AIAA

Doty

Notes on the CCD Driver Board Design:

Clock Regulators

The clock voltage regulators have some difficult requirements:

1. Peak output current is much larger than the average output current. This makes power efficiency difficult to achieve. A large output capacitor is of great benefit here.

2. Clock levels crosstalk (at roughly the 1% level) with the video output. This means that the clocks must behave consistently from pixel to pixel to avoid affecting the video bias. It also means that the clock regulators must be reasonably quiet (super low noise isn't required, but you may not be able to ignore noise performance entirely).

Reset Clock Regulators

The easiest regulators to design were the reset clock regulators (sheet 5). The load capacitance is relatively small (perhaps ~100pf, mostly in cabling), so the current requirement is small (~1nC/pixel), about 40μA DC. Furthermore, it is reasonable to ask that the sequencer program operate this clock in a strictly periodic way: it can be clocked the same every pixel time regardless of the state of the other clocks. This means that there is no problem making it behave consistently from pixel to pixel, and it can therefore have a relatively long time constant.

Capacitors C29 and C30 store the charge to operate the clocks. A withdrawal of 1nC will cause a 10mV drop in the regulator output: this is replenished through R70 or R73 with a time constant of 22μs. These resistors keep the capacitors from destabilizing the op amps that perform the actual regulation.

Serial Clock Regulators

The serial clock regulators are more difficult. The load capacitance is higher (a

few hundred pf), and each regulator serves three drivers. The average load current is thus a large fraction of a milliamp. The clocks cannot be run strictly periodically; they must cease during the parallel transfer. This means that the circuit time constants must be short: otherwise the serial clock levels will change after the parallel transfer, causing a bias shift for the pixels read out immediately after the serial transfer.

A conventional op amp is not the best device for driving capacitive loads with short time constants. Fortunately, an Operational Transconductance Amplifier (OTA) is effective in this application. OTA's are less familiar to engineers than they should be: there are few OTA IC's available, and the only common application for bare OTA's I'm aware of is that they are widely used as variable gain amplifiers in audio equipment. On the other hand, they are very commonly used as building blocks *within* analog IC's like power regulators, sample/hold modules, and charge division ADC's. The big advantage they have in these applications is that they drive capacitive loads well: a conventional op amp is destabilized by capacitive loads, but the "compensation" capacitor for an OTA is its load capacitor, so the more load capacitance you give it, the more stable the OTA will be.

The time constant for a capacitively loaded OTA is GC/g_m, where G is the closed-loop gain, C is the load capacitance, and g_m is the transconductance. To minimize the magnitude of the voltage transient when the clock changes state, we need to use as large a C as possible. This means we need a large g_m. The transconductance g_m is controlled by the amplifier bias current: for the CA3080 OTA, the transconductance is about 20 millisiemens per milliamp of bias current. A really large transconductance thus requires a large bias current and large power dissipation.

We can avoid the large dissipation by putting a transistor current amplifier at the output of the OTA: this increases the transconductance by a factor of β. The time constant for the regulator is thus $GC/\beta g_m$. For example, the positive serial clock regulator (right side of sheet 4) has a bias current of 4.4V/20kohm=220μA, yielding a tranconductance of 4.4mS. The closed loop gain is 1+60.4k/14.7k=5.11. The load capacitance is 0.33μf, and the transistor is a 2N4401 with a typical β of 200. The resulting time constant is 1.9μs.

For stability, any time constants within the loop should be short compared to 1.9µs. The OTA itself is quite fast: its quoted bandwidth of 2Mhz corresponds to a time constant of 60µs (it's not that simple, but in this case it's fast enough to be neglected). The dominant destabilizing factor is the time constant associated with the transistor's "beta cutoff frequency", $f_\beta=f_T/\beta$. At high ouput currents, this effect is dominated by the transistor's transit time, and is nearly independent of the current. For the 2N4401, f_β is approximately 1Mhz in this domain, corresponding to a negligible time constant of 160ns. At low collector currents, however, f_β is dominated by the loading effect of parasitic capacitance and is approximately proportional to collector current. In this domain, the transistor's time constant may be taken as $h_{ie}C_p$. Thus, as the output current declines, the circuit becomes less stable, and will eventually oscillate at a low level. A resistor attached from the base to the emitter of the transistor can stabilize it against this effect: by stealing some of the base current it both limits how large h_{ie} (the input resistance) of the composite circuit can become, and it also reduces the β of the composite circuit at low currents.

