

Advanced VLSI Design
Prof. D. K. Sharma
Department of Electrical Engineering
Indian Institute of Technology- Bombay
Lecture – 16

Interconnect Aware Design: Capacitively Coupled Interconnects

In the last few lectures, we have seen the improvements that we can bring about over the standard buffer insertion technique by using low swing current mode signal. There were two variants that we saw, one of them used inductive termination to provide high pass function which countered the low-pass nature of the wire.

The other technique actually boosted the high-frequency components of the signal before it passed through the wire that was done through dynamic over driving in which at every transition additional drive was added on through a strong driver. This boosts the high-frequency components of the original signal and therefore even after attenuation of these through the wire you get an undistorted signal at the end.

As result we could get very power and energy efficient transmission at high speeds over wires which are essentially low pass. We had seen a comparison of the inductive termination method and dynamic over drive method. There is yet another technique which has lately been suggested as it is worth mentioning this because it does essentially the same thing as the dynamic overdrive but accomplishes it using much simpler components.

After all what we want, we want to dump extra current into the line every time there is a transition and this can be done by putting a capacitor in series with a strong driver. The serious capacitor will make sure that no extra drive is injected when there are no transitions and at every transition extra current going to the line, so it is a much simpler way of boosting the high-frequency components than using the NAND.

NOR combination that we have seen for the dynamic over driver case. So let us have a look at this capacity driving technique, what are its problems and how we overcome these problems.

(Refer Slide Time: 03:00)

Capacitive Peaking

- ▶ Inductive peaking counters the low pass nature of the wire by providing a high pass function at the receiver.
- ▶ Dynamic overdriving provides compensation by boosting high frequency components *before* transmission by providing extra drive during transitions.
- ▶ The same effect can also be achieved by putting a capacitor in series with the driver at the transmitter.
- ▶ However, this causes the DC common mode voltage of the line to be undefined.
- ▶ This can be handled by using a weak driver to set the DC level and to provide low frequency coupling to the line.



As you have seen inductive peaking counters the low pass nature of the wire by providing a high pass function at the receiver. Dynamic over driving provides compensation by boosting high-frequency components before transmission by providing extra drive during transitions. The same effect can also be achieved by putting a capacitor in series with the driver at the transmitter.


However, this causes problems, the DC common mode voltage of the line will now become undefined because you put a capacitor in series. To counter this, we can put a weak driver to set the DC level and to provide low frequency coupling to the line. Otherwise, we might in fact just by putting a capacitor in series go to the other extreme where the drive is much diminished at the low frequency end and indeed the DC drive is removed altogether.

So essentially, a combination of a weak driver and a capacitive driver, would provide very similar functionality to the dynamic overdrive solution. Let us look at the circuit which accomplishes this.

(Refer Slide Time: 04:22)

Capacitive Peaking

- ▶ The grounded gate pMOS at the receiver keeps the line at Vdd when the input is at 1, because the nMOS is OFF.
- ▶ When the input is at 0, the DC value at the line is pulled down by the nMOS. Thus the low frequency swing on the line is defined by the current drawn by the nMOS and the termination resistance provided by the grounded gate pMOS.
- ▶ These two transistors consume static power when the input is at 0. Therefore the current level through these is kept low.

 The graded inverter chain injects a large charge at each transition from 0 to 1 or from 1 to 0.

- ▶ This has the same effect as dynamic over driving.

So if you look at the circuit on the top, essentially the strong driver is now as simple series of increased geometry inverters. This is a standard way of driving high capacitance loads. The inverted digital input noticed the one inverter here, the inverted digital input actually drives an nMOS, which is directly coupled to the line and this is the model of the low pass interconnector.

At the receiver end, we have a grounded gate pMOS transistor which essentially acts as the termination and also pulls up the wire to be ready because it is always on. The gate is grounded, this is the pMOS and therefore it pulls up this line. The receiver is actually a comparator and we will see the details of its working in a little while. Now the ground in the pMOS at the receiver keeps a line at VDD when the input is at one.

Because when the input is at one, this voltage is at zero and therefore the nMOS is off. The nMOS being off but this is of course always on and that establishes a proper and determined DC voltage at this point. This point then floats up close to be ready. When the input is that zero then this point is that one and the nMOS turns on. This pulls the line to a voltage which is lower than earlier.

In fact, the geometry and the current for which this transistor is biased, determines the low-frequency swing at this line because it normally when the input is one, this will flow to be ready. And it will be at a lower voltage because of the pull down provided by this nMOS and the

amount by which it is lower will depend on the amount of current that this transistor draws when the input is zero and consequently this input is one.

The drop is a combination of the current drawn by this and the resistance provided by this transistor. Now these two transistors together will consume static power. When the input is at 0. Therefore, we would like to keep the current level through these as low as possible, after all these are only providing DC and low frequency coupling and a high power is therefore not required through these.

So, we would not like to waste a lot of power as static power in these two transistors by keeping the drive through these any higher than is required. The actual high frequency drive is in fact provided by this chain of inverters, only two are shown here but any even number can be used. These inverters have progressively larger geometries in order to be able to drive a large capacity of load.

This capacitor then provides the capacity peaking only where is a transition at this point is it coupled to the line through this, if the input remains at zero or remains at one then there is no drive through this capacitor. This is exactly like providing a sharp current pulse of positive or negative value to the line during transitions which is what we had done during dynamic overdriving.

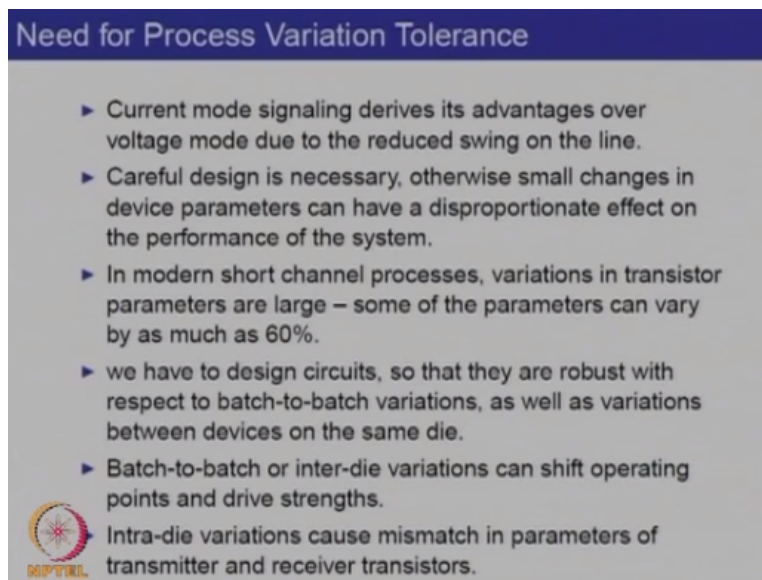
So it is a simpler replacement for the concept of dynamic overriding and now we do not need that NAND, NOR and feedback circuit. So this is actually an attractive solution and it is now being used in addition to the dynamic overdrive solutions that we have seen earlier. There are various pros and cons of this technique versus the dynamic overdrive. We shall see that this receiver is a little harder to design simply because the headroom available is small.

