3.2 Current-Steering DACs

Current-steering DACs are a more common integrated DAC compared to resistor DACs. They are probably the most common bias DAC architecture due to their small size and simplicity. The current-steering DAC replaces the resistor element in the resistor DAC architectures with a MOSFET current element and uses some form of summation of the current elements to produce the result. Sometimes the result needs to be a current such as in integrated bias circuits. This current is then passed to the next stage of a current mirror bias or passes through a MOSFET stack to produce a set of current mirror bias voltages. Voltage-mode DACs convert the current to voltage with a simple resistor or a low-output resistance transimpedance circuit for improved linearity, as shown in Fig. 3.4.

Current source DACs are attractive due to the lower space used by MOSFET current elements and the ability to perform some calibration tricks. The resistor DAC architectures discussed in Section 3.1 can be directly repeated using current sources instead of resistors. This even includes the R-2R ladder DAC.[3]

3.2.1 Unary DACs

Unary current DACs use a single-current element for each quantization step. Unary current DACs are analogous to resistor divider DACs with a resistor element for each LSB. Consequently, these DACs are inherently monotonic, but they consume a large area for medium to high resolutions, just like their resistor divider DAC cousins. In Fig. 3.5, I_u is a single unit current and T_x is the thermometric equivalent of the binary input. A binary input of 00...000 leave all switches open. If we had a binary input of 00...0011, or 3, then T_0, T_1, and T_2 switches would be closed, and the rest of the switches would be open.

3.2.2 Binary DACs

Binary current DACs group current elements into binary multiples that are turned on or off directly with the input bits. This eliminates the decoder required in unary current DACs. However, binary DACs still often use

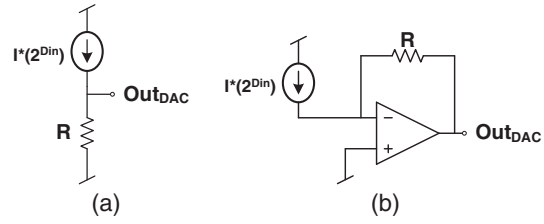

Figure 3.4 (a) Non-inverting and (b) inverting current-steering DAC architecture.

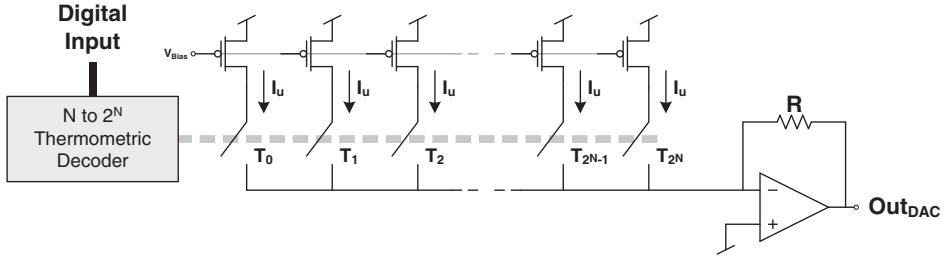

Figure 3.5 Typical unary current divider DAC.

individual unit current sources to make up the larger binary units for matching purposes. For example, the $2 \times I_u$ element would be composed of two identical I_u elements. This results in binary DACs being composed of the same number of current elements as in a unary DAC for equivalent resolution. However, it is entirely possible to simply change the transistor aspect ratios to achieve the binary current multiples instead of exactly repeating the unit current element. This will result in more nonlinearity, but many applications can live with that, in exchange for a much smaller layout area. This technique will produce nonlinearity which is worst between MSB switch points such as 011...11 switching to 100...00.

