

LOGIC SIMULATIONS	2
DIALOGUE	2
LOGIC SIMULATION: LOGIC CELL DRIVE STRENGTH	3
LOGIC SIMULATION: HV LOGIC CELLS	5
LOGIC SIMULATION: MONOSTABLES	7
LOGIC SIMULATIONS: SRAM REGISTER	12
LOGIC SIMULATION: DATA OUTPUT MULTIPLEX	18
LOGIC SIMULATION: SRAM SHIFT REGISTER CELL	20
LOGIC SIMULATION: BI-DIRECTIONAL SRAM CELLS	24
LOGIC SIMULATION: BI-DIRECTIONAL SRAM SHIFT REGISTER	27
LOGIC SIMULATION: LATCH-HOLD CIRCUITS	29
LOGIC SIMULATION: 3-CLOCK SR CELL	32
LOGIC SIMULATION: MASK & CONFIG PROGRAMMING (SHIFT5)	34
CONFIGURATION "SLOW" CLOCKS	36
LOGIC SIMULATION: FAST CONFIG SHIFT REGISTER	38
OVERVIEW: CONFIGURATION PROGRAMMING AND VERIFICATION	40
SIMULATION RESULTS	41
LOGIC SIMULATION: FULL PIXEL 'SLICE' SIMULATION	42
LOGIC SIMULATION: FULL PIXEL SLICE INCLUDING MASKING	45
LOGIC SIMULATION: ROW ENCODER	47
LOGIC SIMULATION: ROW ENCODER	48
LOGIC SIMULATION: DRIVING LONG COLUMN SIGNALS	49
LOGIC SIMULATION: DRIVING LONG COLUMN SIGNALS	50
LOGIC SIMULATION: THREE MASTER-CONTROLLERS + ROW ENCODER	52

LOGIC Simulations

Dialogue

This is a summary of the logic cells for review at IDR.

This is a long document, intending to cover all aspects of the digital control circuits on the ASIC1 test structure, starting with individual logic gates and including mixed-mode simulations incorporating Verilog stimulus, full row-controller custom logic, SRAM memory banks and analog pixels.

Dialog is included after results where appropriate – the contents of this document will be discussed in more detail in the IDR.

Logic simulation: Logic Cell Drive Strength


Simulation top-level = "sim_inverter"

Logic inverters were designed with different drive strengths. Functional logic blocks (2 and 3 input NAND and NOR gates) were designed to the minimum drive strength only. All transistors used are 1.8v parts to reduce power consumption.


Cell variant	x0	x1	x2	x3	x4
Test load	10fF	30fF	55fF	90fF	125fF
Rise time	132p	129p	133p	125p	125p
Fall time	113p	120p	127p	129p	133p
Pmos width	0.8	2.0	3.5	6.0	8.4
Nmos width	0.3	0.8	1.4	2.2	2.9

Above: Tabulated results for typical process corner. Rise and fall time (for these tests) are recorded as the transition time between 15% and 85% of the full-scale 1.8v signal.

Results waveforms



Above: Inverters of each strength variant are simulated together with their corresponding test load. Transient waveforms are shown on the left (50ps per division); DC sweep to determine switching point is shown on the right.



Above: Inverters and logic gates of each strength variant are simulated together with their corresponding test load in all process corners. Transient waveforms are shown on the left (100ps per division); DC sweep to determine switching point is shown on the right: spread approx 200mV.


Logic simulation: HV Logic Cells

Simulation top-level = "sim_inverter"

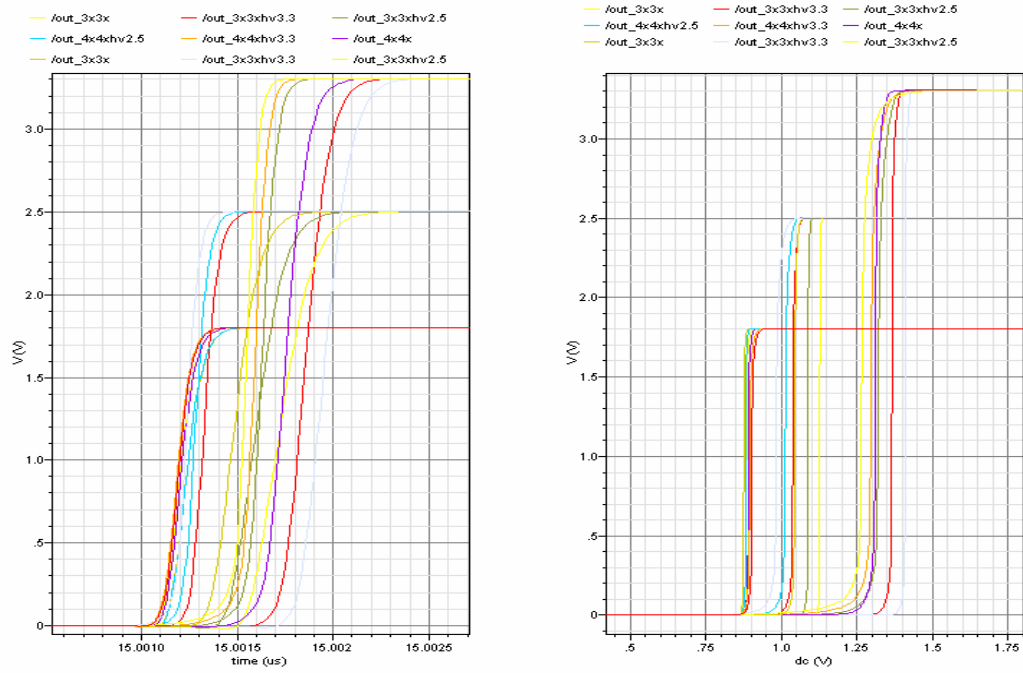
Specifically and only for driving the write transistor of the SRAM cells, inverters for high voltages (>1.8v) were created: The x3_hv and x4_hv strength inverters are copies of their low-voltage equivalents, with the nmos transistor enlarged to the same dimensions of the pmos. The equal transistor sizing lowers the switching point such that these cells could safely be used with a 1.8v logic input (and guarantees that '1' and '0' are both correctly identified, although the noise margin on a '1' will be smaller). Increasing the size of the nmos device means the rise- and fall-times of these devices are compromised (unequal).

Cell variant	x3_hv	x3_hv	X4_hv	X4_hv
Test load	90fF	90fF	125fF	125fF
VDD	2.5	3.3	2.5	3.3
Rise time	235p	200p	110p	111p
Fall time	131p	131p	82p	103p
Pmos width	6.0	6.0	8.4	8.4
Nmos width	6.0	6.0	8.4	8.4

Above: Tabulated data for typical process corner



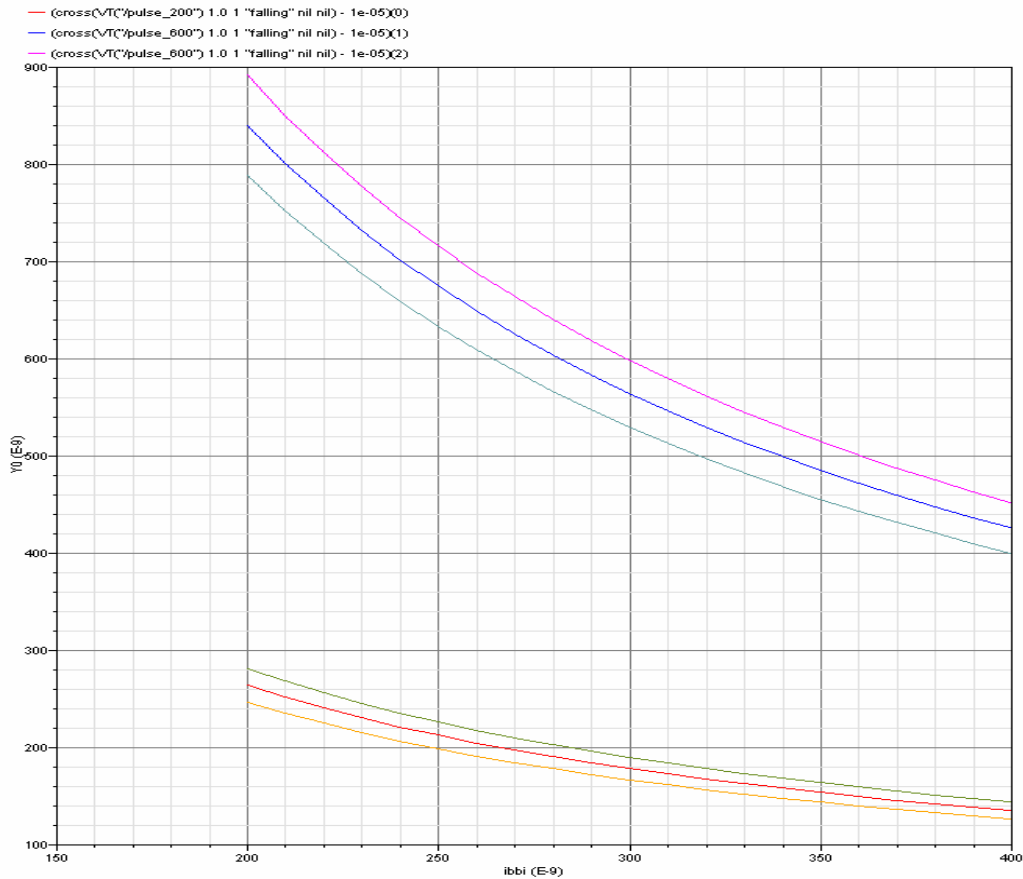
Above: Standard and high-voltage x3 and x4 inverters are simulated with their respective load. 2.5v and 3.3v supplies for the hv parts are simulated to check the switching point is a safe margin from the 1.8v supply. Transient results are shown on the left (500ps per major division), DC sweep shown on the right.



Above: Standard and high-voltage x3 and x4 inverters are simulated with their respective loads at 1.8v, 2.5v and 3.3v supplies in three process corners.

