x,y)
M1
M2
M3
Ibias
Vslc
Vreset
Ack
Iout(x,y)
Vref
Rqst
C

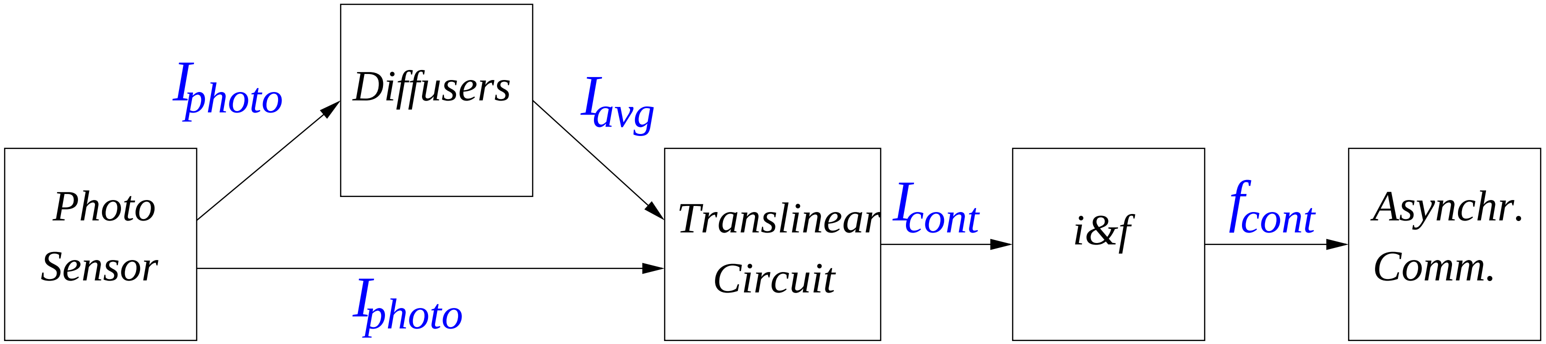

$$f_{cont}(x, y) = \frac{I_{ref}}{C(V_{reset} - V_{ref})} \frac{I_{avg}(x, y)}{I_{photo}(x, y)}$$

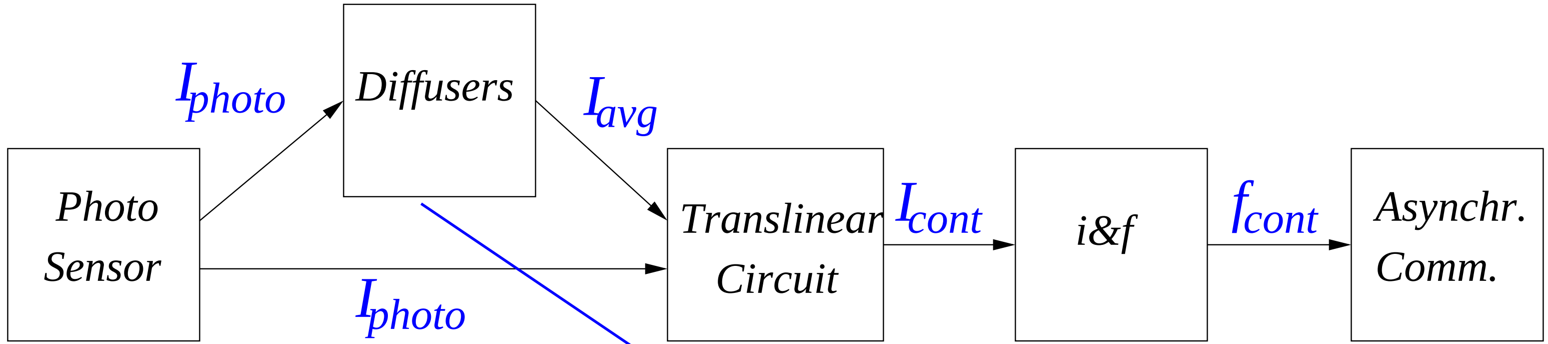

$$\begin{aligned} f_{cont}(x,y) &= \frac{I_{ref}}{C(V_{reset}-V_{ref})}\frac{I_{avg}}{I_{photo}} \times \left(1 + \frac{\Delta I_{ref}}{I_{ref}} - \frac{\Delta V}{V_{reset}-V_{ref}} - \frac{\Delta C}{C} + \frac{\Delta I_{avg}}{I_{avg}} - \frac{\Delta I_{photo}}{I_{photo}} + \Delta_{TL}\right) \\ &= f|_{nominal}(1+\Delta(x,y)) \end{aligned}$$

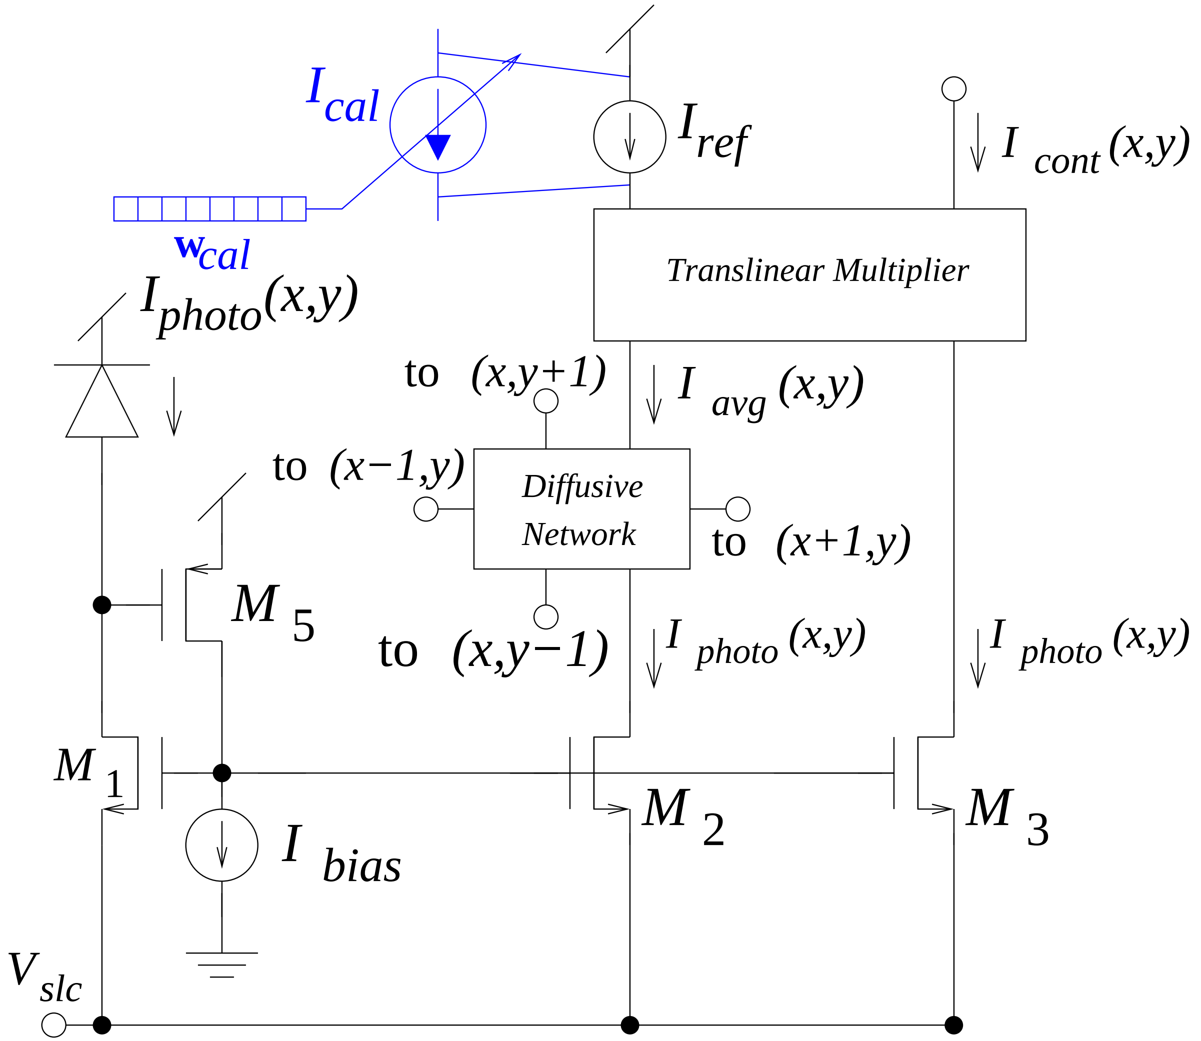

$$f_{cont}(x,y) = \frac{I_{ref}}{C(V_{reset} - V_{ref})} \frac{I_{avg}}{I_{photo}} \times \left(1 + \frac{\Delta I_{ref} + I_{cal}}{I_{ref}} - \frac{\Delta V}{V_{reset} - V_{ref}} - \frac{\Delta C}{C} + \frac{\Delta I_{avg}}{I_{avg}} - \frac{\Delta I_{photo}}{I_{photo}} + \Delta_{TL}\right)$$

$$= f|_{nominal}\left(1 + \Delta(x,y) + \frac{I_{cal}}{I_{ref}}\right)$$

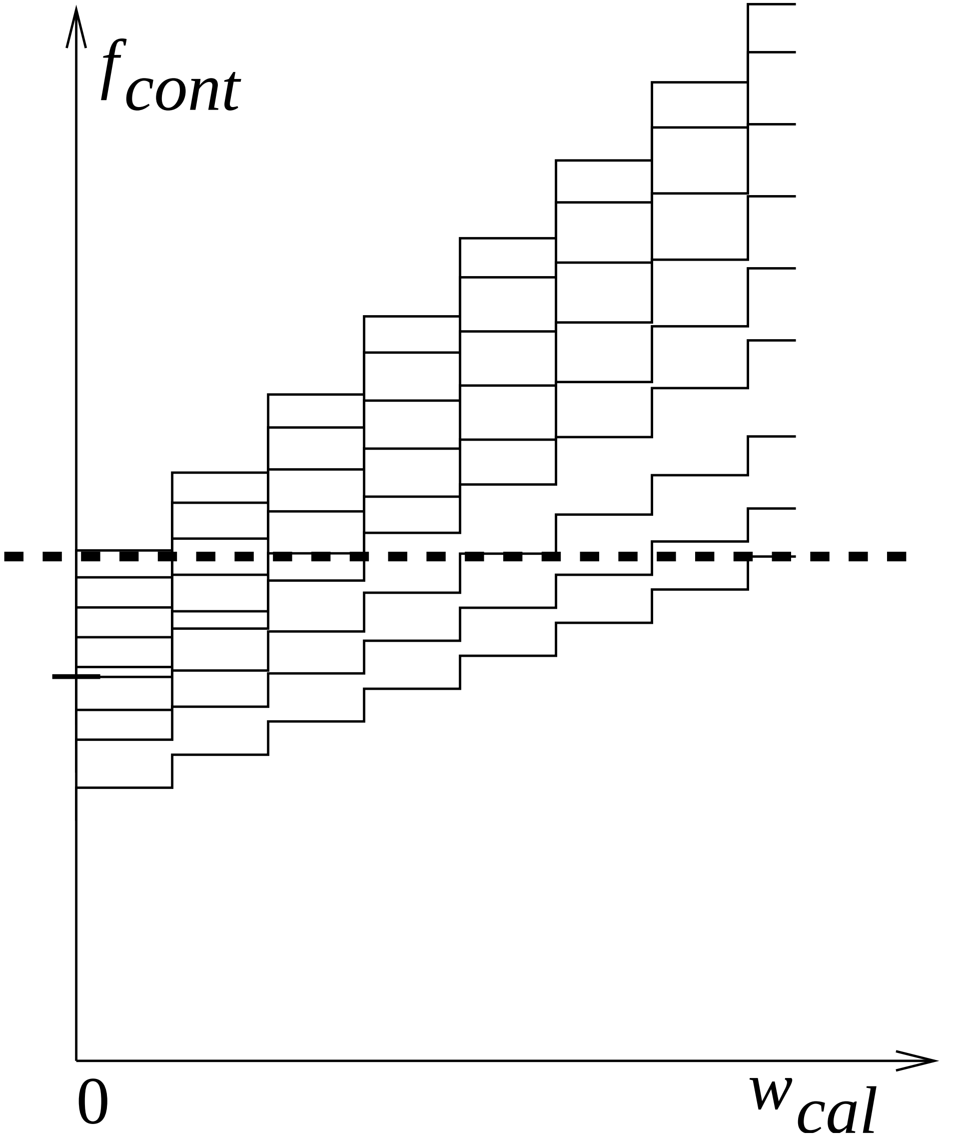

$$f_{cont}(x, y) = \frac{I_{ref}}{C(V_{reset} - V_{ref})} \frac{I_{avg}}{I_{photo}} \times \left(1 + \Delta(x, y) + \frac{I_{cal}(w_{cal})}{I_{ref}}\right)$$

CMOS test prototype in AMS 0.35µm

array size	32 x 32
pixel size	58µm x 56µm
pixel components	104 transistors + 1 capacitor
photodiode quantum efficiency	0.34 @ 450nm
fill factor	3%
pixel current consumption	20*nA* @ 1*keps*, 1*nA* @ standby
matching before calibration (indoor light)	57%
matching after calibration (indoor light)	6.6%
contrast sensitivity	10 Hz / %relative contrast @ 400Hz DC
range of diffusers	~10 pixels
noise standard deviation	~6% fluctuation of spike rate
dark current	~500fA
Handshaking cycle	15ns/ev (shorting Ack and Rqst)

