Low-power CMOS clock drivers

Mircea R. Stan, Wayne P. Burleson
mstan@risky.ecs.umass.edu, burleson@galois.ecs.umass.edu
http://www.ecs.umass.edu/ece/vspgroup/index.html

Abstract— The clock tree of modern synchronous VLSI circuits can consume as much as 50% of their entire power budget. Different methods of decreasing clock power dissipation have been proposed based on low-voltage swings, double-edge triggered flip-flops, gated clocks, etc.

In this paper we propose two types of full-swing low-power CMOS clock drivers. Both are based on a stepped charging and discharging of the clock tree capacitance in order to achieve up to 50% power savings. The first CMOS driver targets single-phase clocking schemes and is based on a quantized adiabatic operation that uses two power supply voltages (Vdd and $Vdd/2$, $Vdd/2$ can be replaced by a tank capacitor). The second CMOS driver proposed targets dual-phase clocking schemes and achieves low-power operation by charge reuse.

The proposed circuits are more than twice larger and slightly slower than standard inverter-chain clock drivers. In this way the circuits present the designer with the usual trade-offs area, speed vs. power dissipation. The theoretical power savings of 50% compared with the inverter-chain driver are smaller for actual circuits (around 25-30% depending on the capacitive load).

The two-step charging and discharging used by both drivers proposed in this paper introduce small discontinuities on the signal edges and for this reason they are not suitable for clocking schemes that are extremely critical to clock transition times.

The proposed low-power CMOS drivers can be used for off-chip or on-chip clock transmission for highly pipelined circuits, systolic arrays, wafer-scale integration synchronous circuits or as I/O drivers for heavily loaded buses whenever low-power operation with full voltage swing is desired and absolute maximum speed is not a must.

Keywords— CMOS clock driver, adiabatic operation, charge recovery, low-power.

I. Introduction

Synchronous circuits need global clock signals that must be broadcast over the entire chip. As feature sizes decrease and chip dimensions increase the relative importance of the clock tree in the overall power consumption increases. For highly pipelined microprocessors like the DEC Alpha the clock power dissipation can be up to 50% of the entire power budget [4]. Similar clock power consumptions can be expected for fine-grain systolic arrays or synchronous wafer-scale integration circuits.

Paper presented at TAU'95. This work was supported in part by NSF grants MIP-9208267 and CDA-9320325.

The authors are with the Department of Electrical and Computer Engineering, University of Massachusetts at Amherst, MA 01003

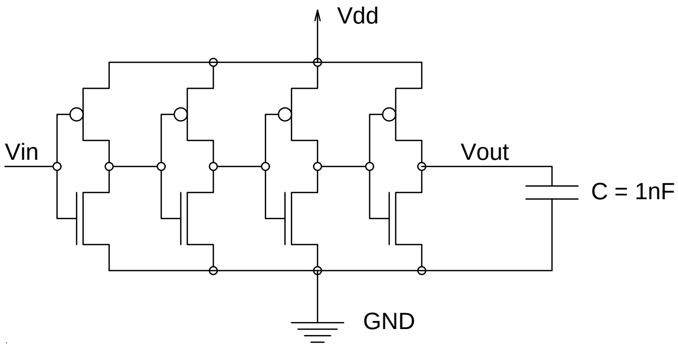

Fig. 1. A chain of inverters for driving the large capacitance of a clock tree (1nF in this case). Each inverter stage is n times larger than the previous one.

The dynamic power dissipated by a CMOS clock node is $P = C \cdot Vdd^2 \cdot f$ where C is the node capacitance, Vdd the power supply voltage and f the frequency. There are many clocking strategies for CMOS circuits [12] which differ in the number of clock phases and loading of the clock node. For single-phase clocking schemes the clock tree can be considered a single electrical node with a large (several nF) capacitance $C_{clk} = C_w + C_g + C_j$ [3] where C_w is the clock tree wiring capacitance, C_g is the gate capacitance of all transistors driven by the clock and C_j is the junction capacitance of the clock driver. For timing purposes because of the ever decreasing feature sizes and increasing die size it becomes necessary to treat the clock tree as a transmission line or as a distributed RC load, but for power dissipation treating the clock node as a lumped capacitance is still a good approximation. Dual-phase clocking schemes have two clock trees with very similar electrical characteristics $C_{phi1} = C_{w1} + C_{g1} + C_{j1}$, $C_{phi2} = C_{w2} + C_{g2} + C_{j2}$.

The usual technique for driving large capacitances is to use a chain of inverters with increasing drive capability as in figure 1. The power dissipated by such a circuit will be dominated by the charging and discharging of C. Many methods have been proposed for decreasing the clock tree power consumption by minimizing one or more of the terms in the power dissipation formula (Vdd, C and f):

- Decreasing Vdd for the entire circuit [2] has a quadratic effect on the overall power dissipation but as a ratio to total power the clock power dissipation remains the same. In order to decrease independently the clock power dissipation the clock voltage swing can be further decreased. Using charge reuse and a clock voltage swing at half Vdd the power dissipated in the clock tree can be decreased by 60-70% [3], [7]. One drawback of using small clock voltage swings is that it

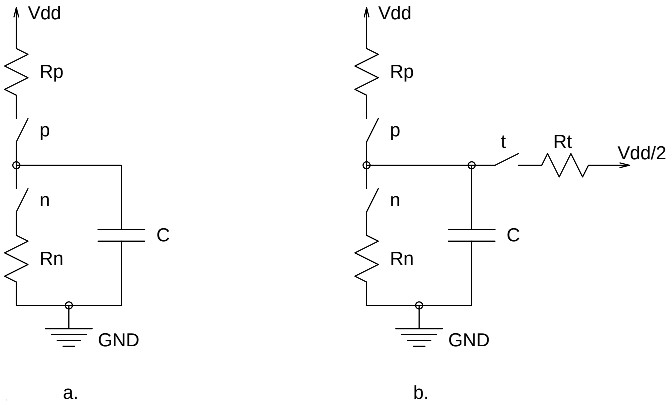

Fig. 2. Switch-level diagrams of the final stage of a conventional CMOS driver (a.) and quantized adiabatic CMOS driver (b.)

requires modified latches and thus an entire redesign of the circuit. Low-power clocks with a full voltage will be necessary for circuits which use a very small Vdd as in [2] and they have the added advantage of not requiring modified latches.