Logic simulation: Monostables

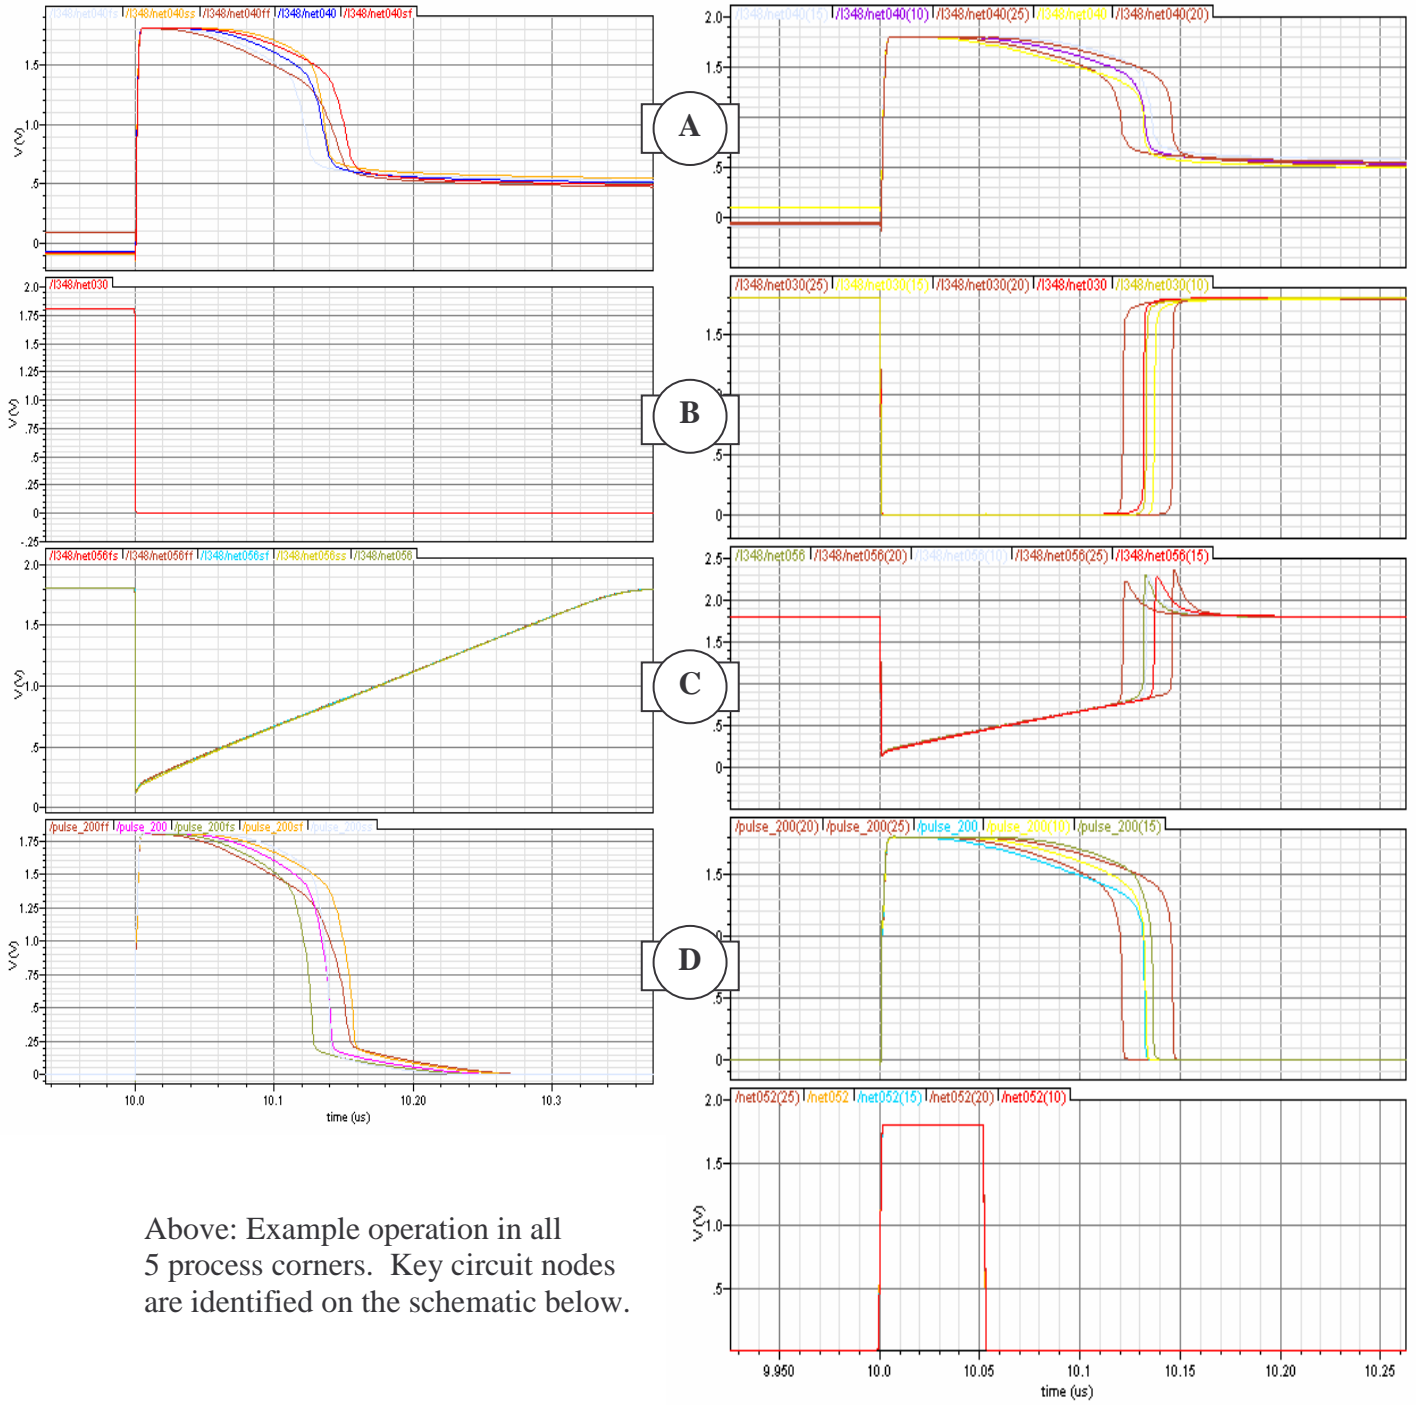
Two monostables were designed with delays in ratio 1:3. These can provide the reset timing for the pre-sample pixel, and also the correct length hit pulse for the preshape pixel. The graph and table below allow the desired timing to be selected by setting the appropriate bias current. The monostables can be triggered with a short or long pulse, the output will be high for the desired time.




Above: Capacitor corners are checked since these will have the dominant effect on the timing. The timing ratio remains consistent, so the external control would allow the circuits to be trimmed to the desired time delays. These circuits can be evaluated on ASIC1 to determine how localised the control of these needs to be.

Bias current	“200ns” monostable output	“600ns” monostable output
210 nA	250 ns	798 ns
265 nA	200 ns	637 ns
360 nA	150 ns	476 ns

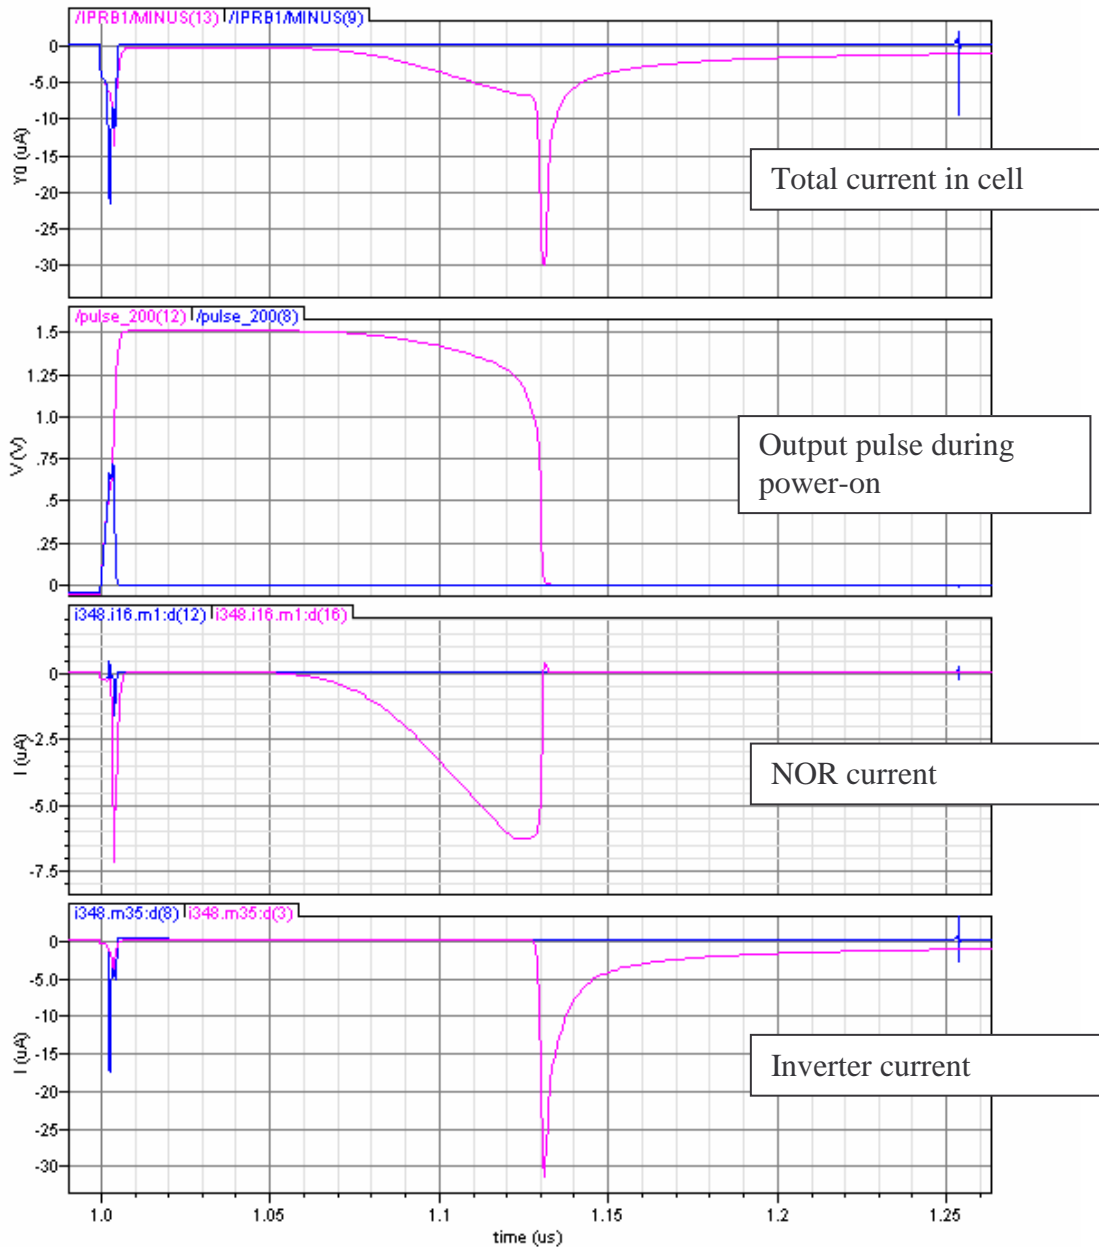
(These circuit simulations have been updated and re-simulated using NW capacitors.)





Above: Example operation in all 5 process corners. Key circuit nodes are identified on the schematic below.




The monostable circuit is designed with the “power-on-reset” input which minimises current consumption at power-up. The power is turned on at time 1us in the simulation. The POR input is held high from time 0 and released after 1.25us.




Above: The currents flowing in the monostable circuit during power-up. The pink traces show the circuit operation without the power-on-reset. The blue show the more satisfactory operation where current flow is minimized.



Above: Monte Carlo mismatch simulations show the likely spread in the pulse duration from the two monostable circuits at one particular setting.




Above: Using a single monostable circuit to generate the POR signal for all the other monostable circuits (external access is also provided for full control).



Logic simulations: SRAM register

The basic SRAM register is illustrated below:



The cross coupled inverters are made from minimum size transistors, and must be overpowered to successfully overwrite their state. Note the readout mechanism is current mode, hence the data column connection is via a pull-down transistor. This ensures the load driven by the inverter cell is kept very small – a large load could also destroy the stored state.

SRAM overdrive strength

It is important to understand the strength required to overpower a SRAM cell – if too many cells were written from the same data (ie the timestamp signal) they might become corrupted.

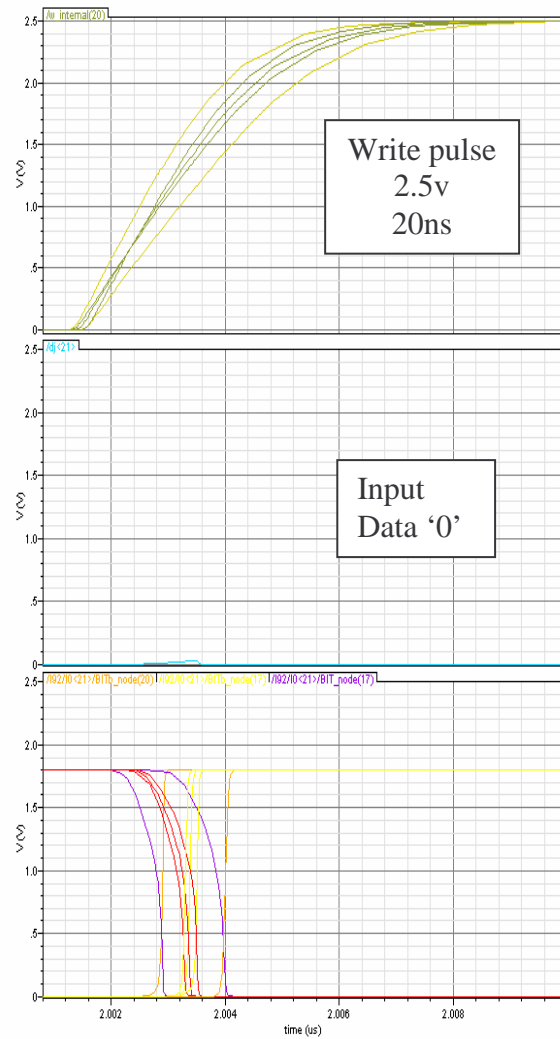
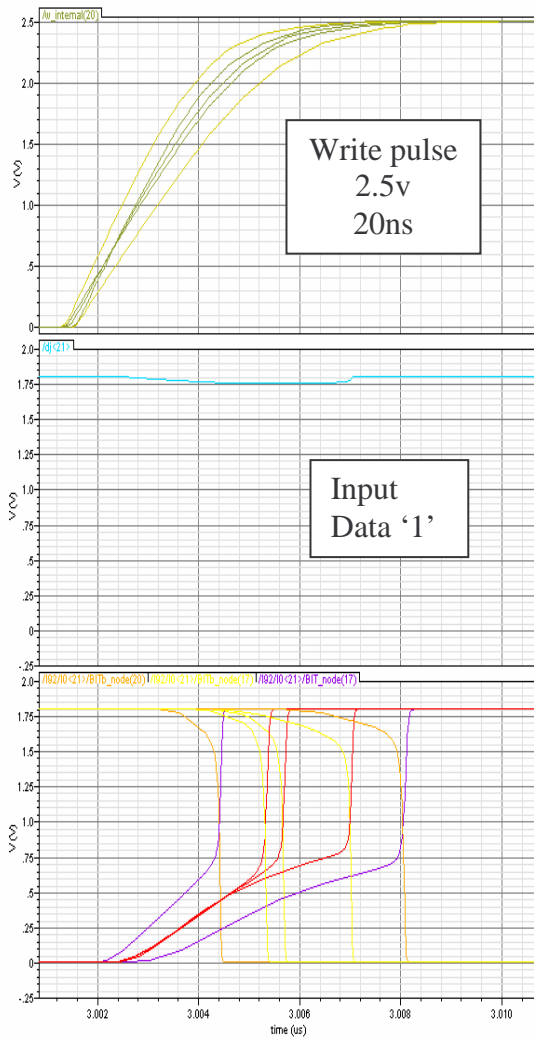
Process corner	Driving cell (SRAM Write signal)	
	x3_hv (2.5v)	x3 (1.8v)
SS	9	0
SF	8	0
TT	10	11
FS	10	10
FF	11	12

Above: Number of active SRAM cells that can be overpowered (successfully written).