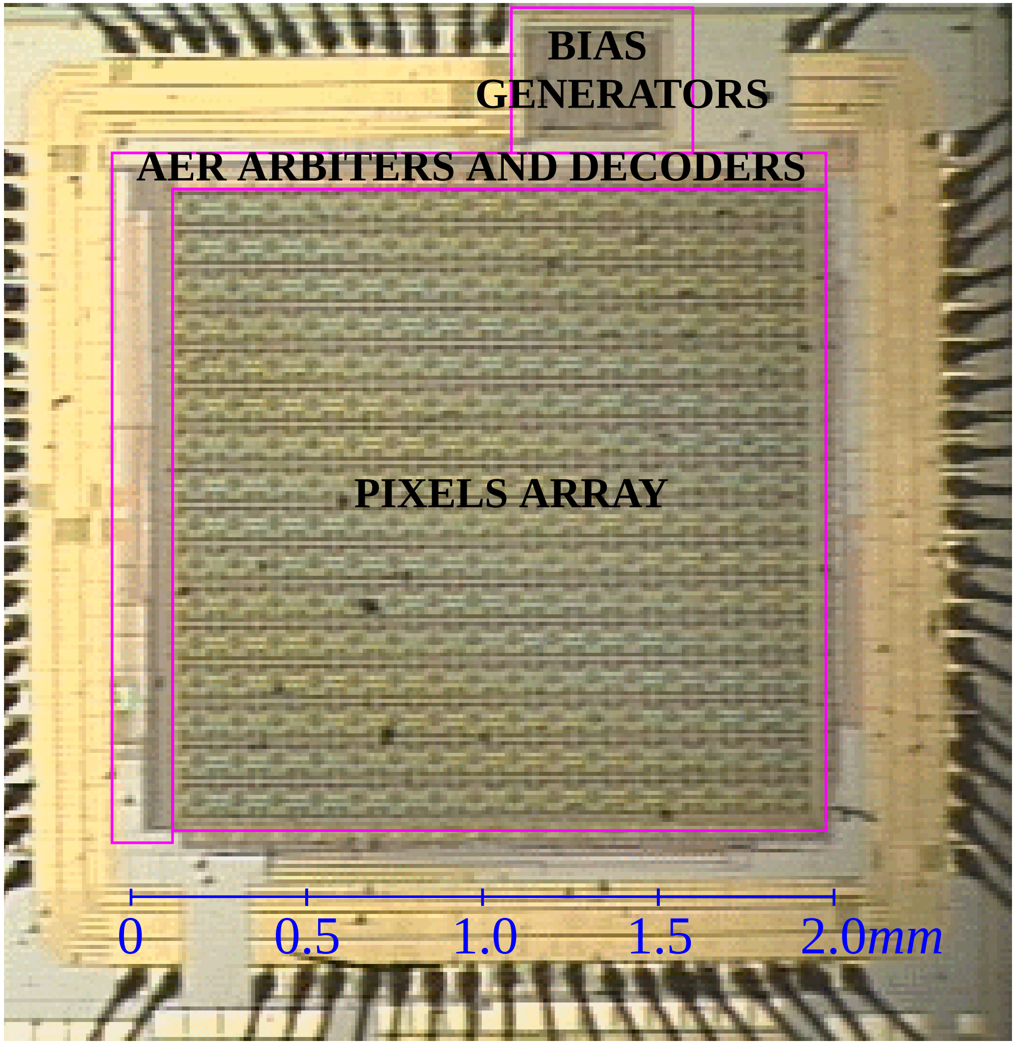

Pixel Layout

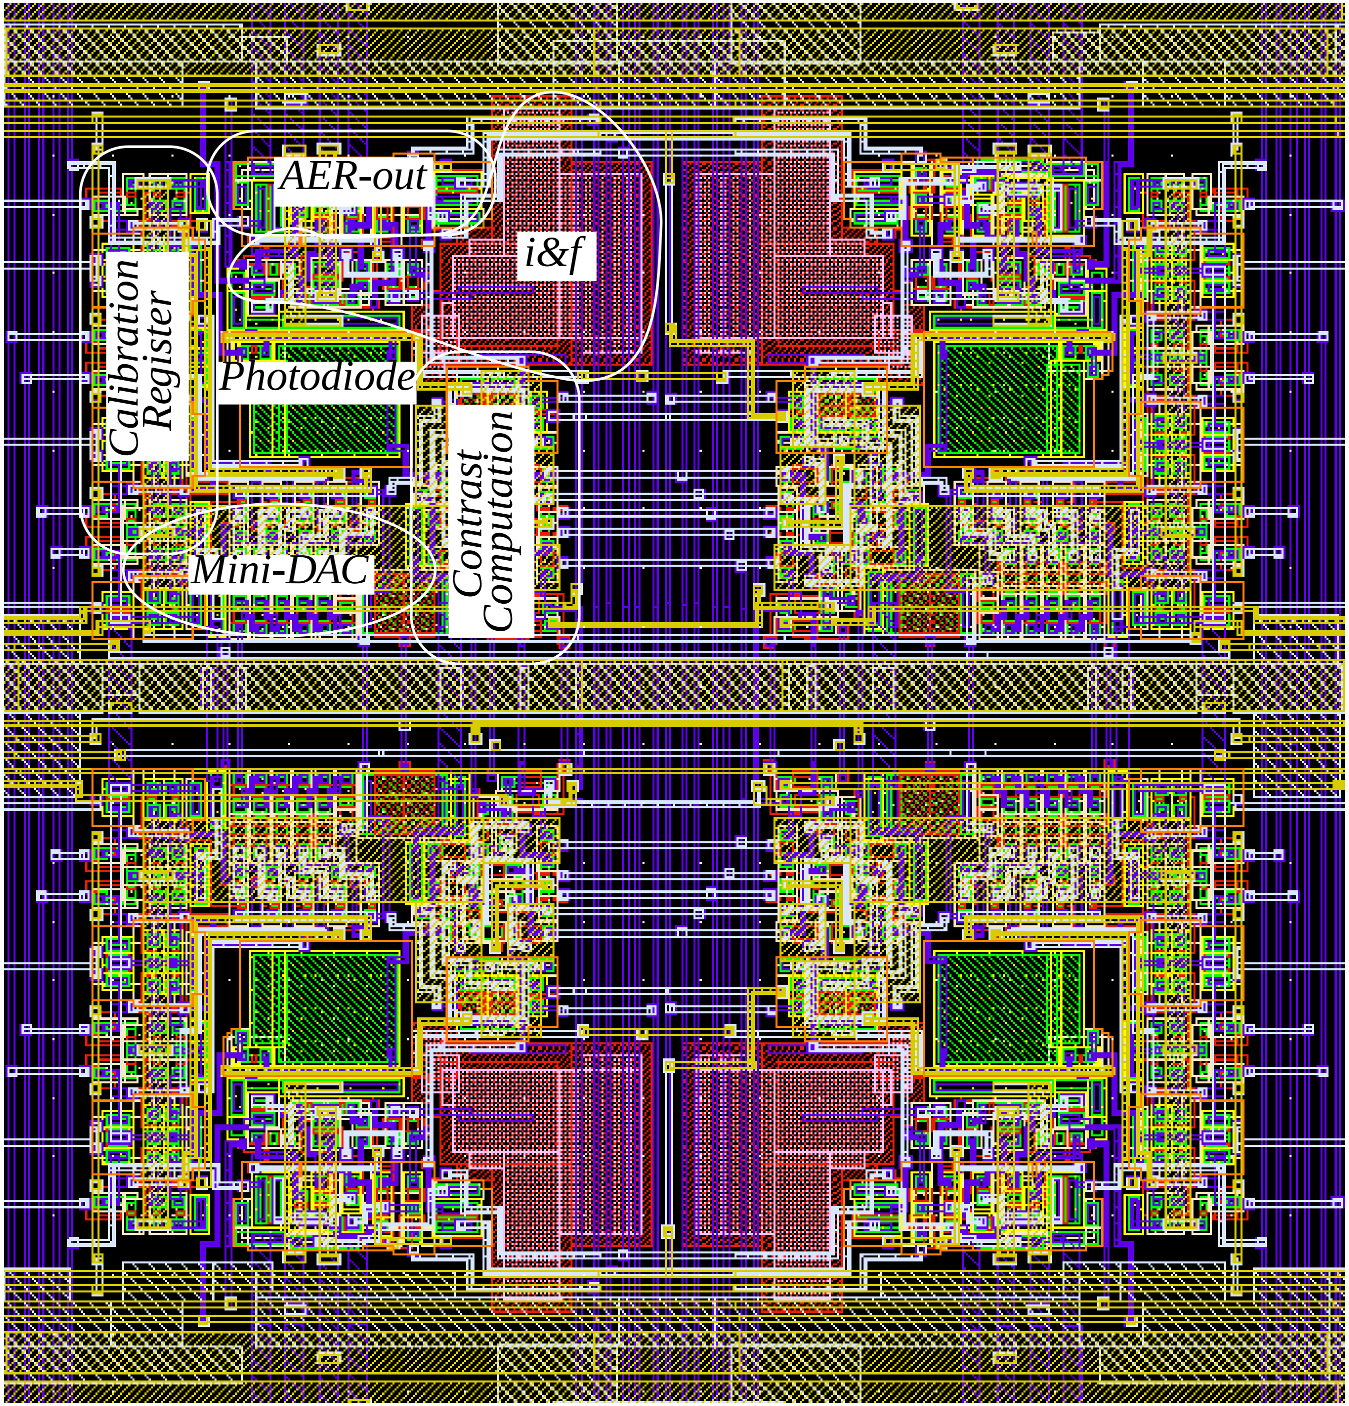

indoor light

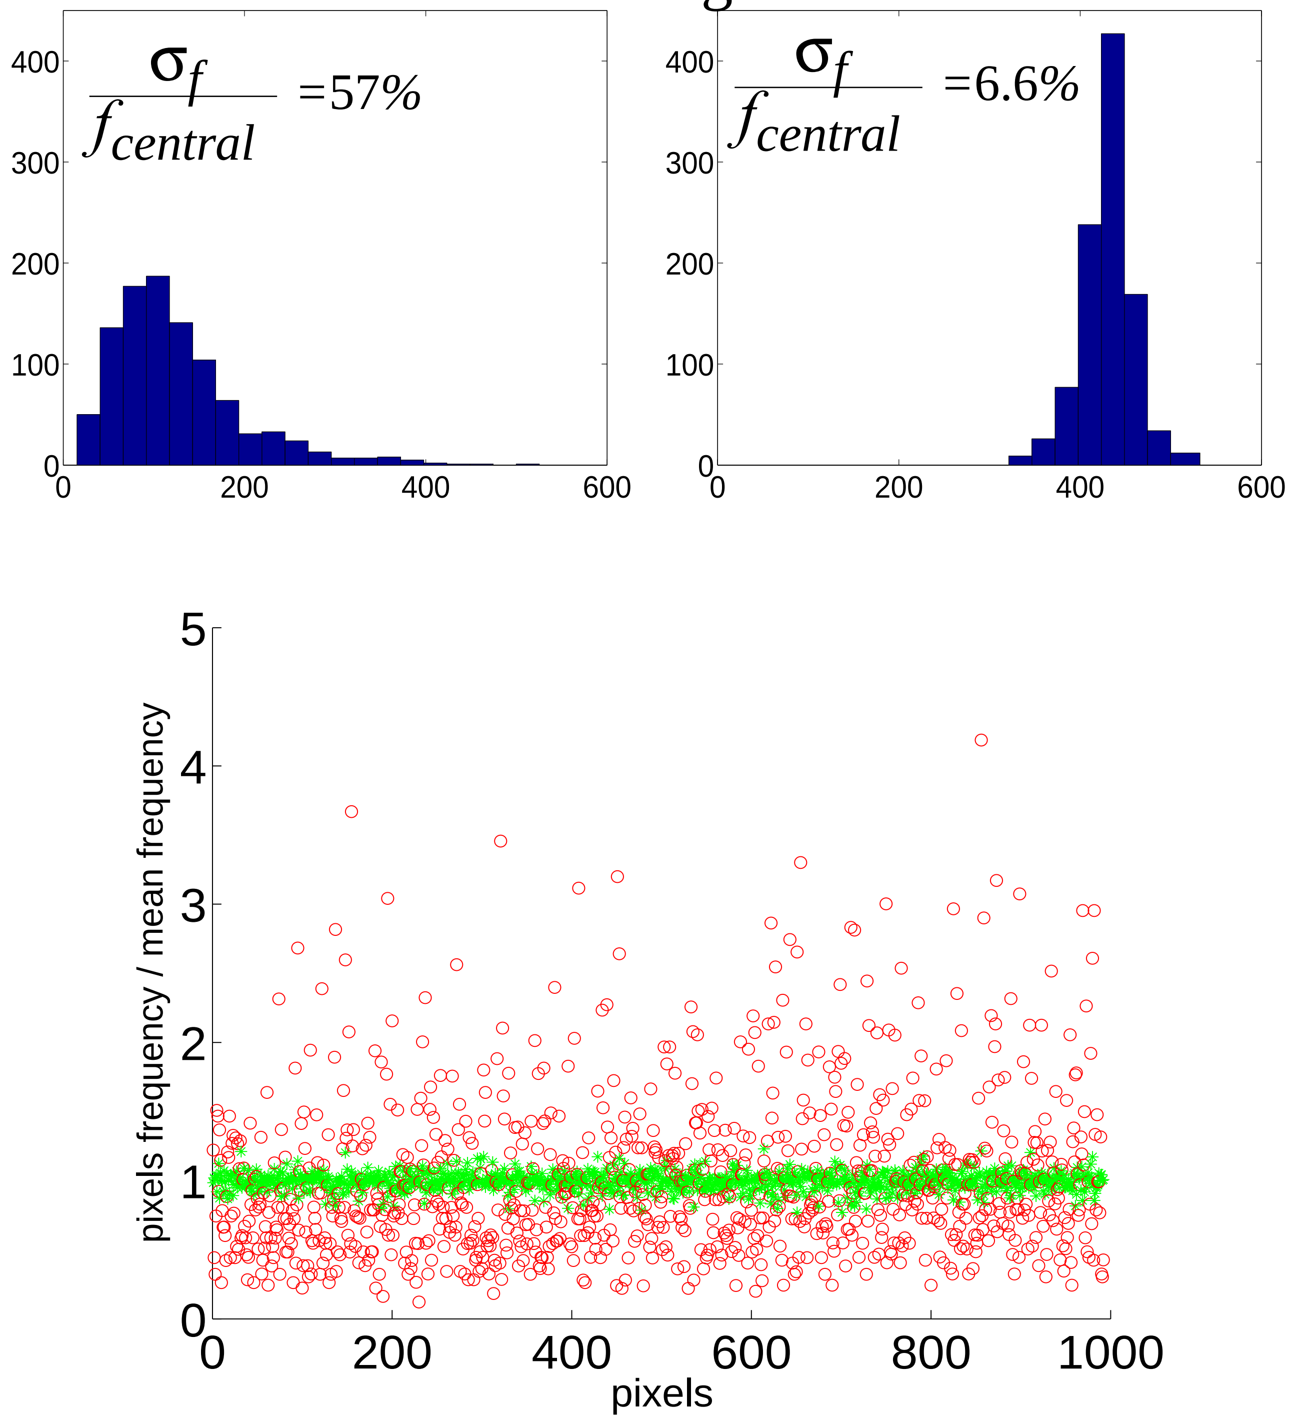

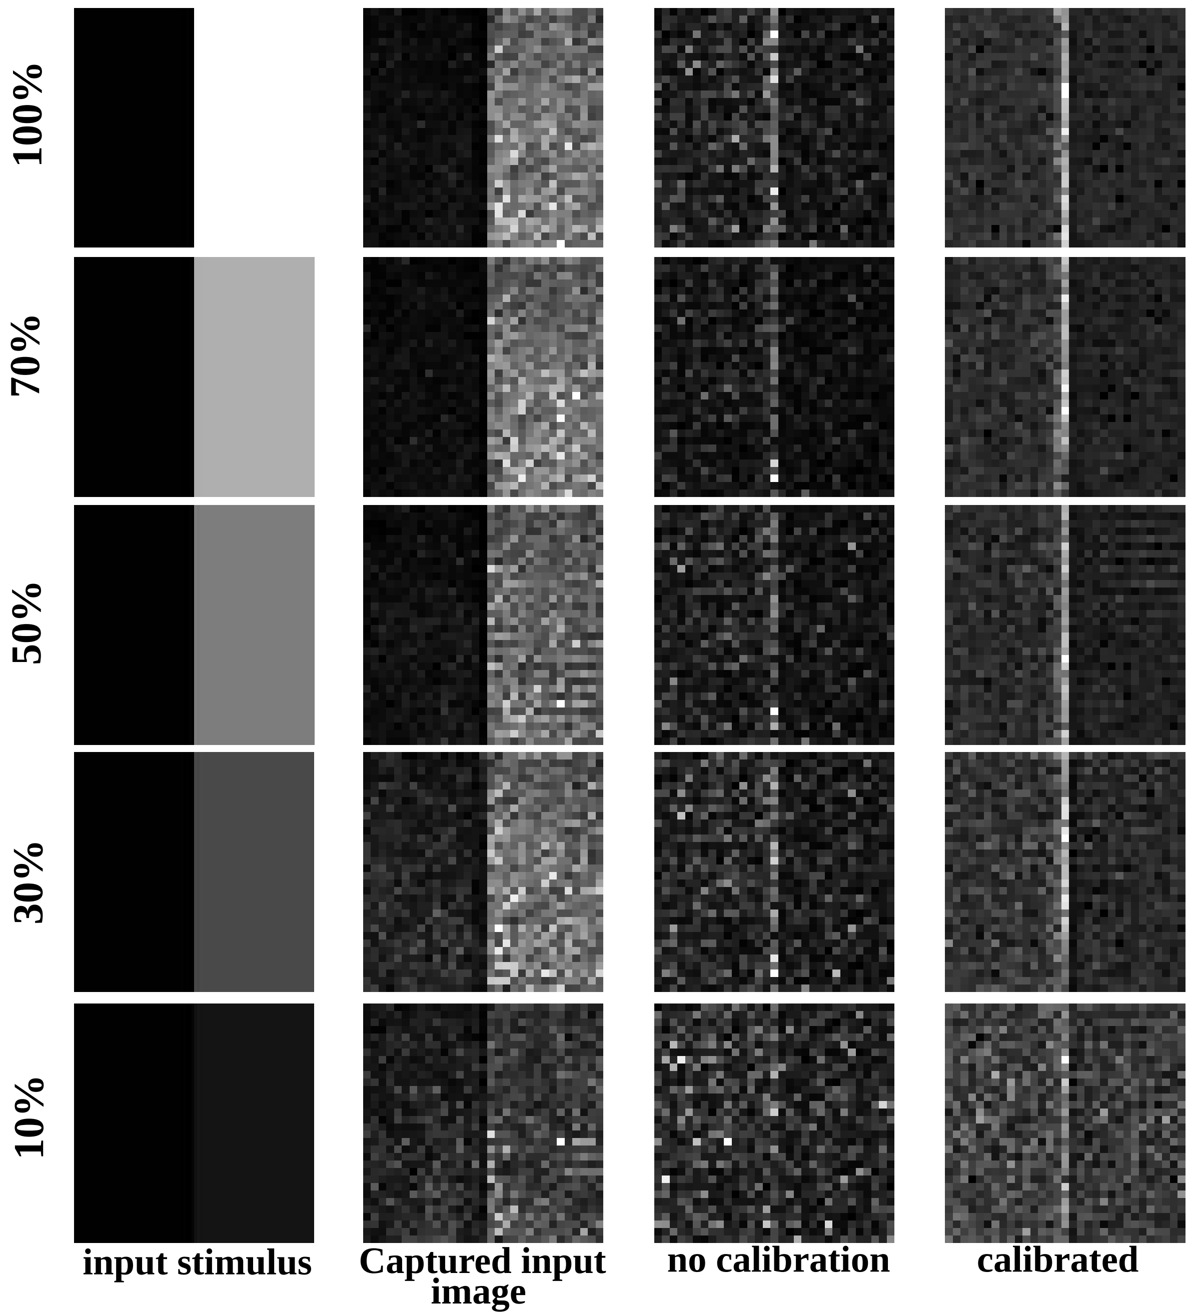
100%
70%
50%
30%
10%
input stimulus
Captured input image
no calibration
calibrated

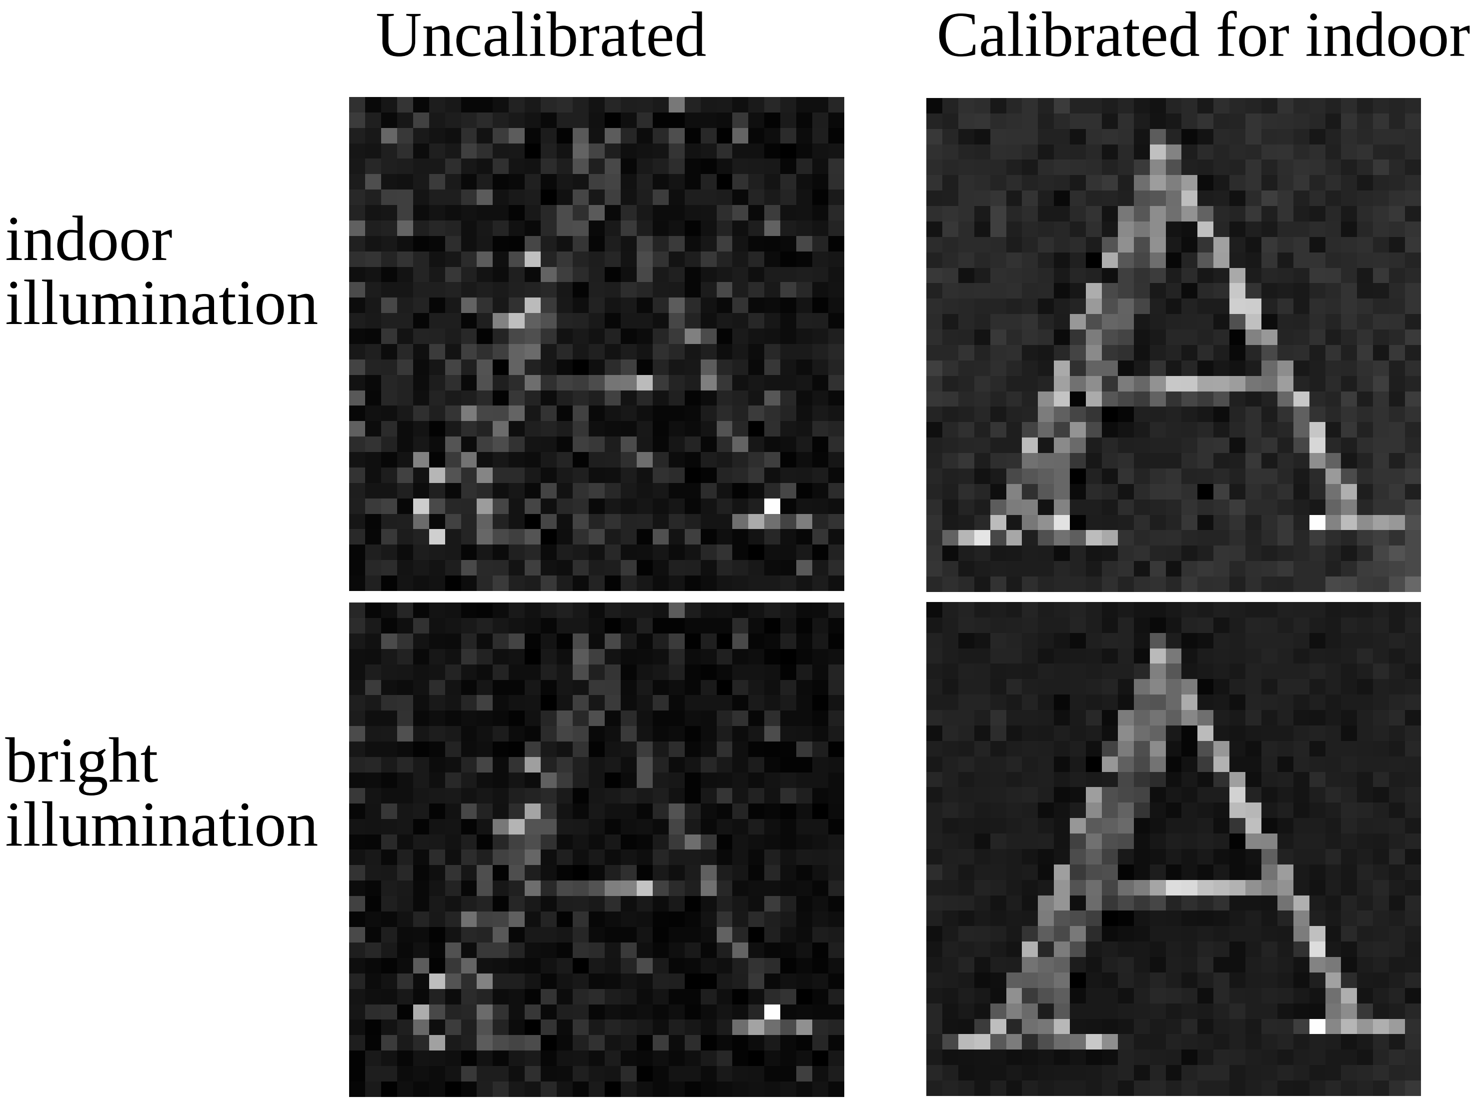
Uncalibrated
Calibrated for indoor
indoor illumination
bright illumination

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems

- Some useful circuits: - calibration

  - LVDS interface

- Some example systems at IMSE: - spatial contrast retina

  - mixed-mode convolution chip

  - fully digital convolution chip

- HW Tools from Sevilla: - some FPGA-based PCBs

  - example use in CAVIAR

- SW Tool: - Behavioral Matlab Simulator

  - Example 1: neocognitron emulation

  - Example 2: texture classification

AER Convolution Chip

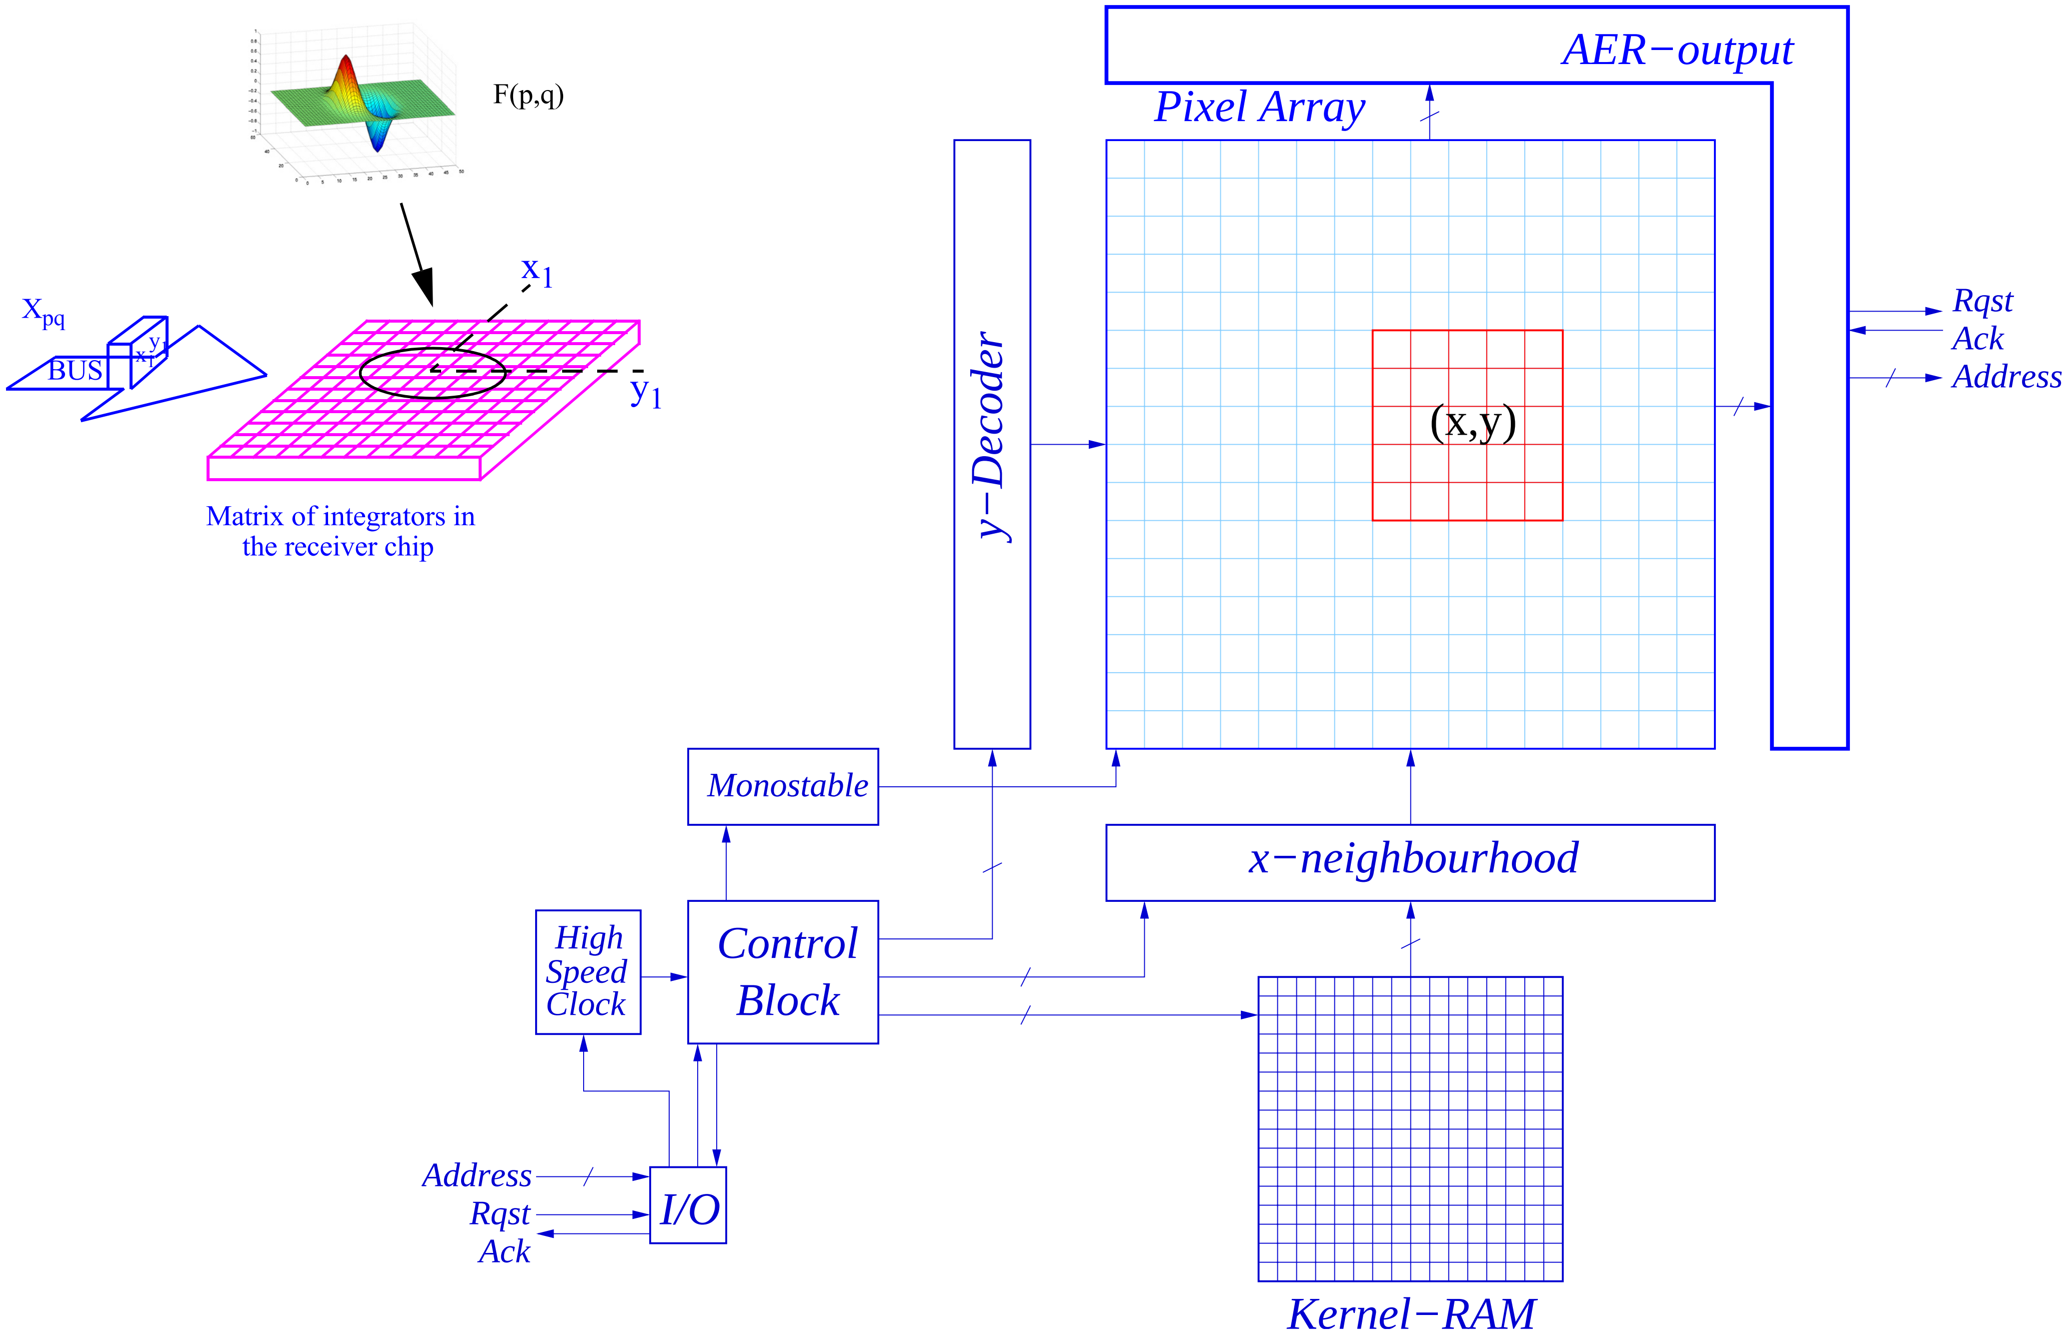

Pixel

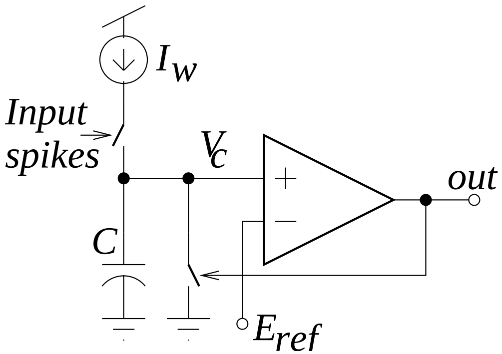
Iw
Input
spikes
Vc
out
C
Eref

(a)

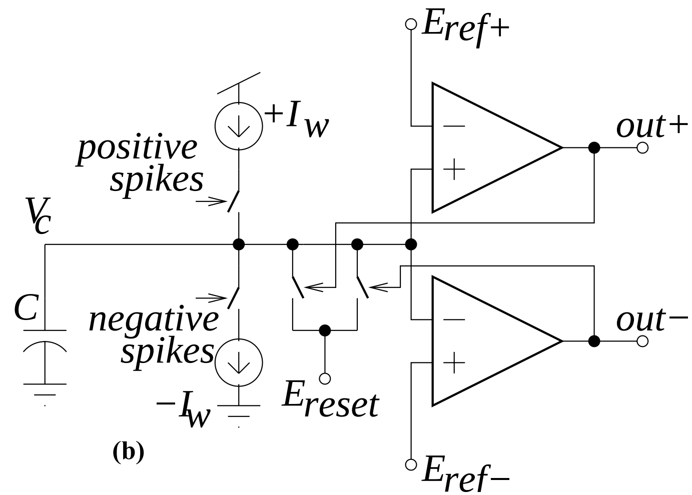
Eref+
+Iw
positive
spikes
out+
Vc
C
negative
spikes
out−
−Iw
Ereset
Eref−

(b)

Pixel

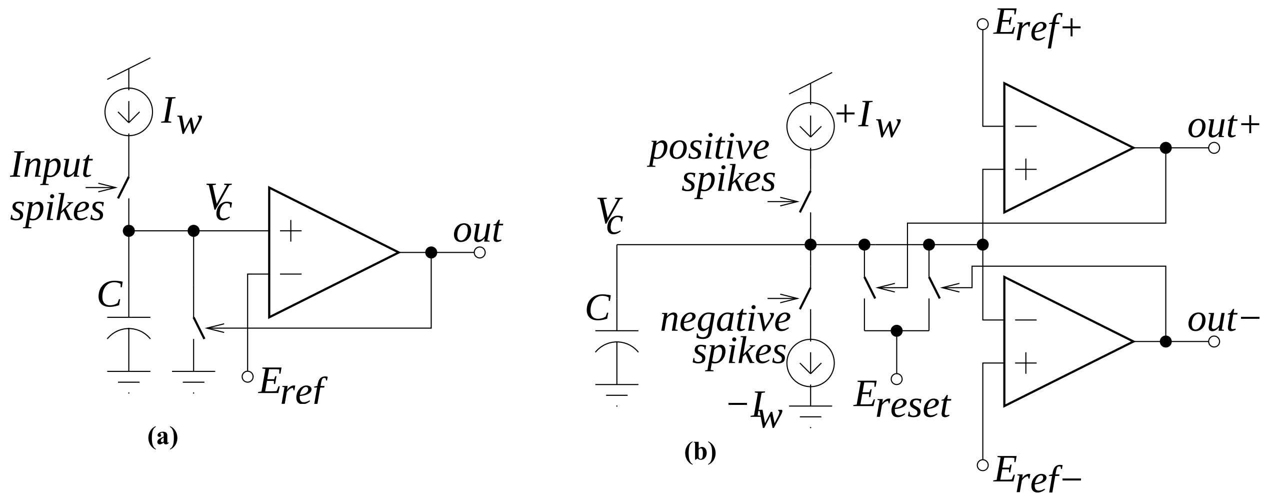

I_w
Input spikes
V_c
out
C
E_{ref}
(a)
E_{ref+}
$+I_w$
positive spikes
out+
V_c
C
negative spikes
$-I_w$
E_{reset}
out−
E_{ref-}
(b)