- The frequency f is generally chosen as high as possible for large throughput but this is contrary to low-power requirements. There are techniques that vary the clock frequency depending on the computational load with a corresponding decrease in power consumption. Other techniques use a double edge triggered flip-flop design [5] which uses half the frequency of a single edge triggered design. If the capacitance is kept constant the power dissipated can be decreased in this way by 50%. At the system level another common technique is to broadcast on the PCB a lower frequency (2-4 times lower) clock and recover the high frequency on-chip by using a PLL.
- In order to decrease the clock tree capacitance C one method is to use a latch design with as few clock loads as possible like the TSPC latch [10] or to use clock gating and to decrease the on-chip wiring capacitance with area pad interconnects [11].

All the above clock power minimization methods are generally orthogonal to one another and can be combined for low-power design. In sections II and III we describe two full-swing clocking methods which theoretically decrease by 50% (around 25-30% for real circuits) the clock power dissipation by using a two-step charging and discharging of the large clock tree capacitances.

A. Principles of operation

One common feature of both drivers is that they achieve low-power operation with a full voltage swing and consequently can directly replace standard clock drivers without any modifications to the clocked circuit. The only requirement for the clocked circuit is that it be tolerant to small discontinuities in the clock edges which practically translate into longer transition times.

The single-phase CMOS clock driver described in section II uses the principle of adiabatic capacitance charging and discharging [1] for achieving low-power operation. The continuous version of adiabatic operation requires fundamental circuit design changes but the discrete version based on stepped charging and discharging [9] operates well within the framework of standard CMOS design. The single-phase driver proposed here uses a two-step quantized adiabatic operation for theoretically achieving 50% power savings over the standard inverter chain. Using more than two steps for charging and discharging can theoretically lead to even better power savings but in practice the extra circuit complexity, the number of additional power supplies and the fact that the savings for each additional step get smaller determines the choice of two-step operation as the best compromise between practical design issues and theoretical power savings.

The dual-phase CMOS driver described in section III uses the principle of charge recovery [8] to achieve low-power operation. Although it appears totally different from the adiabatic principle, charge recovery can be also explained in terms of adiabatic charging and discharging. This is confirmed by the fact that the two-phase driver proposed in III has many similarities with the driver proposed in II. One difference and a main advantage is that the dual-phase driver does not use any extra supply voltage.

II. Two-step quantized adiabatic charging and discharging

A simplified switch-level schematic of a standard clock driver final stage is shown in figure 2 a. Only one of the two switches (p or n) can be closed at a time. The clock cycle can be described in two phases (C initially discharged):

- first phase (clock rising edge) - p is closed with n open and C is charged through the resistor Rp to Vdd. The energy drawn from the power supply for charging C is $W = C \cdot Vdd^2$ where exactly half is dissipated in the resistor Rp and the other half is stored on C.
- second phase (clock falling edge) - n is closed and p open to discharge C to GND and the energy stored on C is dissipated in Rn. No extra energy is drawn from the power supply in this phase. The values of Rp and Rn do not influence the power dissipated.

From the above description there is a clear difference between instantaneous power *consumption* and power *dissipation* and although the two are equal on average sometimes it is more convenient to think in terms of one or the other. Power consumption occurs only when current is actually drawn from the power supply while power dissipation appears whenever there is a nonzero voltage across resistors. Power consumption is generally of interest when looking at battery life for portable devices and for dimensioning GND and Vdd pin counts and wire-widths. Power dissipation is important when dimensioning heat removal devices and for assessing possible heat-related IC failures.

There have been attempts to break the "tyranny of $C \cdot Vdd^2$" CMOS power dissipation by using adiabatic capacitance charging and discharging [1]. Adiabatic operation achieves low-power consumption by always keeping the voltage across resistors small. This generally requires that

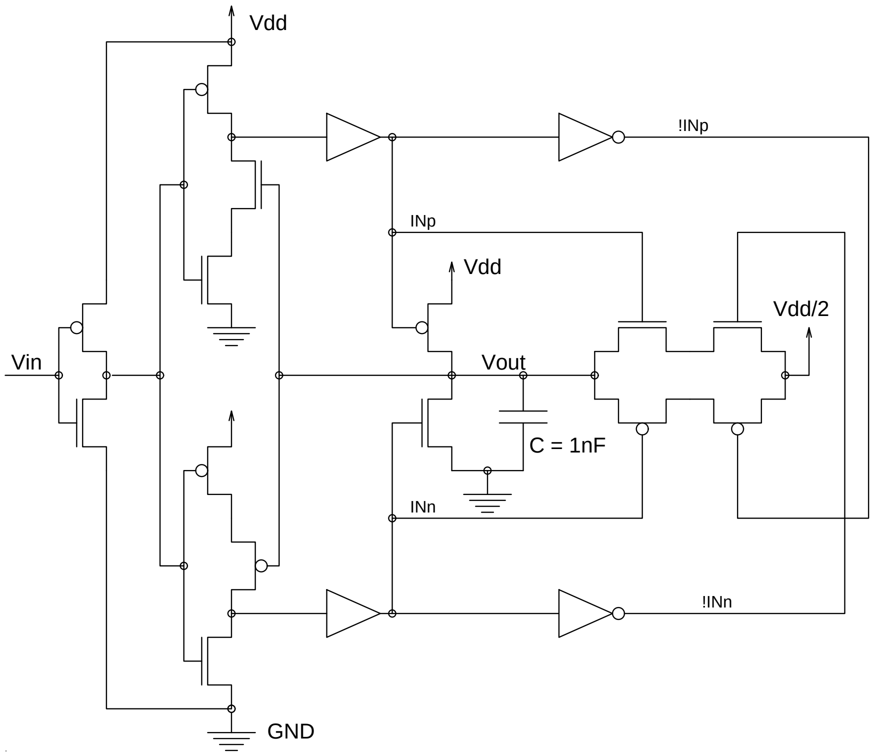

Fig. 3. Schematic of the proposed quantized adiabatic single phase driver.

ramp power supplies be used for charging capacitors and for this reason true adiabatic operation is hard to implement. A practical approximation to continuous adiabatic operation is a quantized version based on stepwise charging and discharging [9]. The first approximation of a continuous ramp is a two-step and we use this theoretical model for designing a practical low-power two-step clock driver.