This data indicates the need to overdrive the write transistor for reliable SRAM operation in all process corners. A high-voltage tolerant transistor must be used for the write transistor (these have a larger minimum size and therefore will make the SRAM layout larger than previously estimated).


This information is also important to ensure timecode buffers are not expected to drive too many rows, maximum 8. Monte-Carlo simulations have shown that a driver cell of weaker strength (ie x2) is insufficient and results in corrupted data.

Example operation




Above: Internal 2.5v write signal successfully overwrites the SRAM (the internal BIT and BITb nodes) with the input state.

Right: Monte carlo simulation (N=250) shows 4 failed write attempts. The write transistor is the critical component for this circuit action and has probably been down-sized in the extreme corners of the monte-carlo variations.



The write transistor can be driven with up to 3.3v in the event that SRAM writing fails. The write transistor can be enlarged slightly in the layout, which will also help to guard against this rare write failure. Updating just the write voltage to 3.3v the monte-carlo simulation is run again:

Right: Monte carlo simulation (N=250) shows 100% success with the write transistor driven at 3.3v.

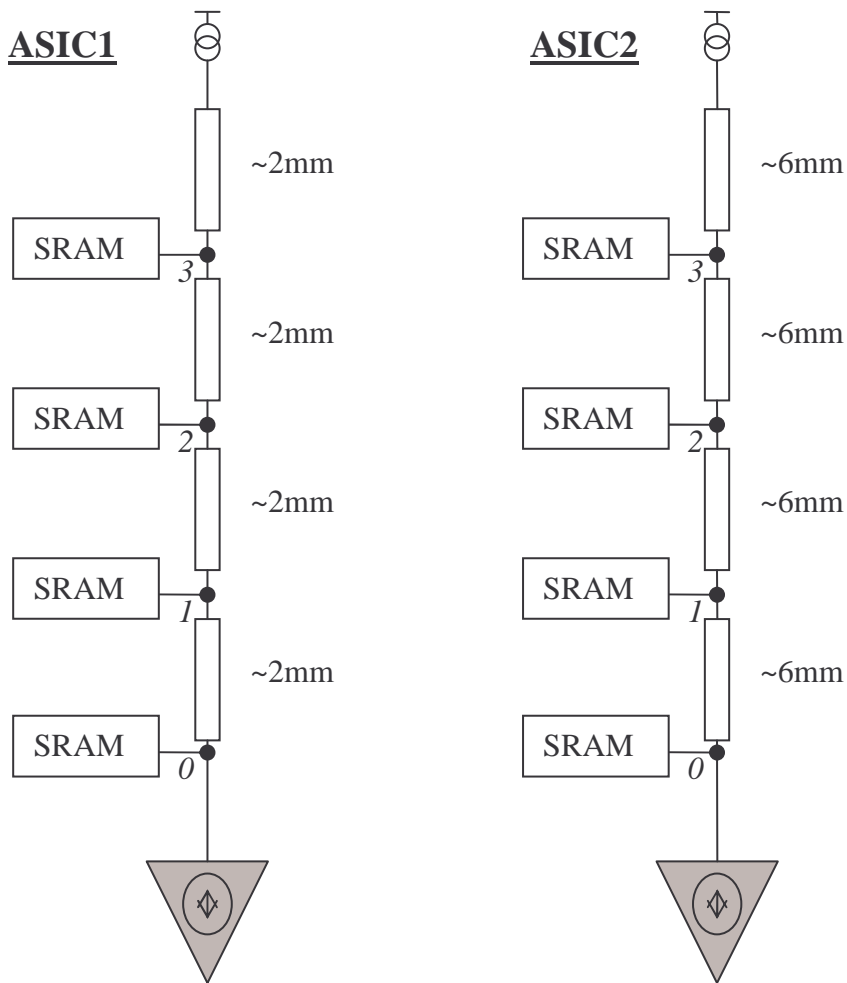


Logic Simulation: Data Sense

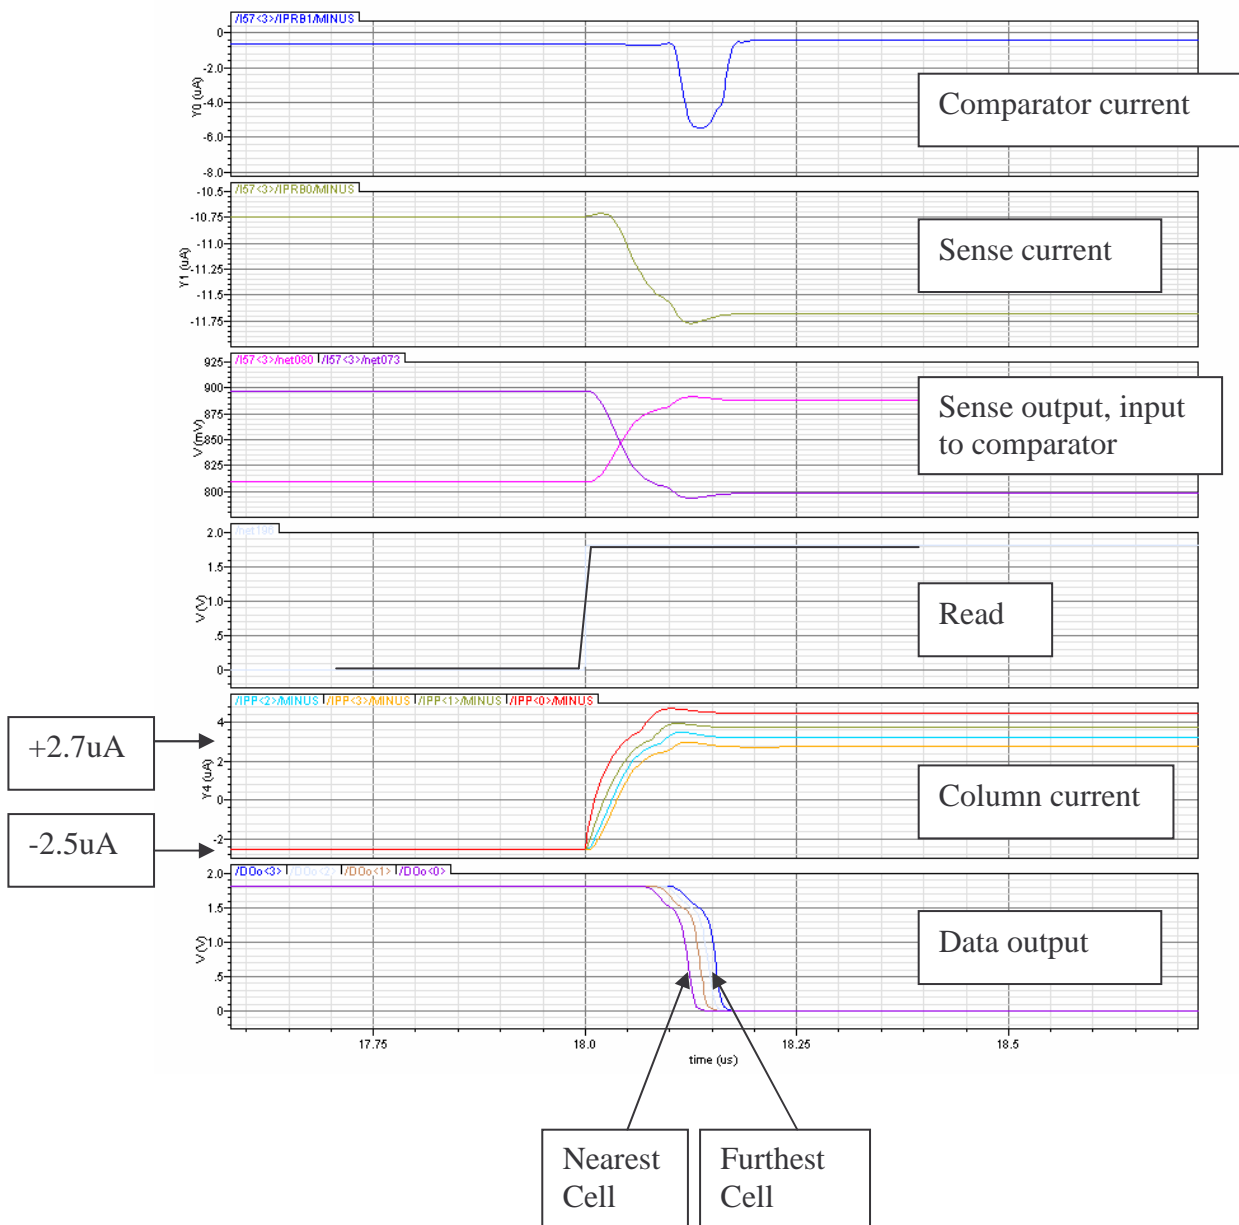
The circuit is operated in reduced current mode as detailed in the table below:

Data Column	Static current bias	10uA
	Mirrored 4:1 onto line as	2.5uA
SRAM cell	Current sink for 'zero'	5uA
	Current sink for 'one'	0
Data Sense	3.3v bias ref (pad)	130uA
	3.3v bias ref (local)	13uA
Inverter (comp)	Switching current limited to	5uA
	Total in sense amplifier	15.4uA