And this line actually swings from VDD to a few millivolts below VDD, therefore essentially amplifying this voltage, this low voltage swing, is not that easy because we do not have much headroom for the transistors connected to this line and this comparator therefore is harder to

resign. In fact, our group at IIT Bombay has worked on this technique also and modified this technique so that the resting voltage at this point is brought closer to be ready by two.


So that efficient comparators can be designed easily here without adding too much to the static power consumption. That combined with a worst-case data sequence technique actually optimizes the behavior of this capacitive peaking and gives very efficient data transmission on long wires.

(Refer Slide Time: 10:20)



Need for Process Variation Tolerance

- ▶ Current mode signaling derives its advantages over voltage mode due to the reduced swing on the line.
- ▶ Careful design is necessary, otherwise small changes in device parameters can have a disproportionate effect on the performance of the system.
- ▶ In modern short channel processes, variations in transistor parameters are large – some of the parameters can vary by as much as 60%.
- ▶ we have to design circuits, so that they are robust with respect to batch-to-batch variations, as well as variations between devices on the same die.
- ▶ Batch-to-batch or inter-die variations can shift operating points and drive strengths.
- ▶ Intra-die variations cause mismatch in parameters of transmitter and receiver transistors.

 NPTEL

However, all these techniques have brought in efficiency by reducing the swing on the line. This means that our design has to be very careful otherwise if there are small changes in device parameters which will always happen, can have a disproportionate effect on the performance of the system.

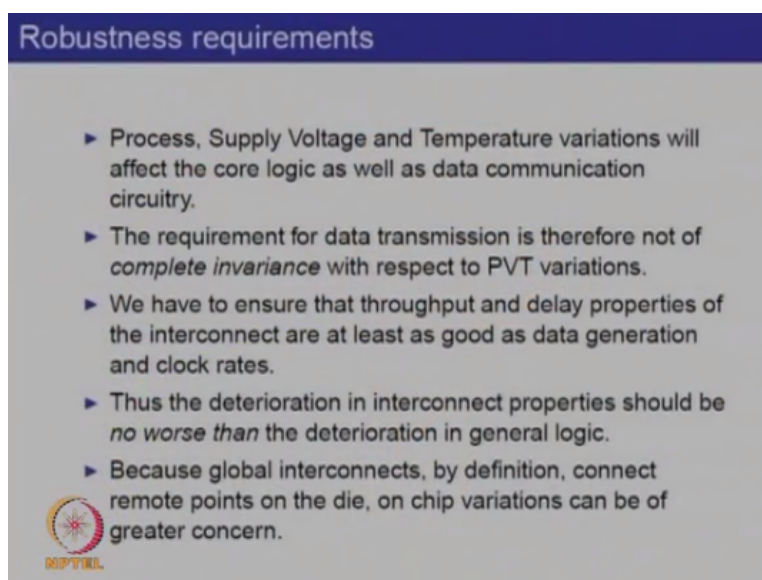
In the voltage mode this signal is swinging all the way from ground to be ready and therefore small changes in V_T , etc., of the transistors do not have such a big effect. However, in modern short channel processes variations in transistor parameters are large, sum of the parameters can vary by as much as 60%. And therefore we have to design circuits so that their robust with respect to batch to batch variations as well as variations between devices on the same die.

What will these variations do, they can in fact change the operating points and the strengths of the driver connected. Therefore, we have to design our techniques such that they are practical and robust with respect to such expected variations. Intra-die variations are also important because after all we are talking of long wires, therefore it stands to reason that the transmitter and receiver are in different parts of the chip which are separated by a longest distance.

As a result, the transistor parameters at the transmitter end and those at the receiver end will not be identical. So, there are two kinds of variations which worry us and therefore the design of interconnector where design must take this into account, one of these, is a batch to batch variation, that means if for one run the scheme works if we are not careful in our design in another run where the values of V_T and mobility and so on are different for the NP channel transistors.


The operating point may shift and because your swings are extremely small it is not guarantee that your scheme will work as well anymore. This is one problem. The other problems are that the transmitter and receiver are in different parts of the chip and there may be mismatch on the same chip in the same run and this mismatch can cause malfunctions. So therefore we need to have a style of design which takes these variations into account and permits our circuits to keep working in spite of these dynamic variations. So what are our robustness requirements.

(Refer Slide Time: 13:29)



Robustness requirements

- ▶ Process, Supply Voltage and Temperature variations will affect the core logic as well as data communication circuitry.
- ▶ The requirement for data transmission is therefore not of *complete invariance* with respect to PVT variations.
- ▶ We have to ensure that throughput and delay properties of the interconnect are at least as good as data generation and clock rates.
- ▶ Thus the deterioration in interconnect properties should be *no worse than* the deterioration in general logic.
- ▶ Because global interconnects, by definition, connect remote points on the die, on chip variations can be of greater concern.



We are saying that the process supply, voltage and temperature variations popularly known as PVT variations will affect the core logic as well as data communication circuit, it is not only the interconnect which will slow down in fact the rate at which we generate the data will also change. Therefore, the requirement for data transmission is not complete invariance with respect to be PVT variation that is not our robustness requirement.

We just have to ensure that throughput and delay properties of the interconnect are at least as good as the data generations and clock rates. If we land in a slow version of the transistor parameters, then the data generation and the clock that it can support will also come down. What we have to ensure is that the deterioration in interconnect properties should be no worse than the deterioration in general logic.

Because of global interconnects by definition, these connect remote points of the die and on-chip variations must also be accounted for.

(Refer Slide Time: 14:44)

Effect of common mode voltage mismatch


Ideal

Transmitter Receiver

Misaligned

V_{cm-Rx}

- ▶ In case of ideal match, small fluctuations in line voltage are converted to rail to rail swing by the receiver.
- ▶ If, however, the mismatch is large, the small swing on the line may be completely ignored by the receiver.
- ▶ It is important, therefore, that the amount of swing on the line is much more than the mismatch in common mode voltages.
- ▶ But high swing will cause power dissipation.
- ▶ It is better to have smart bias circuits, which will reduce mismatch and the need for a large swing.



Let me just give a simple example of how this local variation can be of concern. The batch to batch variation in transistor parameter variations with runs is easy enough to understand because after all a slow circuit when VTs are high, mobilities are low will not keep up to the speed requirements. However, it is a much more softer nuance to understand why local variations

should worry us so much and why is there to worry more in case of low swing techniques that we have been describing.