A clock driver using two-step charging and discharging (see figure 2 b.) can theoretically save 50% power compared with a conventional driver. The energy consumed in one period by a conventional driver will be denoted by $W = C \cdot Vdd^2$. This implicitly assumes that the power consumption is dominated by charging and discharging C. In the proposed step driver circuit there are three switches (p, n and t) and two power supply voltages (Vdd and $Vdd/2$). Two of the switches (n and p) are for statically driving the output to LO and HI while the t switch is used only in a transient manner. Only one of the 3 switches must be closed at a time in order to avoid static power dissipation. Four phases, two static and two transient, will explain how this circuit works (C initially discharged):

- phase one (transient, clock rising edge to $Vdd/2$) - t is closed and in this way C is charged to $Vdd/2$. $W/8$ is dissipated in Rt and another $W/8$ is stored on C ($W = C \cdot Vdd^2$). A total of $W/4$ is drawn in this phase from the $Vdd/2$ power supply.
- phase two (static, clock rising edge to Vdd) - when the output reaches $Vdd/2$ t gets open and p closed such that C is charged from $Vdd/2$ to Vdd. Another $W/8$ is dissipated in Rp while C stores an additional $3W/8$ with the total energy stored on C being $W/8+3W/8 = W/2$ as for the conventional driver. The energy drawn from the Vdd power supply in this second phase is $W/2$.
- phase three (transient, clock falling edge to $Vdd/2$) - t is again closed for discharging C to $Vdd/2$. $W/8$ is dissipated in Rt while $W/4$ is returned to $Vdd/2$. C will have $W/8$ stored at the end of phase three.
- phase four (static, clock falling edge to GND) - t is opened and n closed when $Vout$ reaches $Vdd/2$ and C is discharged to GND while $W/8$ is dissipated in Rn.

From the above simplified analysis it can be seen that only $W/2$ energy is drawn from the Vdd power supply (in the second phase) while the $W/4$ energy drawn from $Vdd/2$ in phase one is actually returned to $Vdd/2$ in phase three. If $Vdd/2$ can both supply and accept current (e.g. rechargeable battery) then the overall energy drawn from the $Vdd/2$ power supply is zero and theoretically the power savings for the step driver are 50% compared with the conventional case. Even better is to replace the $Vdd/2$ power supply with a tank capacitor since the voltage on the tank capacitor will automatically converge to $Vdd/2$. In order to provide a convenient initial condition for the tank capacitor a simple circuit with two Zener diodes in series will suffice. If the voltage Vz of the Zener diodes is slightly larger than $Vdd/2$ there will be no DC current flowing and the midpoint will be guaranteed to be between $Vdd - Vz$ and Vz. For a Vdd = 5V and a Vz = 3V the midpoint will be between 2V and 3V which is good for the correct functioning of the step driver.

There are many challenges in actually implementing a circuit using the above ideas. The power savings are likely

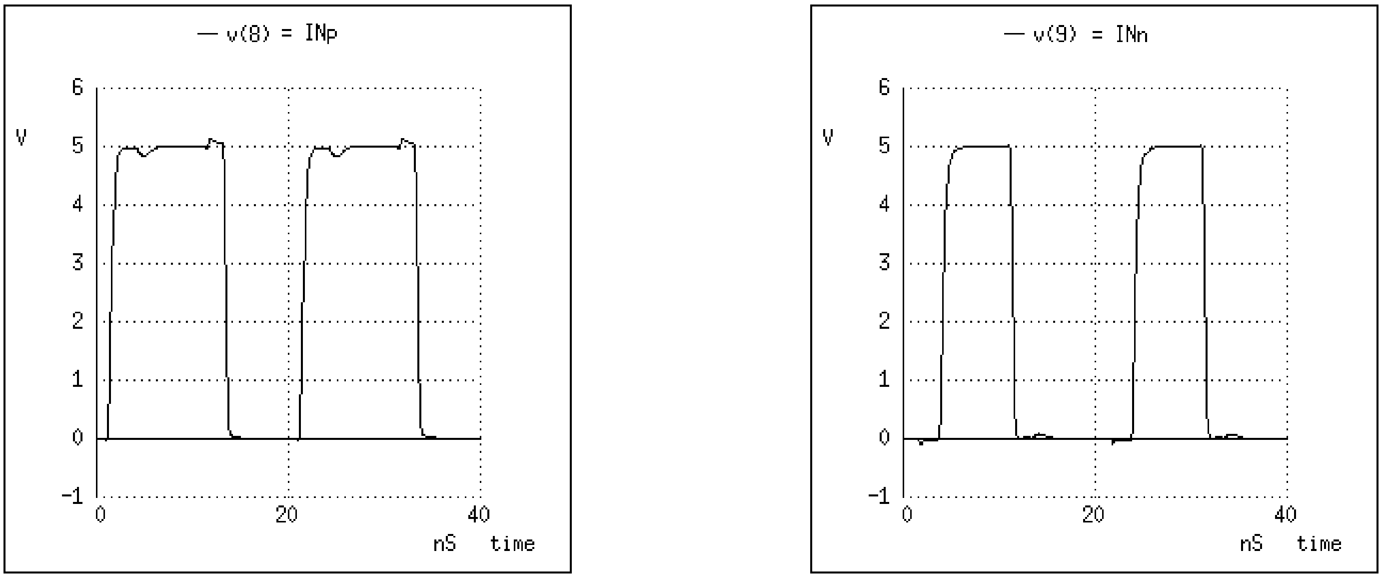

Fig. 4. SPICE simulation of the INp signal generated by a dynamic NAND and INn signal generated by a reverse dynamic NOR.