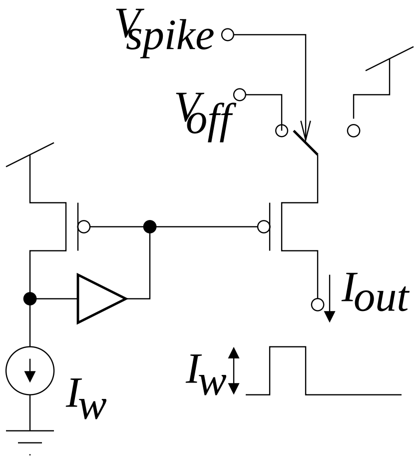

V_{spike}
V_{off}
I_{out}
I_w
I_w

Pixel

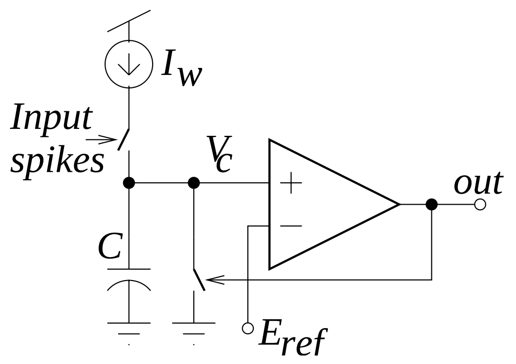

I_w
Input spikes
V_c
out
C
E_{ref}

(a)

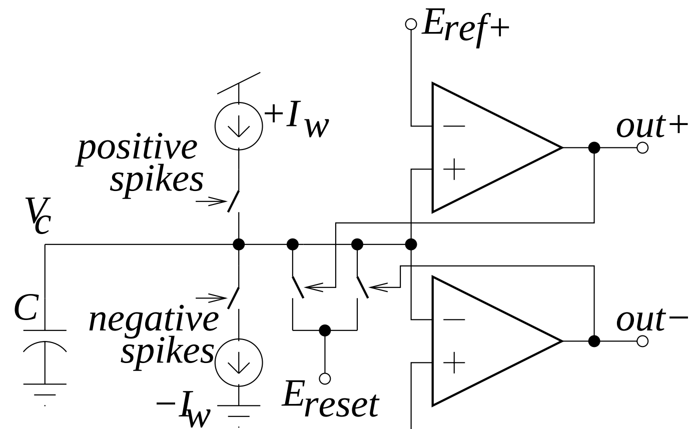

E_{ref+}
$+I_w$
positive spikes
out+
V_c
C
negative spikes
$-I_w$
E_{reset}
out−
E_{ref-}

(b)

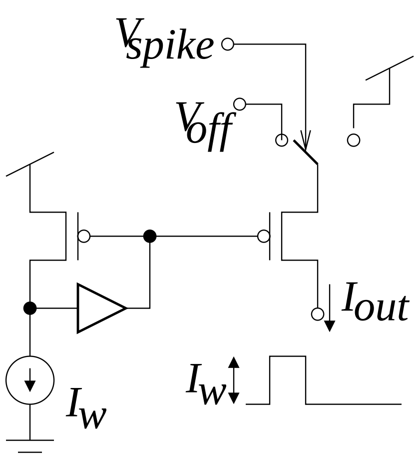

V_{spike}
V_{off}
I_{out}
I_w

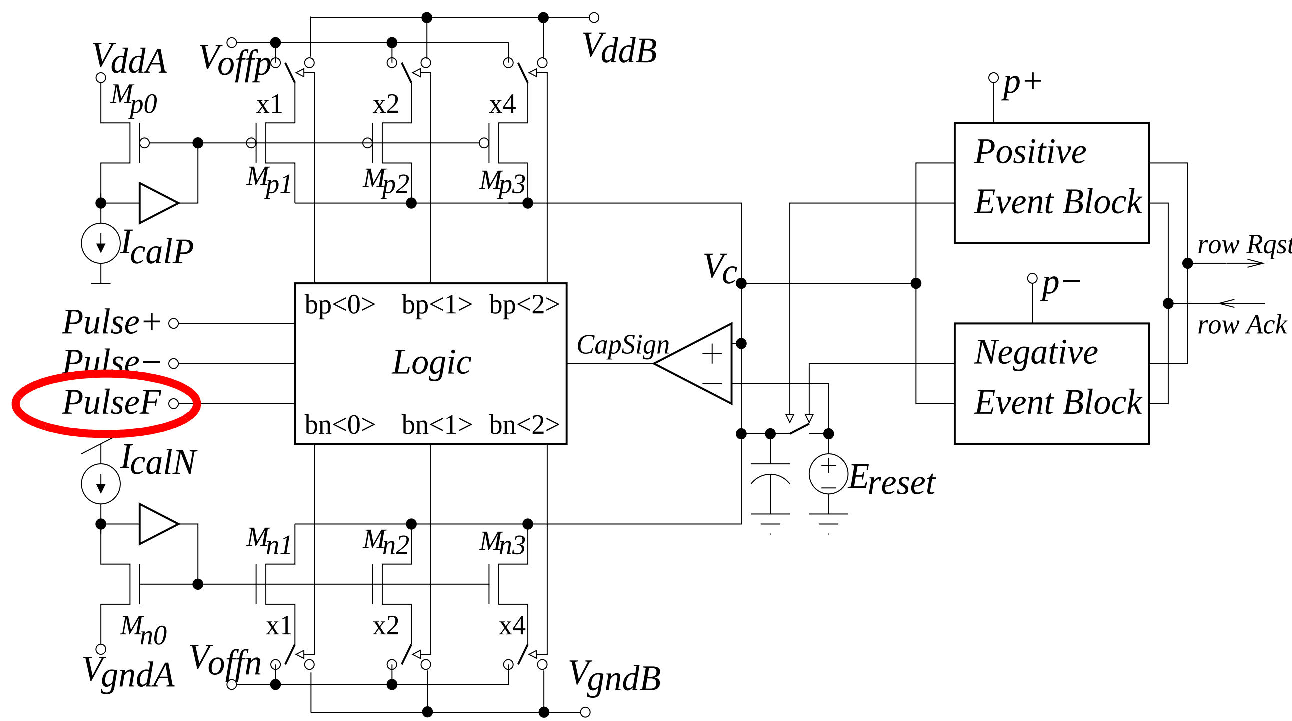

VddA
Voffp
VddB
M_{p0}
x1
x2
x4
M_{p1}
M_{p2}
M_{p3}
I_{calP}
p+
Positive Event Block
V_c
row Rqst
p−
row Ack
Pulse+
Pulse−
PulseF
bp<0> bp<1> bp<2>
Logic
bn<0> bn<1> bn<2>
CapSign
Negative Event Block
I_{calN}
Ereset
M_{n1}
M_{n2}
M_{n3}
M_{n0}
VgndA
Voffn
VgndB

Pixel Calibration

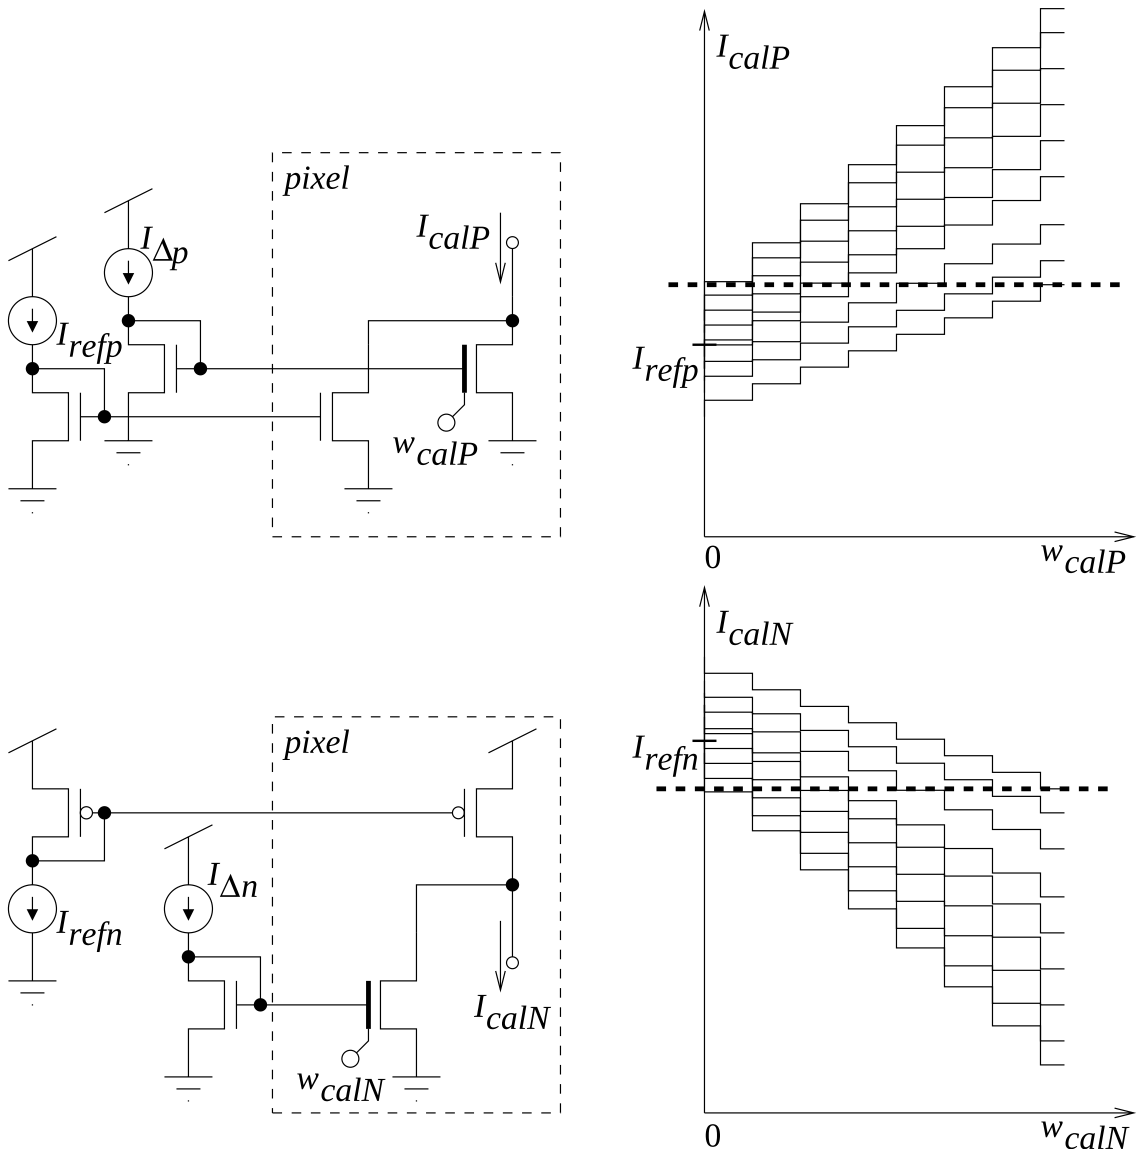

- pixel current pulses may range from ~1pA to ~1µA

Fabricated Chip

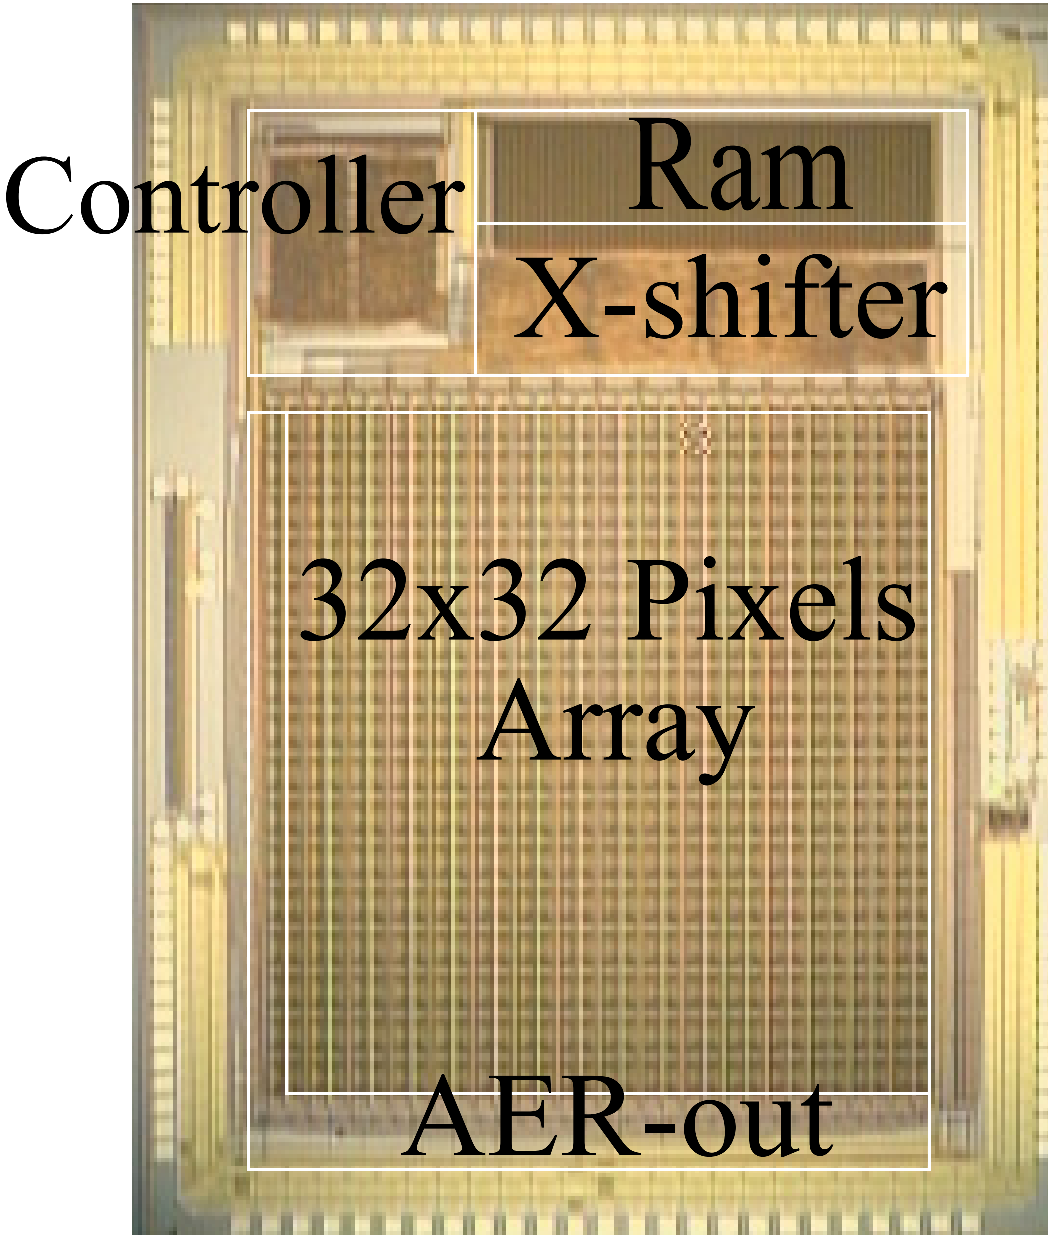

Pixel Layout

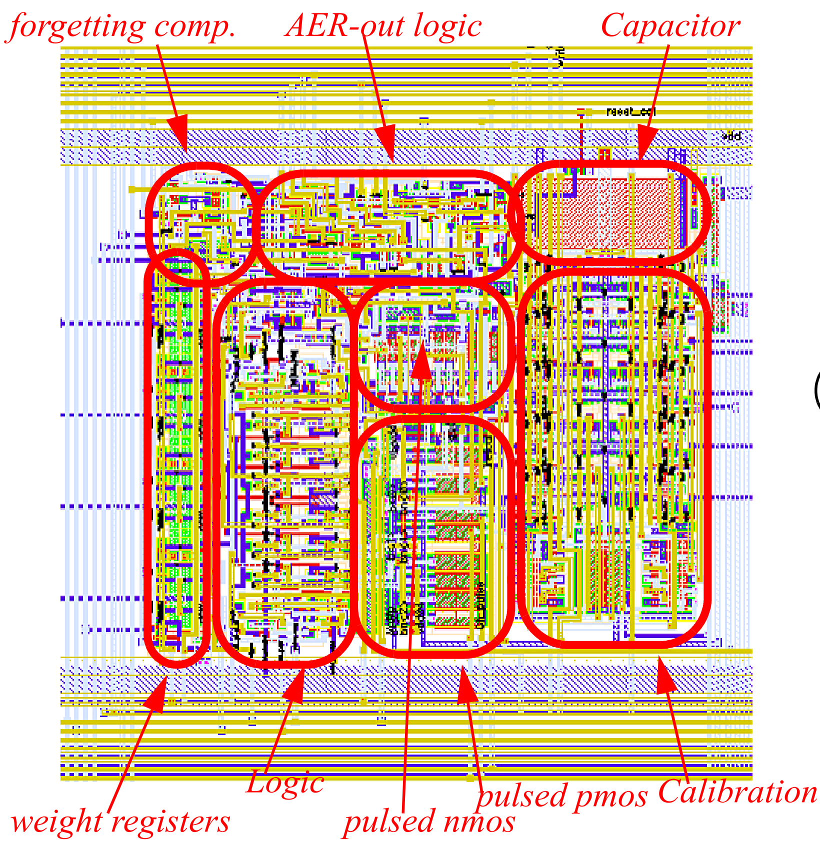

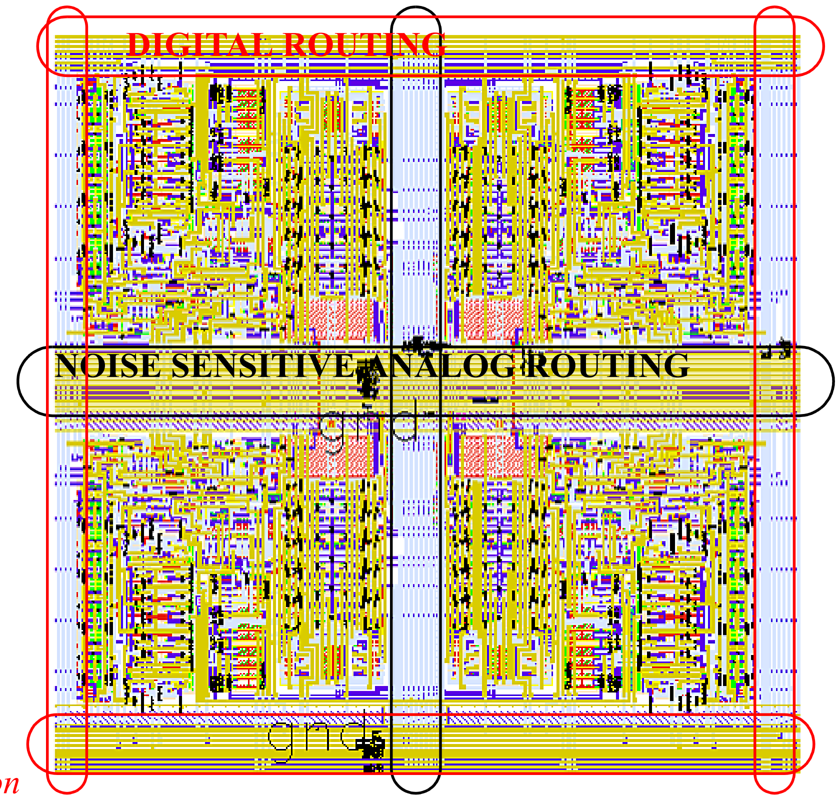

- pixel size 90µm x 90µm
- 364 transistors + 1 capacitor
- kernel weight resolution: 4 bit
- calibration register resolution: 5 bit
- interpixel mismatch (after calibration): < 2%