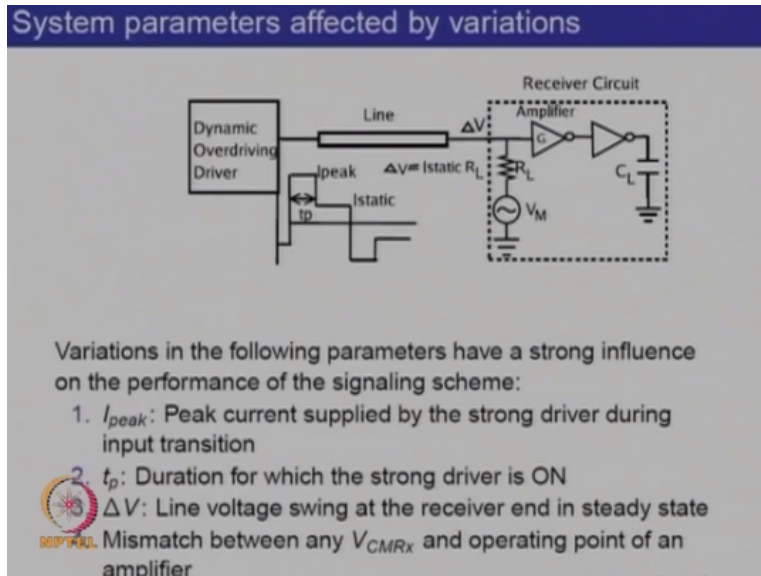
Consider this case and here, at the receiver, we are trying to resolve this small swing around a common mode voltage into a full-blown ground to be ready swing which will then be used by the receiver. Now, if the switching threshold of the receiver is exactly aligned with the common mode voltage as driven by the transmitter then we do not have a problem. Designing an amplifier which will take this small swing and amplify it to a full rail to rail swing with a comparator is not very difficult.

But let us say that because of parametric variations and mismatch between the transmitter and receiver, the common mode voltage at the transmitter is slightly below, in fact, it is so low compared to the common mode voltage at the receiver that even at the highest level of the swing at the transmitter, it remains below the resolution threshold of the receiver comparator. As a result, while we have a healthy swing around this common mode voltage at the transmitter.

This entire signal whether high or low is below the threshold of zero at the receiver and as a result, the receiver will be stuck in zero. Exactly, the same thing happens if the common mode voltage is much higher than this one level and because this swing is rather low, relatively small mismatches between the transmitter and receiver can lead to problems. There is no problem for the rail to rail swing of the buffer inserted technique

So while we have come up with a better technique which is energy-efficient it brings in its own requirement of robustness of design which we must be aware.

(Refer Slide Time: 17:41)



So to analyze this, we have essentially a somewhat idealized model of either dynamic overdriving or the capacity drive case in which we apply an enhanced drive for a short time, this could be because of capacitive coupling or because of dynamic overdrive Nand, NOR combination and then maintain the line at a low drive. Similarly, when there is a 1 to 0 transition, we give it a large boost for a short time and then maintain it at low drive.


At the receiver end, we have a reference voltage, V_M , which is the switching voltage of this inverter at the receiver and a terminating resistance. This amplifier has a high gain provided this line is kept at V_M and finally this drives a buffer which drives a load capacitance. There are various parameters of the transmitter and receiver, which will affect the robustness of this solution.

The value of I_{peak} is the peak current supported by the strong driver during input transition that is the I_P value. T_P is the duration for which the strong driver is on and ΔV is the line voltage swing at the receiver end, so as a result of this drive shape in current at the transmitter, we shall get a ΔV at the receiver and after it has passed through the low pass line, and finally the mismatch between the common mode voltage seen at the receiver at the operating point of this transmitter.

So these are the various parameters which will affect the robustness of our design. The scheme with feedback which we had described which has essentially a feedback inverter which stops the drive when the line at the transmitter end reaches a one or zero, has a particular problem. The reason for that is that this sensor inverter which turns off the drive is at the transmitter end and this inverter which is very similar.

It transfers the low swing voltage to rail to rail swing is at the receiver. These two might not match, in that case, if the mismatches to large, we may have a problem.

(Refer Slide Time: 20:26)



Let us look at this case and let us say that the common mode voltage at the transmitter end and the common mode voltage at the receiver end have a certain mismatch. Now what happens is that because the receiver tries to maintain the line at this voltage, the sensing at the transmitter goes completely a rille. Consider the case here, let us say that the line was resting at one and we are trying to pull it down to 0.

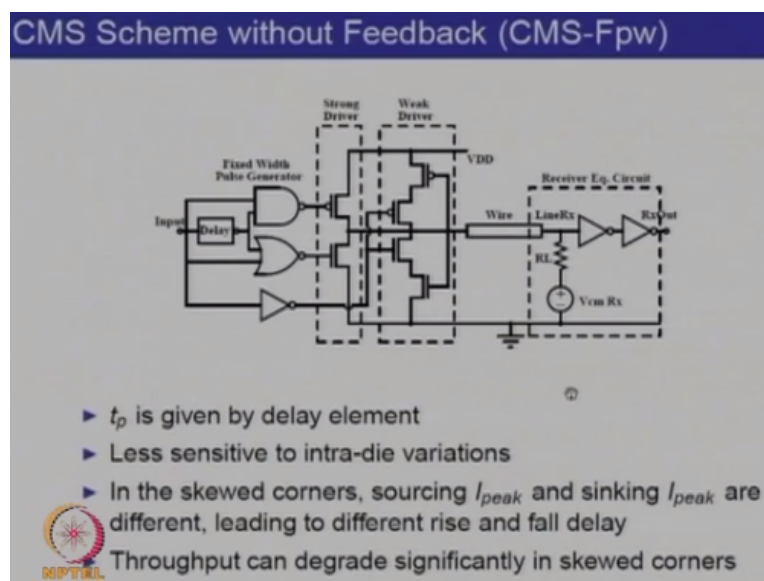
As we pull it down to 0, the feedback inverter after the transition is complete turns the strong driver off. However, the common mode voltage at which this turning off occurs, is much lower than the receiver common mode voltage. As a result, the voltage here goes to a voltage much lower during than the receiver common mode voltage.

And therefore as soon as a strong driver turns off, the receiver start charging this line up because the receiver is trying to keep it $V_{CM\ RX}$, as soon as it reaches a certain voltage, the feedback inverter at the transmitter thinks that this voltage is too high and turns a strong driver on again. Notice, the input has no transition at all, however, because of this feedback there is this back and forth between the receiver and the transmitter, when the strong driver turns off, the receiver start charging it to its common mode voltage.

This common mode voltage is too high for the feedback inverter at the transmitter and that again turns the transmitter on which then takes it down to the low value which it sees as the appropriate low voltage. When that voltage is reached, the strong driver turns off. As soon as a strong driver turns off, the receiver tries to take the voltage to its common mode voltage.