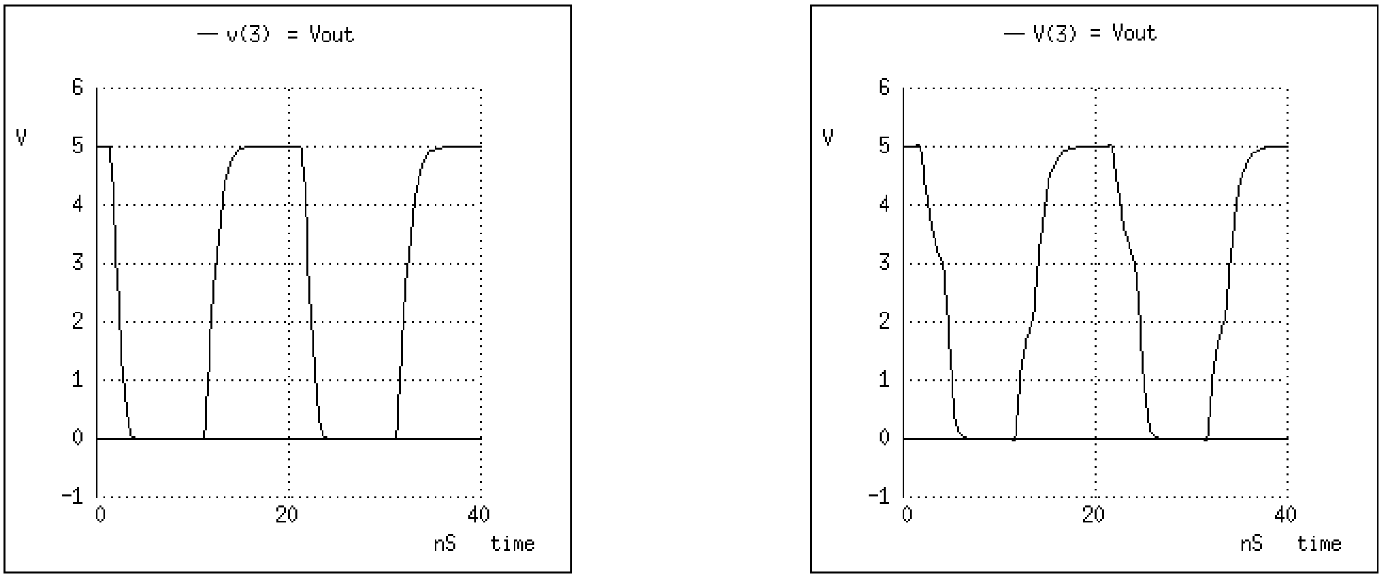

Fig. 5. SPICE simulation of the output of the step driver (right) compared with the output of the conventional driver.

to be less than 50% because in a real circuit there are other sources of power consumption besides the load C. The main difficulty is generating the various signals needed for closing and opening the p, n, and t switches at the proper time. If phases one or three are too long there will be noticeable discontinuities on the rising and falling edges of the clock and the rise and fall times will be unnecessarily large. For a clock signal this is unacceptable and for this reason using a state machine for generating the signals as proposed in [9] is not feasible. Furthermore a state machine would consume a lot of power itself. The next section will describe a simple circuit that uses a feedback from the output node in order to open and close the switches at the proper times. A major advantage of using this feedback is that the circuit behavior adapts itself to the output load C. If C is large phases one and three will be longer and this will let C charge to $Vdd/2$ before t is opened, if C is small the phases will be shorter.

A. *Low power single-phase clock driver using two-step adiabatic charging*

A schematic of a two-step driver with quantized adiabatic operation is shown in figure 3. Because there are three switches operating alternatively it is no longer possible to use a single chain of inverters in order to drive the final stage. The proposed circuit (see figure 3) has two separate chains, one for the final pMOS and another for the final nMOS.

The most important part of this circuit is the feedback from the output to the two NAND and NOR gates at the

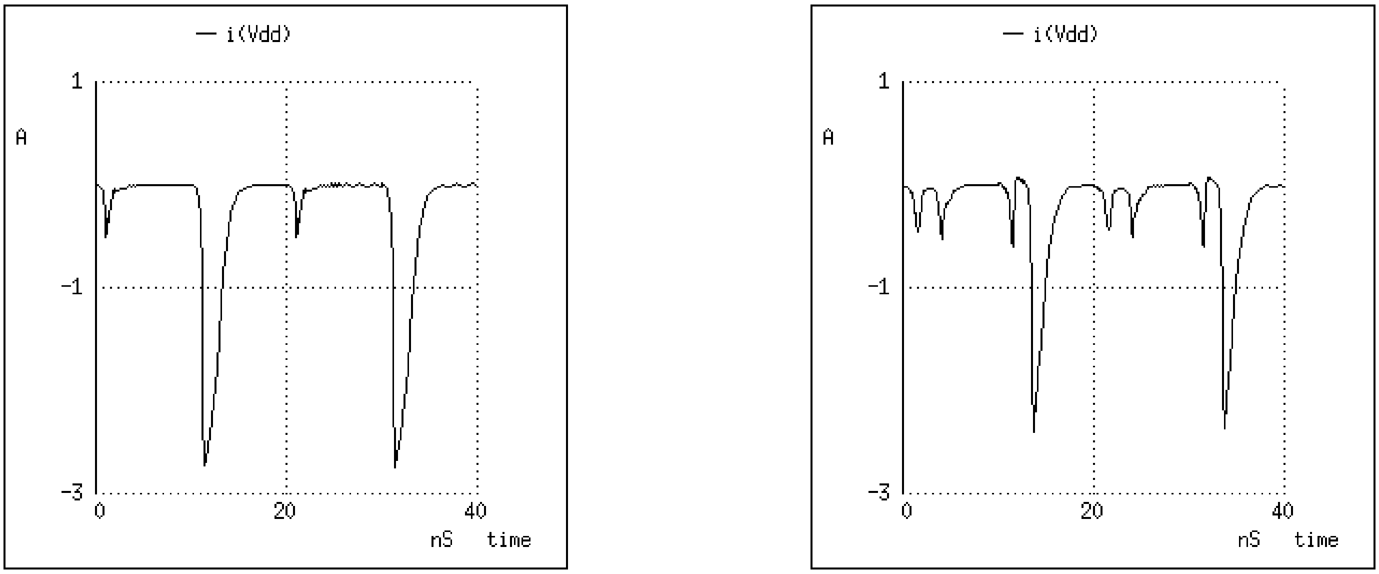

Fig. 6. SPICE simulation of the current consumed by the step driver (right) compared with the current of the conventional driver. The current has a negative polarity according to SPICE conventions.