The circuit effectively operates in two modes: in the active mode the transistor is on, and the time constant is 1.9µs. In the passive mode, the transistor is off, and the ouput current flows through the 100k base-emitter resistor. In this case, β is effectively 1, so the time constant is 380µs. The circuit is quite stable at the extremes, but the crossover point between the two modes requires some analysis. At the crossover point the transistor h_{ie} is 100k, the effective h_{ie} is 50k (two 100k's in parallel) and the effective β is 100 (if you change the ouput current from the OTA by a small amount, half of it bypasses the transistor, reducing the β by a factor of two). If we assume C_p is 30pf, the transistor's time constant is 1.5µs, comfortably below the circuit's overall time constant of 3.8µs.

The crossover between the two modes occurs when the OTA's output current is just barely enough to generate the transistor's V_{be} voltage through the 100k resistor. The output voltage difference between this point and the point at which the OTA's output current is zero is GV_{be}/g_mR. For V_{be}=0.5V and R=100k, this voltage is about 6mV. Thus, if the DC output current is zero, the transient reponse shows two time constants: it initially has a time constant of 1.9µs, but

when the output voltage is within 6mV of its final value, the time connstant increases to 380μs.

A regulation uncertainty of 6mV is enough to be noticeable at the CCD output: each mV of change in one of the serial clock levels will produce a video bias shift of about 1e-. Fortunately, the active output current is sufficiently high that while the serials are actually clocking the circuit never enters the passive mode, and even the drain due to a single clock edge is sufficient to turn on the transistor. Thus, when clocking, the circuit forgets its history very rapidly, and yields consistent clock levels from cycle to cycle. This behavior is confirmed by circuit simulations.

Note that if the output voltage more than ~0.5V from ground, the feedback network's load is sufficient to keep the transistor turned on, so quiescence is only an issue for the lowest possible output voltages.

Parallel Clock Regulators

The parallel clock regulators are similar to the serials, but the loads they are designed to drive are much larger: the parallel clock capacitance will be in the tens of nanofarads. The regulator output capacitance has thus been increased to 2.2μf, and the OTA bias current is 700μA. The active mode time constant is thus about 4μs, and the quiescent mode time constant is about 800μs. The transition from the active to the quiescent state occurs only 2mV from the zero output current equilibrium.

The parallel clock level doesn't directly affect the video bias, but a change in the parallel clock level can change the video bias throgh "substrate bounce". This is essentially a ground loop on the CCD chip: the current flow due to the change of the clock level changes the substrate potential because the substrate isn't perfectly grounded. Our differential video input stage is relatively immune to this effect, so we shouldn't have to worry too much. The relatively short (compared to previous MIT CCD systems) parallel clock regulator time constant means that the regulator settles rapidly to the point where the slope of its output is negligible. If it reaches quiescence, the time constant lengthens, but the rate of change of the output in quiescence is at most 2.5μV/μs. This slope would be barely noticeable

even if it were directly coupled to the video amplifier's input; in practice it should be completely negligible.