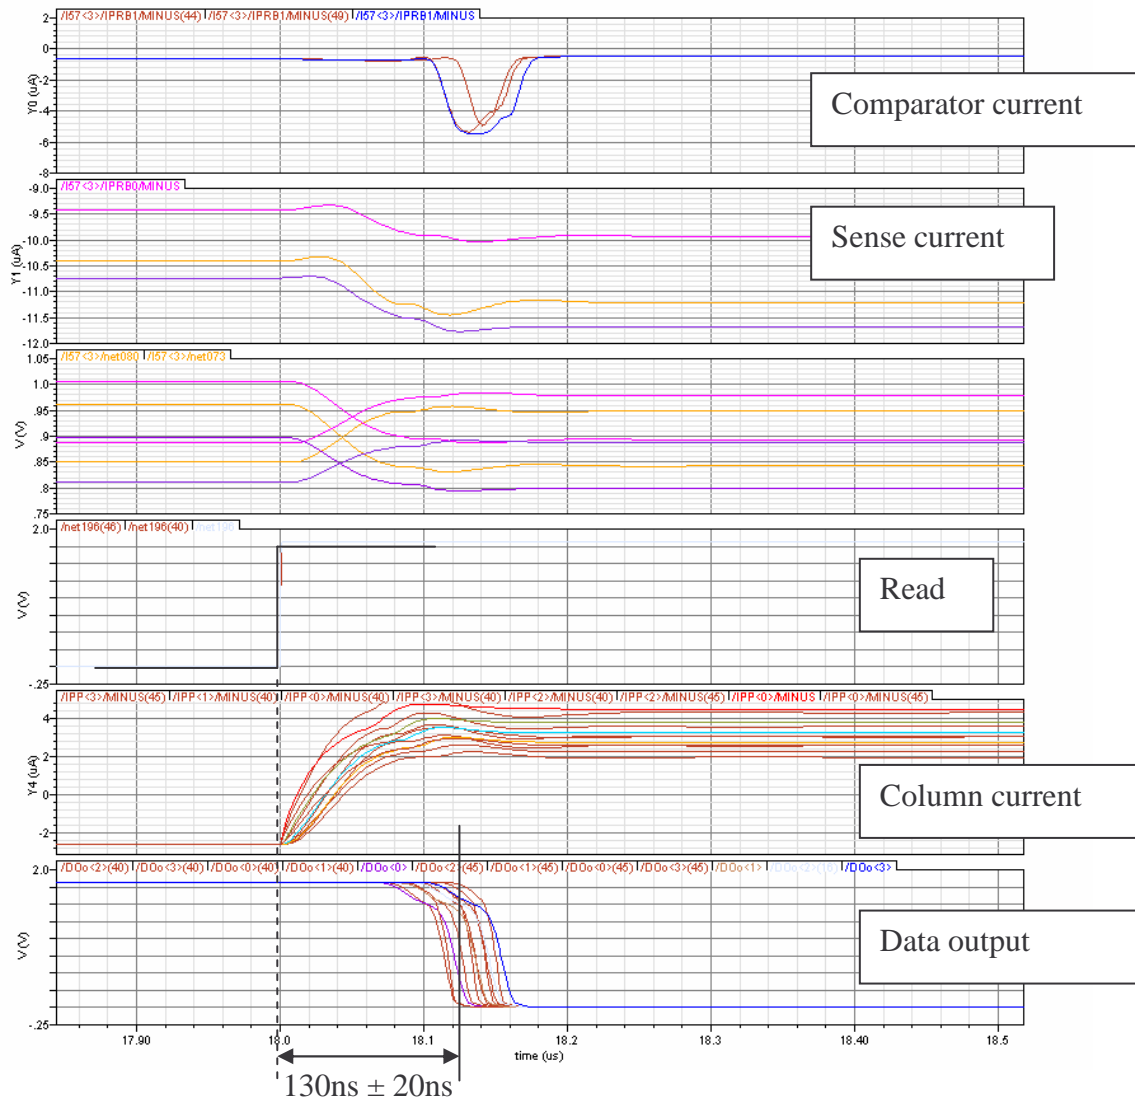
The circuit is simulated with models for full line resistance and capacitance. The full reticle size has been verified, although this will not be implemented until ASIC2. Results for ASIC1 are also checked, although this performance will be faster due to the smaller distances involved.



Results Waveforms




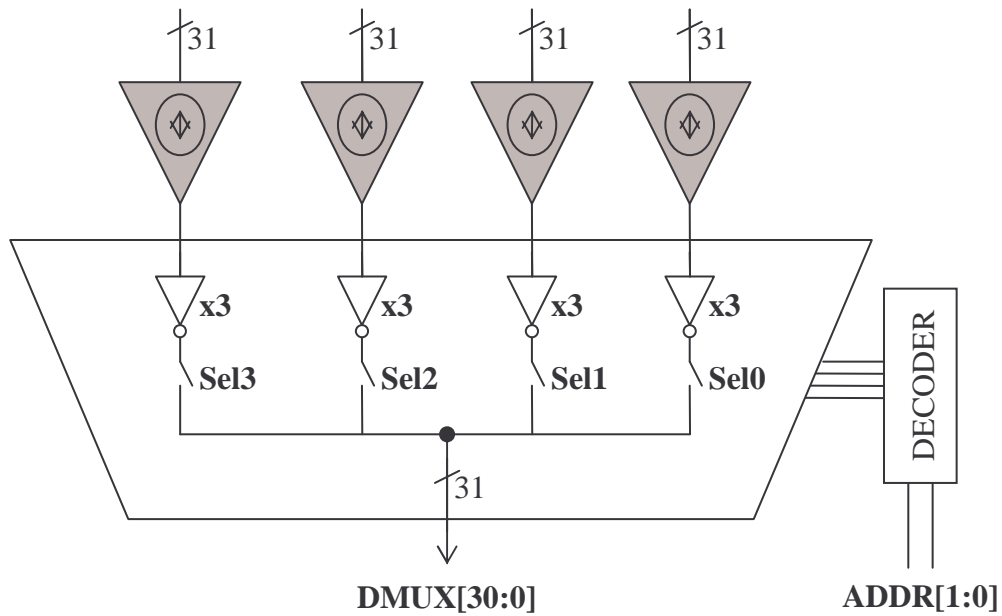
Above: Key circuit nodes are shown to demonstrate typical circuit operation. There is little difference between the read times for near and far cells.



Above: Process corners simulated to show the variation in data sense times. Allowing 200ns for each read permits readout at 5Mhz.


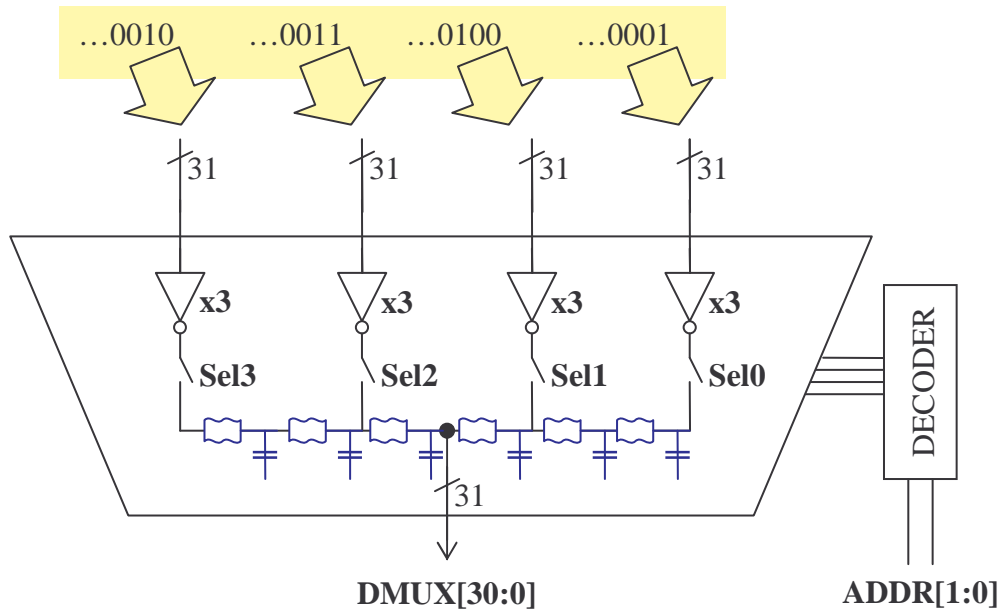
Logic Simulation: Data output multiplex

The diagram below shows how the data sense amplifiers are multiplexed to a single 31bit parallel output



Above: Internal select lines as decoded from 2-bit Gray code ADDR input.


The circuit is primed with test data values at the multiplexer inputs. The multiplex address are cycled through the four states. Only the four LSBs are plotted for clarity. This simulation must demonstrate the x3 inverters can drive the common output line including RC loading.

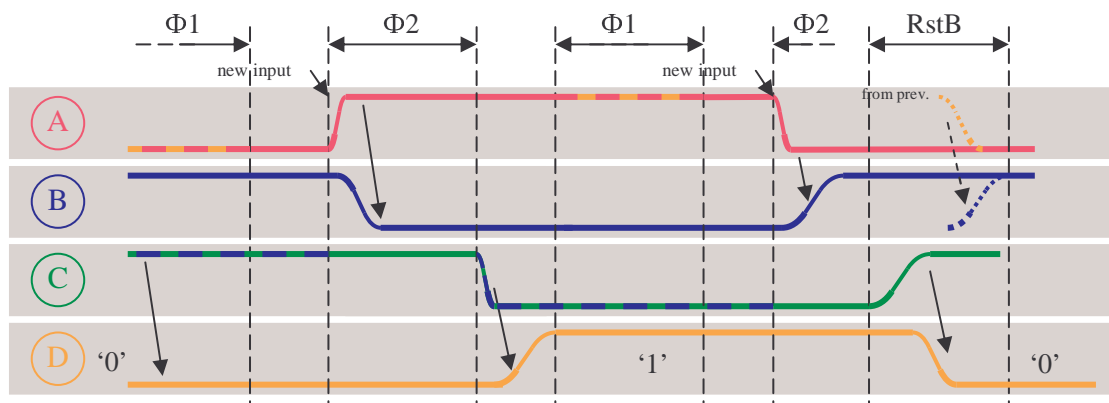


Above: Annotated digital data output from the multiplexer. Sel à Data delay is approx 25ns with the estimated line loadings (acceptable for 10Mhz performance).

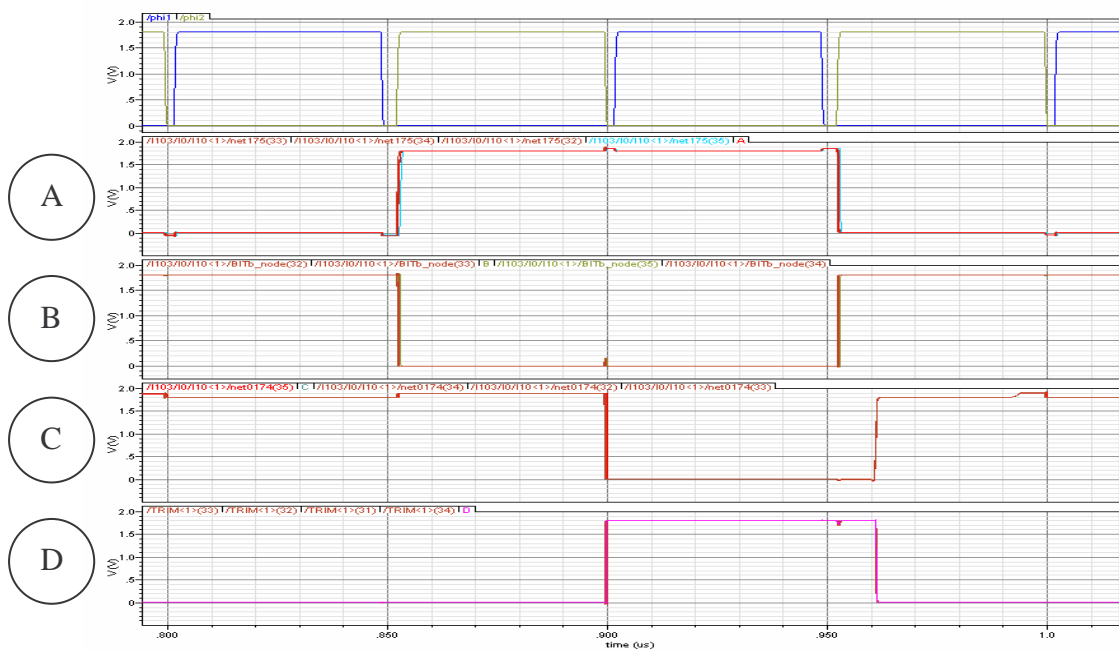
Logic Simulation: SRAM shift register cell

The SRAM shift register cell and waveform timings are illustrated below.