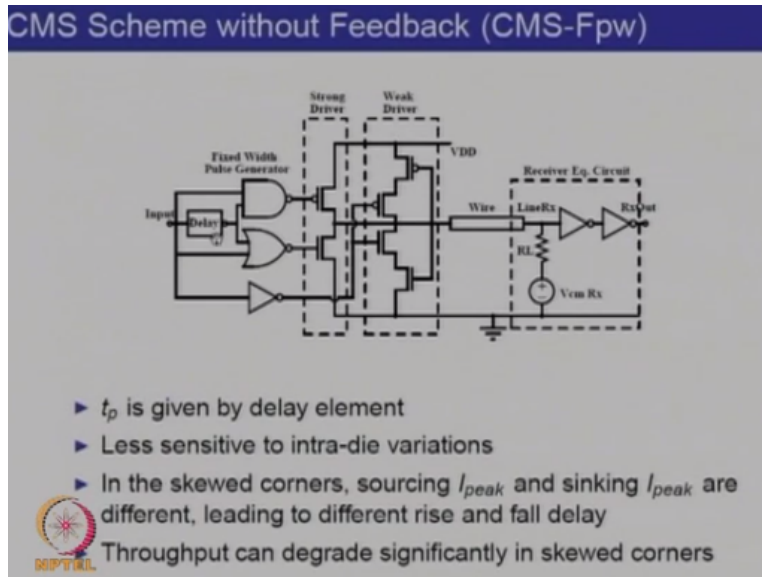
And because these two are not the same, you get an intermittent turning on and off of the strong driver, which reduces the average swing of the line which can cause robustness problems.

(Refer Slide Time: 23:28)



These problems are removed by a technique which we have advocated and which is affixed pulse width driver and this gets rid of the feedback noticed that the circuit is not feedback and the drive is now provided for a fixed delay that means the strong driver is not turned off on sensing the line, we are not sensing the line any more. The strong driver is turned off after a delay which is process dependent.

(Refer Slide Time: 24:07)



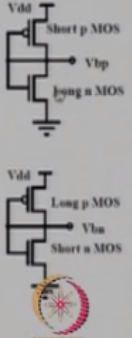
We would like to minimize this process dependence and this we have done in some work done at IIT Bombay by developing bias which actually senses the current process. This whole thing depends on a short channel pMOS and long channel nMOS or a long channel pMOS and a short channel nMOS. This system relies on the following fact that the short channel transistors have a much higher variation than the longer transistor.

So, consider this because this short nMOS will vary with the process where the long pMOS will not at least to 0 level, therefore it sends more or less a process independent current through this diode connected nMOS and as a result of this output tracks V_{TN} . If V_{TN} is higher, this voltage also becomes higher and that corrects the bias for transistor parameter variations. Exactly the same thing happens for this pMOS.


(Refer Slide Time: 25:27)

Minimizing Process Dependence

To minimize process dependence, we need smart bias circuits which sense the process corner and adjust the bias to compensate for variations.



- ▶ Long Channel transistors show relatively less variation with process compared to Short Channel transistors in the same process.
- ▶ We can make use of this difference to design a bias generator which senses the process corner and tries to increase the transistor current in the slow corners and to decrease it in the fast corners.
- ▶ Simple bias generators using inverters with input and output shorted and which use this feature are shown here.




So using such auto bias circuits, we have developed a system in which the drive through this is in fact corrected for process variation and also it does not use any feedback. By combining these two techniques in fact we have been able to come up with a very robust technique.

(Refer Slide Time: 25:54)

Effect of Intra-die Process Variations

Mismatch in VM of inverter can be up to 40 mV. ¹. For VM-mismatch of 40 mV

CMS system	Percentage Degradation	
	Delay	Throughput
CMS-Fb	25	33
CMS-Fpw	10	14
CMS-Bias	4	9.5



We have simulated these techniques and we find that the degradation for the scheme that we have suggested, these are the three schemes, this is the current mode scheme with feedback. This is a current mode scheme with fixed pulse width and this is the current mode scheme with the smart bias which I have just now described.


We find that the degradation and the mismatch is much reduced in case of delay the percentage degradation can be as much as 25% for the feedback is, 10% for a fixed pulse width case and only 4% when we combine fixed pulse width with a smart bias case. Similarly, the throughput degrades by about 33% in case of feedback about 14% in case of fixed pulse width but only 9.5% when we use this technique which combines the fixed pulse width with a smart bias generation.

So by using essentially good VLSI design techniques, it is possible to meet the robustness requirements so that the current mode solution can in fact become practical.

(Refer Slide Time: 27:20)

Effect of Inter-die Process Variations			
Signaling System/ Logic Circuit	Percentage Degradation		
	SS	SNFP	FNSP
CMS-Fb	17.5	5.7	2.9
CMS-Fpw	32	33.6	34.9
CMS-Bias	18.75	8.2	7.14
Voltage Mode	27	< 1	2.8
Ring Oscillator Freq	23	2.88	3

- ▶ Interconnects with CMS-Fpw scheme become the bottleneck in overall performance of the chip in skewed corners
- ▶ Degradation in the throughput of the proposed scheme in the skewed corners is around 7% which is less than that in CMS-Fpw scheme



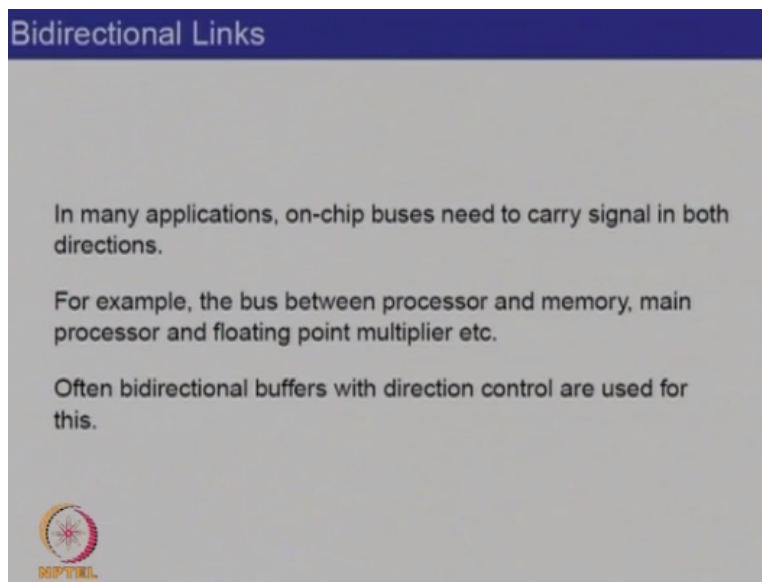
You can see that the ring oscillator frequency degrades by about the same order, 23% here and the one with bias is much smaller than the ring oscillator frequency due to process variation, essentially what it means is that the ring oscillator frequency will determine the digital rate of generation of data. If this degrades by 23% then as long as we degrade by less than 23%, everything is fine.