input. This feedback does not significantly load the output node (small devices in the NAND and NOR) and allows the same circuit to work reliably with different loads. When the load is large $Vout$ will take longer to transition and this will be reflected in how fast the NAND and NOR transition. The behavior of the NAND and NOR gates can be seen from their SPICE simulated outputs in figure 4. It was chosen that the NAND and NOR be dynamic for several reasons:

- to decrease the load on the output,
- to lower the overlap current consumption,
- to work with $Vdd/2$ input voltage swings. The NAND and NOR gates have to transition for inputs at $Vdd/2$ and this does not happen for static gates unless their transistors are asymmetrically sized.

A comparison of the output of the proposed driver with the output of a standard inverter chain can be seen in figure 5. The current drawn by the low-power driver from the Vdd power supply compared with the standard inverter chain can be seen in figure 6. It can be seen that both the average and the peak current are lower for the step driver.

III. Two-step dual phase clock driver using charge recovery

A very promising technique for achieving low-power operation in CMOS is charge recovery. The basic idea is to redirect some (as much as possible) of the charge stored on capacitors that are to be discharged to those capacitors that need to be charged. This recovery of the charge translates directly into power savings because the recovered charge is not drawn from the power supply.

Charge recovery can be explained in terms of stepwise adiabatic charging and discharging and can be similarly performed in two or more steps. For a dual-phase clock driver a two-step charge recovery operation achieves 50% theoretical power savings.

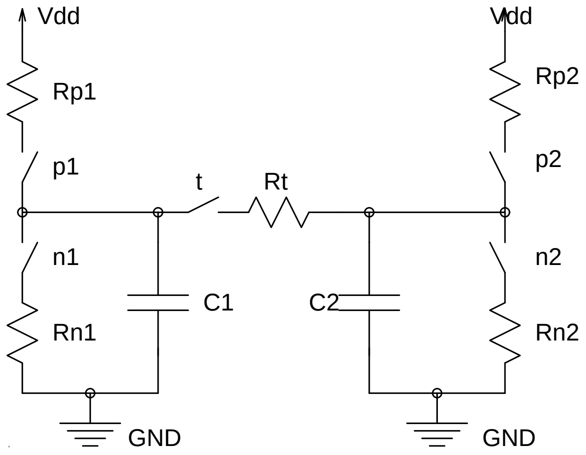

Fig. 7. Switch level diagram of the dual phase driver with charge recovery

A dual-phase clocking scheme uses two clock phases of opposite polarities. A standard dual-phase clock driver will use two drivers similar to the single-phase driver in figure 2 a. and will consume twice the power $2W = 2 \cdot C \cdot Vdd^2$. The new scheme proposed here is close to the databus charge recovery technique proposed in [8] with the important difference that taking advantage of the clock's deterministic nature significantly simplifies the circuit. A switch level diagram of the dual phase clock driver can be seen in fig. 7. Notice that this time there is no need for an extra power supply.

This driver has many similarities with the single phase step driver. There are 5 switches and four phases, two static and two transient, that explain how this circuit works (initially $C1$ is discharged and $C2$ charged):

- phase one (transient, PHI1 rising edge and PHI2 falling edge) - t is closed and in this way $C1$ and $C2$ share the charge originally on $C2$. In this way half of the