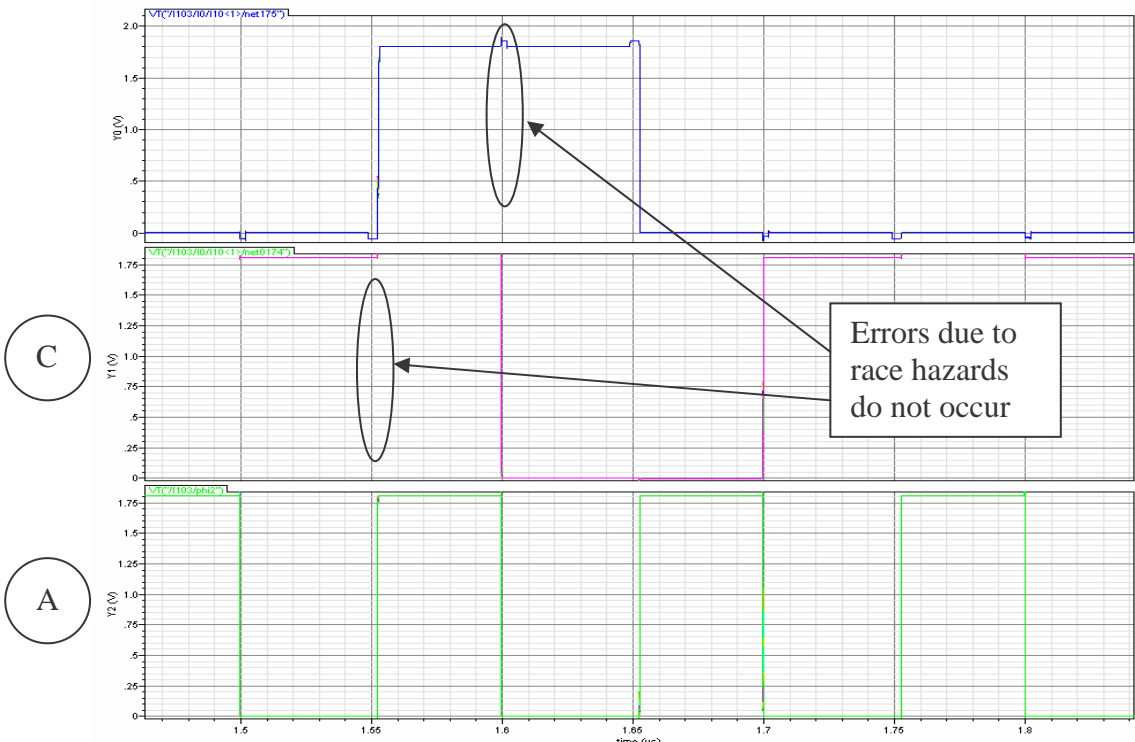





- During $\phi 1$ the SRAM cell holds its own state with two cross-coupled inverters.
- As $\phi 2$ switches on, node A is updated with the new input state (connected from previous SR cell through a $\phi 2$ switch). Node B updates accordingly. At the same time switch $\phi 2b$ switches off, thus allowing the cell to drive its current state at the output without corruption from the new input.
 - RACE CONDITION
 - Node B is updated with the new value at node A. Node A must not be allowed to propagate through the inverter and $\phi 2b$ switch to the node C (where it would corrupt the cell's previous data).
 - The undesirable signal path includes an extra inverter with respect to the desirable signal sequence, therefore the race condition will result in the desired outcome.
- During $\phi 2$ the SRAM cell drives its stored state to the next cell (connected to node D through a $\phi 2$ switch). The stored state is held on node C, comprising parasitic and gate capacitances of the inverter.
- As $\phi 2$ switches off the cell becomes isolated from its neighbour, and switch $\phi 2b$ closes to update node C with the new state.
 - RACE CONDITION
 - Node C is updated with the new value at node B. Node C must not be allowed to propagate through the inverter and $\phi 2$ switch to the output node and into the next cell (where it would corrupt data at node A).
 - The undesirable signal path includes an extra inverter and $\phi 2$ switch with respect to the desirable signal sequence, therefore the race condition will result in the desired outcome.
 - An intermediate value may be sampled on the parasitic capacitance between two SR cells – this would not affect the correct operation of the SR cell.
- As $\phi 1$ switches on the cell completes the internal feedback loop and the new state is stored indefinitely (whilst powered).



Above: Corner simulations showing the internal nodes in the SR cell. The race conditions at both edges of phi2 behave correctly.




Above: Mismatch Monte Carlo (100 runs) to check performance is robust in race conditions.



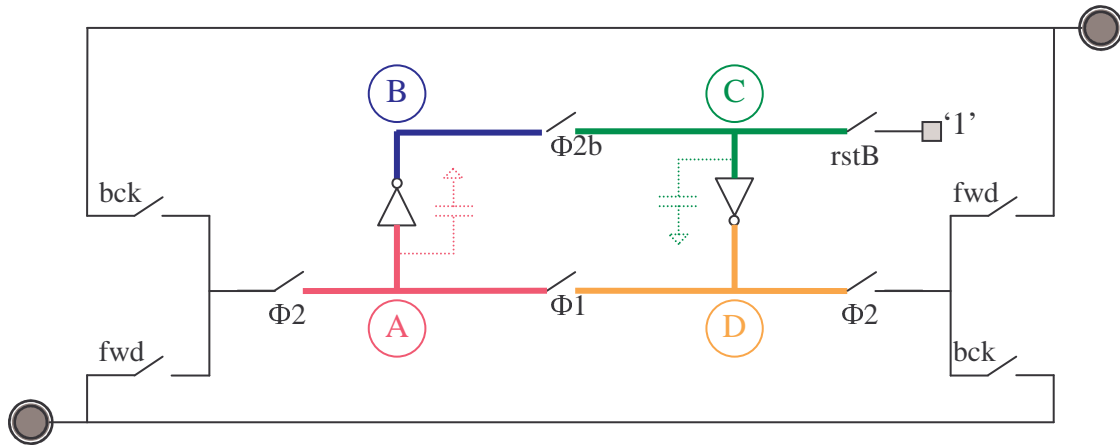
Above: SRAM reset is shown in detail. Bit 12 of the shift register is currently active; the reset sets this to zero along with all other bits in the shift register; This sets the local OR evaluation to true (signal SHIFT_IN_CHK; also generates the “Done” signal for readout control) which sets the input to the first register (bit 19) high. Thus the shift register is reset to its initial state, 1000000000000000000000.

Right: Process corner variations from above.

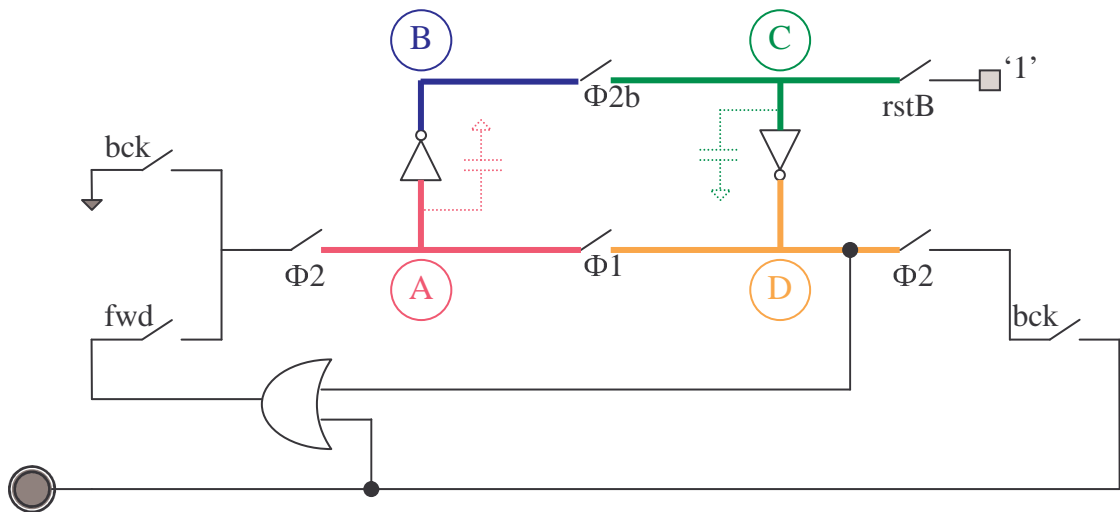


Logic simulation: Bi-Directional SRAM cells


The standard SRAM shift register cell is adapted to make it bi-directional as illustrated below.




The last cell must be modified such that it will hold the token rather than lose it if clocked again (robust operation in overflow conditions) and also ensure that a '0' is driven as the cell input when in backwards mode.




The “bi-directional” and “endstop” shift register cells are arranged to form a 20-element shift register. A logical NOR generates the SR input in the fwd mode such that either after reset or one forward clock cycle when empty, the register is reset to the default condition 10000000000000000000.




Results Waveforms



Above: Typical operation involving 5 consecutive hits (ie all in the same 42-pixel section). These simulations are run at $\sim 50\text{MHz}$ to check operation is possible at the target 150ns bunch-crossing rate.



Above: The overflow condition is verified – the register is clocked for 25 hits. The circuit can be seen to hold the token at position 0 and then assert each enable line for all the registers back to 19.



Above: Example operation at 50Mhz in all 5 process corners with R-C line models included. 8ns pulses separated by 2ns are intact after a full 84 * 50 micron length. SRAM operates correctly.


Logic simulation: Bi-Directional SRAM shift register

Simulation top-level = "sim_sram_shift20"


The bidirectional shift register generates the read and write pointers for the SRAM memory bank.

In normal operation the register is clocked once each register that must be written. During readout mode the register is clocked once in backward direction to initialise, then combinational logic enables each row to access each of its valid registers sequentially until the column read is complete.


If >19 forward ("write") clocks are received the register outputs a "full" wired-or signal, and will correctly read back the first 19 registers that were written.



Above: Typical operation, showing 5 clocks in the forward direction, then 25 clocks in the reverse direction, of which the first 5 act upon the shift register and replay those registers that were written-to. The shift register thereafter ignores any further phi2 clocks and sets the "Done" signal (combinational NOR of all 20 bits).




Above: Overflow condition: 25 hits are simulated, note that reg 0 remains active beyond the 20th hit – this is used to generate the overflow signal, and ensures that when clocked backwards the shift register has not lost its token so that readout of the other 19 registers is achieved (as illustrated).



Above: Correct operation verified in process corners SS,TT,FF.


Logic simulation: Latch-Hold circuits

The structure below uses a ping-pong architecture to sample the hit signals on both rising and falling edges of the Hold signal. This means the sampling “Hold” signal can run at the ~6Mhz.




The “hit” states are stored on the gate capacitance of the second inverters. This capacitance is small, but the state must be held for only 150ns.


Results Waveforms



Above: Typical operation is demonstrated in all 5 process corners. A small injection of charge is seen on the floating storage node due to the adjoined switch.



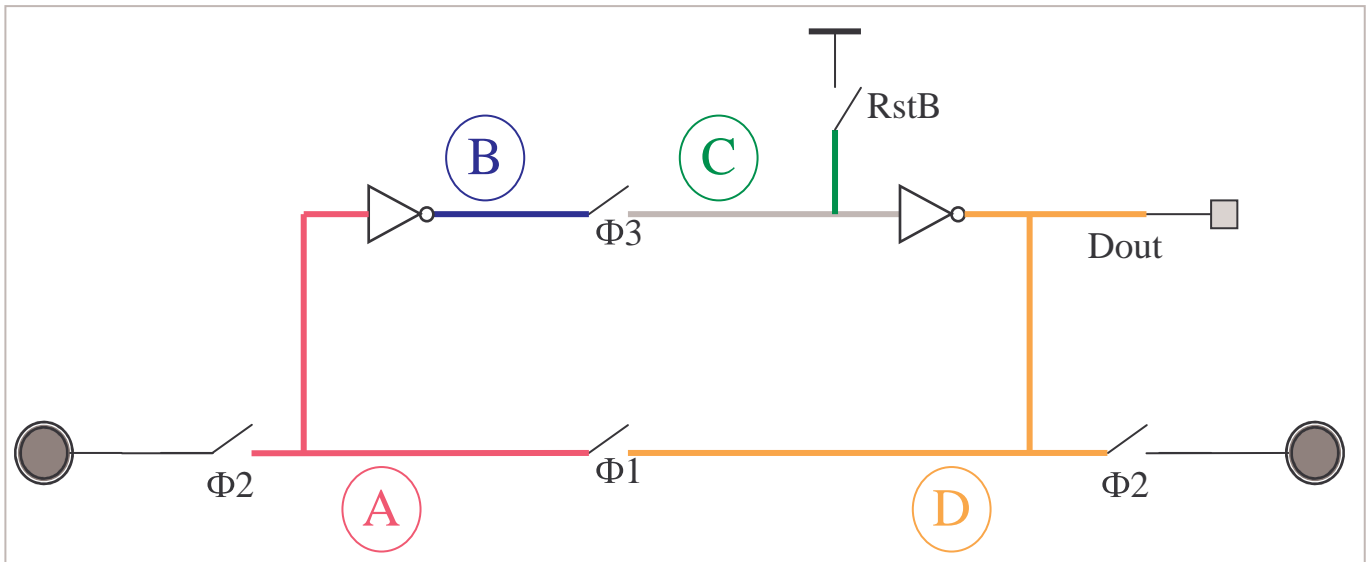
Above: Simulation in all process corners is shown indicate the likely lifetime of this small storage node. Taking 250us as a safe interval, the hold signal must be clocked at a minimum of 4kHz to avoid charge loss effects.




Above: The hold signal is clocked at 300ns (ie 150ns sampling rate). A parasitic extraction of a preliminary layout is simulated. The charge injection effect seems to be reduced when parasitic capacitances are included.

Logic simulation: 3-clock SR cell

The SR cell can be driven with 3 clocks to completely eliminate the race condition. This is most useful for mask programming where long clock lines will be driven slowly – in these circumstances the true complement needed for successful operation of the 2-clock SR cell would require local clock buffers in every pixel – which would draw excessive current.