And we noticed that the voltage mode fails to meet this requirement, so does the current mode scheme with a fixed pulse width, whereas the scheme that we have suggested and the scheme with feedback, they can meet the requirement for the process variation, noticed this is not local variation, this process variation. However, the feedback circuit is not so graceful as we had just

seen in case of on-chip variation between transmitter and receiver, this table is for global process variations and for that the feedback scheme is not very good.

For local variations, the fixed pulse width scheme is okay but this is not very good but the scheme that we have suggested which is the current node scheme with smart bias that meets the requirement in both cases. I think we will skip this to go to a bidirectional link. Now notice that bidirectional links are very important, we have talked about this earlier and we need to have a scheme which will permit bidirectional transmission of data.

(Refer Slide Time: 29:22)

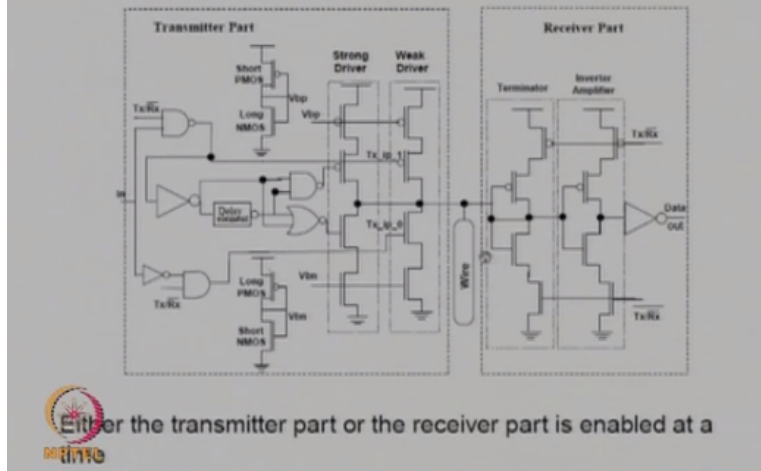


This can be done in voltage buffer mode by using back-to-back connected tri-state buffers, where exactly one of these is activated. However, as we had seen that this leads to problems, first of all the delay of a bidirectional repeater is more than that of unidirectional buffer because of the loading and a direction control signal is required by each repeaters and if there is a buss.

Then the direction control signal is loaded by a large number of such transistors in parallel and the buffers carrying the direction control signal are heavily loaded and they consume additional power. So, we need a repeater less signalling scheme.

(Refer Slide Time: 30:21)


Proposed Current-Mode Transceiver



This can be done in the current mode bidirectional link. Essentially, we have across the line transmitter as well as a receiver connected to the end of the line, noticed that nothing needs to be connected in the middle of the wire, at either end of the wire, we have a transmitter and a receiver. Obviously, the transmitter and receiver must have the information of who is to transmit and who is to receive, in any bidirectional scheme that is the case.

Since that information is available, we can use the to turn on either the transmitter or the receiver at either end of the wire and thus achieve bidirectional transmission quite easily. This is possible because there is not actively circuitry in the middle of the wire.

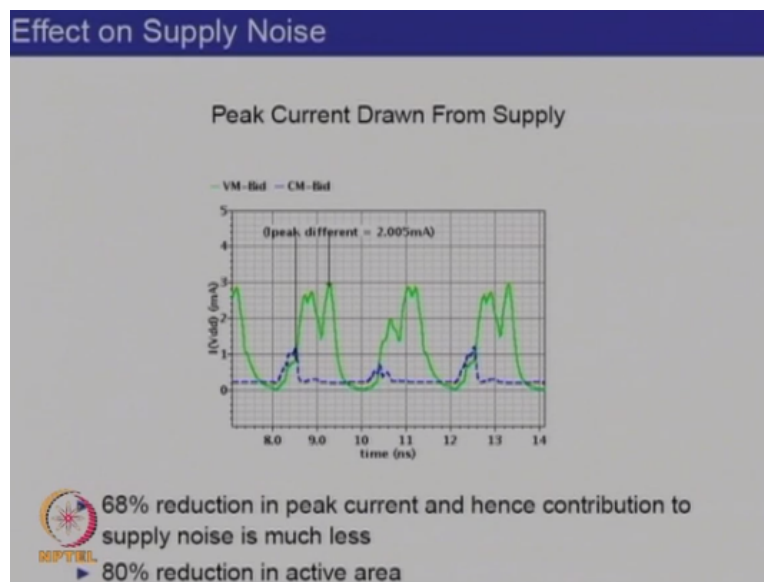
(Refer Slide Time: 31:13)



As a result, we see that as of here, we plot the regions in which the current more bidirectional drivers consume less power than the voltage mode bidirectional power. The plot is data rate versus line length and for all combinations where the current mode consumes less power, we have this shaded region that mean for this line length and this data rate and beyond the current mode will consume less power.

And you will notice that most of the useful range is covered by that region for example line lengths greater than 2 mm and data rates which are say a few hundred megabits per second, for all such combinations whatever the line length, whatever the data rate the current mode consumes less power compared to the voltage mode and this is a very important point because these are in fact the robust designs that I have described just a little while ago.

(Refer Slide Time: 32:31)



There is one additional advantage that current mode has and this is the power drawn from the supply. The voltage mode buffers draw huge amount of power from the supply and as a result cause spikes on the supply voltage. This is a source of additional noise to the entire system, because current mode draws less power from the power supply. The spikes that it generates on the supply voltage are much smaller.

And as a result, the noise level injected is much smaller, consequently a current board interconnect runs quieter than a voltage mode interconnect. So, therefore we are talking of a 60%

reduction in peak current and hence contribution to supply noise is smaller and 80% reduction in activity area. Therefore, for bidirectional data transmission, current mode is indeed extremely attractive.

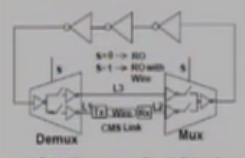
While many of these ideas therefore sound quite attractive, we would like to show that they work in silicon under practical cases and there is one problem which presents itself when we contemplate doing this and that problem is the following. The overall delays of wires of any practical length are quite small. These are of the order of a nanosecond or less, measuring such delays is not an easy thing.

And if you couple it through a pad and bring it out to external instruments which might present loads of the order of picofarad, it extremely difficult to demonstrate which of the techniques that we are talking about is in fact faster. Therefore, we need to develop test circuits which will allow us to compare the performance of various suggested schemes on chip itself.

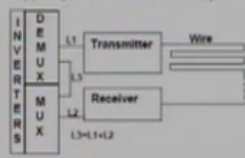
And the output of these test chips, test circuits should be such that it is either DC or some low-frequency which can be brought out from the chip easily and can be measured using inexpensive instruments which illustrate this by only a few representative circuits.