kernel

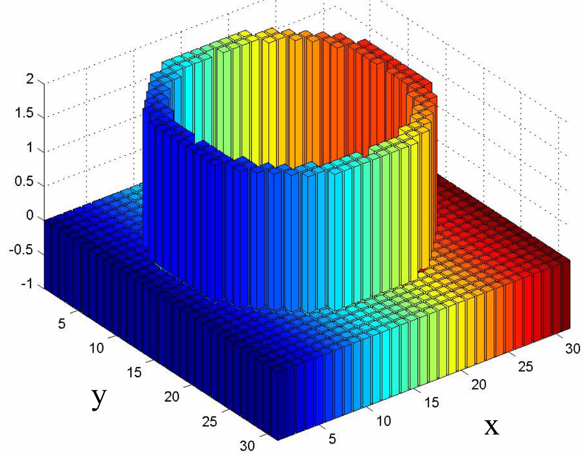

Input A

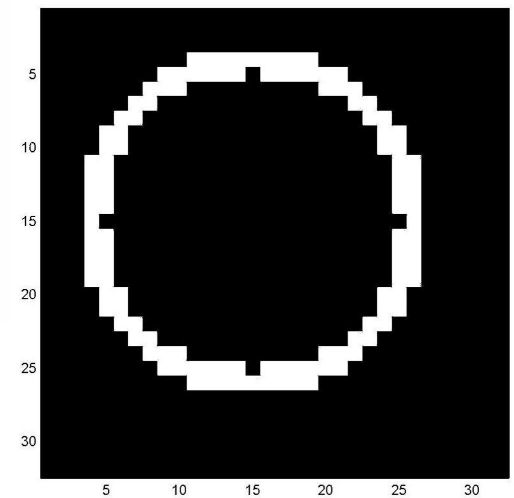

Output 1

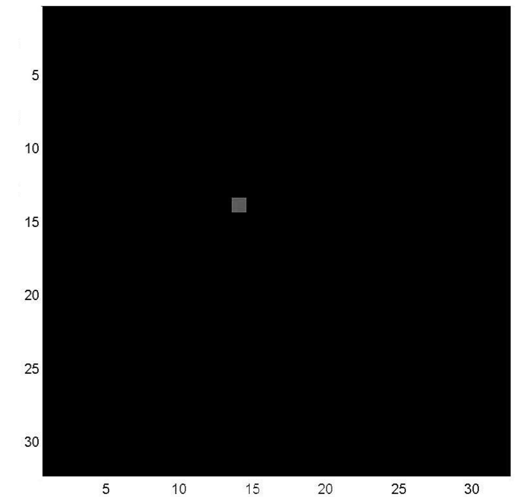

Input B

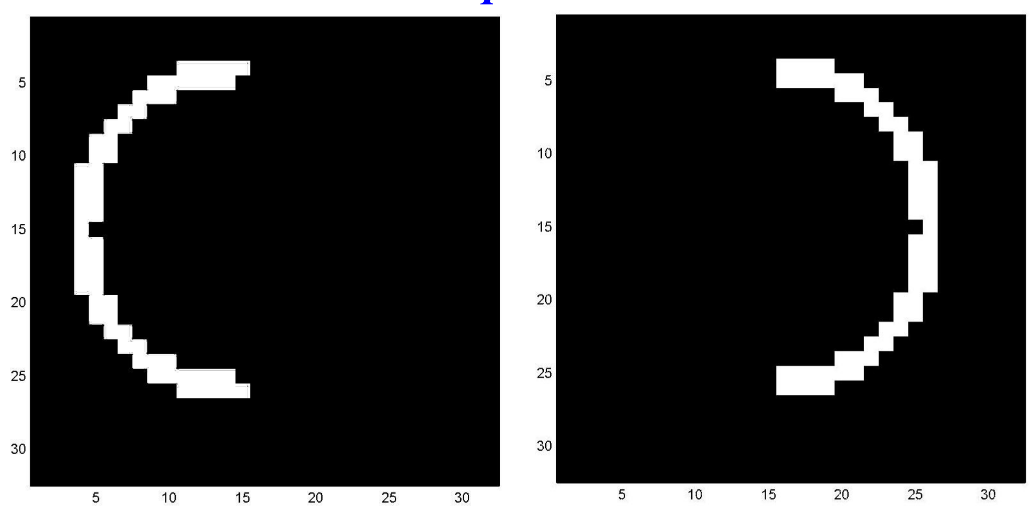

Output 2

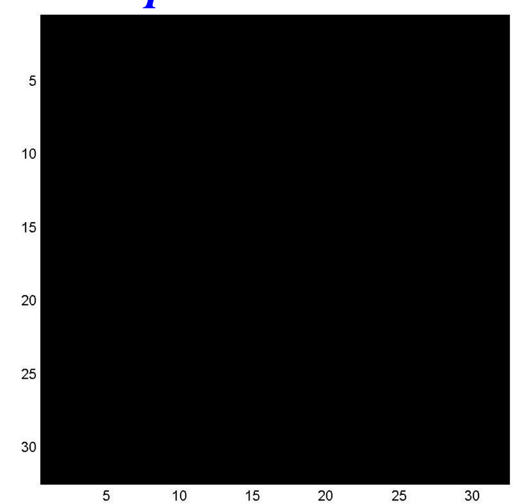

PCB with 4 32x32 Conv. Chips + Event Routing

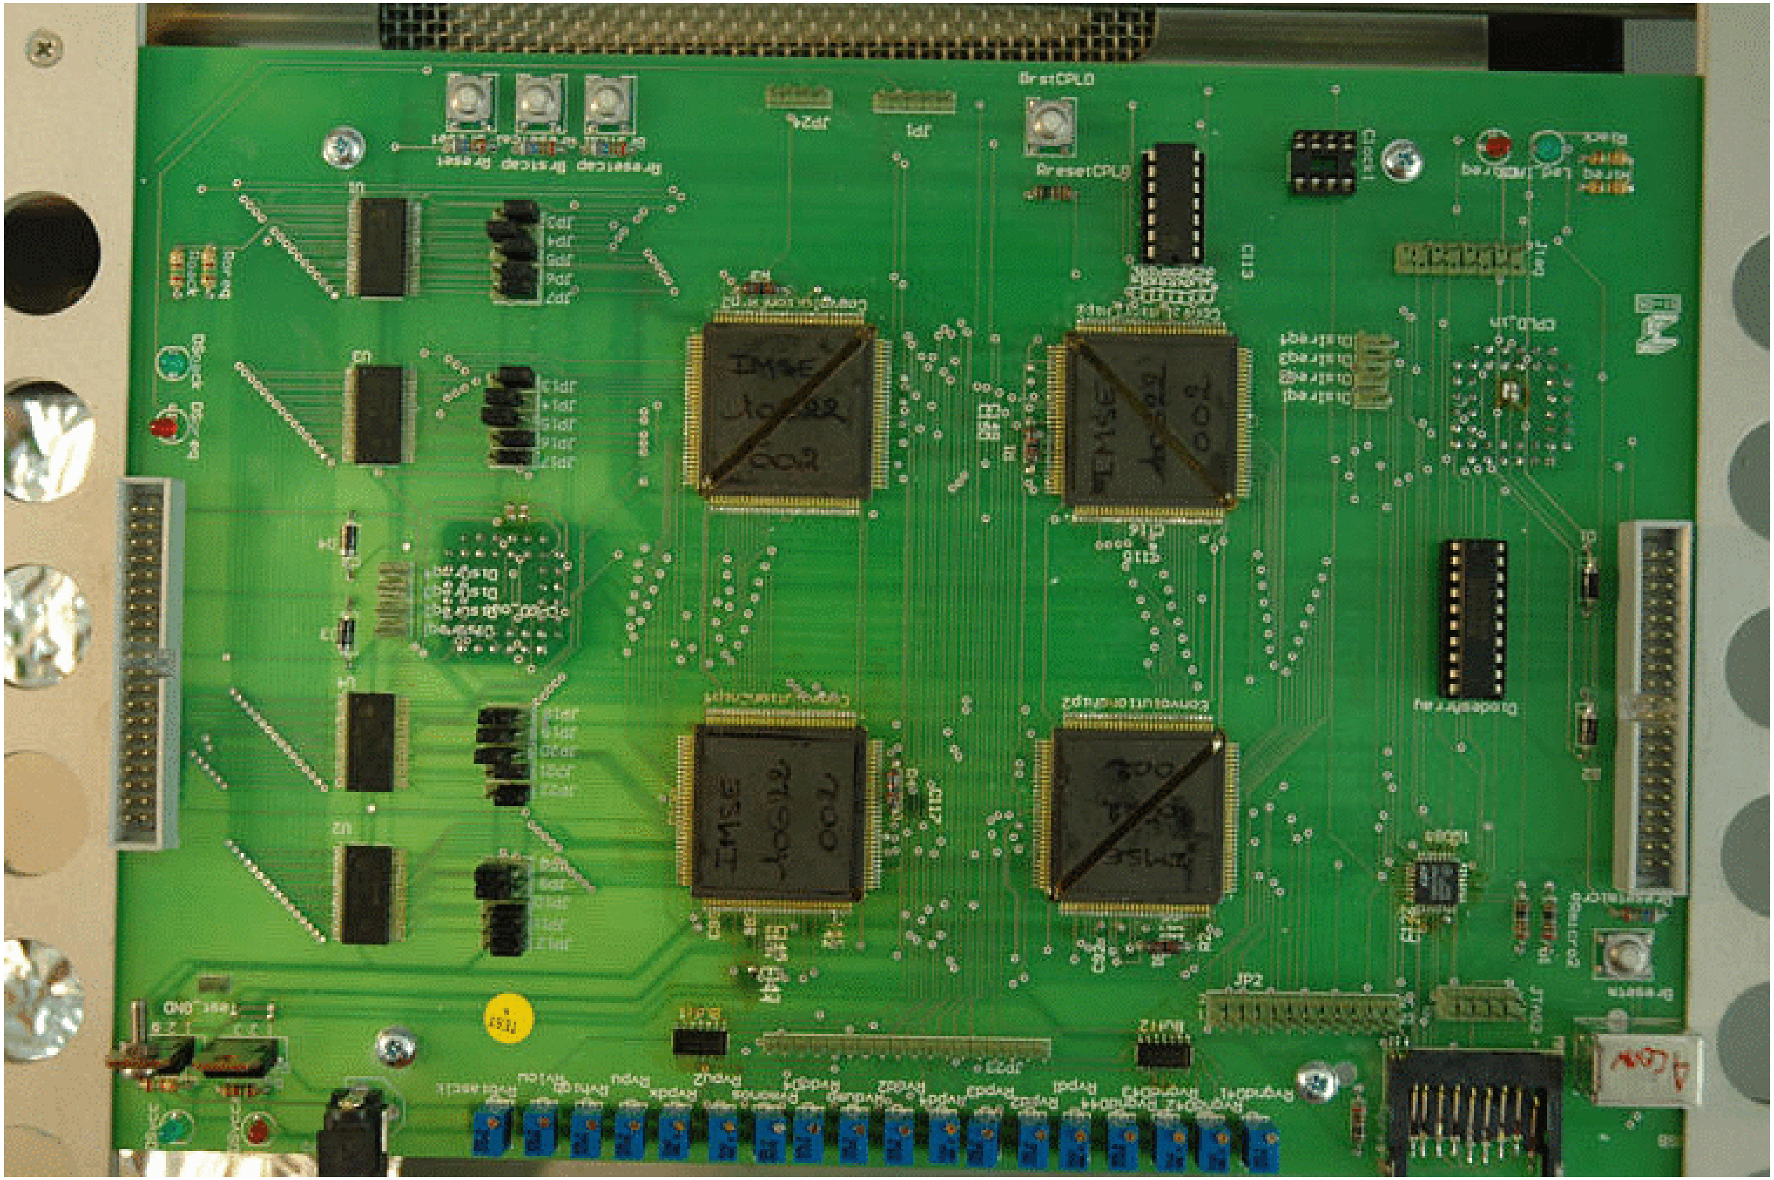

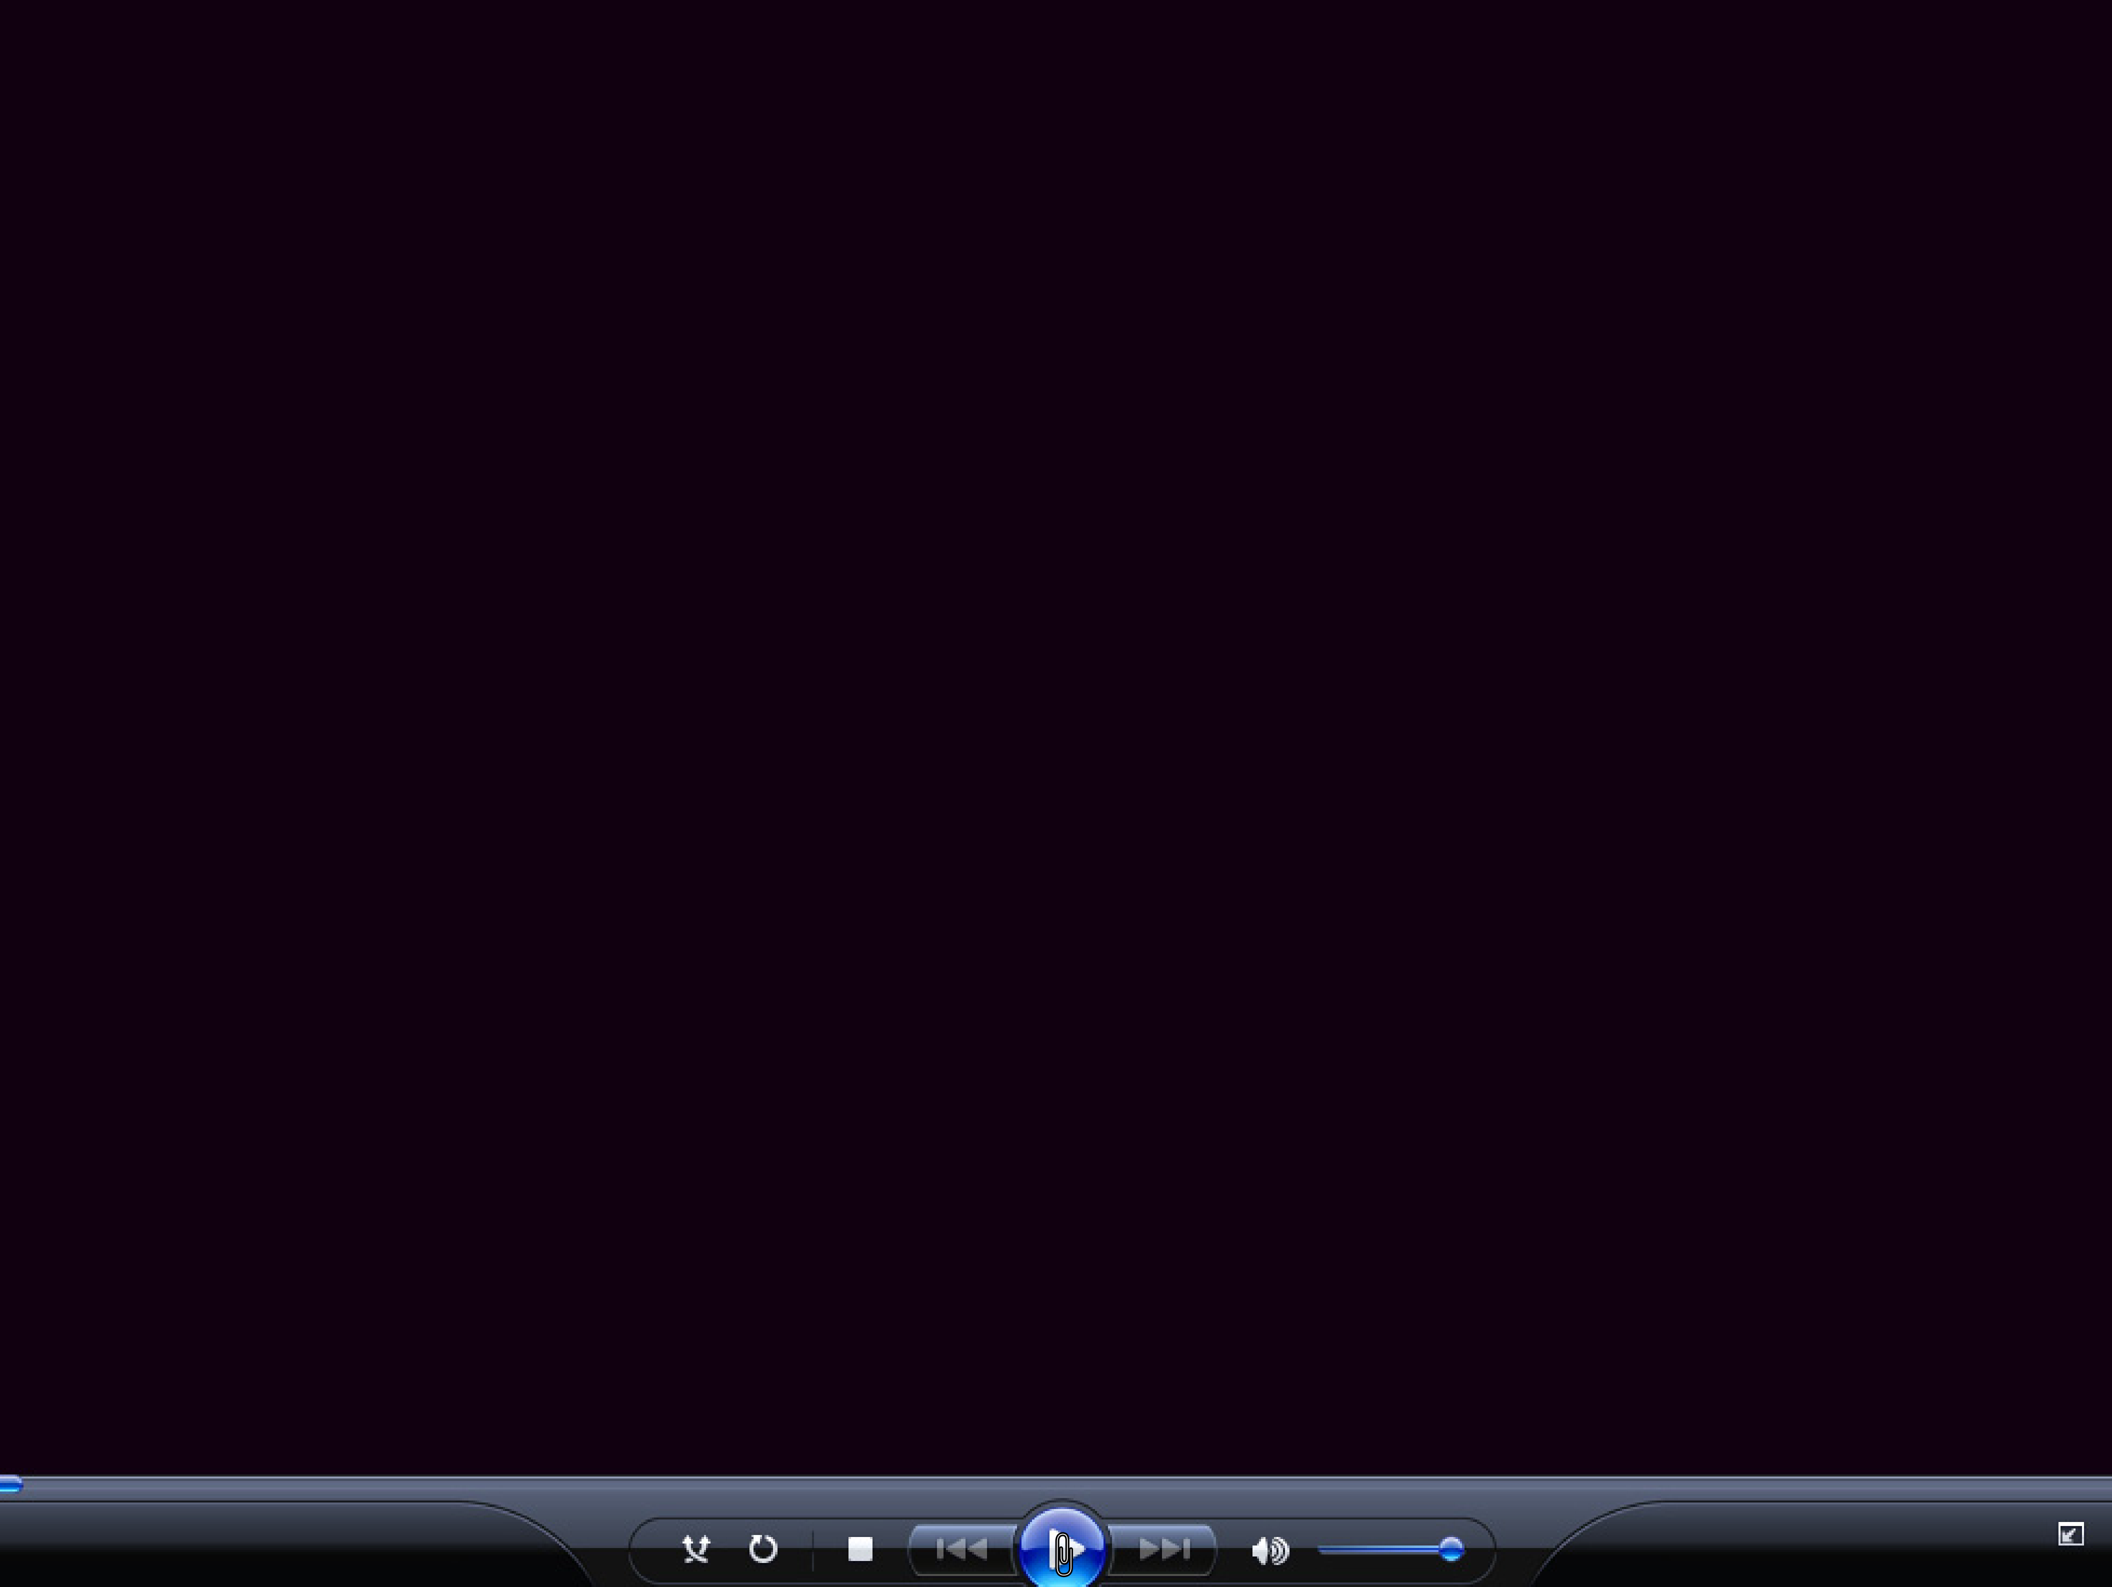

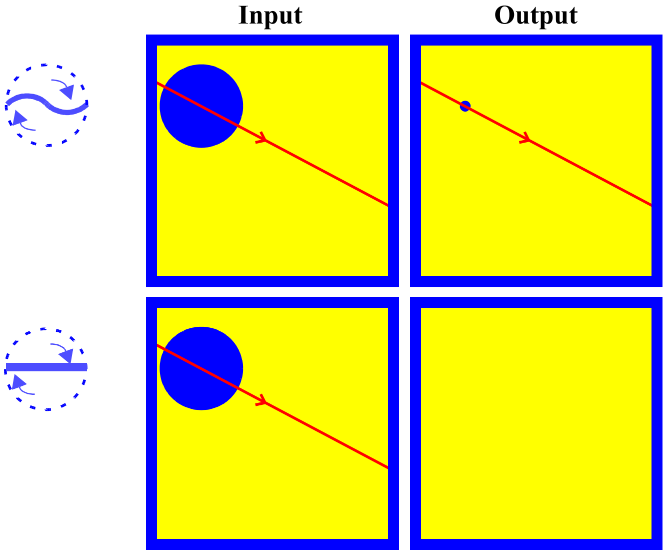
Input
Output

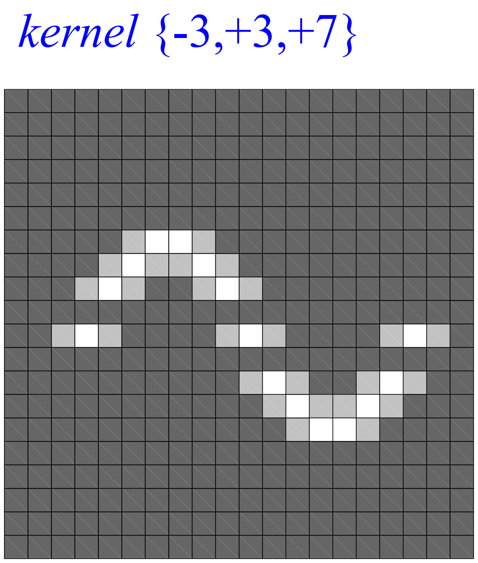
kernel {-3,+3,+7}

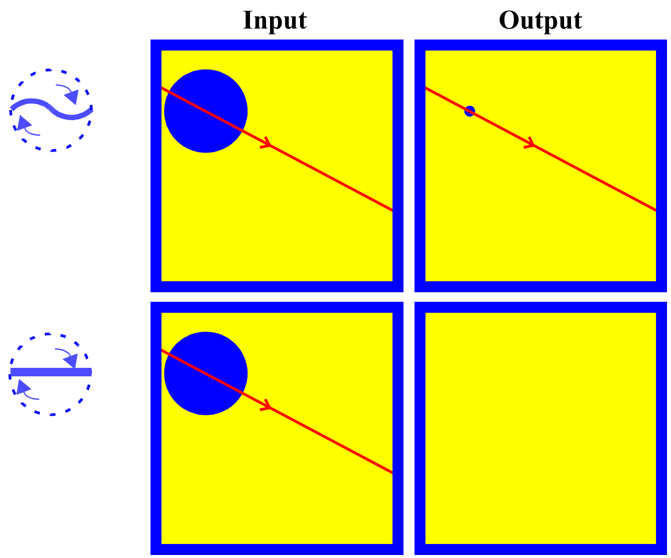
Input
Output

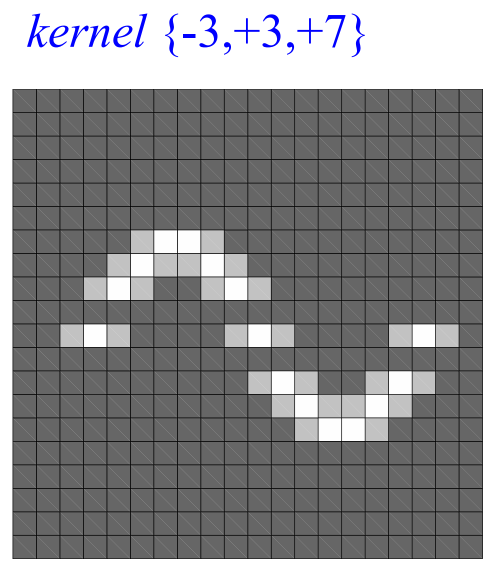
kernel {-3,+3,+7}

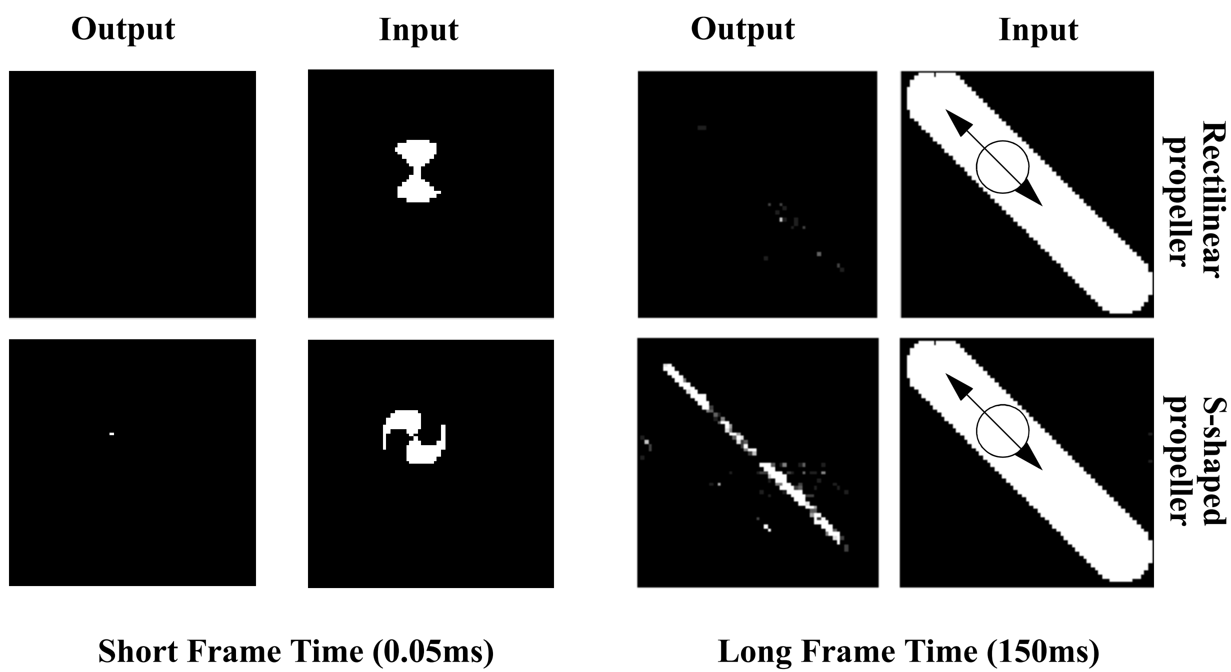
Output
Input
Output
Input
Rectilinear propeller
S-shaped propeller
Short Frame Time (0.05ms)
Long Frame Time (150ms)

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems
- Some useful circuits:
  - calibration
  - LVDS interface
- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip
- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR
- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

Fully Digital Convolution Chip (I): Architecture

- Array of pixels (Digital)
- Random Access Memory (Kernel programmed)
- Horizontal Shift Block
- 2's complement
- Synchronous controller (input communication)
- AER block (output communication)
- Configuration registers

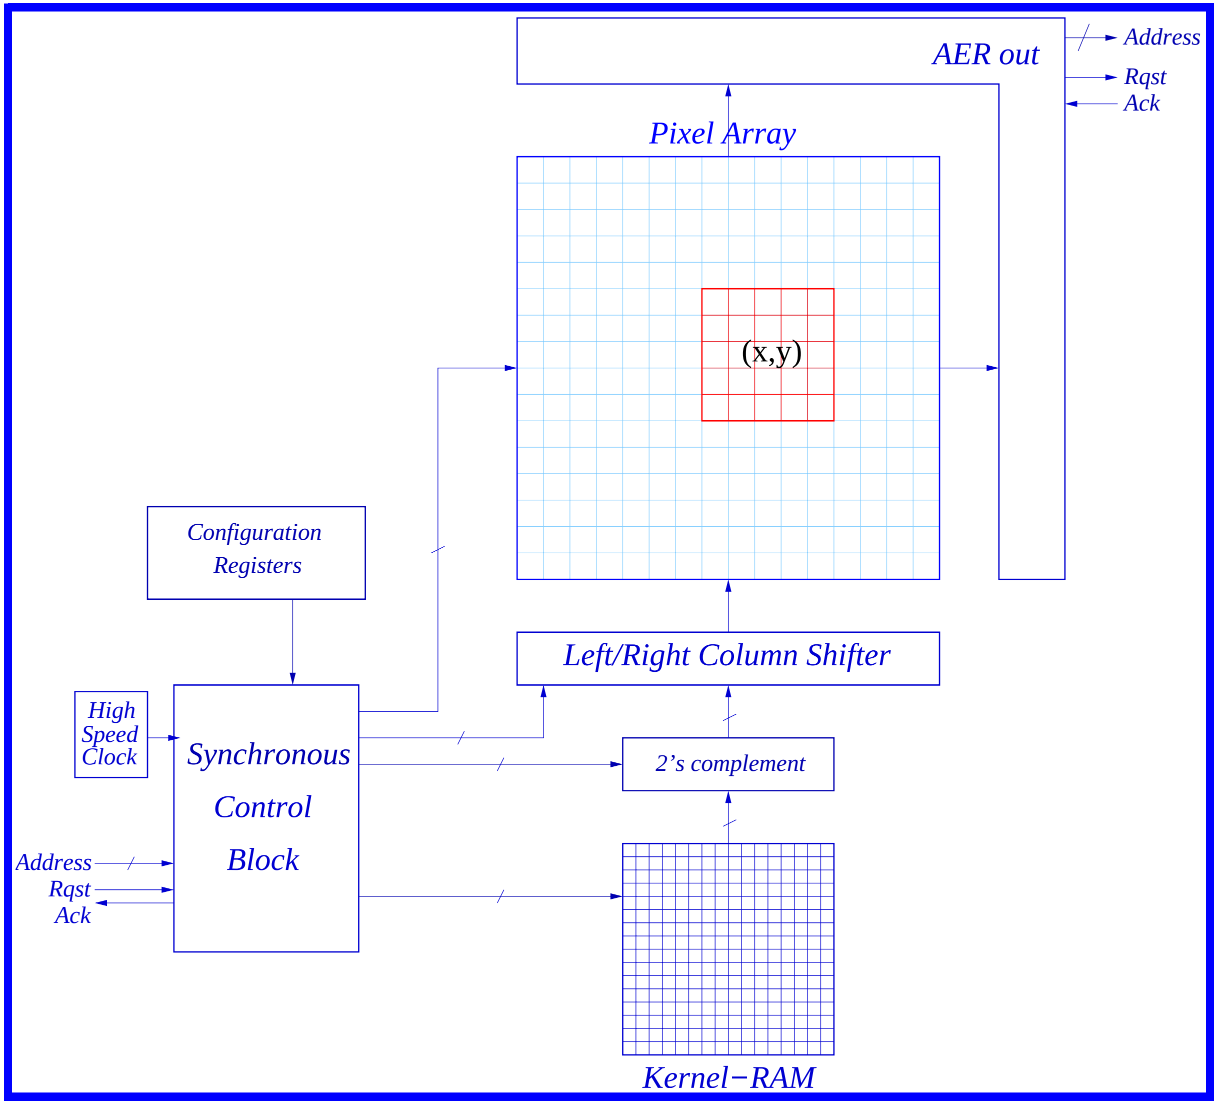

Fully Digital Convolution Chip (II): The Pixel

- Arithmetic unit
- Accumulator (18 bits)
- Comparator
- AER output communication