Results Waveforms




Above: Slow operation (100kHz) is verified in all process corners to ensure charge leakage from transistor gates does not degrade performance at such low speeds. Clock edges are loaded to achieve 50ns rise/fall times for realistic simulation.

Logic simulation: Mask & Config programming (shift5)

Each pixel contains a 5-bit SRAM shift register, arranged such that each column forms a long shift register. Global clocks are distributed along each row. Pixels buffer and generate the compliment clocks internally to ensure clean clock waveforms for the SRAM shift register cells.

Pixel operation



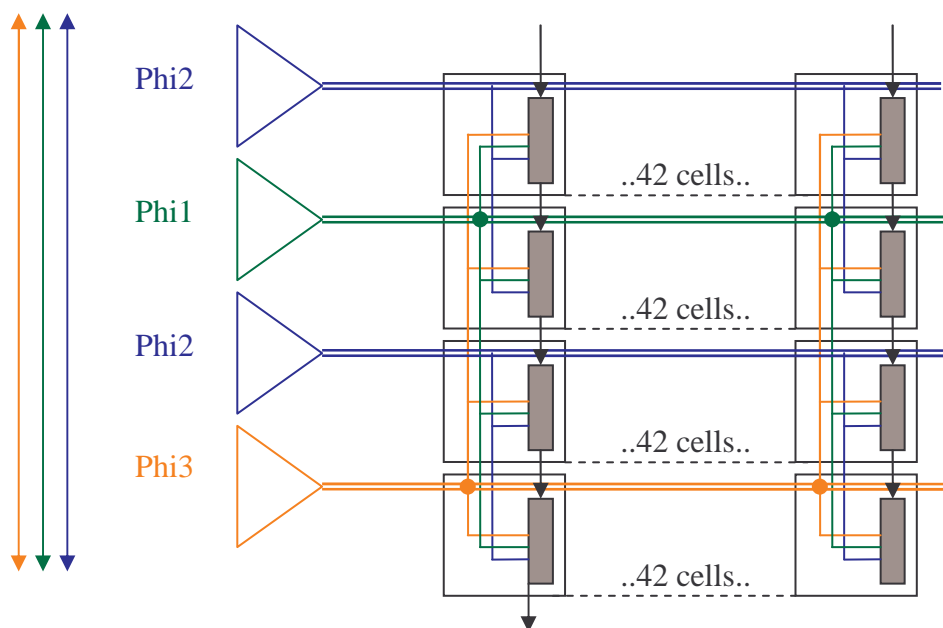
Above: Example waveforms showing shift register functionality using the 3 phase clocking scheme.

Reset occurs during phi2. Input data is latched on the falling edge of phi2. In the example above a single 'one' is shifted through the SR. In the real system, data is shifted in as TRIM-LSB first with the MASK as the final bit. The full string of 5-bit codes in this order are shifted through the column.

Configuration “SLOW” clocks


Clock buffering is summarised below. Clock loads are made equal.

	Pixels driven	Rows that share	SRAMs per pixel	Gates per SRAM using clk	Total load (# gates)
Phi1	42	4	5	2	1680
Phi2	42	2	5	4	1680
Phi2	42	2	5	4	1680
Phi3	42	4	5	2	1680



Each 4-row section is repeated throughout the pixel array. Global clocks run vertically to feed the large (x4) row buffers.

Results waveforms




Above: Config clocks along one line of 42 pixels. Rise times of $\sim 25\text{ns}$ are consistent between process corners. Switching current (per buffer) is $\sim 300\mu\text{A}$.

Estimated maximum switching current during programming of ASIC1 is 50mA, although for slow clock rates this will average and smooth to a small fraction of this peak figure.

Logic Simulation: Fast Config Shift Register

The input and readback fast shift registers are formed from the 2-phase SRAM shift register cells. Each cell has its own local clock buffering which permits high speed operation, but at higher power.

The input and readback shift registers are directly connected as illustrated below and simulated at 100Mhz.




Results Waveforms

(see next page).

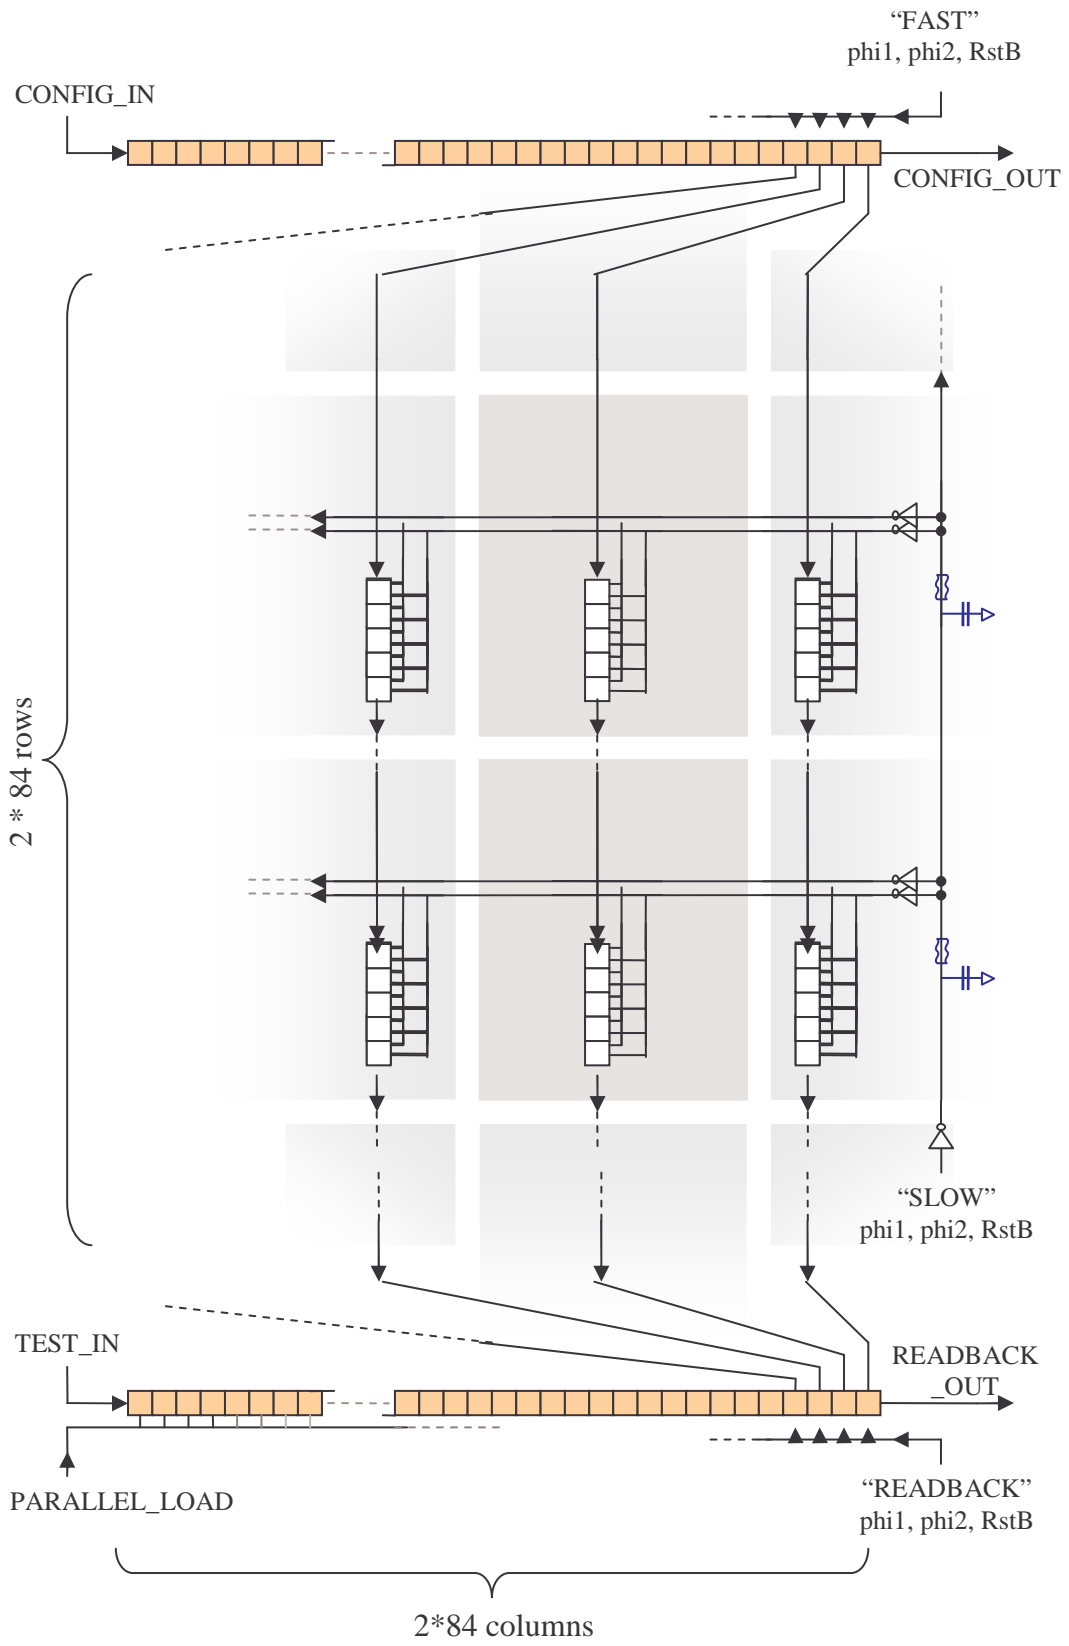
100Mhz clock signals are seen to be intact and non-overlapping at both ends of the row.

Current peak ~3mA for clock edges.

Note that a single cycle delay is seen in the readback shift register data due to the parallel load architecture.



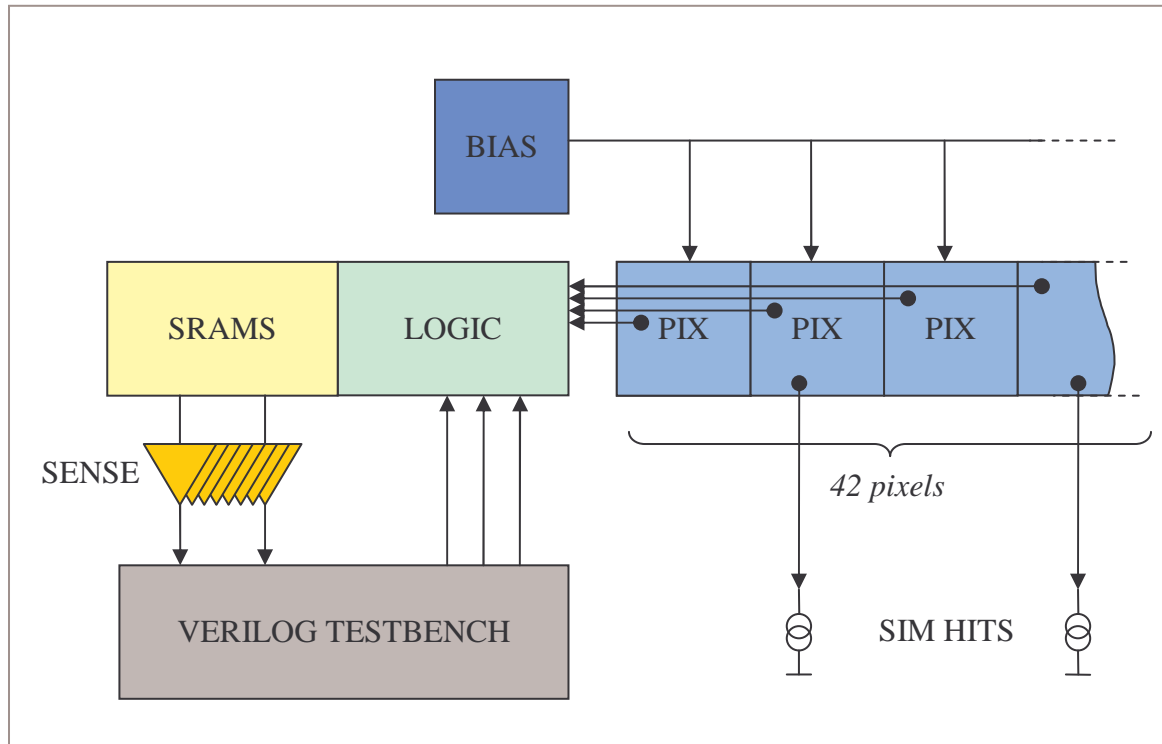
Overview: Configuration programming and Verification