(Refer Slide Time: 35:02)

Time to Frequency Conversion




(a) Delay Measurement Circuit: Principle



(b) Delay Measurement with CMS Link: Floorplan

- ▶ Transmission gates were used to implement switches.
- ▶ Multiplexer(demultiplexer) are designed so that delays for both possible paths through the mux/demux pair are the same.
- ▶ The floor plan of the circuit is such that the beginning and the end of the long interconnect are close to each other.
- ▶ Therefore when the short path L3 is chosen, the total delay corresponds to the delay in inverters, mux/demux etc.



Consider this suggestion what we have here is the multiplex circuit and we have a choice the mux and demux ensures that they the ring oscillator here, and this ring oscillator will oscillate at a particular frequency. The frequency of oscillation will depend on the total loop delay, now what we can do is that in one of the arms of this mux/demux pair, we can put the transmitter wire and the receiver of the suggested scheme.

The other is a dead short. Now we measure the delay using the dead short first. This measure the delay through these inverters, this mux and this demux. Apart from that the delay of this short wire which provides the shorting part L3, is also included. The other option in the other position of the mux/demux what we have are these approach lines L1 and L2 and apart from L1 and L2, we had the transmitter, the long wire over which we are measuring the data rate and power etc., and receiver.

It is laid out in such a way that the transmitter and the wire loops back to the same region, this is a shorting wire of length L3, L1 is the length of the approach wire to the transmitter and L2 is approach wire from the receiver. When the mux and demux is in this position, we ensure that $L1 + L2$ is the same as L3. As a result, the total delay which is common for the two cases includes the entire delay of this path, that is common and because L3 is equal to $L1 + L2$ this delay is also the same.

Therefore, if you take the difference of the two delays, then it measures accurately the delay through, so essentially what we have done is that we have converted the measurement of very short delays to a measurement of oscillation frequency. We have this ring oscillator, we put the mux/demux in the L3 position and because of the much smaller delay through L3, this ring oscillator oscillates at a much higher frequency.

Using a low-frequency signal, we now switch this mux/demux to take the lower path, when it takes the lower part, then the delay of the transmitter and receiver is included in the path and therefore the ring oscillator still oscillates but at a much lower frequency. These 2 frequencies are indicative of the delays in the two cases and if you take the difference of the two delays in that

case all the common delays cancel out, leaving only the delay of the path which we want to measure.

(Refer Slide Time: 38:39)

(a) Delay Measurement Circuit: Principle

(b) Delay Measurement with CMS Link: Floorplan

- ▶ We first measure the frequency of oscillation choosing the short wire path between the demux and mux.
- ▶ This gives the delay of the measurement circuit except for the system under test.
- ▶ We now select the interconnect system whose delay we want to measure and find the frequency again.

$$Delay = 0.5 \left\{ \frac{1}{f_{RO}} - \frac{1}{f_{system}} \right\}$$

MPTEL

So therefore by measuring the frequency in the two cases which is as simple as this formula so therefore the net delay of the transmitter plus wire plus receiver is given simply by this which is $1/f_{ro}$ which is the ring oscillator $-1/f$ system. So system is when it is shorted and this is the ring oscillator with the transmitter receiver in. Notice that this system is much higher.

(Refer Slide Time: 39:13)

Time to Frequency Conversion: Accuracy

To assess the accuracy of the scheme, we simulated the whole circuit, for different line lengths up to 14 mm in a 180 nm process.

- ▶ The delay through the interconnect scheme was noted from the simulation results. We call this the "Simulated Delay"
- ▶ The delay was also calculated by the formula:

$$0.5 \left\{ \frac{1}{f_{RO}} - \frac{1}{f_{system}} \right\}$$

We call this the "Calculated Delay"

- ▶ These results were tabulated to assess the expected accuracy from this test scheme.

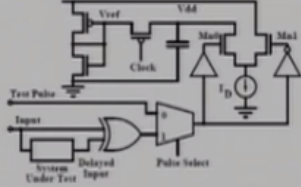
MPTEL

This was assessed by doing first simulations, in simulations of course, we can see the delay and we can also see the frequencies, so when we simulate the circuit, we look at the frequencies and

compute the delay using this formula and compare it to the delay which we see from assimilation transient simulation case and we find that the percentage error is very, very small.

(Refer Slide Time: 39:44)


Time to Voltage Conversion



- ▶ Capacitor C is pre-charged to peak value during the negative phase of the clock.
- ▶ It is then discharged for a time equal to the delay through the system.

$Delay = \frac{C\Delta V}{I} = k\Delta V$

Value of k is found experimentally using a calibration pulse of known duration.

 NIPTRIL

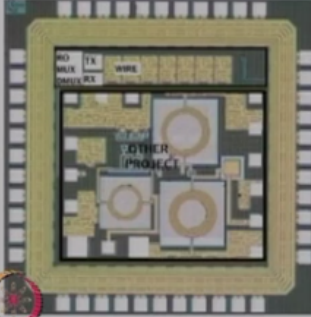
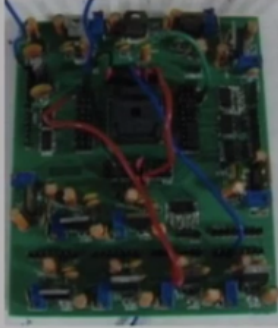
Similarly, by using a time to voltage conversion in which the application at the transmitter of the digital bit starts the charging of a capacitor and the arrival of the bit stops the charging through a current source, we can convert this delay to a DC voltage and this voltage can be read from the outside circuits. So this essentially points out that there are circuits which are possible. These circuits can be put on the same chip as the interconnect.


And by using these circuits, we can actually make very small differences in delay and power visible through signals like frequency and DC voltages which are very easy to measure often. So we actually implemented these various schemes on silicon and used these measurement circuits on chip. The high frequency of the ring oscillator was in fact scaled-down by a factor of 32 to 64 to come down to a level of frequencies, where we can measure it easily using inexpensive frequency meter.

(Refer Slide Time: 41:16)

Current-Mode Signaling Test Chip

- ▶ 1.5mm × 1.5mm chip fabricated in 180nm MM/RF process
- ▶ 44-pin die packaged in QFN56 package








So this is a chip that we actually made, this is this is a photograph not a diagram and the transmitter receiver all the wires and all the circuit are here. We built an external test jig which provides all the voltage, control signals, trial singles and so on and the whole die was packaged in a 44-pin QFN package.

(Refer Slide Time: 41:50)


Measurement Results

(Frequency measured using a 6-digit frequency counter)

Signaling Scheme	Delay (ns)	Energy (pJ)	EDP (pJ×ns)	Measured at Data Rate (Mbps)
Voltage Mode	1.191	4.54	5.328	371
CMS-Fb	1.006	1.52	1.52	400
CMS-Bias	0.938	0.851	0.799	621

The proposed circuit offers 22% improvement in delay and 85% improvement in EDP over voltage-mode scheme.