Unary and binary current DACs are often used together. Typically, unary DACs are used for the MSB current elements because of their inherent monotonicity. Binary DACs are used for the LSB elements because of their much smaller size when created with weighted transistors. This concept is nicely illustrated in a publication by Horsky[4] for a monotonic 16-bit DAC designed for harsh environment sensor applications. Horsky utilized a seven-bit MSB unary current DAC followed by a five-bit binary current DAC. An astute reader will quickly realize that this does not equal 16 bits. The remaining resolution was cleverly realized by using a single LSB current element with pulse width modulation (PWM) to keep the unit current of the unary and binary DACs at a reasonably low level and maintain a reasonable layout area. The current cycling of this PWM LSB element is smoothed out with a large capacitor or with the sensor capacitance itself. The PWM element makes up the four LSBs of the DAC, and the remaining 16th bit is the MSB, which is a simple current polarity bit. Using all four of these approaches together makes Horsky's publication a very good illustration of the flexibility of the current DAC architecture.

3.2.3 Dynamic calibration of DAC current elements for high resolution

Current DACs are compatible with matching algorithms to improve their linearity and still achieve a relatively small layout area. DAC elements can be trimmed (this is also true for resistor DACs), but this becomes unwieldy for

large transistor arrays and is an expensive way to improve linearity. Dynamic current calibration is a more automated way to improve linearity of a current DAC. Dynamic current calibration relies on the principle of a stored floating voltage bias on a capacitor at the gate of each current element. This floating voltage is set with a reference current that is switched through each unit current element during a calibration cycle while each current element transistor is diode connected, meaning the gate is connected to the drain. The gate of the current element is biased according to the reference current, then the gate-to-drain connection opens, and the bias voltage floats on the gate of the current element until it is refreshed during the next calibration cycle. Non-idealities such as charge injection offset are overcome with various schemes. In 1989, Groeneveld et al.[5] achieved a ± 1 INL on a 16-bit current-dividing DAC using this dynamic matching technique. Groeneveld reduced sensitivity to charge injection by dynamically matching a 10% trim current source instead of the entire element, reducing the impact of non-idealities accordingly.

3.2.4 Switching current elements for improved speed

Current DACs are slowed down when their current elements are completely shut off, as would happen in the basic structure shown in Fig. 3.6. This slowdown is caused by the increased time it takes to charge up and settle the parasitic capacitances in each MOSFET current element. Most current DACs that require higher speed simply switch each current element from a low-impedance current dump node into the circuit. This keeps the current source on and properly biased which, in turn, greatly speeds up the DAC switching and settling. The switch clock edges must be closely aligned but non-overlapping in this type of scheme to prevent shorting the current dump node to the output. However, the switching clock non-overlapping time must not be too long because the parasitic capacitance of the current source will charge up and shut off the current element, eliminating the speed improvement. Current-steering DACs used in high-speed ADCs usually require this approach.

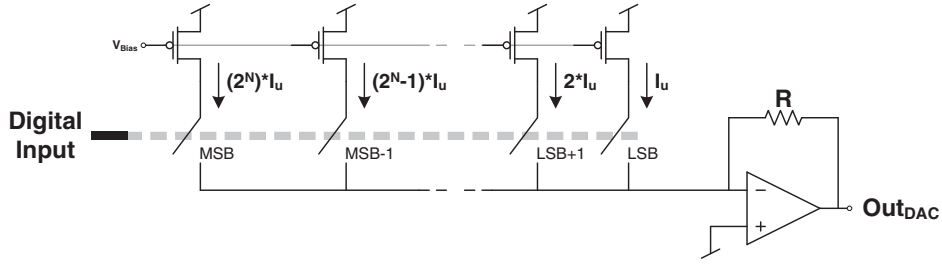

Figure 3.6 Typical binary-weighted current DAC.

3.3 Switched-Capacitor DACs

We have discussed voltage-mode DACs (resistor DACs) and current-mode DACs, and now it is time to move to charge-mode, switched-capacitor DACs. Switched-capacitor DACs are the most popular DAC for ADC architectures. They can be switched into configurations to realize many different functions, making them good for implementing various types of mathematical algorithms in addition to the DAC operation. Capacitor-based DACs leverage highly controlled oxides and doping available in modern CMOS to achieve good capacitor matching with small-area devices. One of the biggest benefits of capacitor DACs is that capacitor arrays use no quiescent current, unlike resistor or current arrays. Switched-capacitor DACs trade clock cycles and time for reduced components versus resolution.