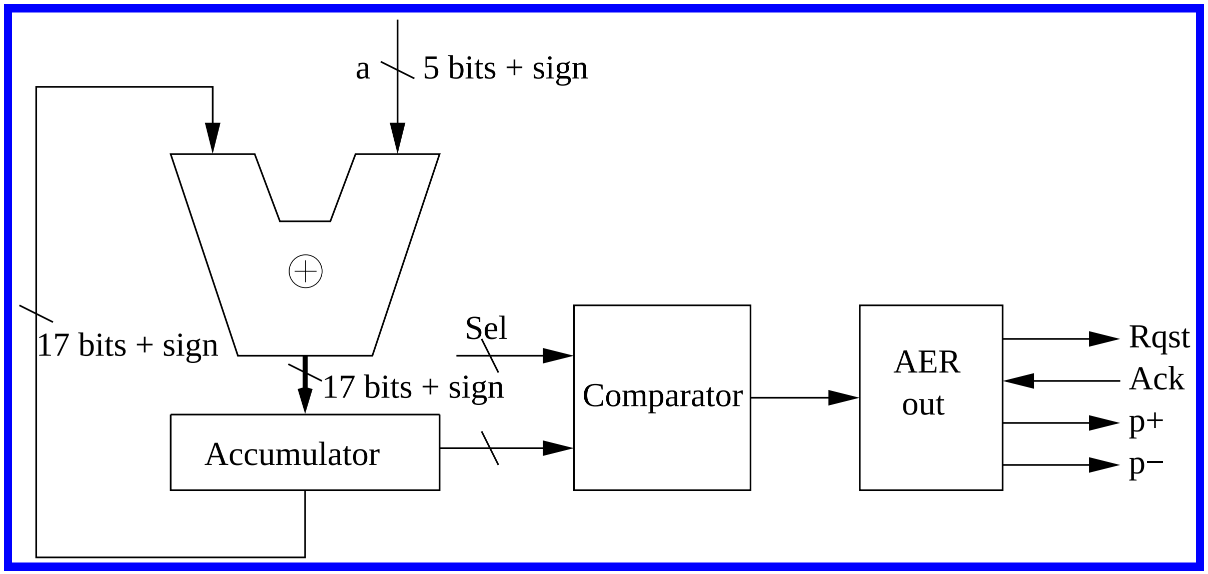

Fully Digital Convolution Chip (III): Layout

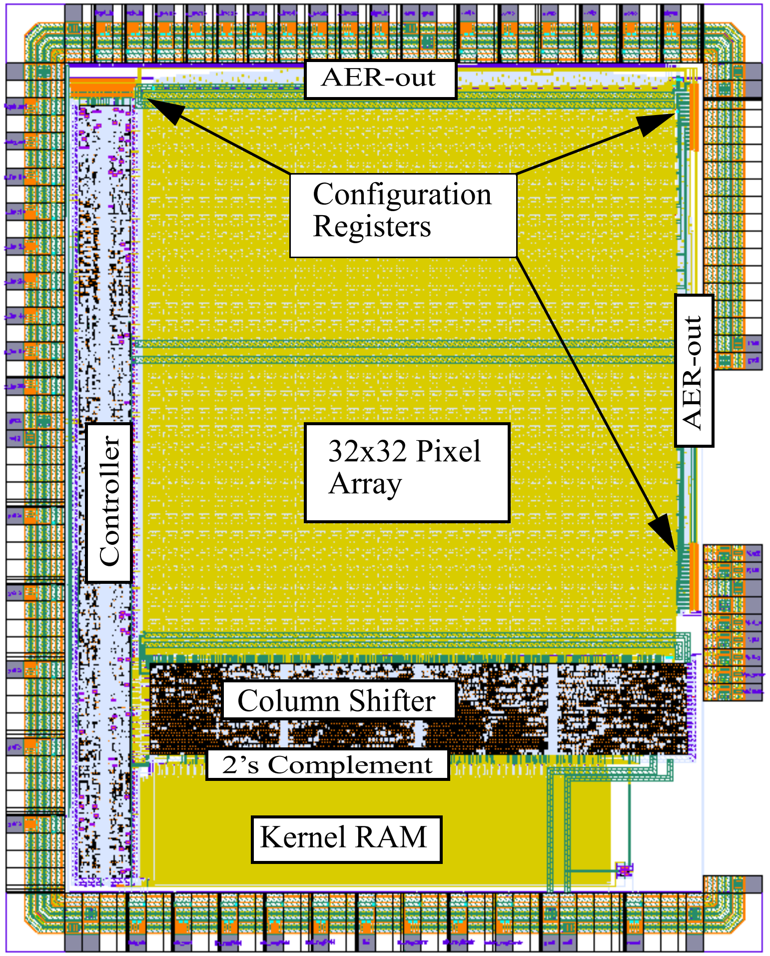

Chip Size: 5.4mm x 4.3mm

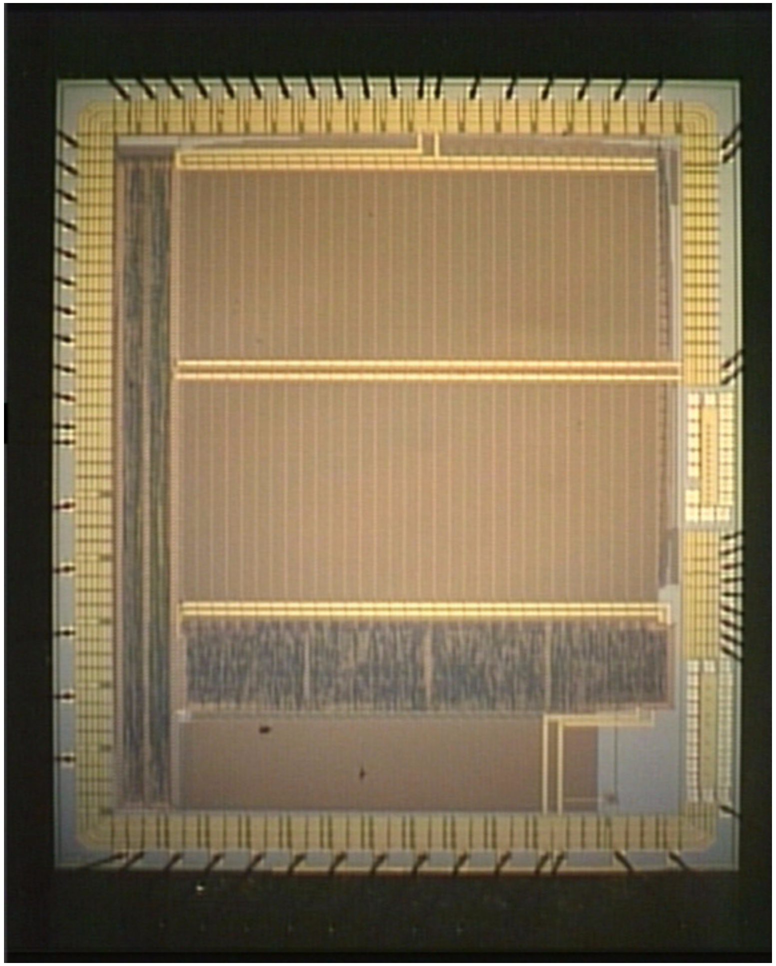

Photograph of the fabricated chip

Experimental Results (I): Single Chip Configuration

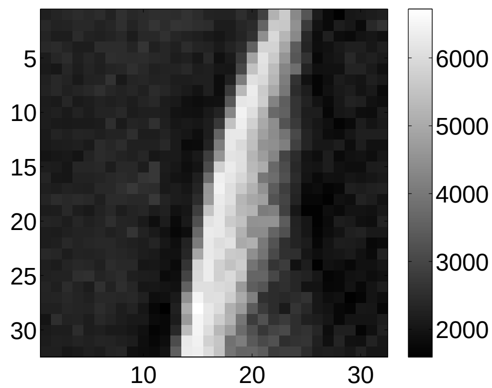

Input image

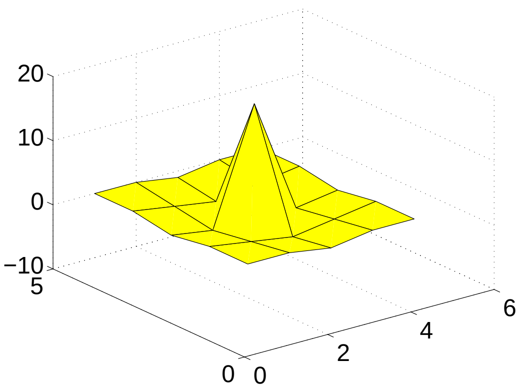

High-Pass Kernel

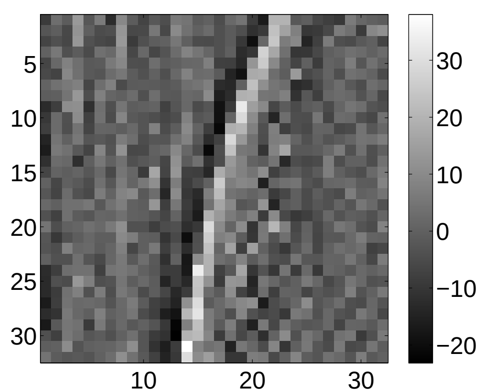

Measured output

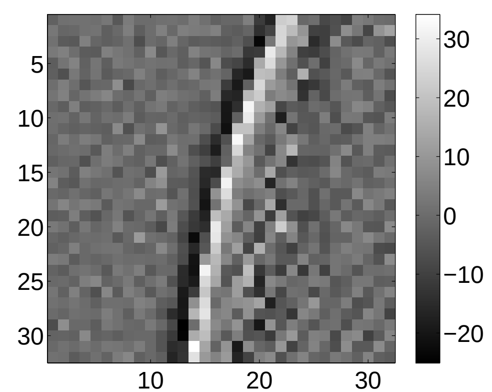

Ideal output

Experimental Results (II): Multichip Configuration

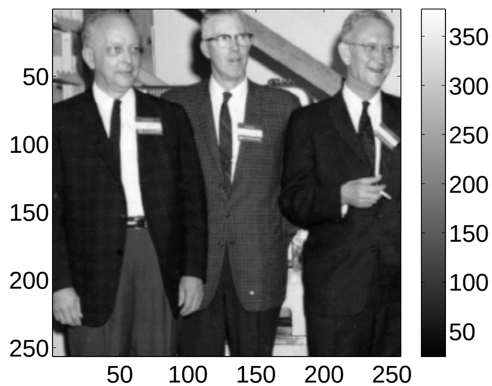

Input image

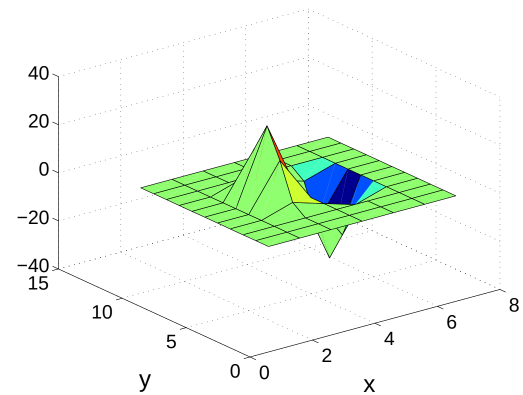

Gabor vertical edge-extraction

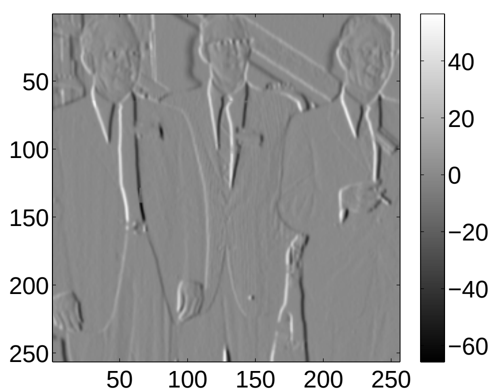

Measured output

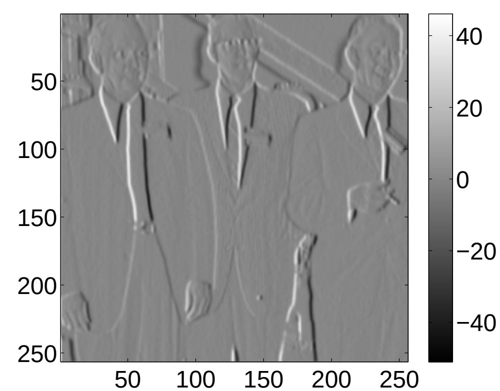

Ideal output

Experimental Results (III): Multichip Configuration

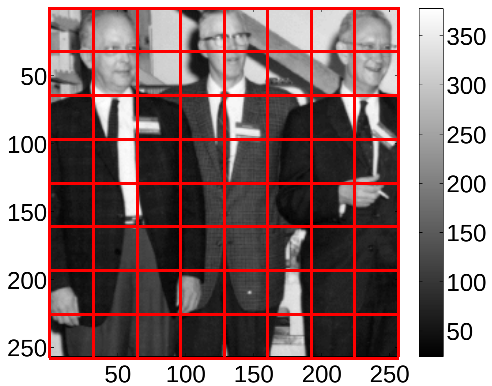

Input image

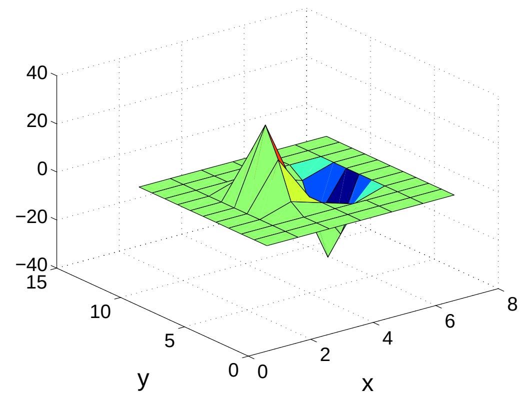

Gabor vertical edge-extraction

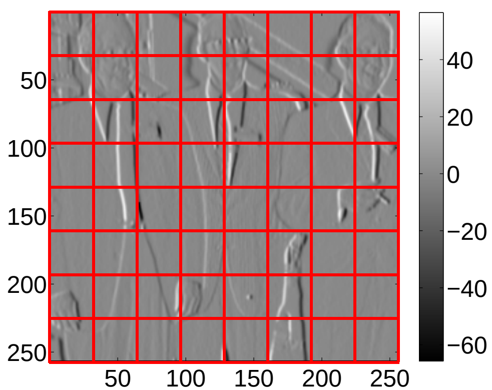

Measured output

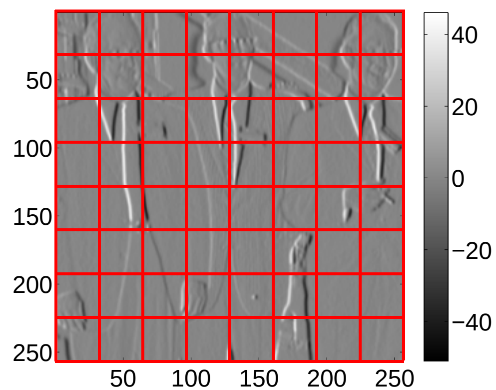

Ideal output

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems
- Some useful circuits:
  - calibration
  - LVDS interface
- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip
- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR
- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

PCI-AER

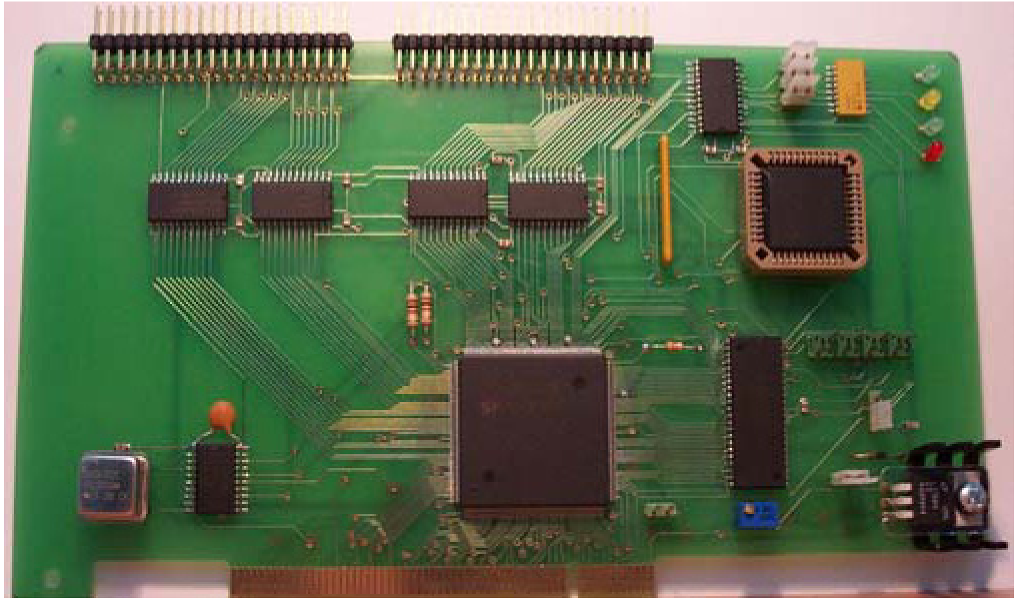

- sequences AEs from the computer out to the AER port
- transforms video-frames to AER in real-time (uses rate coding)
- captures and timestamps AEs from the AER port into the computer
- peak rate 15Meps, sustained rate 10Meps.
- FPGA: Spartan-II
- performs bus mastering

USB-AER

- USB connection to computer
- sequencer: either frame-AER or recorded AE-player (up to 25Meps)
- monitor: either AER-frame or timestamping and data-logging (up to 25Meps)
- data logging/playing up to 512 Kevents (very useful for multi-lab experiments)
- mapper (stand-alone mode) 25Meps; mappings from 1-1 up to 1-8
- VGA output
- firmware loaded through USB or MMC/SD card

Splitter/Merger

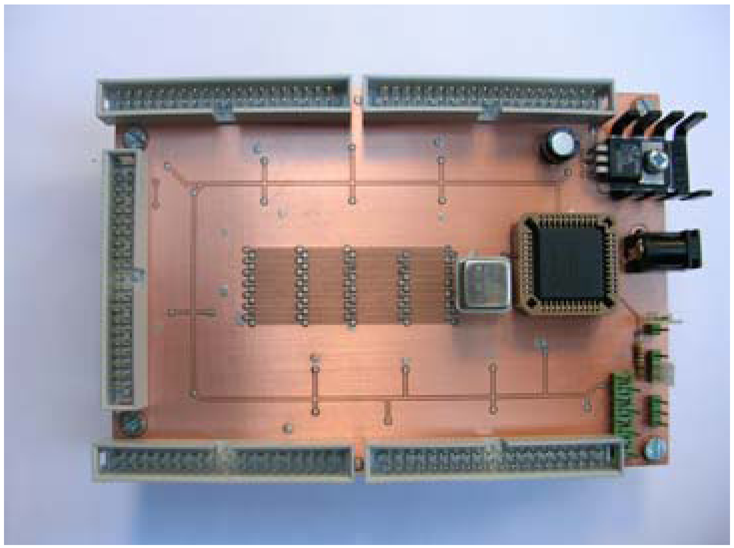
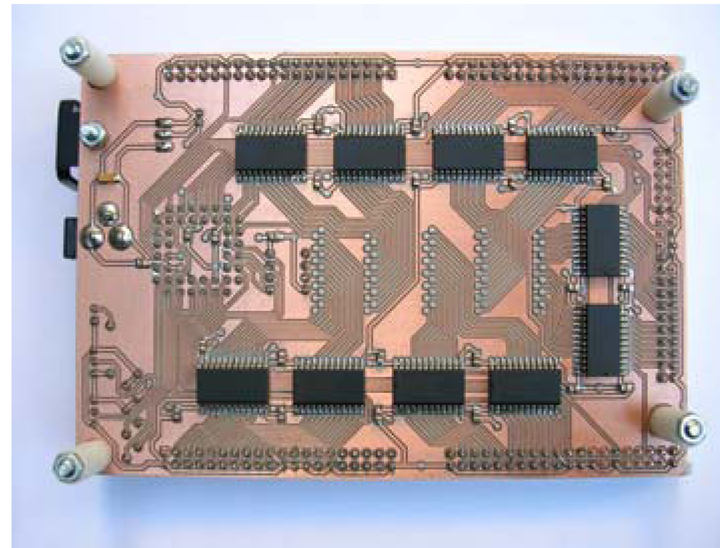

- uses a CPLD as the communication center
- splitter: 1 to 4
- merger: 4 to 1
- reconfigured by jumpers
- delay introduced: 20ns

USB2AER

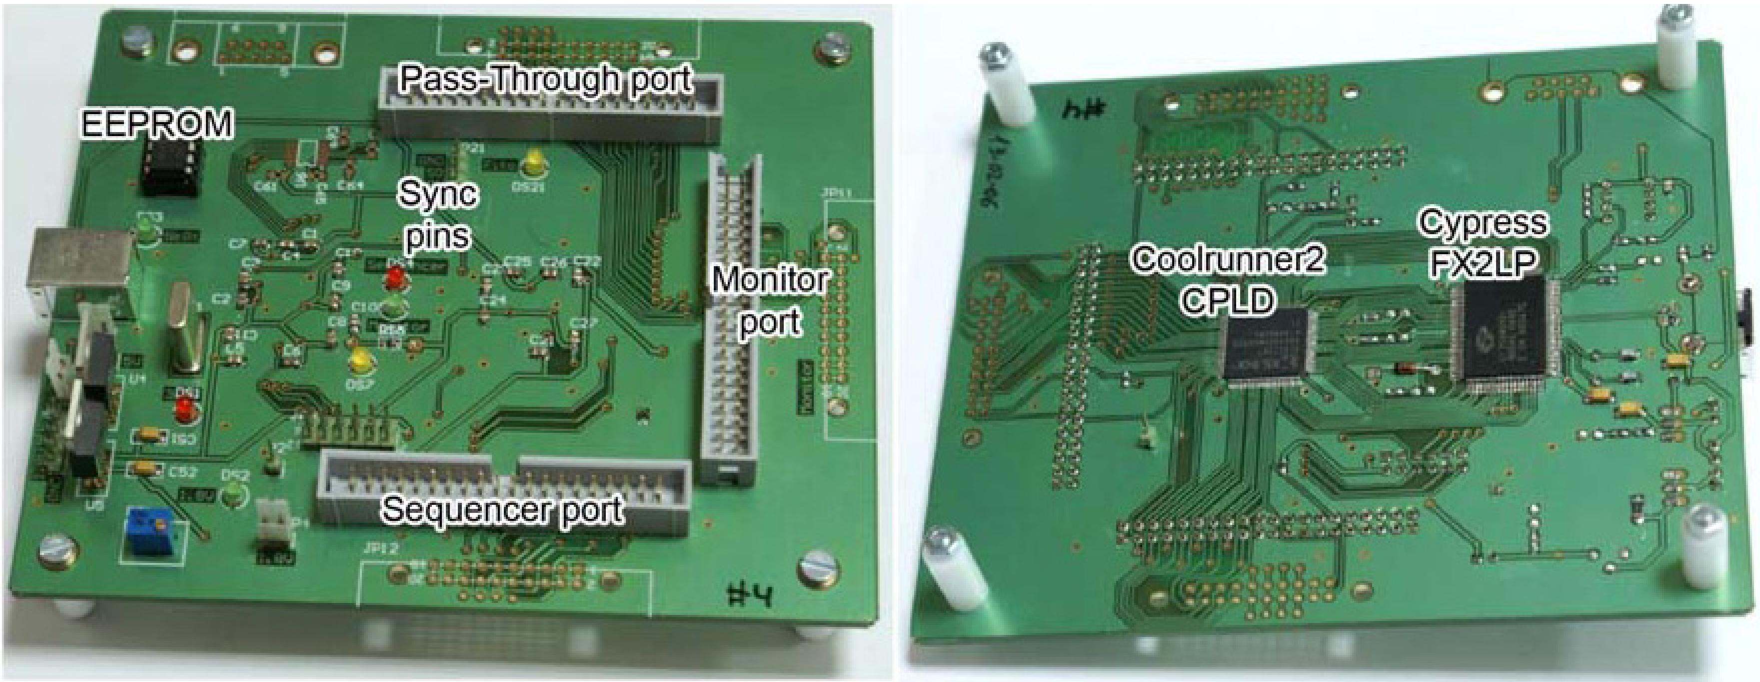

- uses high speed USB2 (up to 6Meps between AER-port and computer)
- only functionalities: AE monitor & AE sequencer (AER-port to/from computer-USB2)
- monitoring & sequencing can be simultaneous
- no FPGA, just a CPLD (timestamping) and a microcontroller (for USB traffic management)
- USB powered
- compatible with jAER viewers and Matlab