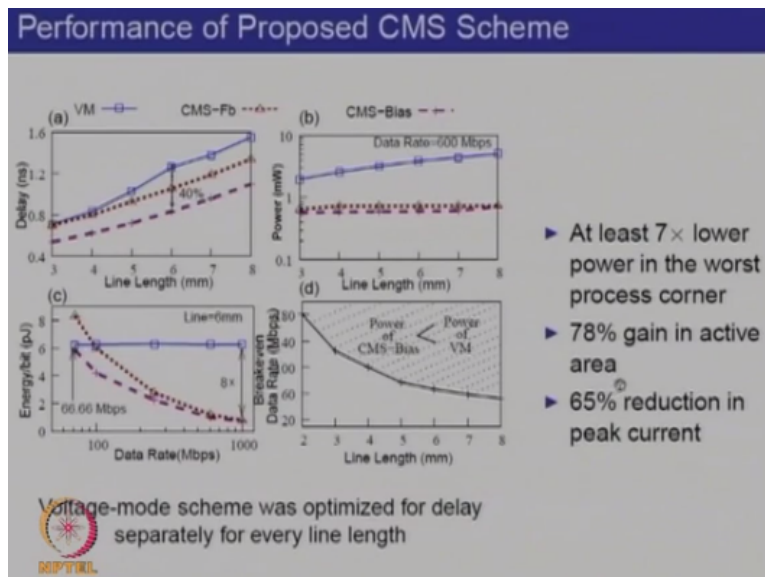


Using this, we measured the actual delay, the power and the energy used by the three schemes and looking at the data rate and the measurements, we can see that the proposed circuit which is the CMS bias, remember this is the circuit, which counters both the batch to batch process variation as well as the transmitter to receiver on chip parametric variations.

So, using CMS bias, we can see that we get about 22% improvement in delay and as much as 85% improvement in the energy delay product over the voltage mode scheme. This establishes the fact that this scheme is much superior to the widely used buffer insertion and at the same time is practical against process variations and on chip variation. Therefore, it is possible that in future circuits, interconnect aware design will make use of circuits of this class.

Remember, this has the advantage that the general design style of the digital circuits which constitutes most of the complexity of the VLSI design, does not change at all. The bits are still rail to rail. They are the conventional voltage mode bits. It is only the transmission of these signals which is now being reduced to current mode.

(Refer Slide Time: 43:40)

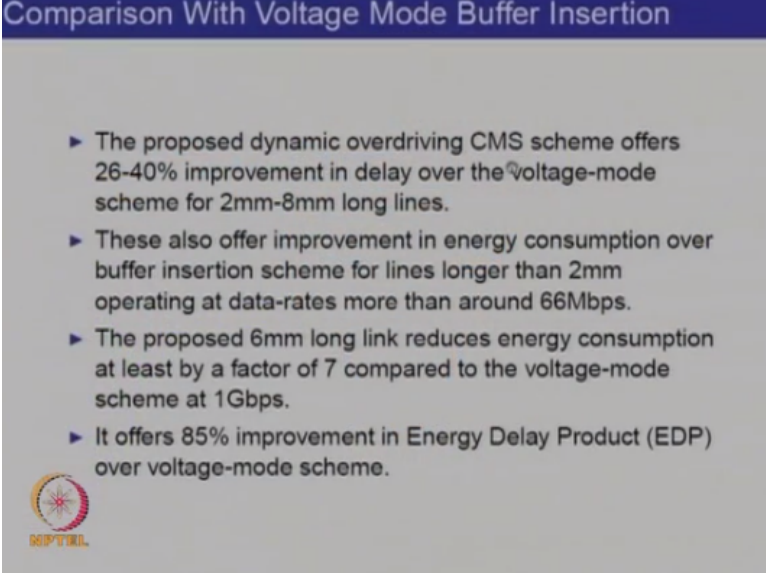


So essentially just to summarize the behavior, there is at least seven times lower power in the worst-case process corner, 78% gain in active area. this is the area on silicon and 65% reduction in the peak current, which then translates to generation of lower supply noise. Another factor which must be pointed out is that the voltage inserted buffers had to be redesigned for every wire length.

If the wire length changes that the placement and sizing of the buffer inverters has to be changed. On the other hand, the current mode signal is very robust, it is designed once and for all and remains unchanged for all wire lengths.


This is an advantage because then you can put it in a library and then not worry, whatever the length of the wire the same component is pulled out and then used.

(Refer Slide Time: 44:47)



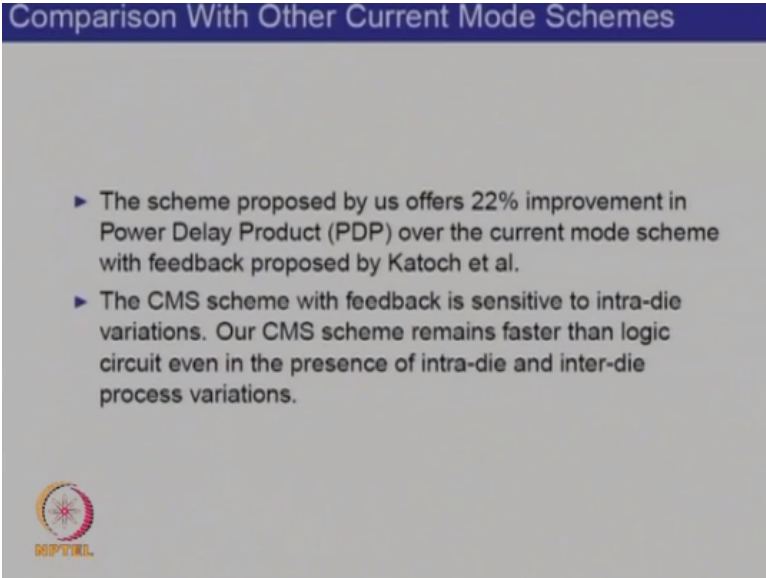
Comparison With Voltage Mode Buffer Insertion

- ▶ The proposed dynamic overdriving CMS scheme offers 26-40% improvement in delay over the voltage-mode scheme for 2mm-8mm long lines.
- ▶ These also offer improvement in energy consumption over buffer insertion scheme for lines longer than 2mm operating at data-rates more than around 66Mbps.
- ▶ The proposed 6mm long link reduces energy consumption at least by a factor of 7 compared to the voltage-mode scheme at 1Gbps.
- ▶ It offers 85% improvement in Energy Delay Product (EDP) over voltage-mode scheme.




The proposed dynamic overriding CMS scheme and by the proposed scheme, I mean the one that we have proposed which corrects for robustness from batch to batch and on chip variation, using a smart bias circuit offers 26 to 40% improvement in delay for 2 mm to 8 mm long lines but also compared to other schemes, it offers a substantial improvement in the energy delay product.