Simulation results

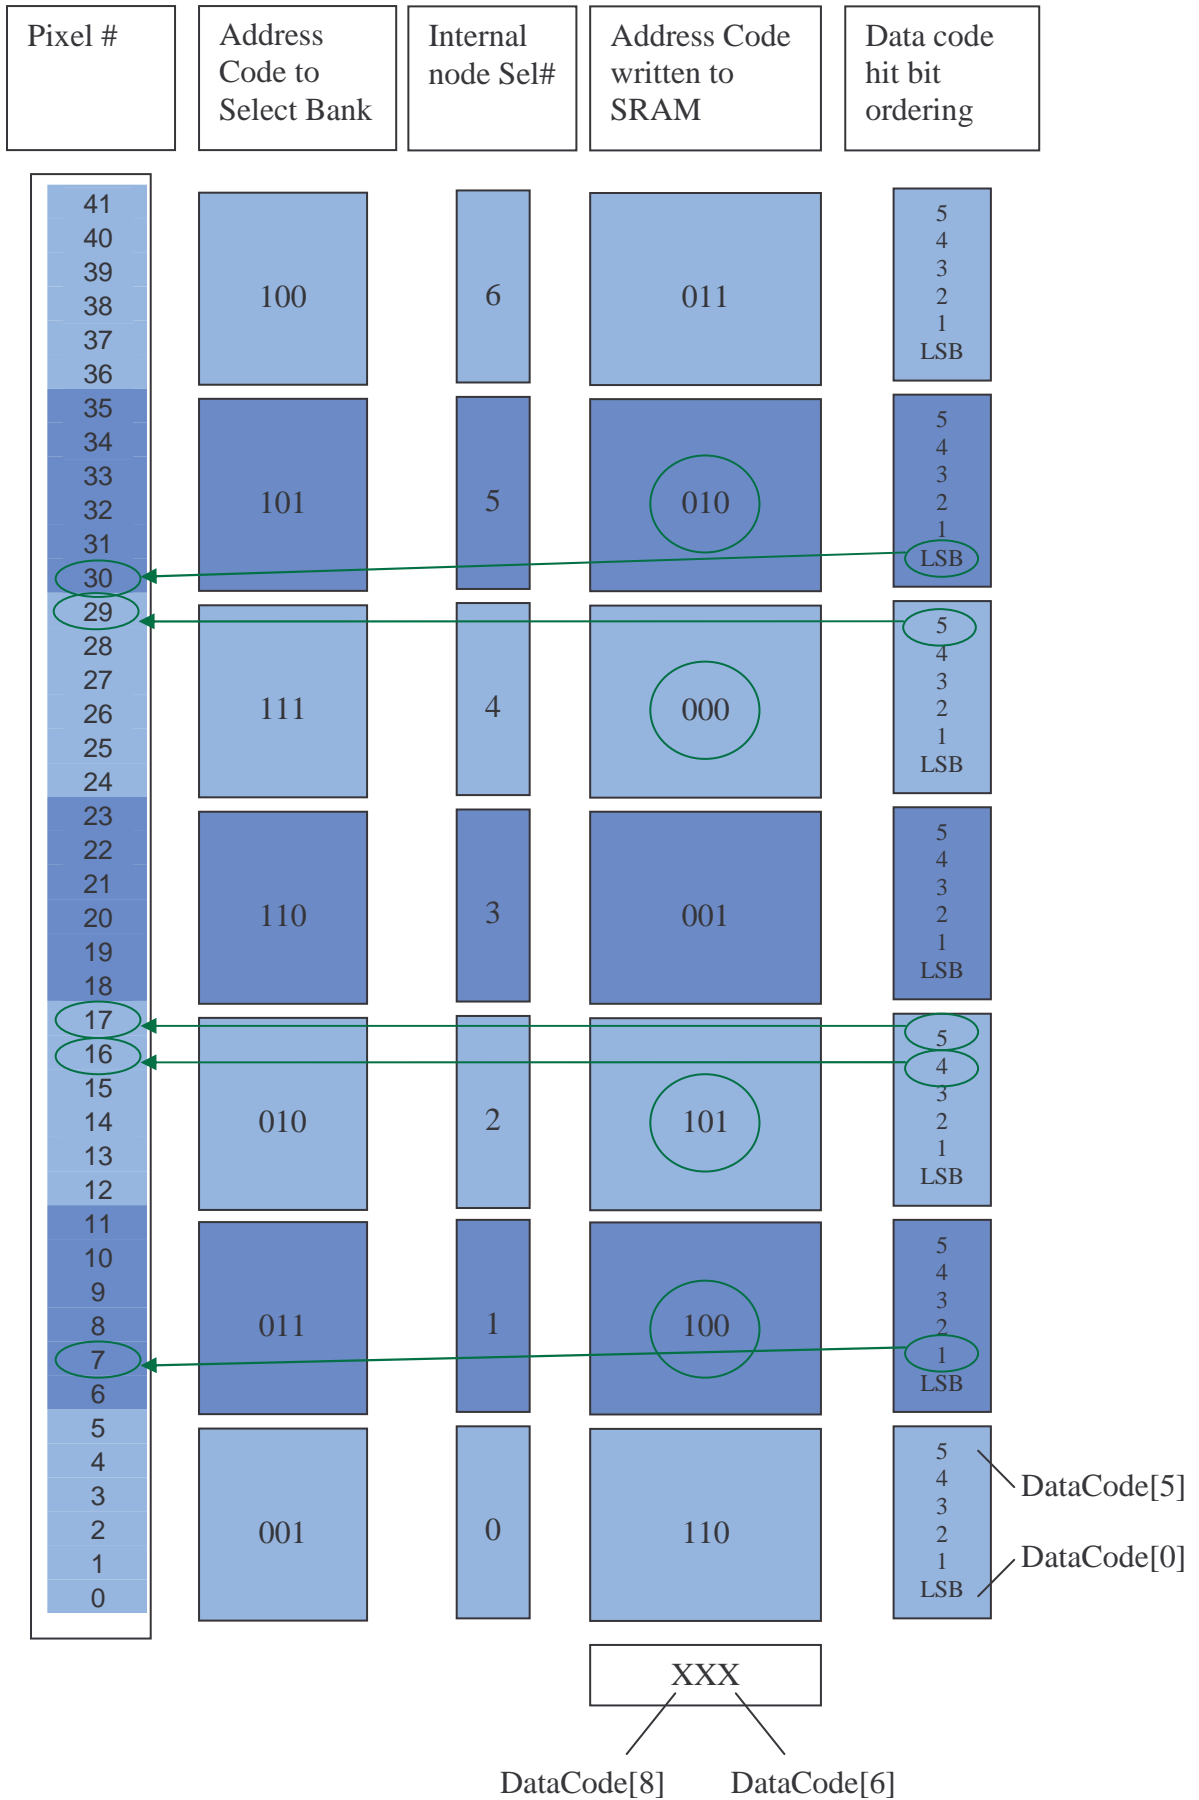
Mixed mode verification simulation implements a fast SR, a few lines of pixel SR cells and a fast “READBACK” SR. The pattern shifted out should match that shifted in.

Logic simulation: Full pixel 'slice' simulation

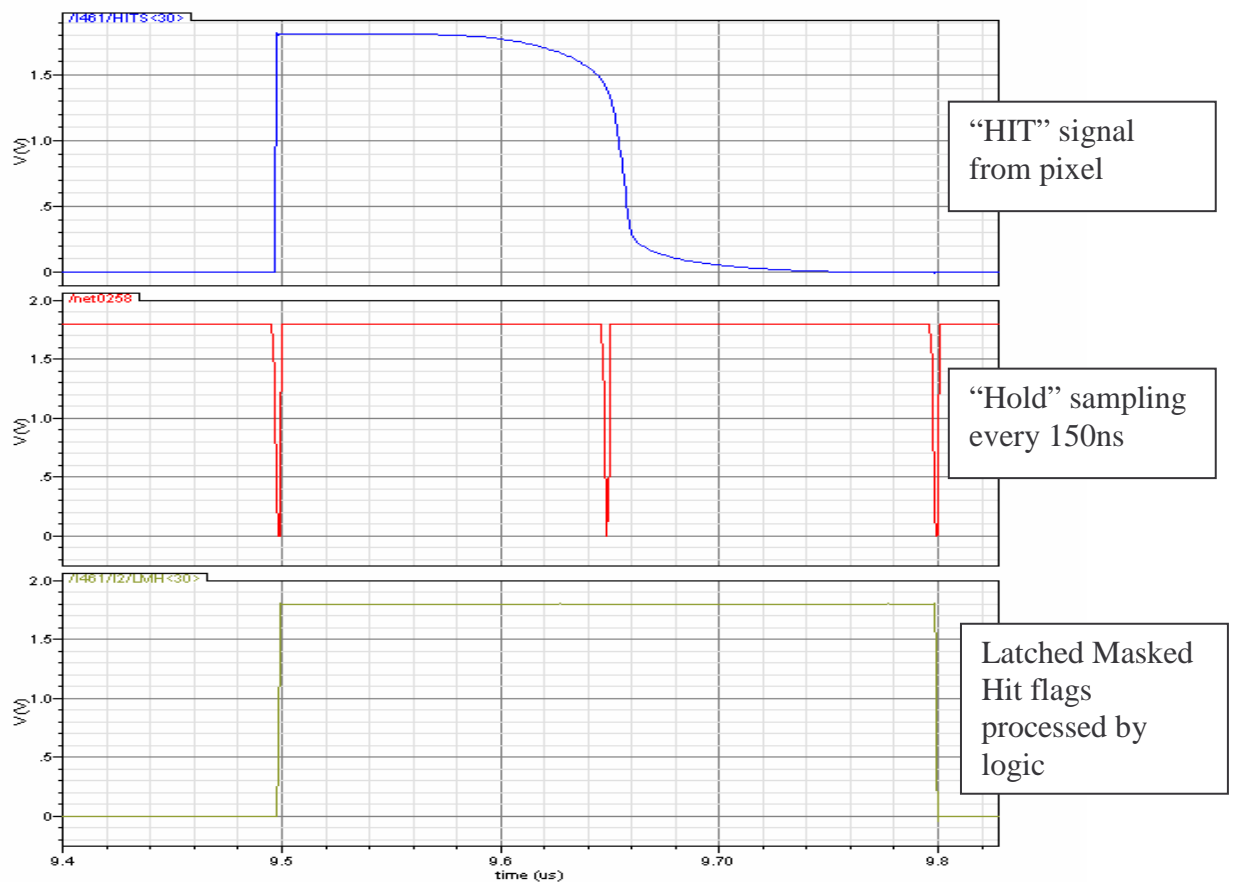


The diagram above depicts the slice simulation including 42 pixels, the master row controller, 19 SRAM registers and data sense amplifiers. This simulation omits any mask or trim programming to reduce the simulation time (all channels are enabled).

The testbench initialises the logic and waits for the analog to settle/power-up. The logic is then sequenced as per bunch-train operation, and simulated hits are applied to different pixels at different times. The logic detects and stores the hit location and time, which is then read out and displayed in the simulation log after the bunch train is complete (16 crossings).



The double hit seen in the simulation deserves further investigation:



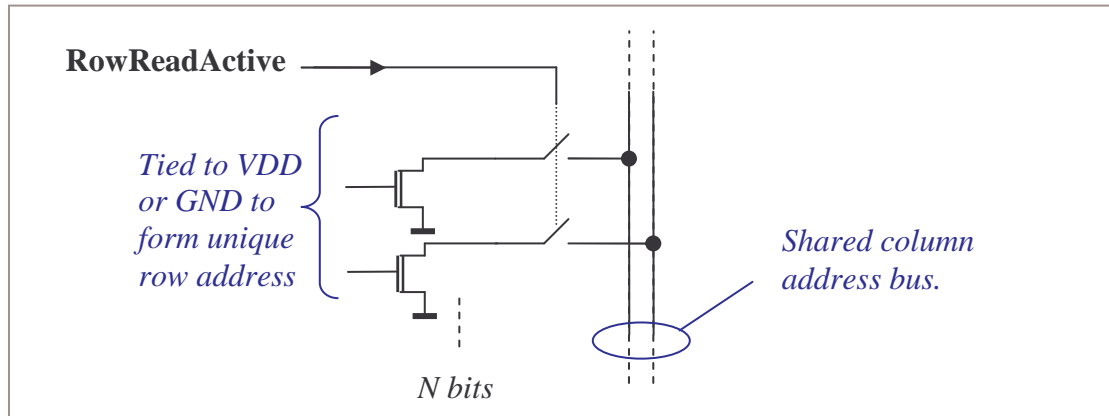
Looking at the internal signals for pixel 30 shows that the rising edge of the hit signal occurs at approximately the same time as the "hold" sampling. The duration of the hit signal is measured to be 157ns, which means it is sampled as active (high) on two consecutive occasions. This shows how a double-hit might occur.