AER-Robot

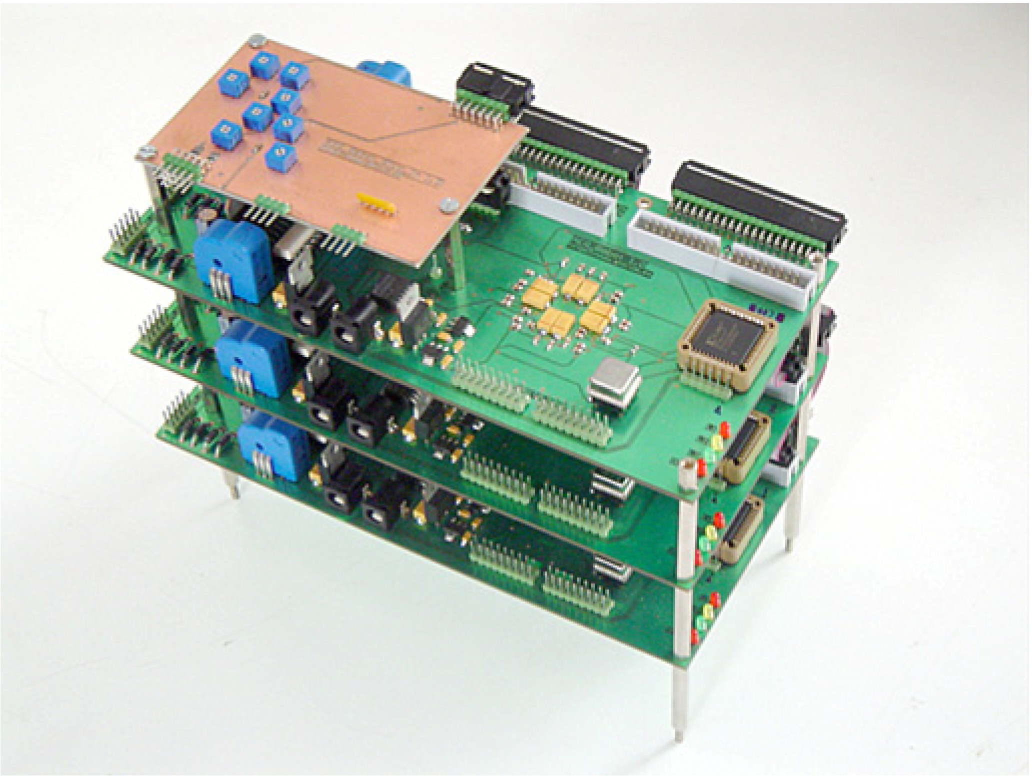

- for controlling motors directly from an AER bus
- each PCB has 4 motor connectors

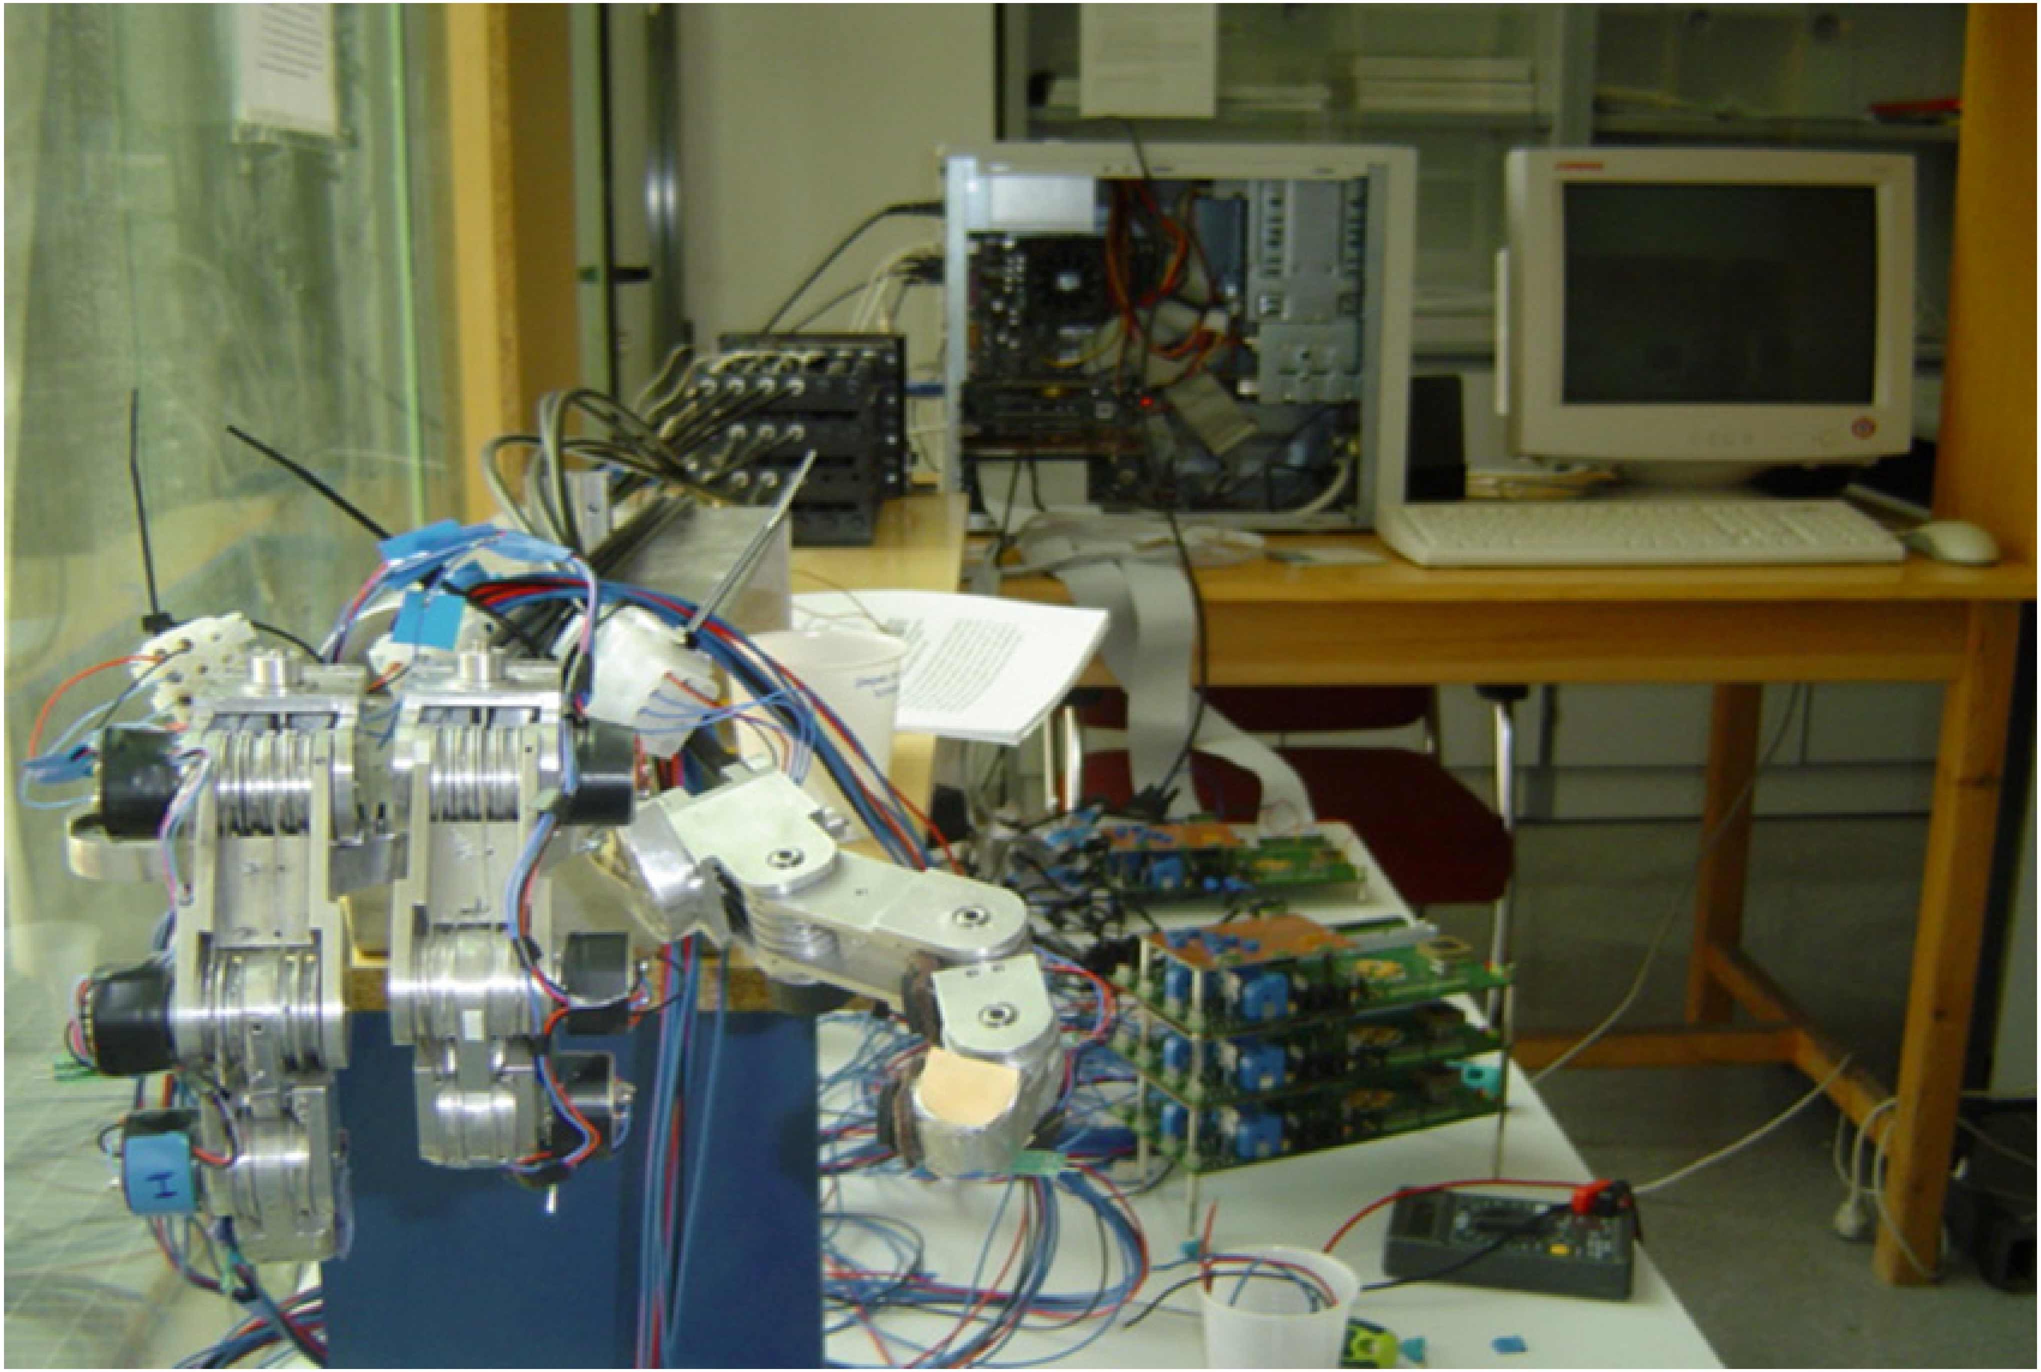

The CAVIAR Vision System

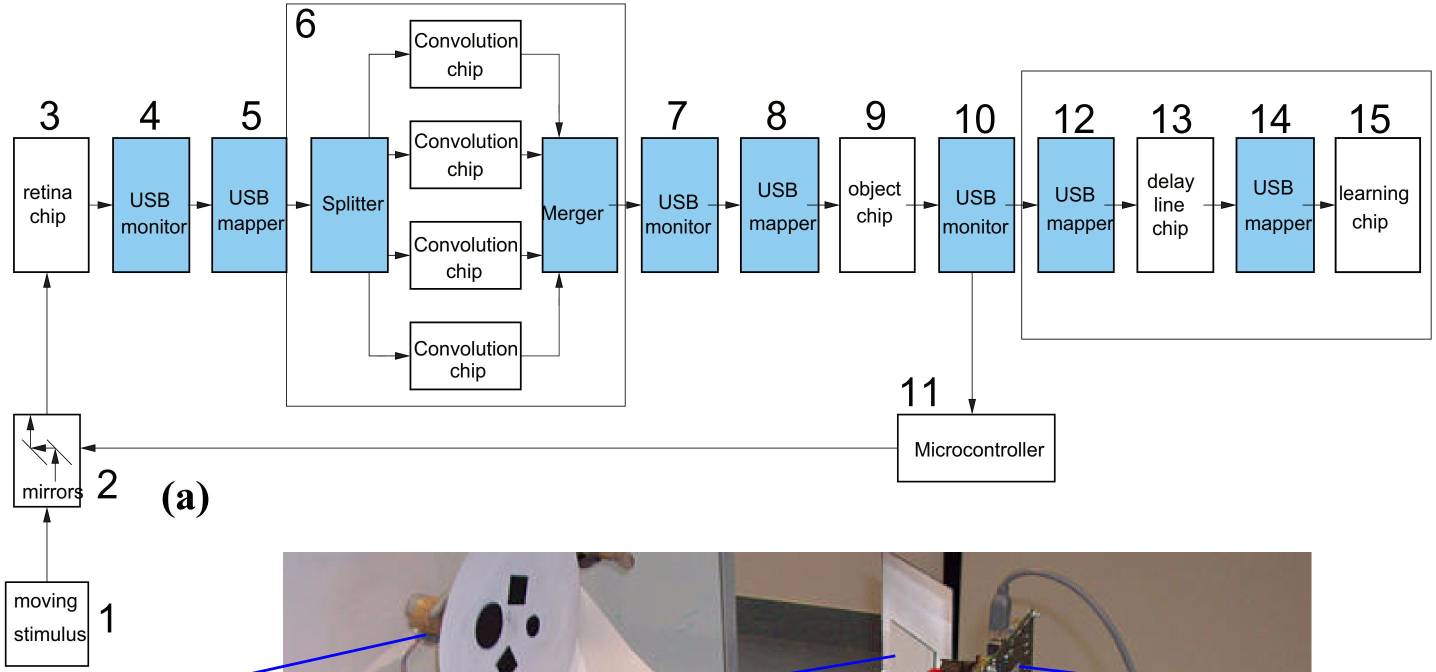

(a)

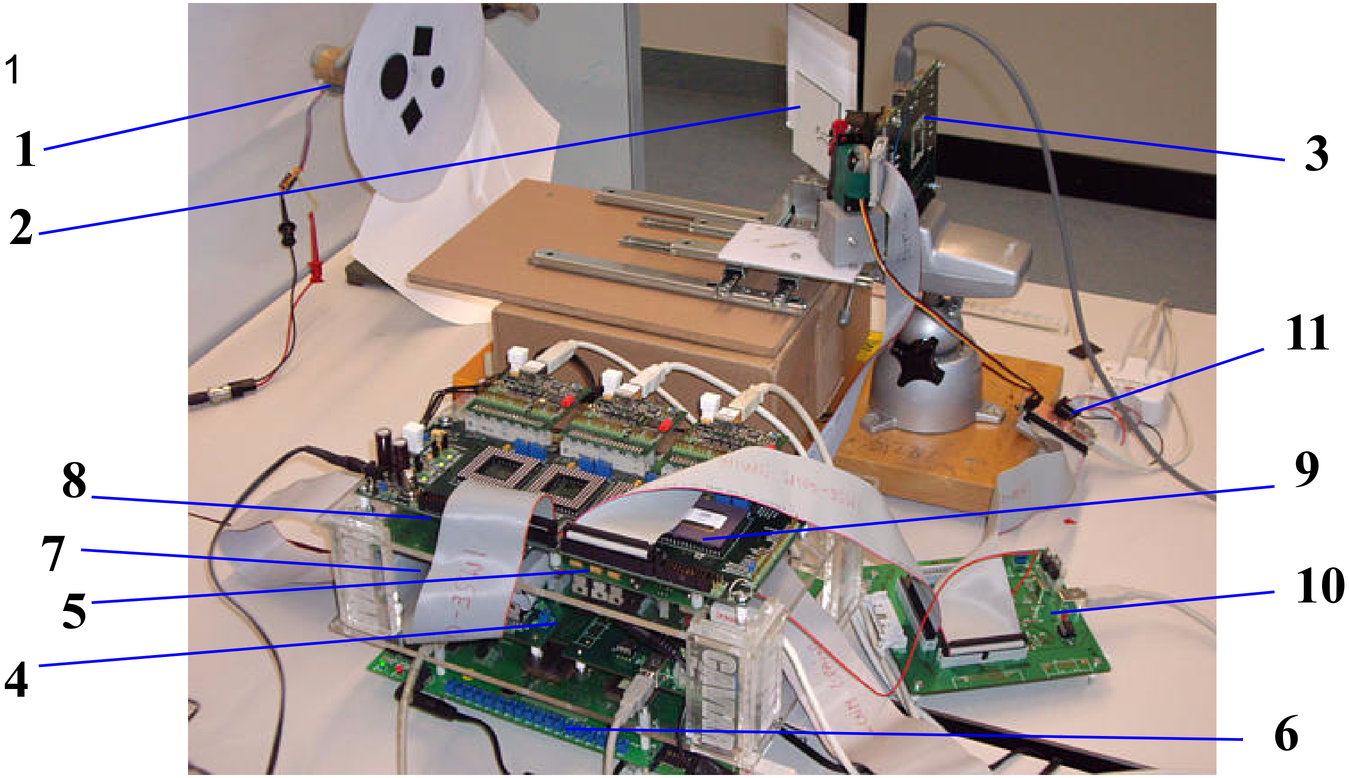

(b)

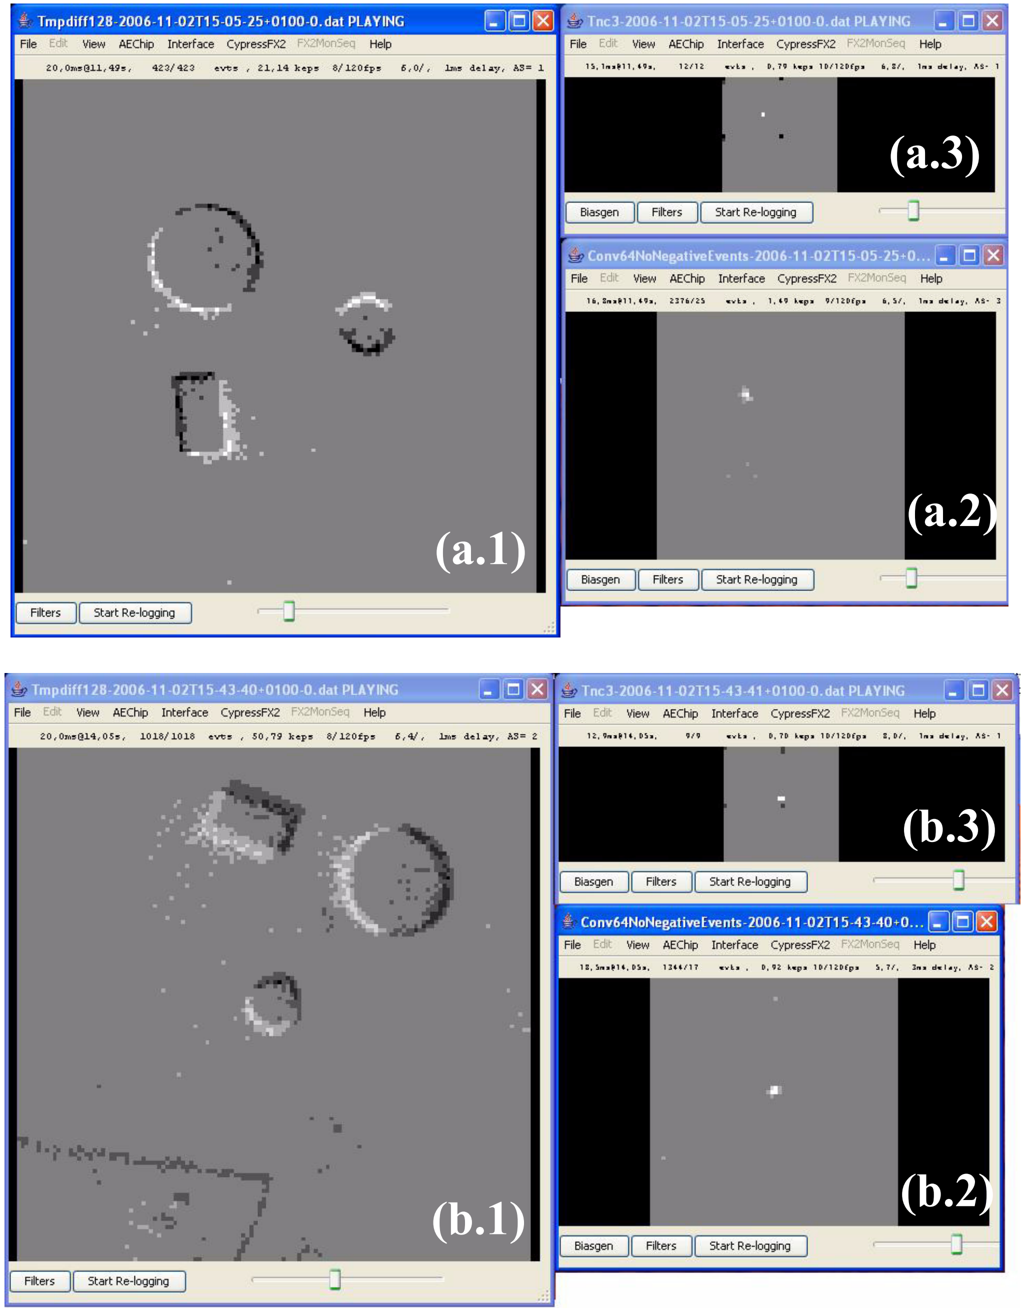
Tmpdiff128-2006-11-02T15-05-25+0100-0.dat PLAYING
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Filters
Start Re-logging
(a.1)
Tnc3-2006-11-02T15-05-25+0100-0.dat PLAYING
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Biasgen
Filters
Start Re-logging
(a.3)
Conv64NoNegativeEvents-2006-11-02T15-05-25+0...
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Biasgen
Filters
Start Re-logging
(a.2)
Tmpdiff128-2006-11-02T15-43-40+0100-0.dat PLAYING
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Filters
Start Re-logging
(b.1)
Tnc3-2006-11-02T15-43-41+0100-0.dat PLAYING
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Biasgen
Filters
Start Re-logging
(b.3)
Conv64NoNegativeEvents-2006-11-02T15-43-40+0...
File Edit View AEChip Interface CypressFX2 FX2MonSeq Help
Biasgen
Filters
Start Re-logging
(b.2)