(Refer Slide Time: 45:22)



Comparison With Other Current Mode Schemes

- ▶ The scheme proposed by us offers 22% improvement in Power Delay Product (PDP) over the current mode scheme with feedback proposed by Katoch et al.
- ▶ The CMS scheme with feedback is sensitive to intra-die variations. Our CMS scheme remains faster than logic circuit even in the presence of intra-die and inter-die process variations.

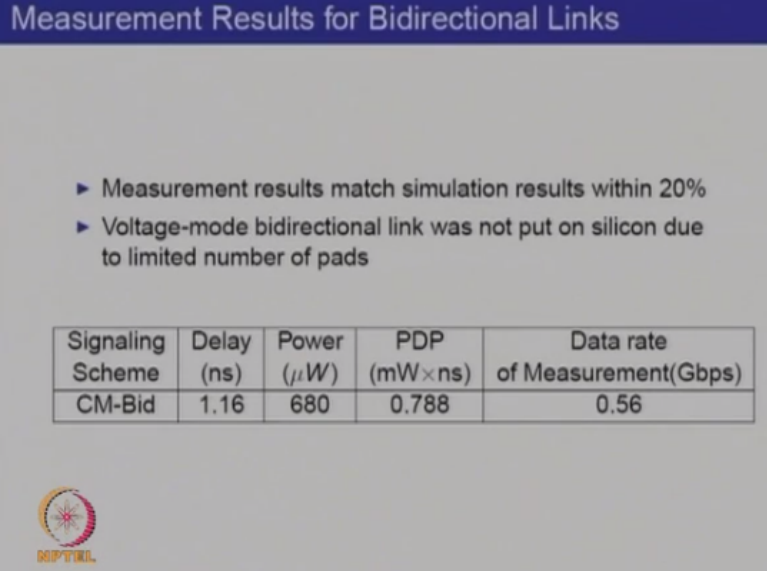


Compared to other current mode schemes like the one with feedback, there is 22% improvement in power delay product which is much smaller of course with voltage buffer. All current mode

schemes perform much better than voltage buffer schemes. So 22% improvement is over the other current mode scheme. A factor of sudden improvement over the voltage mode scheme and the CMS team with feedback is sensitive to intra-die variation.

Whereas, the current node scheme with smart buyers remains faster than logic circuits even in the presence of intra-die and inter-die process variations.


(Refer Slide Time: 46:05)



Measurement Results for Bidirectional Links

- ▶ Measurement results match simulation results within 20%
- ▶ Voltage-mode bidirectional link was not put on silicon due to limited number of pads

Signaling Scheme	Delay (ns)	Power (μW)	PDP ($mW \times ns$)	Data rate of Measurement (Gbps)
CM-Bid	1.16	680	0.788	0.56



We have also made measurements with bidirectional links and we noticed that current mode bidirectional links offer very small delays and small consumption of power compared to traditional voltage mode buffering scheme and in this because simulation showed that the performances are not even comparable, we did not actually compare these two on silicon.

(Refer Slide Time: 46:39)

Simulation with Matched Model Parameters

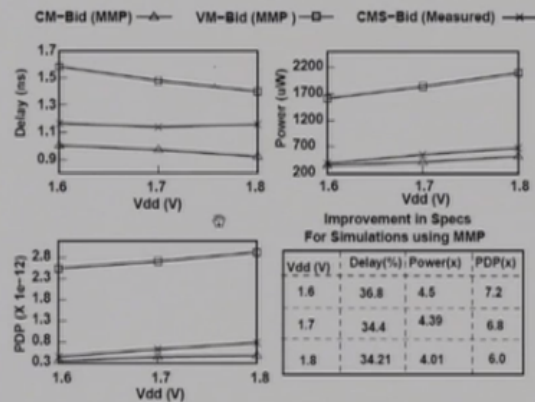
Parameters	TT	Measured	MMP	% Match
I_{satn} (mA)	6.23	6.44	6.43	99.8
I_{satp} (mA)	2.40	2.22	2.28	97.3
V_{tn} (mV)	501	510	506	99.2
V_{tp} (mV)	494	493	499	98.8
I_{offn} (pA)	75	170	120	82.4
I_{offp} (pA)	80	48	58	80.5
$I_{dsn}/I_{dsp}@V_{gs}$		$I_{ds} - V_{gs}$ points		
$I_{dsn}@0.9$ (μA)	66.6	65	66.4	97.85
$I_{dsp}@0.9$ (μA)	76.2	70	67.5	96.45
$I_{dsn}@1.2$ (μA)	154.4	150	145	96.67
$I_{dsp}@1.2$ (μA)	191	170	172	98.82
$I_{dsn}@1.8$ (μA)	347	330	317	96
$I_{dsp}@1.8$ (μA)	491	440	452	97.27



We actually did an extraction of transistor parameters from some extra patterns that we had put on the chip.

(Refer Slide Time: 46:44)


Measurement Results and Simulation Results with MMP



And then we can show that if we use the transistor parameters which occurred on the exact run on which we have made measurements.

(Refer Slide Time: 46:53)

Conclusion

- ▶ Global interconnects form a major bottleneck for performance of digital system at scaled down technology.
 - ▶ Use of current mode signaling is promising to remove this bottleneck.
 - ▶ Through simulation, circuit fabrication and actual measurements, we have demonstrated that current mode signaling has overwhelming advantages over the currently used voltage mode buffer insertion schemes.
 - ▶ We have demonstrated that the particular configuration suggested by us for a current mode scheme is superior to other current mode schemes.
 - ▶ Our scheme is robust with respect to batch to batch parametric variations and to on chip parametric variation.
-  Therefore we assert that it is a practical option for use in modern systems for implementing both unidirectional and bidirectional data links.

Then, we can reproduce the results that we measure. So in conclusion, we can say that global interconnects form a major bottleneck for performance of a digital system at scaled-down technology. Use of current mode signalling is promising to remove this bottleneck through simulation circuit fabrication and actual measurements on silicon. We have demonstrated that current mode signalling has overwhelming advantages over the currently used voltage mode buffer insertion schemes.

We have demonstrated that the particular configurations suggested by us for a current mode scheme is superior even to other current mode scheme at this particular configuration has apart from a fixed width overdriving pulse. A biasing scheme which controls the amount of current dumped by the overdriving in a process independent and variation independent way. Our scheme is robust with respect to batch to batch parameter variations and to on-chip parametric variation.

And therefore it is a practical option for use in modern systems for implementing both unidirectional and bidirectional data links. With this we bring this discussion on current mode and voltage aware data links to an end. So, essentially what it means is that the interconnect wires which are not even considered important earlier have become performance limiters and very careful design has to be done.

The widely used methods are running out of power now and fortunately new schemes which combine mixed signal design with VLSI design which can give interconnect aware design and they can continue to boost the performance of integrated circuits, as we scale down the dimensions at least for the foreseeable future and it is the use of these techniques which will result in interconnect aware designs of tomorrow.

We will bring a discussion of interconnect aware design to a close with this lecture.