Whilst it is interesting that this did occur in simulation, we consider the probability of this occurring to be low, since it depends on so many variables (some of which are artificially generated in the simulation). A full analysis would involve much analysis from many device and circuit simulations and is probably impractical! (Discussion welcome)

If double-hits are seen during testing, this might indicate that the biasing of the monostables is too high. In this case the "mso_bias" that controls the timing of the hit pulse can be adjusted until double-hits cease to appear. Note however that adjusting the bias too low would make the "HIT" signal pulse much less than 150ns, in which case hits could be dropped altogether if a similar set of timing circumstances occur.

Logic Simulation: Row Encoder

The basic structure of the row encoder is illustrated below. The function of this block is to report at the column base the row address of the cell that is currently being read. To minimise re-design effort for ASIC2 the full 9 bits are implemented, although only 7 are only required for the ASIC1 test structure. The row addresses are assigned in a GRAYCODE order.

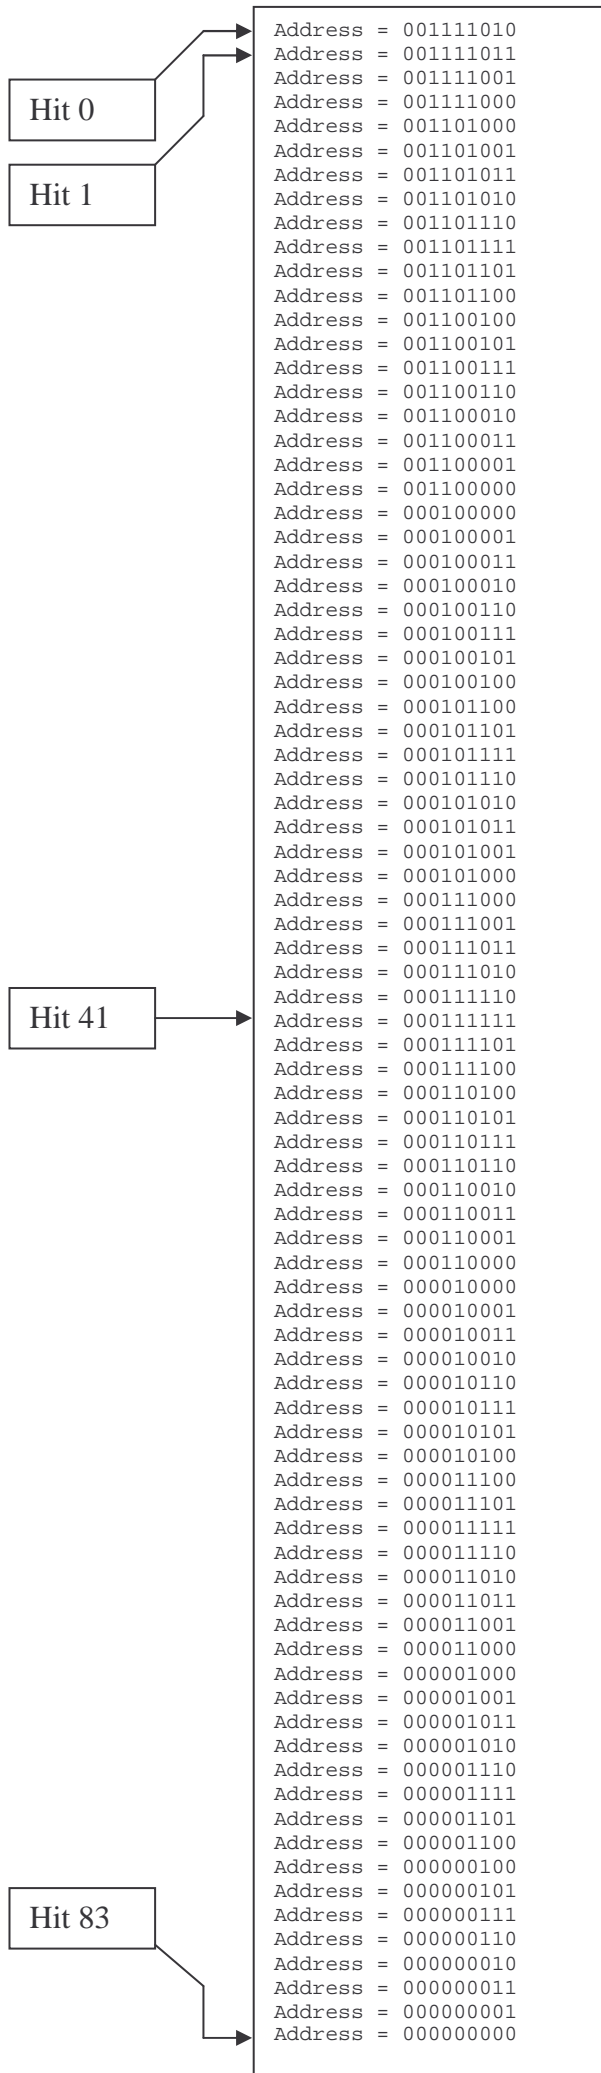


The active cell drives the column bus with its own unique code. The pull-to-ground scheme is compatible with the sense amplifiers used to read the SRAM cells.

Simulations

A row-encoder of 84 cells is simulated.

A Verilog stimulus is used to individually excite each of the 84 enable (RowReadActive) signals and print to screen the address code that is seen at the base of the column. This text file is imported into excel and processed to check that no repeated codes are seen in the 84-code set.




Left: Mixed-mode simulation output log excerpt.

Importing this data into excel, converting to decimal and sorting confirms there are no repeated codes.

Note that the bus ordering is such that the enables input bus' MSB corresponds to the address 0.

Below: Each enable signal is raised in turn to generate the data output shown on the left.

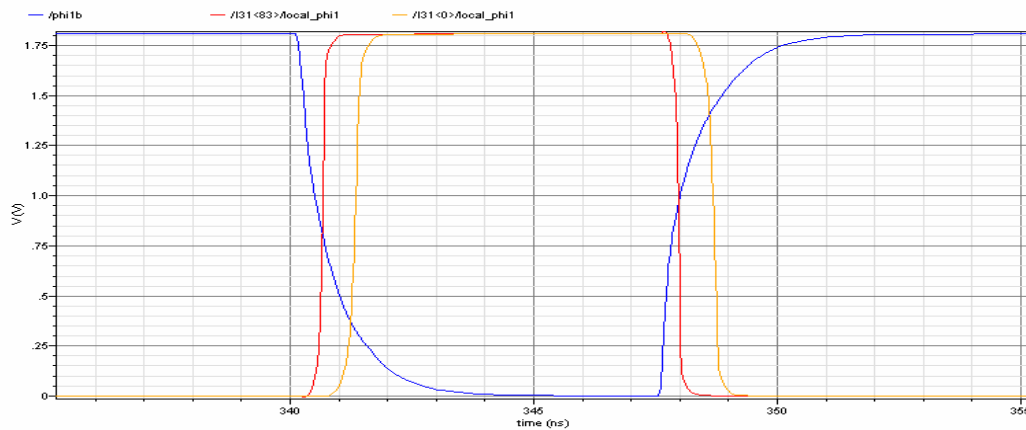


The condition when more than one enable signal is raised will cause conflict on the address bus and draw additional current. This scenario should be avoided – since the enable inputs are derived from the readout logic where only one cell is accessed at any one time this condition should be assured.

Logic Simulation: Driving Long Column Signals


The master-controller logic is stripped of all functionality and only the input gates are left. This dummy block is instantiated 84 times with RC line models and driven with x4 strength buffers to check the clock and control signals can be operated at the required speeds.

Results Waveforms




Above: Internal to each row, phi1 clock edges differ by <1ns between the top and bottom cells in the 84-high column. The blue trace shows the true clock signal on the shared column line that distributes the clock. The internal buffering within each row ensures the local clock edges are fast and true compliment.

The other control signals (fwd, hold, init etc) are buffered with the same x4 buffer, and loaded in the same way as the clocks (a single input port on an inverter/logic cell) at each row. This simulation is taken as verification for the buffering and operating speed of all these control signals.



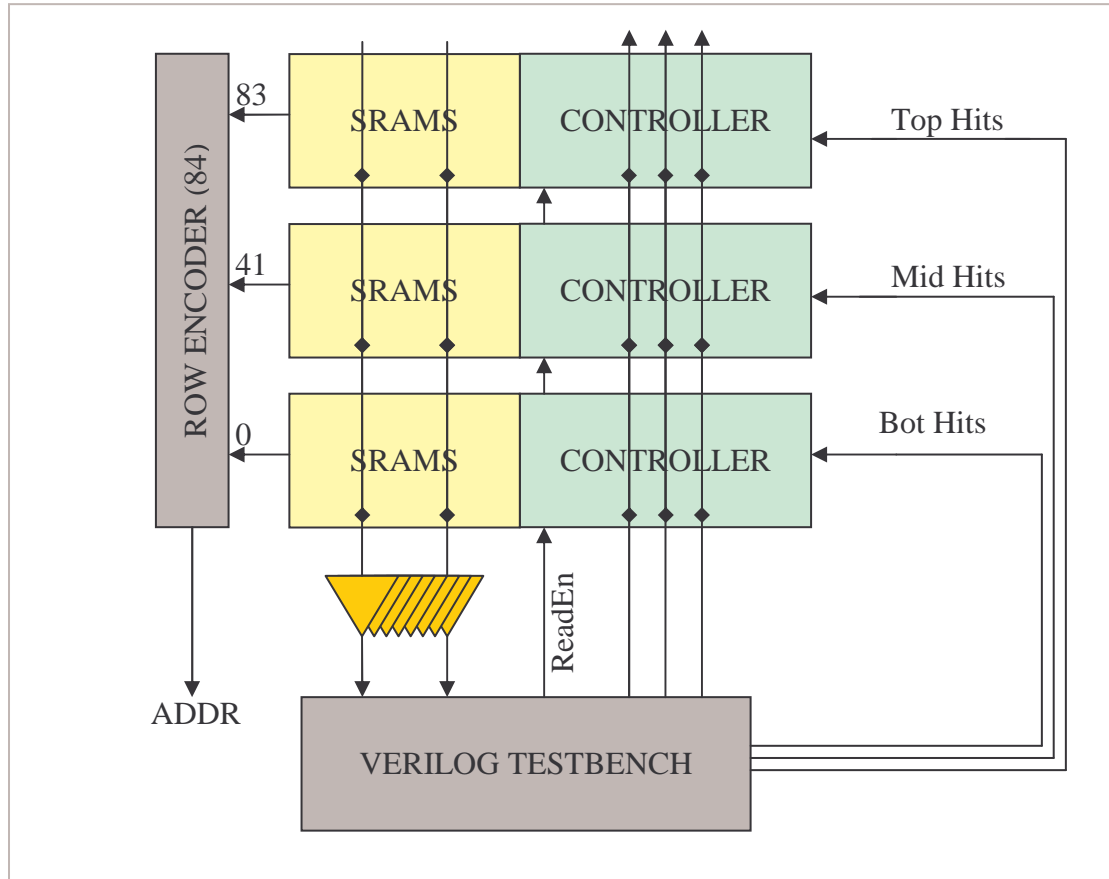
Above: The wired-or used to indicate overflow is seen to respond quickly on the occurrence of overflow, but is much slower to respond once the overflow has cleared. This is primarily a debug feature and therefore operation speed is unimportant; functionality is verified.




Above: The read-enable signal propagates from the bottom to the top of the column (since the dummy cells are all un-hit). The 17ns delay is the typical time taken for the read token to pass through all 84 cells this varies by up to 3ns in process corners. This sets a realistic maximum readout clocking rate of 25Mhz (assuming the token is collected and reset for a clock cycle after every 84-pixel stretch.) This operating speed is more than enough for the data rates of this project, especially if columns are operated in parallel.

Logic Simulation: Three Master-Controllers + Row Encoder

The diagram below depicts the system simulation to check integration of several key circuit blocks in a logic column formation. 17 hits are applied in 3 time divisions.



Results



Above: Inverted row addresses during readout from 'Bot', 'Mid' & 'Top' memories.
Below: Verilog simulation log excerpt with colour-coded annotations to aid reading.