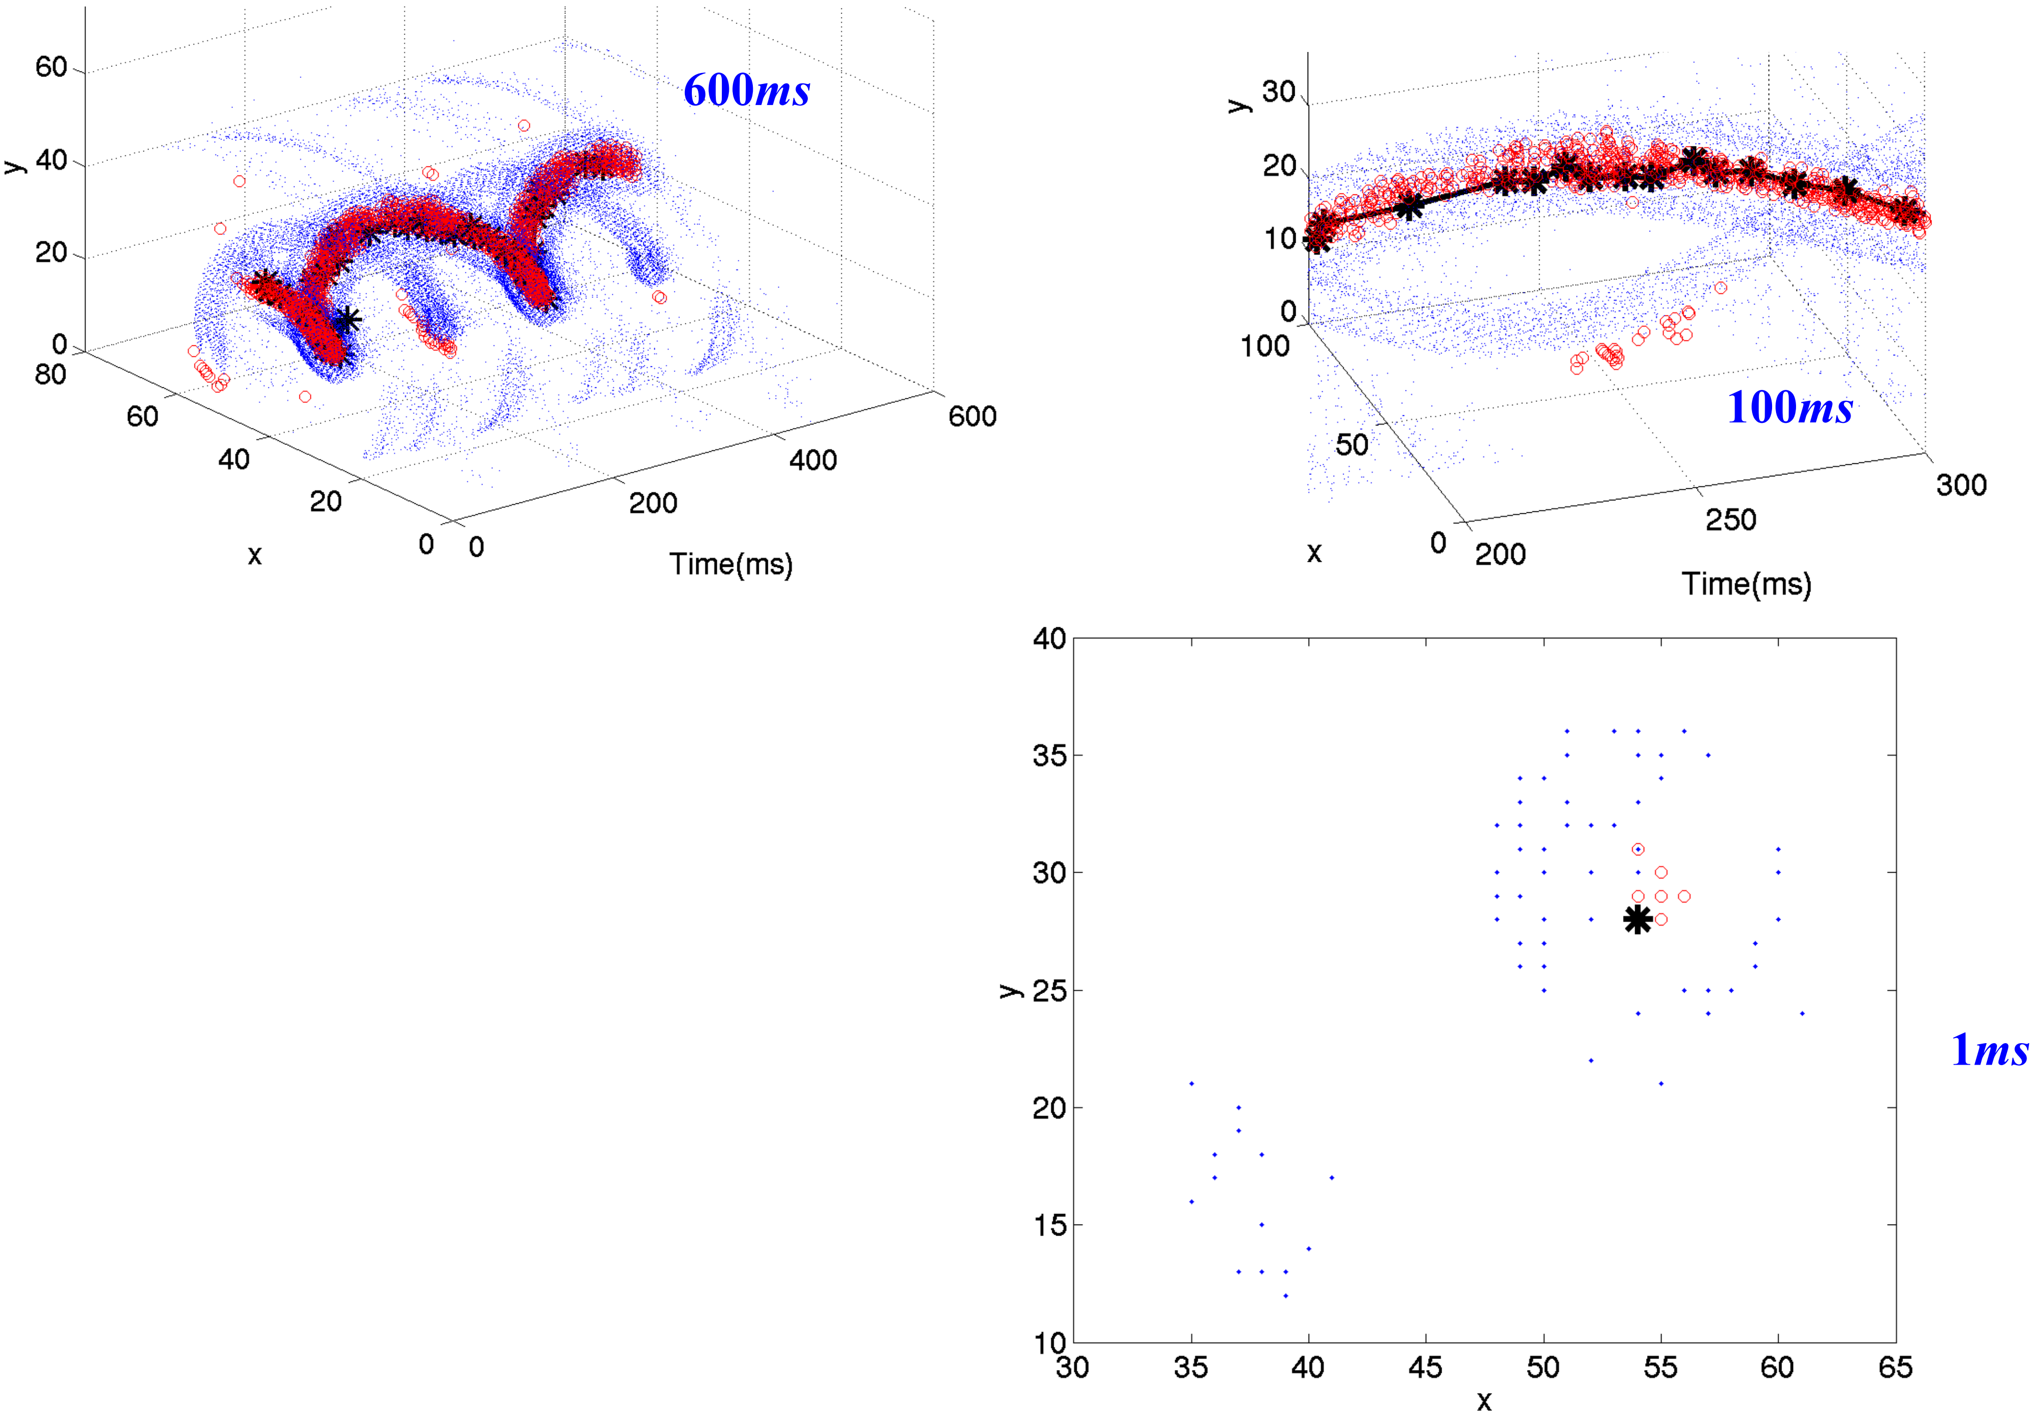
600ms
Time(ms)
x
y
100ms
Time(ms)
x
y
1ms
x
y

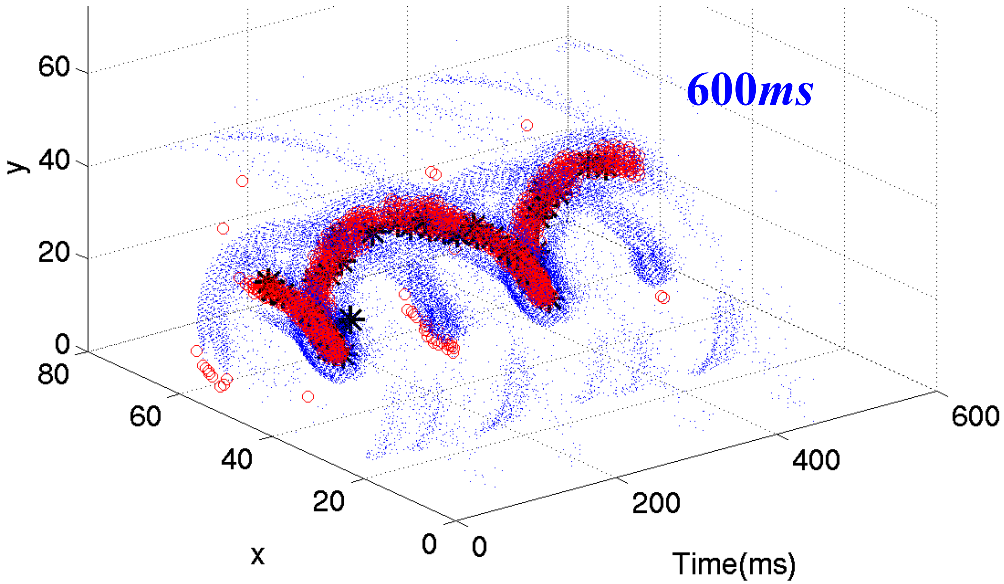

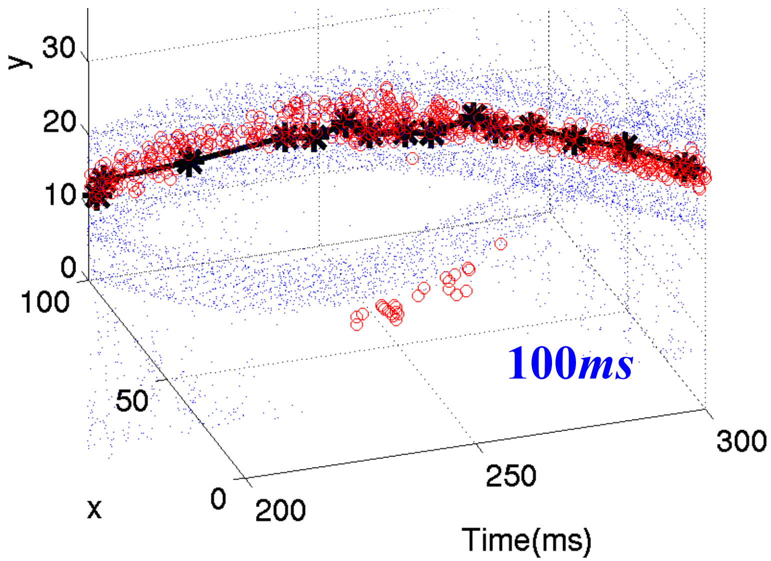

- 4-layer system
- 45k neurons
- up to 5M synapses
- 12Geps
- 1-3ms latency for tracking
- scalable w/o performance degration

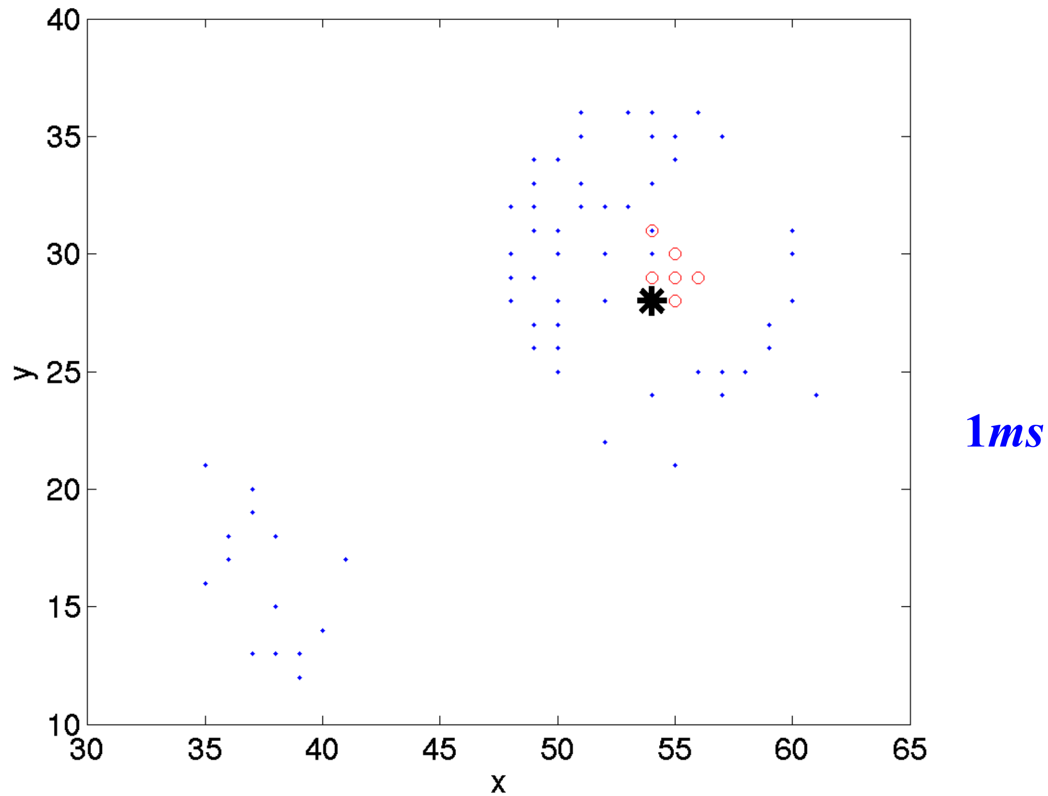

Latency Measurement

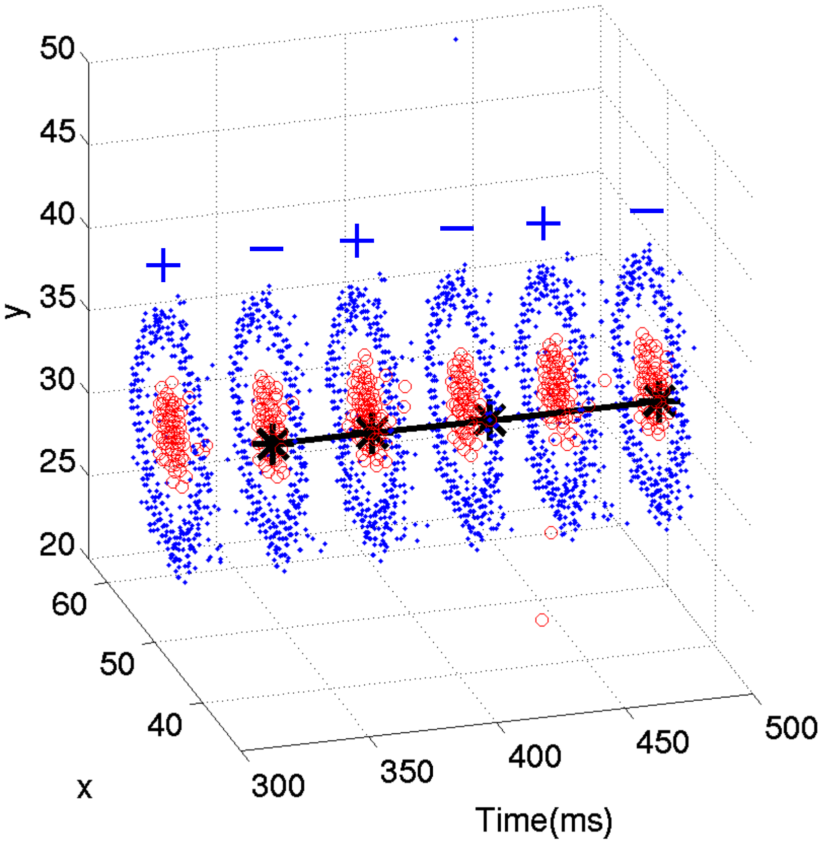

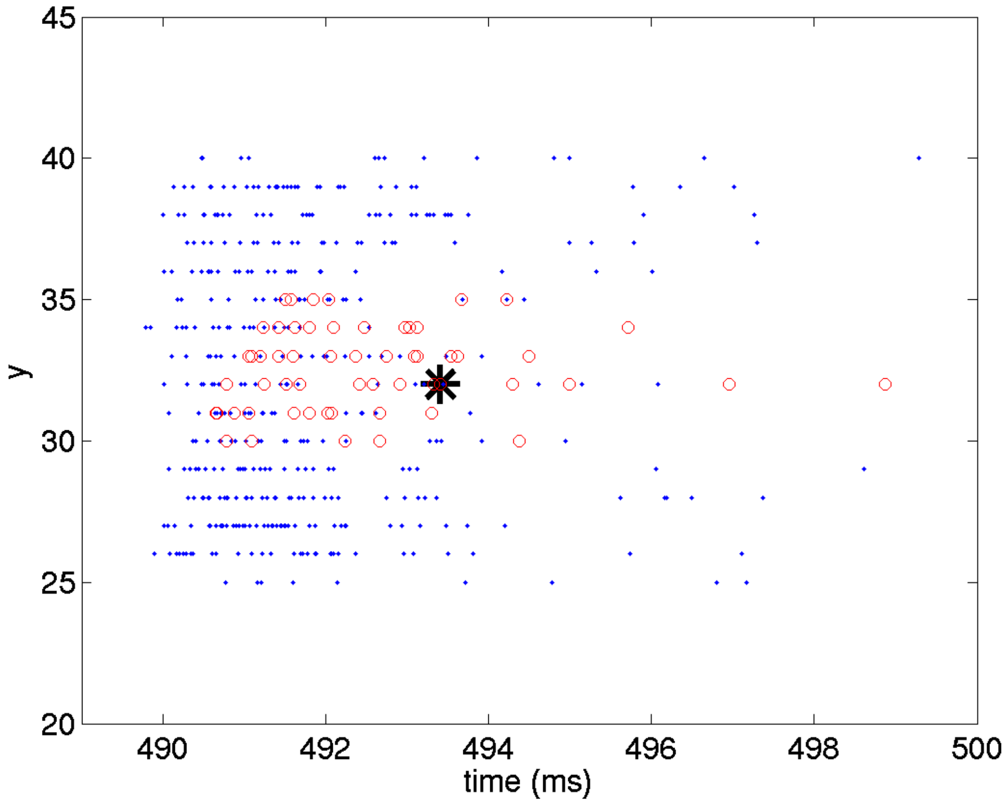

Outline

- Introduction: AER, a technology for building large scalable neuromorphic systems

- Some useful circuits:
  - calibration
  - LVDS interface

- Some example systems at IMSE:
  - spatial contrast retina
  - mixed-mode convolution chip
  - fully digital convolution chip

- HW Tools from Sevilla:
  - some FPGA-based PCBs
  - example use in CAVIAR

- SW Tool:
  - Behavioral Matlab Simulator
  - Example 1: neocognitron emulation
  - Example 2: texture classification

CORTICAL TISSUE

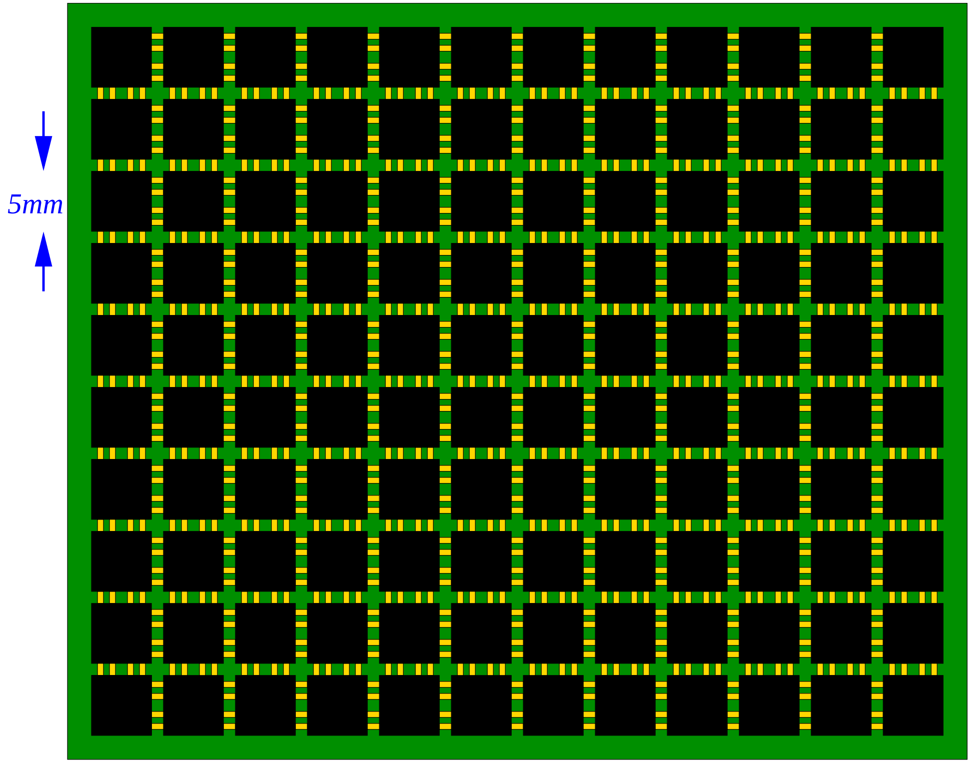

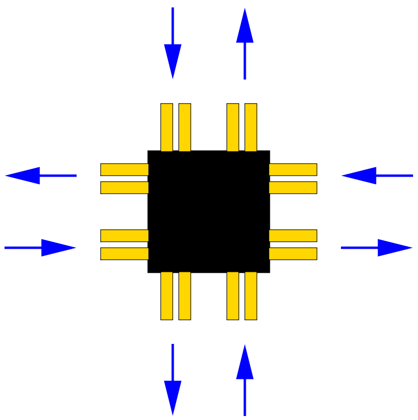

- Potentially Very High Computational Power: 2Mneurons, 32Gsynapses, 238Tconn/sec
- ¿How to reconfigure?
- ¿What hierarchies and structures?
- ¿What kernels?
- We need theories for implementing desired functionalities (hopefully before the HW is available)

MATLAB based AER Behavioral Simulator

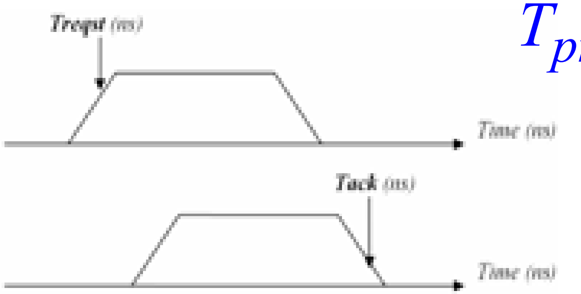

$T_{preRqst}$	T_{Rqst}	T_{Ack}	x	y	$sign$
0	0	10	45	29	1
20800	20800	20810	44	29	1
156250	156300	156310	43	29	-1
214250	214250	214260	44	30	1
250000	250000	250010	45	29	1
291600	291650	291660	44	29	-1
291650	291700	291710	43	30	1
399750	399750	399760	45	32	1
399800	399800	399810	44	42	1
399850	399850	399860	43	28	-1
399900	399900	399910	23	40	1
399950	399950	399960	9	38	1
400000	0	-1	2	41	1
400050	0	-1	3	42	-1
400100	0	-1	23	5	1
400150	0	-1	26	32	1
400200	0	-1	44	28	1
400250	0	-1	45	30	1
400300	0	-1	45	34	-1
500000	0	-1	45	29	1

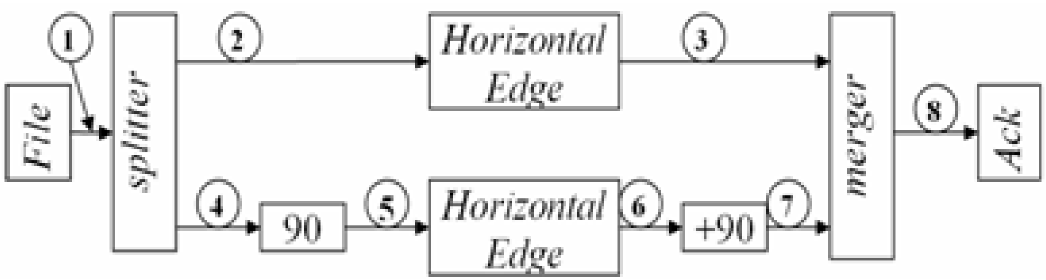


```
%First, we declare sources to the system
%          SOURCES    SOURCES DATA
sources    {1}        {data1}

%Next, we declare priorities
priorities {0.9,0.8,0.7,0.6,0.5,0.4,0.3,0.2}

%Next, we declare blocks
%NAME          IN-CHANNELS    OUT-CHANNELS    PARAMS       STATES
splitter       {1}            {2,4}           {params1}    {state1}
h_sobel        {2}            {3}             {params2}    {state2}
imrotate90     {4}            {5}             {params3}    {state3}
h_sobel        {5}            {6}             {params4}    {state4}
imrotate90     {6}            {7}             {params5}    {state5}
merger         {3,7}          {8}             {params6}    {state6}
ack            {8}            {}              {params7}    {state7}
```