Notes on the CCD Video Board Design:

Preamplifier

The video preamplifier is an unusual differential amplifier:

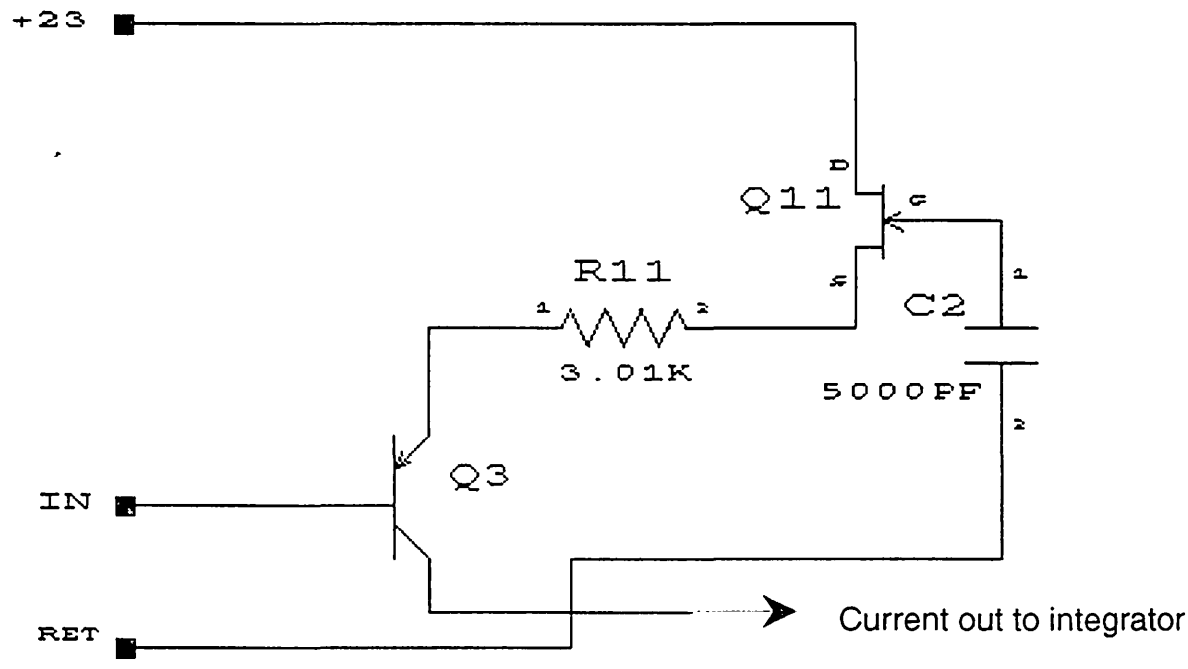

The usual differential amplifier contains two identical transistors coupled together by their emitters (or sources), with input voltage applied differentially to the bases (or gates) and output current taken fromone or both collectors (or drains).

The circuit above is somewhat similar: Q3 and Q11 are coupled together emitter to source, the input volage is applied differentially between the base of Q3 and the gate of Q11, and the output current is taken from the collector of Q3. As with a conventional differential amplifier, the conductance from the emitter of Q3 to the source of Q11 determines the transconductance: this conductance is dominated by R11, with the larger transconductances of the transistors as small corrections. Unlike the conventional circuit, the operating current flows in series

through the two transistors, so no separate bias circuit is needed. Bias current is controlled by the charge on C2.

The DC offset of this circuit is much larger (and harder to define) than that of a conventional differential amplifier, but that doesn't matter here, because the DC offset of the CCD output is even larger (~15V!). A "clamp" circuit, described below, controls the charge on C2 in such a way that the "zero signal" output level of the CCD produces a nominal output current (0.5 mA) at the collector of Q3.

Integrator

The output current from the preamplifier is converted to a voltage by a "dual slope" current integrator constructed of discrete transistors:

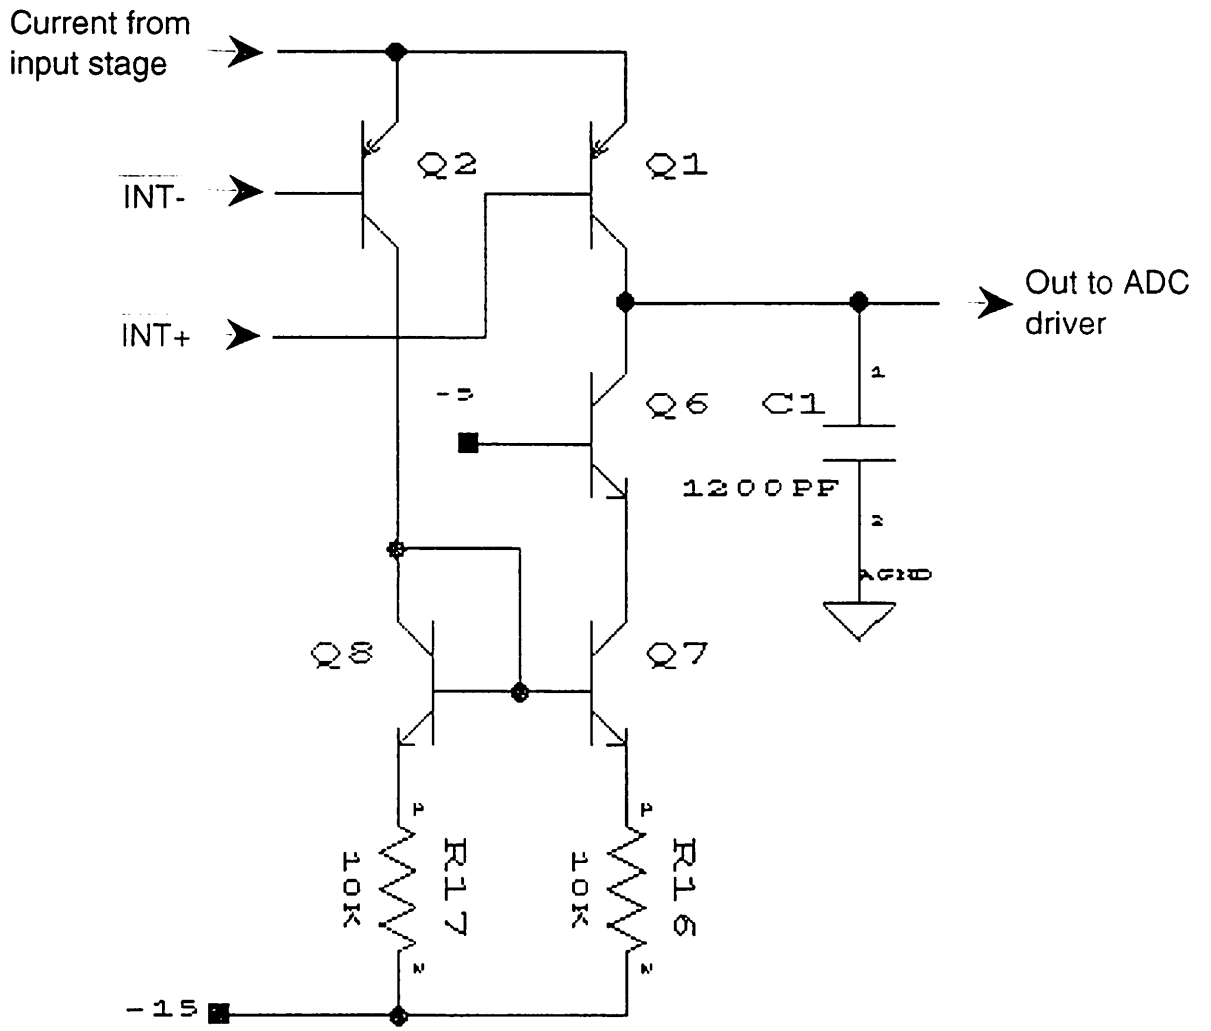

Q1 and Q2 are switches that control which branch of the integrator the current goes to. After the CCD output amplifier is reset, the sequencer will assert (pull down) INT-/, turning on Q2. This directs the current through a current mirror (a current amplifier with a gain of -1) formed by Q8 and Q7. This pulls charge from C1. After a period of time, nominally 8 µs, the sequencer will deassert INT-/, turning off Q2, and assert INT+/, turning on Q1. At the same time, the sequencer

lowers S3-OR, pushing charge into the CCD output node. The current, which now includes a contribution proportional to the charge at the output node of the CCD, now charges C1. After another 8 µs (nominal), the sequencer will deassert INT+/ (to learn where the current goes, see below). C1 now contains an amount of charge determined by the difference of the average current during the two integration intervals.