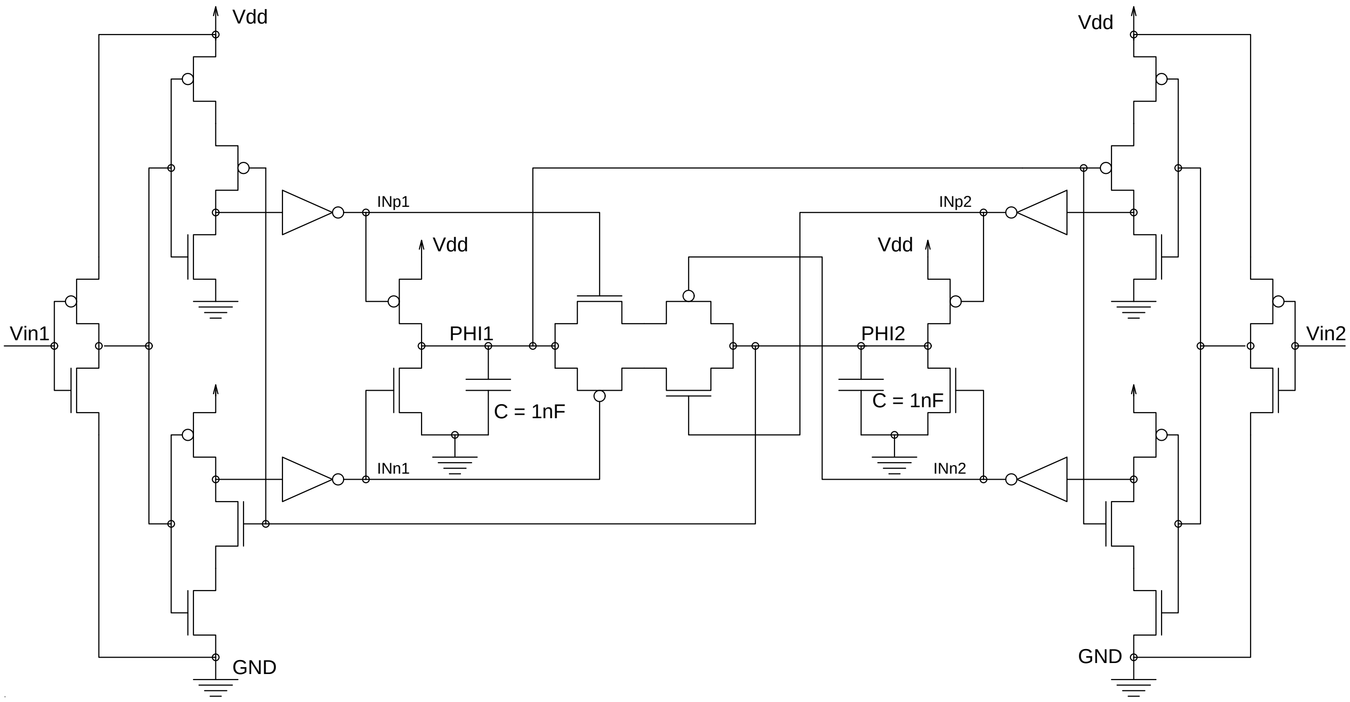

Fig. 8. Schematic of the two-phase driver with charge recovery.

charge on $C2$ is recovered. No power is drawn in this phase from the power supply.

- phase two (static) - starts when the outputs reach $Vdd/2$ by having t open and $p1$ and $n2$ closed such that $C1$ is charged to Vdd and $C2$ is discharged to GND. The energy drawn from the Vdd power supply in this second phase is $W/2$.
- phase three (transient) - t is closed with all the other switches open for sharing the charge on $C1$. Half of the charge on $C1$ is thus recovered.
- phase four (static) - t is opened and $n1$ and $p2$ closed. $C2$ is charged to Vdd and $C1$ is discharged to GND. The energy drawn from the Vdd power supply in this phase is $W/2$.

From the above simplified analysis it can be seen that W energy is drawn from the Vdd power supply for this circuit as opposed to $2W$ for the standard dual inverter chain driver.

A. *Low power dual-phase clock driver using two-step charge recovery*

A schematic of a two-phase driver with charge recovery is shown in figure 8. The signals driving the final pMOS and nMOS transistors are similar to the corresponding signals for the single phase driver (see fig. 9). The main difference is in the feedbacks from the output to the NAND and NOR which now are cross-coupled (the output of PHI1 drives the gates for PHI2 and vice-versa). This was necessary in order to make sure that the circuit will function correctly independent of initial conditions on $C1$ and $C2$. This also resulted in changes in the polarity of some signals and the reversed role of the NAND and NOR gates.

The output of the dual phase driver with charge recovery can be seen in fig. 10. A comparison of the current drawn from the power supply by the driver with charge recovery compared with a standard inverter chain driver can be seen in figure 11 from which it can be seen that both the average and the peak current for the step driver are lower.

It should be noted that although in the theoretical simplified analysis of the circuit it was assumed that the initial conditions were with $C1$ discharged and $C2$ charged this is not required for the correct functioning of the circuit. The capacitors can be in any initial condition and after only one clock cycle they get charged in sync. Figure 12 shows a simulation where both capacitors were initially discharged.

Conclusions and future work

Trying to minimize the clock power dissipation of a CMOS circuit is appealing for several reasons:

- the power dissipated by the clock represents a large percentage of the total power dissipation. Savings in the clock power dissipation will have then a large impact on the overall power dissipation.
- the clock circuit is principially simple and it makes sense to spend extra design effort for optimizing it.
- the clock signal is deterministic. For this reason techniques that only work with some probability on general purpose logic circuits will work in a deterministic way for a clock driver.

In this paper we described practical implementations with simulation results of two CMOS clock driver circuits that have similar operation characteristics although they use different principles for achieving low-power operation. The single-phase clock driver uses a two-step adiabatic charging and requires an extra power supply or tank capacitor. The dual-phase clock driver uses a two-step charge

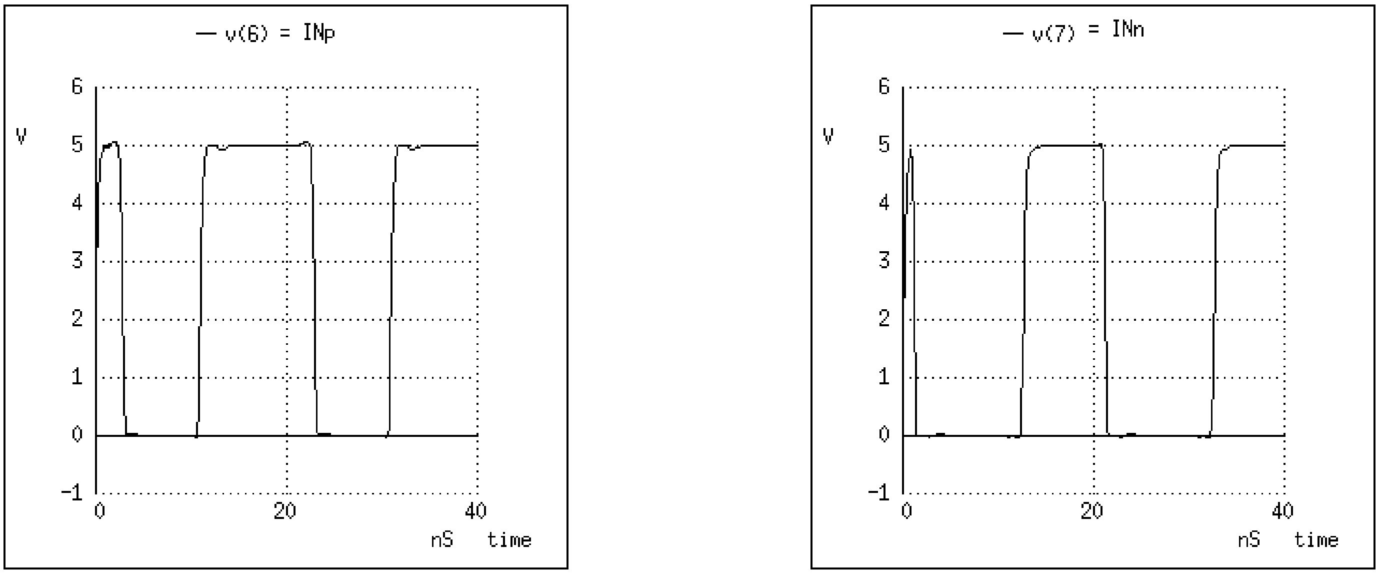

Fig. 9. SPICE simulation of the INp1 and INn1 signals that drive the phase 1 driver. They are very similar to the corresponding signals for the single phase driver