3.3.1 Capacitive divider

The basic capacitor divider DAC is quite similar to the other binary DACs discussed in the previous sections. It is made up of binary-weighted capacitors that are switched in or out, depending on the binary input, as shown in Fig. 3.7. The difference is that the capacitor divider DAC uses two clock phases. The first phase is used to reset the capacitor array, and the second phase is used to perform the charge-sharing DAC operation.

The capacitor divider DAC in Fig. 3.7 will experience charge sharing with amplifier input drivers in the unity gain buffering amplifier. This charge sharing can severely degrade linearity as the input to the buffer amp swings across the FSR. A more common capacitive divider DAC technique, shown in Fig. 3.8, uses the high gain of an op-amp to hold the common node steady and eliminate this problem. This DAC inverts the output, but the inverted output can easily be accounted for by using inverted digital input. The small size and binary-weighted behavior of this type of DAC make it attractive for use as the feedback DAC in successive approximation ADCs (discussed in Section 4.5), which are common for column-level conversion in image sensors.

Capacitor divider DACs are common in image sensors. They are an important part of the architecture of many successive approximation ADCs.

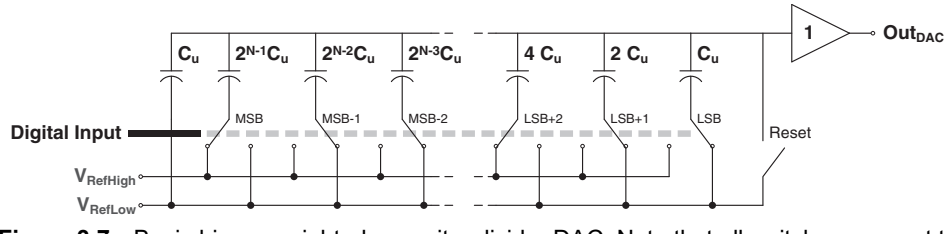

Figure 3.7 Basic binary-weighted capacitor divider DAC. Note that all switches connect to V_{RefLow} during reset.

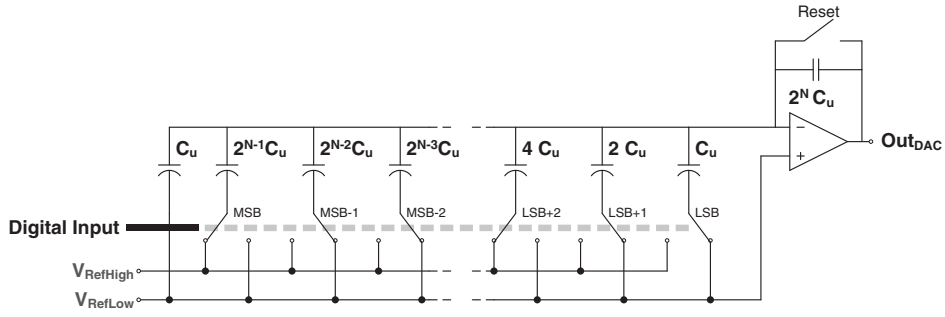

Figure 3.8 More typical binary-weighted capacitor divider DAC maintains a steady common node with a high-gain op-amp and feedback. Note that all switches connect to V_{RefLow} during reset.

They are often combined with an input sampling scheme to serve as programmable gain amplifiers to adjust image sensor signal chain gain and offset.

3.3.2 Charge redistribution 2-cap DACs

One problem with the binary-weighted DACs is the size and complexity of the capacitor arrays for resolutions exceeding just a few bits. This may be fine for DACs located on the periphery of a sensor array, but DACs with medium to high resolution inside a single image-sensing column, or even inside a single pixel may not be able to tolerate these large arrays. Instead of a large array of capacitors, we can perform conversion using only two capacitors switched in time to share charge, as shown in Fig. 3.9. Suarez, Gray, and Hodges[6] introduced this algorithm in 1975 and achieved 8 bits of resolution with the technology of the time. The charge redistribution 2-cap DAC is one of the first types of DACs employed in column-level analog-to-digital conversion by JPL.[7]

Additional switching of the capacitor array to that shown in Fig. 3.9(a) allows the two capacitors to interchange during each sample to obtain a resultant voltage that is almost independent of capacitor mismatch. This allows smaller capacitors to be used for the conversion of much higher resolutions at the cost of extra switch time and a little more complexity.