This circuit thus implements the "differential averager" signal processing function, known to be optimal when the CCD output noise is white, and near optimal for a $1/f$ CCD output noise spectrum.

Cascode transistor Q6 corrects a subtle defect in the current mirror: because the mirror transistors must be forward biased to function, there is a certain amount of charge which must be present in the mirror before it will start operation. This charge is provided by the first pixel in a sequence; as I recommend that the video processing be performed every pixel time whether it is needed or not, this charge effectively comes from the first pixel clocked in any continuous sequence of frames. This charge is held in the base-emitter capacitances of the mirror transistors. The trouble comes when capacitor C1 is discharged: if Q6 is omitted, the collector voltage on Q7 goes down, pulling charge out of the mirror through the collector -base capacitance of Q7. This charge must be made up during the next pixel, and it thus causes the level of a pixel (which controls the size of the step when C1 is discharged, and therefore the charge removed from the mirror) to affect the level of the next pixel. Q6 prevents this: when its collector voltage changes, the resulting charge transfer is to the -5V supply (which won't move much for a couple of picocoulombs).

Current mirrors are often made without emitter resistors, but in this case they are needed. Resistors R16 and R17 help balance the mirror, that is, help keep its gain near -1, but they serve a more important function as well. The large transconctances of the mirror transistors lead to large noise currents. R16 and R17 reduce the effective transconductances of the transistors by ~200, drastically reducing current noise.

To avoid visual clutter, I haven't shown the voltage dividers that reduce the logic

swings of INT+/ and INT-/ from 0-5V to 4-5V. The reduced swing is necessary to keep the switch transistors in their active regions when the output voltage becomes positive. Reducing the logic swing also minimizes crosstalk from the analog to the digital parts of the system.

Integrator Reset

The integration capacitor C1 is discharged by the following circuit:

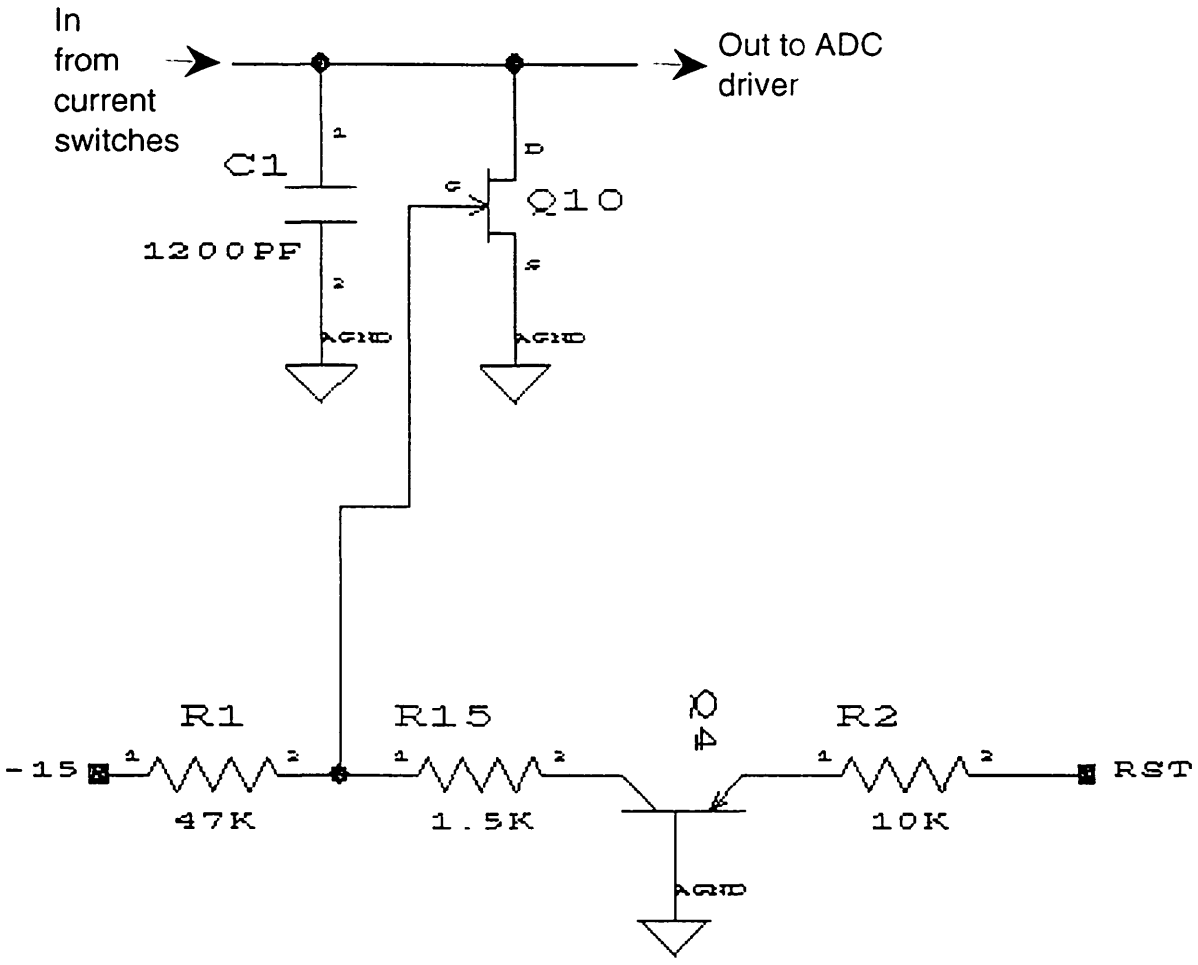

Transistor Q4 acts as a level translator, converting the 0-5V logic swing to a -15-0V swing at the gate of the reset switch transistor Q10. R15 is needed to prevent forward biasing of the gate of Q10; when Q4 is on, its collector voltage saturates at about 1 diode drop above ground. If R15 were absent, a gate-source current would flow through Q10, preventing it from fully discharging C1.

Video Baseline Clamp

To function properly, the video preamplifier must be biased in such a way that it conducts an appropriate current at zero signal level. Also, the current output from the preamplifier must go to some harmless sink when the integrator is inactive. These functions are performed by the following circuit:

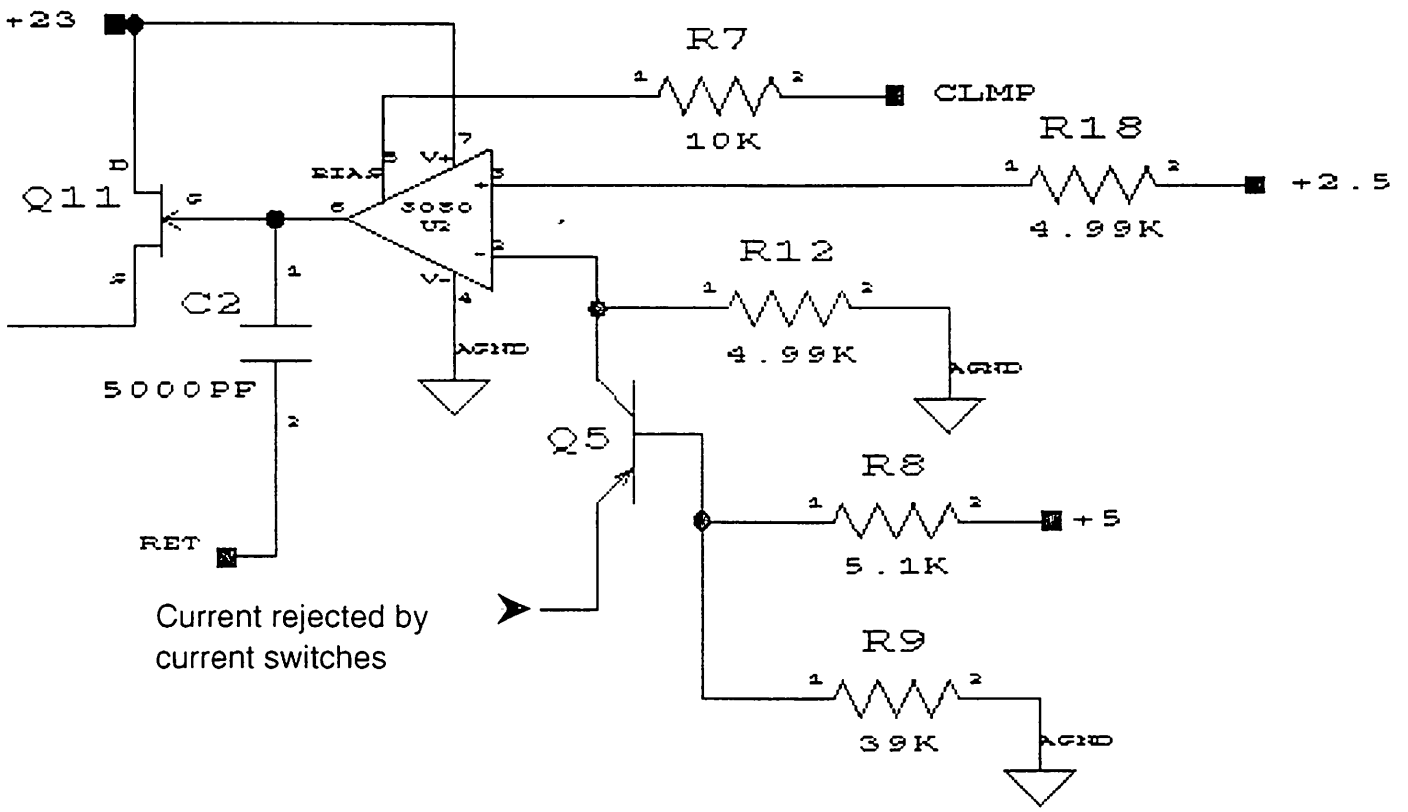

Q5's base is biased to about 4.5V, or roughly halfway between the logic levels for the switch transistors Q1 and Q2 described previously. Thus, if either Q1 or Q2 is turned on, essentially no current flows through Q5 (the base voltage difference is ~0.5V which is >>kT/q (~26 mV)). On the other hand, if both Q1 and Q2 are turned off, all of the current from the preamplifier will pass through Q5. This current will develop a voltage across R12; at the nominal preamplifier bias level of 0.5 mA, the voltage will be 2.5V.

After CCD output stage reset, but before video integration, the sequencer will assert CLMP, turning operational transconductance amplifier U2 on. U2 is connected in such a way that it drives the charge on C2 in the direction that will

make the voltage on R12 closer to 5v (the coupling is through the preamplifier, only patially shown in this drawing, see the preamplifier description above).

R18 fixes another subtle problem: when U2 is turned on by asserting CLMP, there is a transient common mode current pulse at the inputs. If the impedances at the inputs are different, the amplifier sees this pulse as a differential voltage and responds to it. Once the transient passes, the circuit behaves properly, but the effect is large enough to increase the settling time.

ADC Driver

The ADC driver amplifier is a gated current feedback op amp:

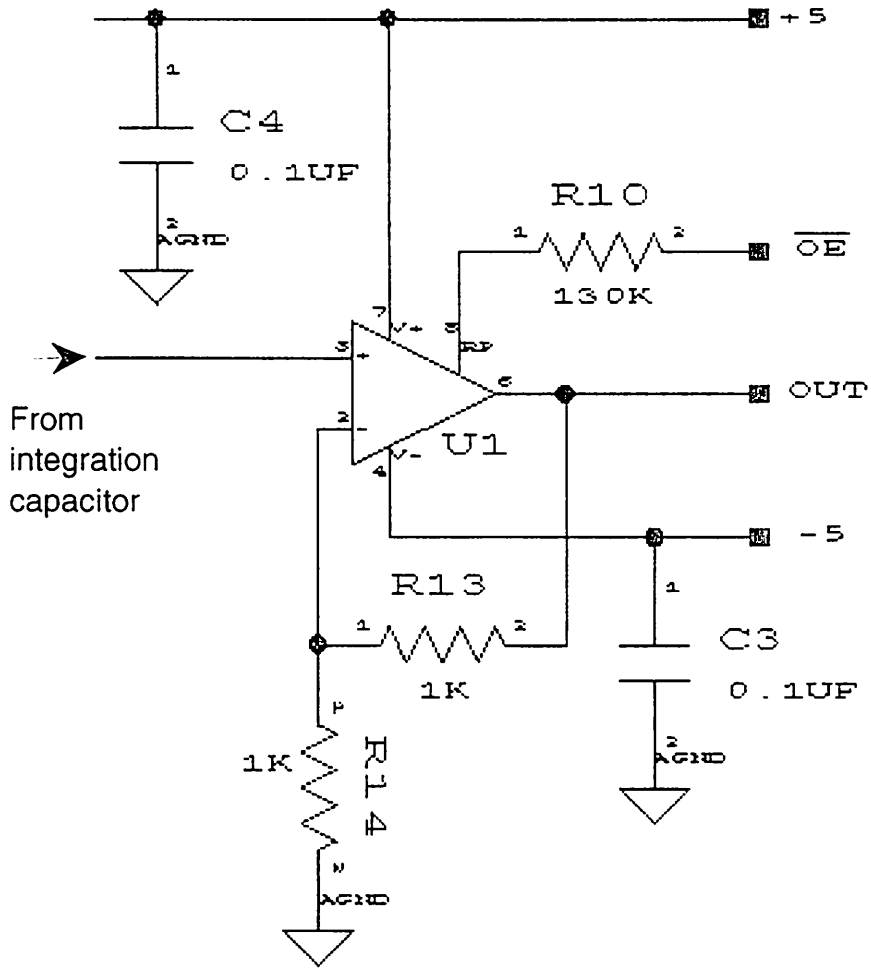

U1 cannot overdrive the ADC because its ±5V supplies are the same supplies that power the ADC. The amplifier is gated on only when the ADC is tracking its input, thus reducing power consumption and reducing integrator errors due to the load represented by the op amp's input.

Parts Selection

Noise and nonlinearity are minimized by choosing Q11 to have the maximum possible transconductance. Peculiarly, by this criterion the J106 is the best transistor I could find, even though it is primarily intended as a switch, and Q11 is an amplifier!

For symmetry, the key virtue for Q10 is low drain-gate capacitance, making a UHF amplifier transistor, the MMBTJ310, the best choice, even though Q10 is a switch! The reason is that the value of the -15V supply affects the video bias level directly via capacitive coupling through Q10.

For the bipolar transistors, the selection criteria are more conventional. Transistor base currents steal from the flows through the switches and the mirror, causing gain errors. Base currents are also noisy. Therefore, we choose high β transistors (sometimes designated as "low noise" amplifier transistors) to minimize the base current. Q4 is a saturated switch, but it need not be very fast, and a common base saturated switch tends to be fast even if the transistor isn't optimized for saturated switching, so for simplicity we use a high β, low noise transistor here too. Chosen transistors are the NPN BCX70K and the PNP BCX71K.

For U1, we choose the astonishing CLC505, a low power current feedback op amp. The CLC505 is much quieter than a low power FET input op amp, and much faster than a low power bipolar voltage feedback op amp. It may also be gated off for additional power savings.

U2 is the classic CA3080A OTA, a relic of the 60's: it's slow by modern standards, but good enough to do the job here (it is a curious feature of modern analog technology that OTA's are rather common within complex analog function IC's, but one has very few choices if one needs a "bare" OTA).

The integration capacitor C1 needs to have a low voltage coefficient, a low temperature coefficient, and a low hysteresis. A ceramic chip capacitor with an "NPO" or "C0G" characteristic should be good enough (fancy dielectrics like Teflon aren't needed), but don't use other sorts of ceramics.