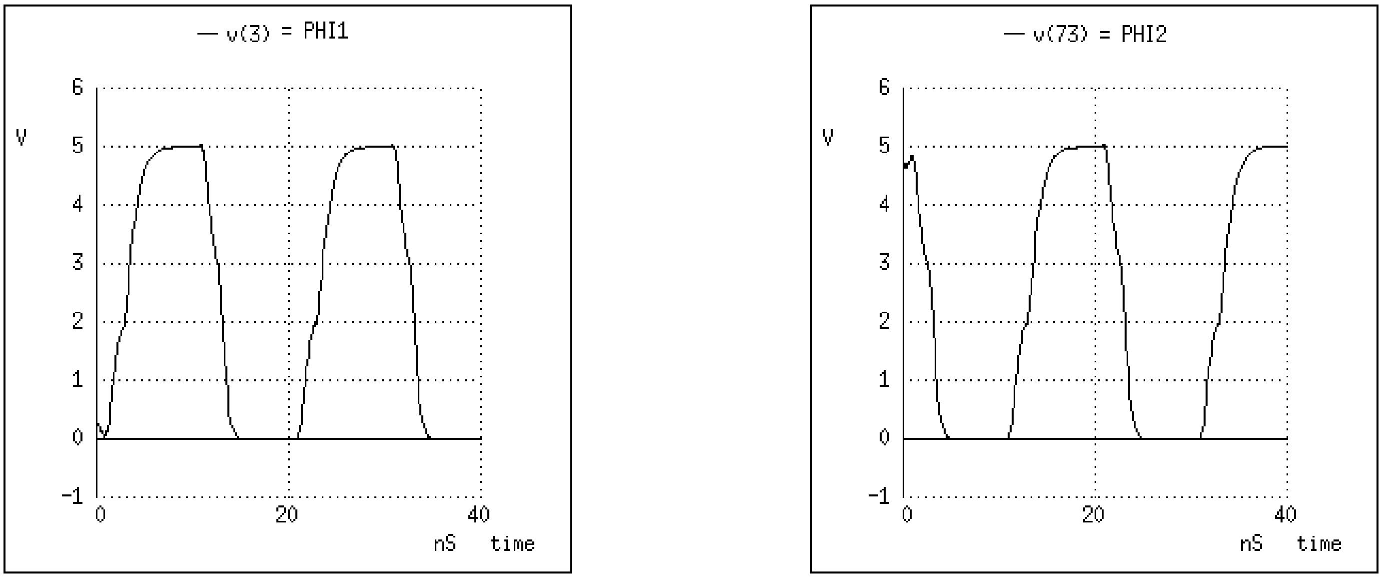

Fig. 10. SPICE simulation of the output of the dual step driver with charge recovery.

recovery scheme. Both drivers exhibit a theoretical 50% (25-30% for actual circuits) power savings compared with standard inverter-chain clock drivers but they are significantly more complex. The actual power was determined with a technique described in [6] and the power savings for the current implementations were found to be around 25% for the single-phase and 30% for the dual-phase driver.

The circuits can be used for on-chip or off-chip clock or data transmission for highly capacitive loads when the larger complexity and slightly slower operation are not detrimental. Further work could refine transistor sizes or use BiCMOS circuits in order to decrease the driver area and minimize clock edge discontinuities. Further work is also needed to determine which types of latches and flip-flops work with these drivers and which don't.

The circuits where not layed-out or fabricated but it was tried to use realistic SPICE files. SPICE level 3 models of a 0.8 micron HP process available through MOSIS where used and the transistors where described with their AD, AS, PD and PS parameters in order to take diffusion capacitances into account.

References

[1] W. C. Athas, L. J. Svensson, J. G. Koller, N. Tzrtzanis, E. Chou "A Framework for Practical Low-Power Digital CMOS Systems using Adiabatic Switching Principles", *International Workshop on Low Power Design*, pp. 189-194, Napa Valley, Apr. 24-27, 1994.

[2] A. P. Chandrakasan, S. Sheng, R. W. Brodersen, "Low-Power CMOS Digital Design", *IEEE Journal of Solid-State Circuits*, pp. 473-484, April 1992.