A good example of a two-cap DAC with this capacitor mismatch compensation can be seen in Lu et al.[8] Their 2008 study concluded that up to 20 bits of resolution can be achieved with capacitors matching only to 0.1%, although they acknowledge that noise will be the real limit. This matching technique sacrifices some conversion speed, but the authors suggest that this algorithm is only needed on a few of the MSB switching cycles that most impact linearity in this architecture. The simple routine shown in Fig. 3.9,

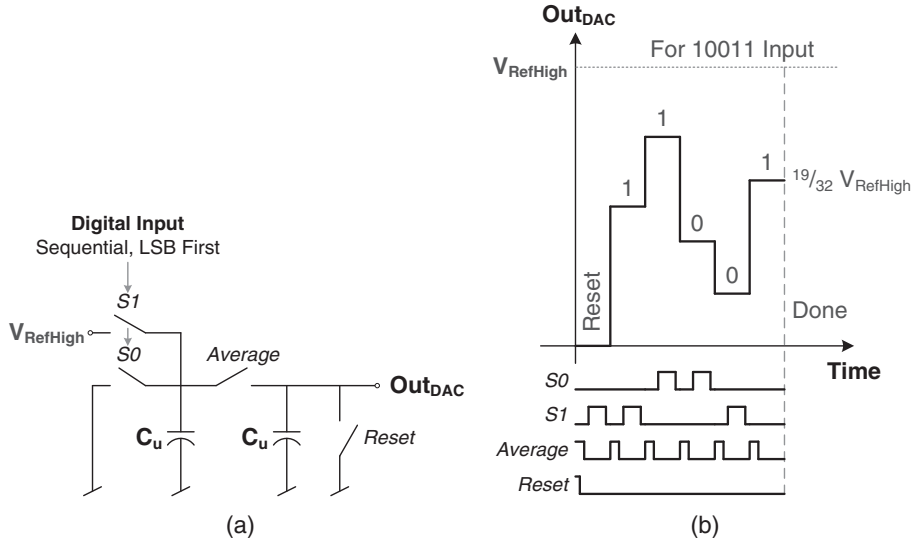

Figure 3.9 (a) Two-capacitor DAC and (b) the resultant output for 10011 (19) input word.

neglecting the reset, may be used on the remaining LSBs to speed up the remainder of the conversion cycle.

The small size and capacitor mismatch reduction schemes of cyclic two-capacitor DACs make them a good choice as subcomponents for column-level or even pixel-level ADC architectures on image sensors. They have also been proven for other column-parallel solutions such as read-in integrated circuits for driving scene generator chips and display array driver DACs. For example, in 2005 Bell[9] published a good paper utilizing 2-cap DACs to drive an LCD chip. Bell's read-in IC system utilized 820 column-parallel 2-cap DACs with resolutions of 10 or 12 bits. Each DAC required 15 μs for conversion, and two DACs fit within a column pitch of 45 μm.

3.4 The Specialized Multiplying DAC

Pipeline, algorithmic, and two-step ADCs will usually utilize a specialized multiplying DAC (MDAC) circuit that performs multiple operations in addition to the digital-to-analog conversion. These operations are performed in two clock phases, as illustrated in Fig. 3.10. During clock phase 1, the MDAC samples an analog input. During clock phase 2, the MDAC switches over and performs digital-to-analog conversion on the sampled input, and subtracts the DAC output from the analog input sampled during phase 1, producing a residue. This residue voltage is then gained up and held on the output for the next stage to sample. We will discuss this type of DAC architecture in much more detail in Chapter 5.