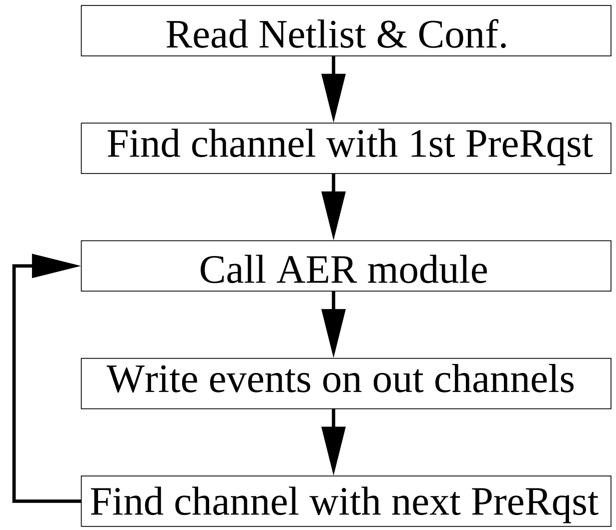

Multi-Chips Multi-Layer Processing Systems

Neocognitron & Convolution Neural Networks

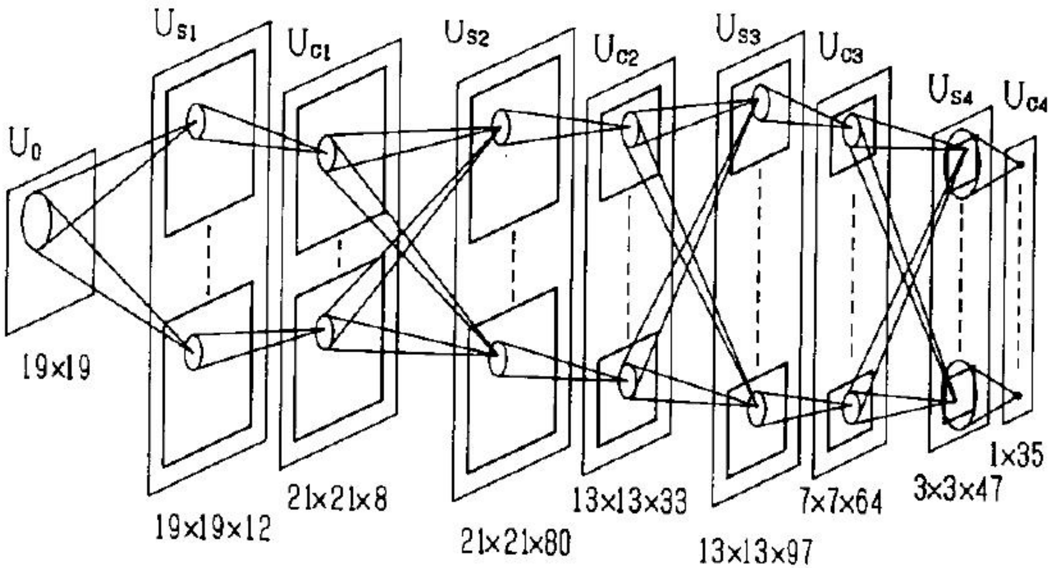

8 layers, 376 Convolutions

K. Fukushima
1969

Applied to handwritten character recognition

Example: Simplified Neocognitron

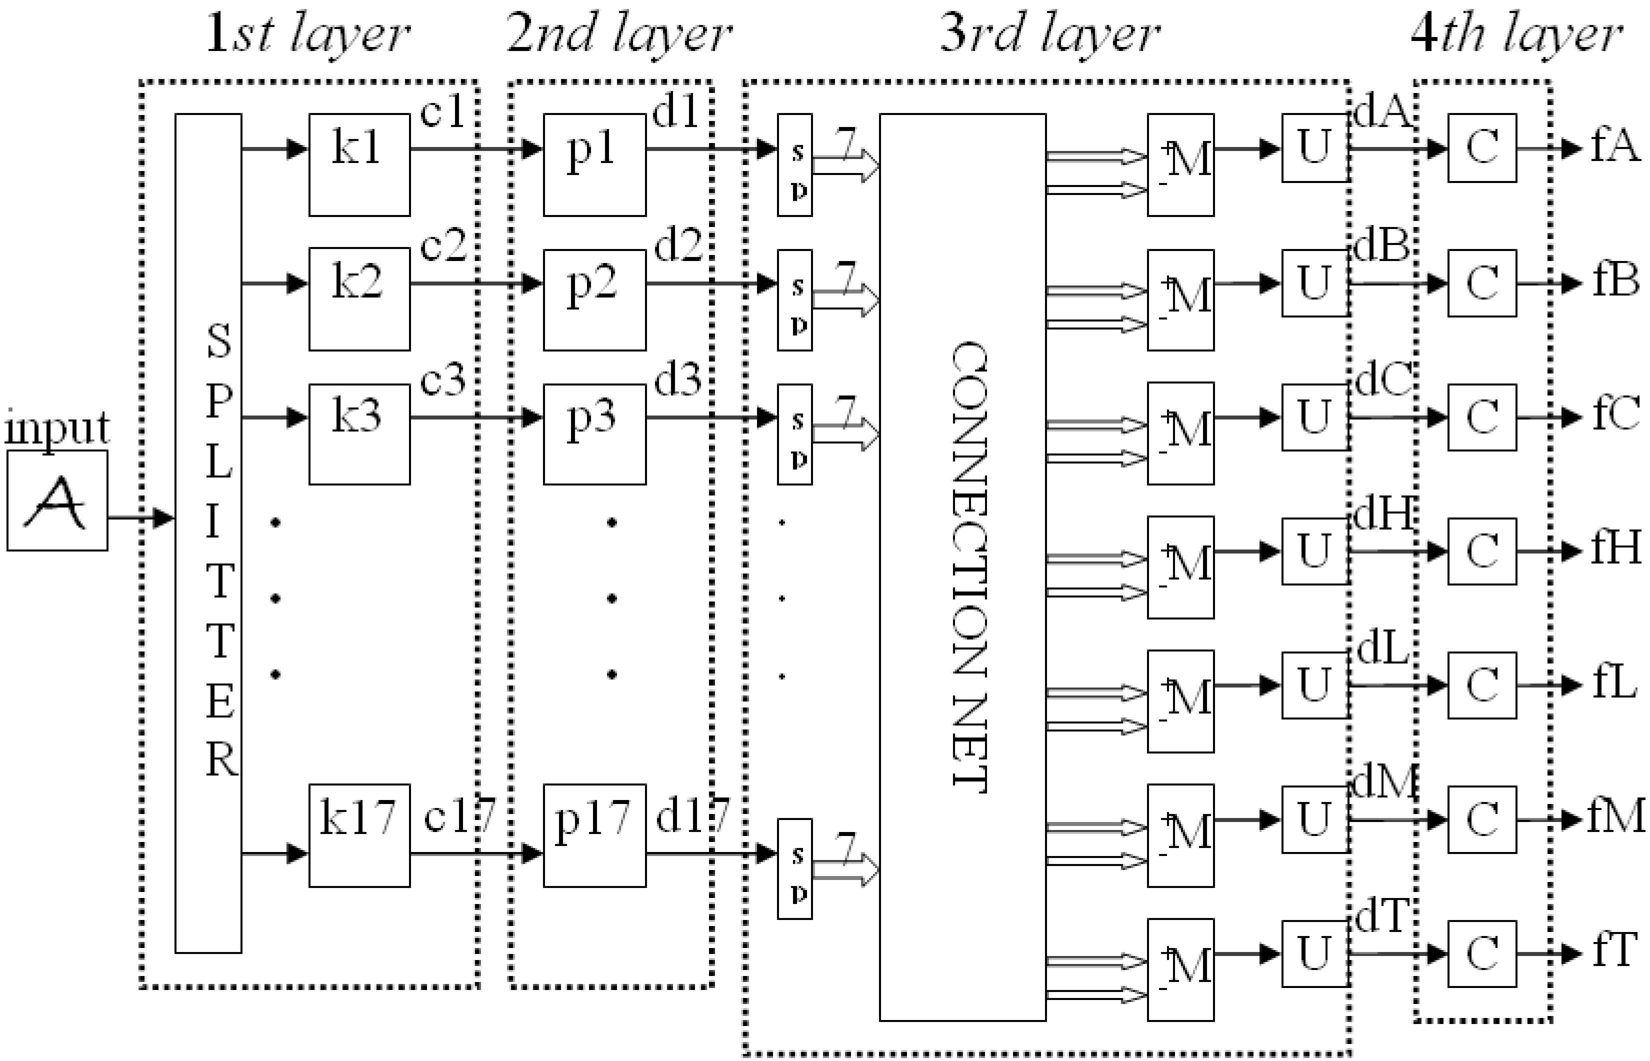

- 4 layers, 68 convolution modules
- inputs 16x16 b&w pixels
- 7 output categories

Large kernels

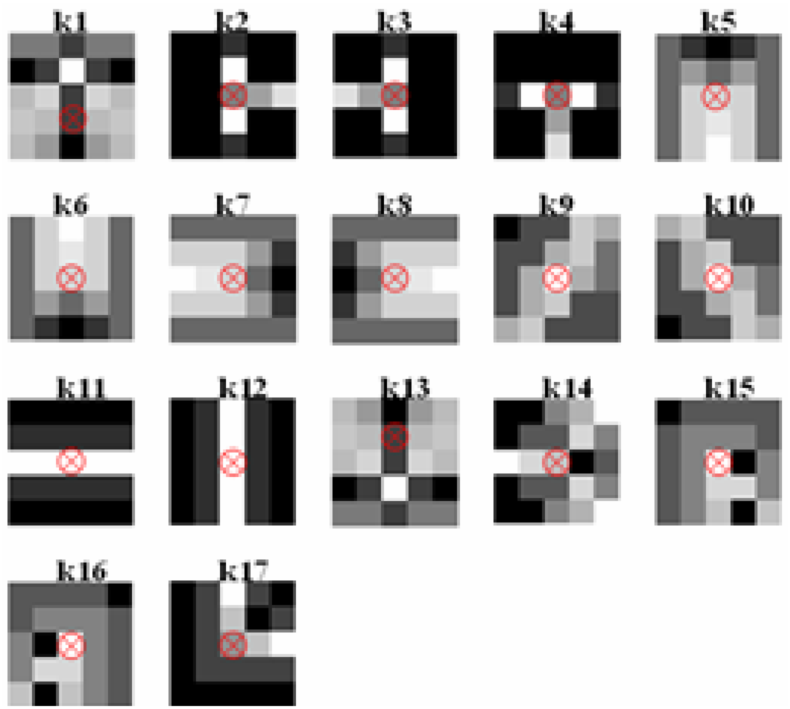

Layer 1

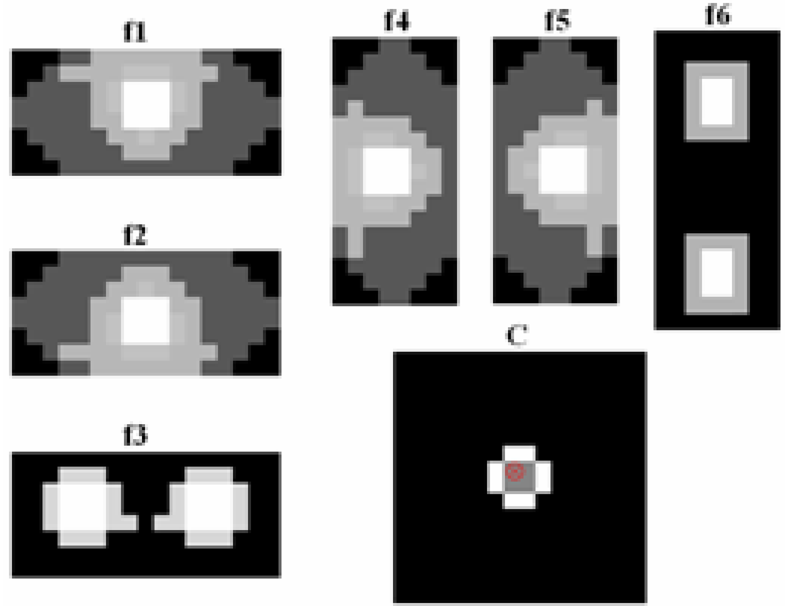

Layers 2 & 4

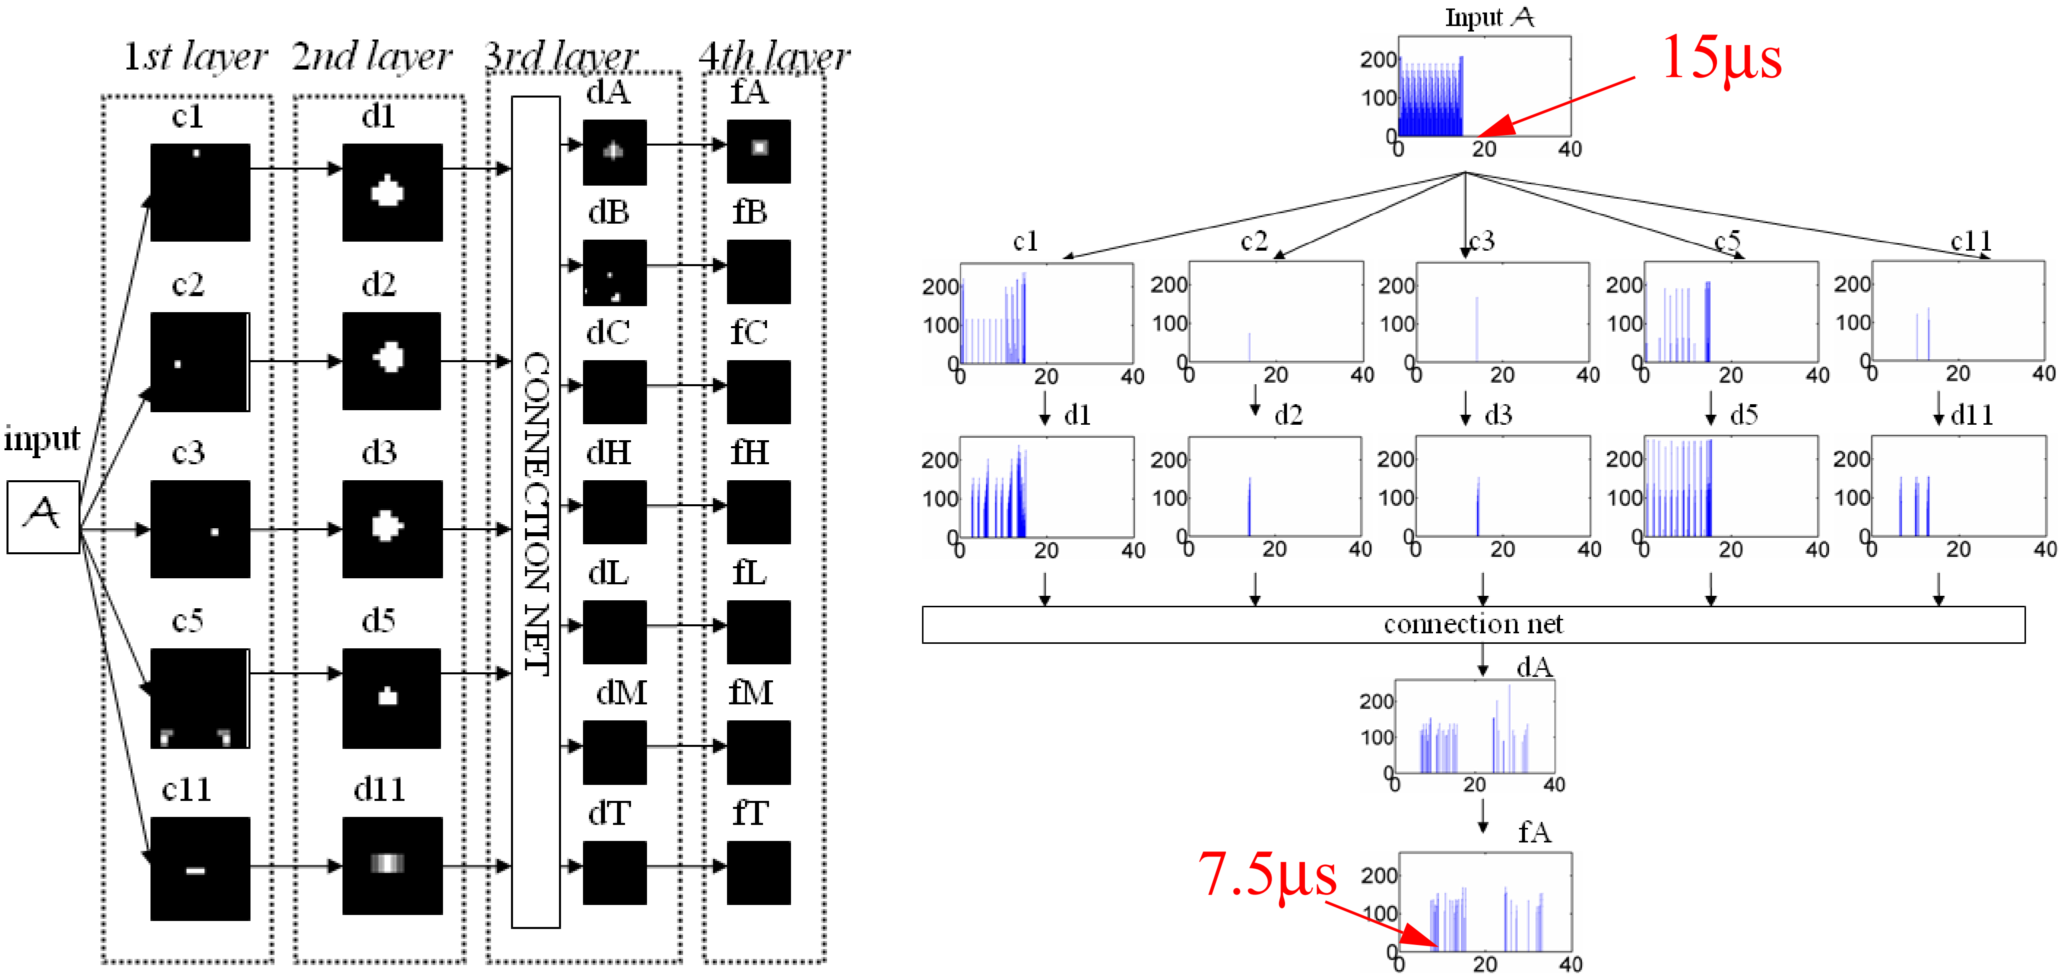

1st layer
2nd layer
3rd layer
4th layer
input
c1
c2
c3
c5
c11
d1
d2
d3
d5
d11
CONNECTION NET
dA
dB
dC
dH
dL
dM
dT
fA
fB
fC
fH
fL
fM
fT
Input A
15μs
connection net
7.5μs

Texture Classification

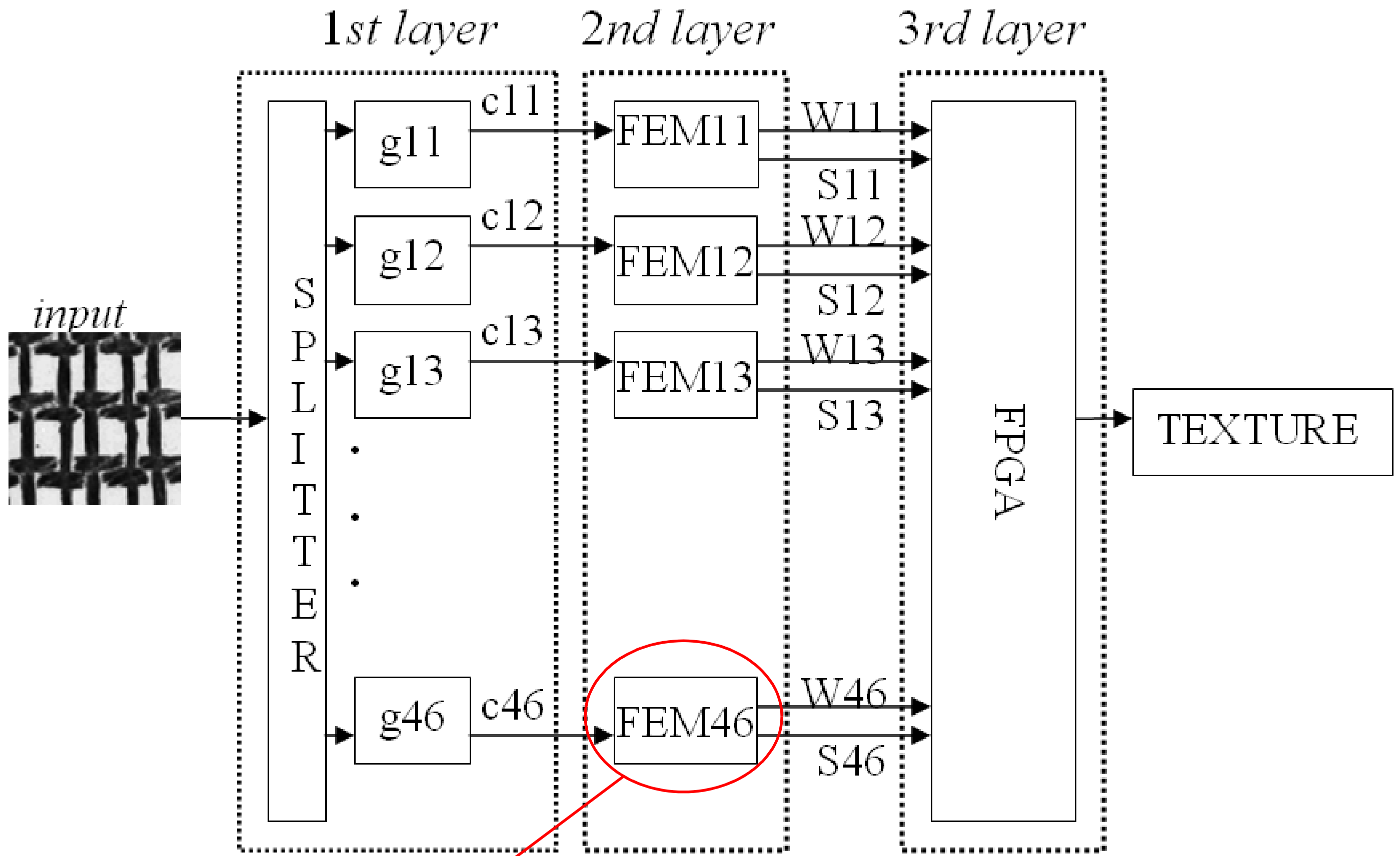

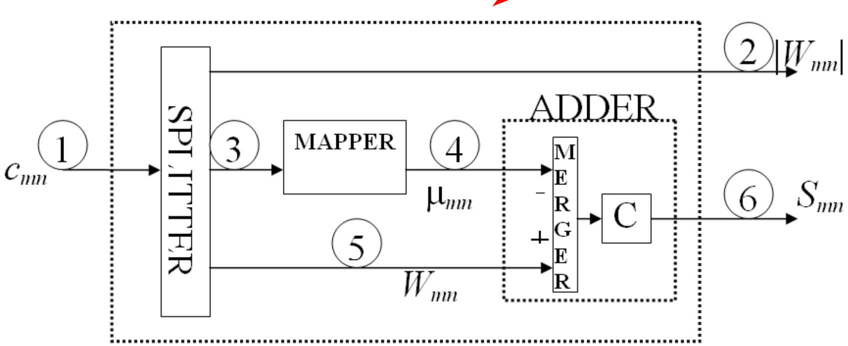

- input 90x90 pixels
- 48 convolutions
- kernel sizes up to 50x50

Detecting People & displaying using jAER

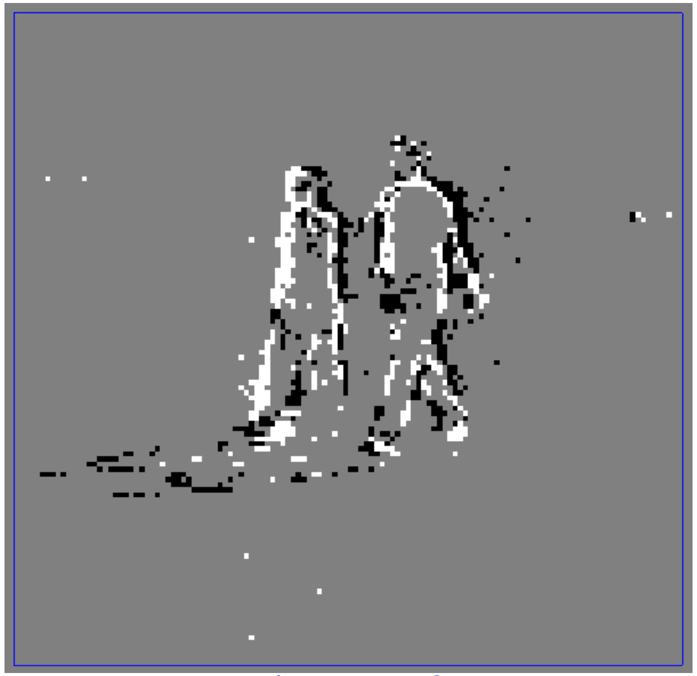

raw input from Tobi's temporal contral DVS

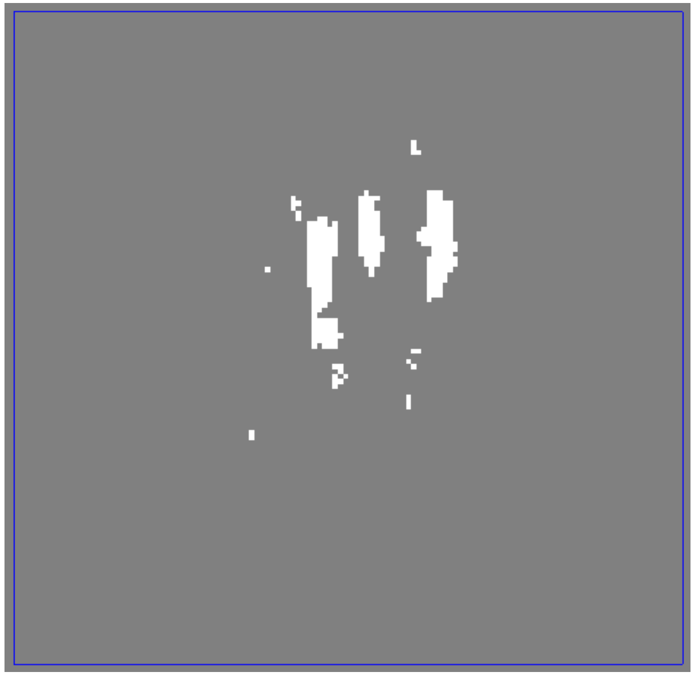

Vertical Gabor Filter 7x7 kernel

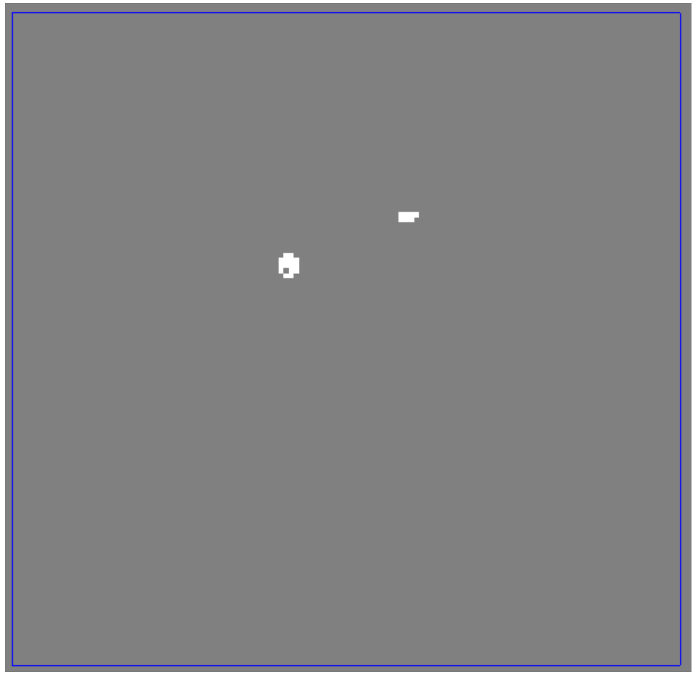

crude template matching 19x36 kernel

Conclusions

- AER has high potential for building complex neurocortical hierarchies.
- A variety of AER sensors are available.
- With present day technology it is feasible to build programmable & reconfigurable “Cortical Tissues” with millions of neurons, billions of synapses, and Tconn/sec.
- We need to develop knowledge for configuring, programming, and training optimally such systems.