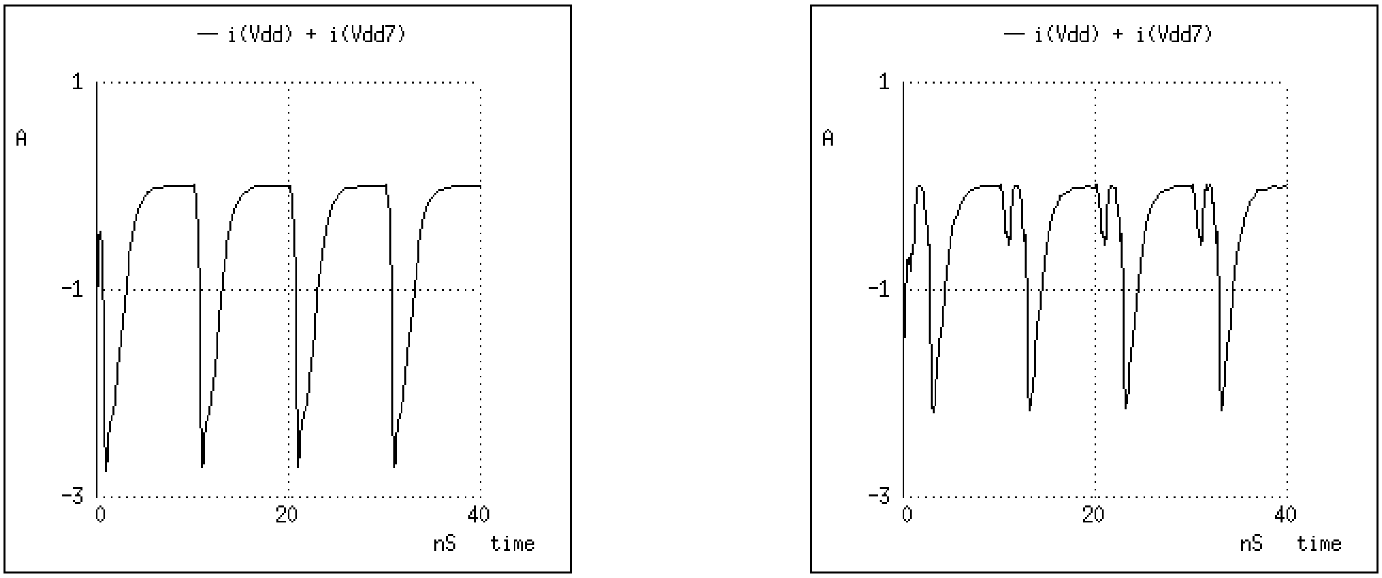

Fig. 11. SPICE simulation of the current consumed by the driver with charge recovery (right) compared with the current of a conventional dual inverter-chain driver.

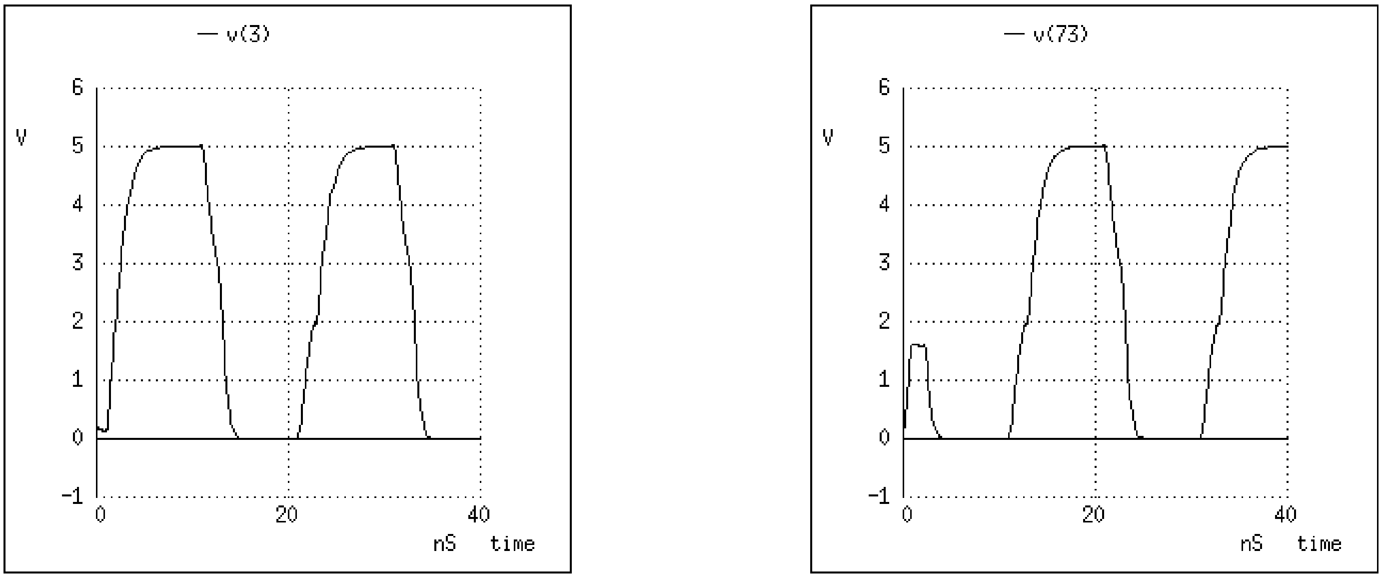

Fig. 12. The output of the two-phase driver when both capacitors are initially discharged.

[3] E. De Man, M. Schobinger, "Power Dissipation in the Clock System of highly pipelined ULSI CMOS Circuits", *International Workshop on Low Power Design*, pp. 133-138, Napa Valley, Apr. 24-27, 1994.

[4] D. Dobberpuhl et al. "A 200-MHz 64-bit Dual-Issue CMOS Microprocessor", *IEEE Journal of Solid-State Circuits*, pp. 1555-1567, Nov. 1992.

[5] R. Hossain, L. D. Wronski, A. Albicki "Low Power Design Using Double Edge Triggered Flip-Flops", *IEEE Transactions on VLSI Systems*, pp. 261-265, June 1994.

[6] S. M. Kang "Accurate Simulation of Power Dissipation in VLSI Circuits", *IEEE Journal of Solid-State Circuits*, pp. 889-891, Oct. 1986.

[7] H. Kojima, S. Tanaka, K. Sasaki "Half-Swing Clocking Scheme for 75% Power Saving in Clocking Circuitry", *Symposium on VLSI Circuits*, pp. 23-24, 1994.

[8] K. Y. Khoo, A. N. Willson, "Charge recovery on a Databus", *International Symposium on Low Power Design*, pp. 185-189, Dana Point, CA, Apr. 23-26, 1995.

[9] L. J. Svensson, J. G. Koller, "Adiabatic Charging without Inductors", *International Workshop on Low Power Design*, pp. 159-164, Napa Valley, Apr. 24-27, 1994.

[10] J. Yuan, C. Svensson "High-speed CMOS Circuit Technique", *IEEE Journal of Solid-State Circuits*, pp. 62-70, Feb. 1989.

[11] Q. Zhu, J. G. Xi, W. W.-M. Dai, R. Shukla, "Low Power Clock Distribution Based on Area Pad Interconnect for MCM",*International Workshop on Low Power Design*, pp. 87-92, Napa Valley, Apr. 24-27, 1994.

[12] N. Weste, K. Eshraghian, *Principles of CMOS VLSI Design, A Systems Perspective*, Addison-Wesley Publishing Company, 1993.