

IC Op-Amps Through the Ages

1.0 Introduction

The operational amplifier concept emerged from extensive development of electronic analog computers in the 1940s. Operational amplifiers get their name from their ability to perform mathematical operations such as summation, integration and differentiation. With the addition of logarithmic amplifiers and feedback, one may perform multiplication and division as well. These abilities allow op-amp circuits to simulate differential equations, such as those describing the trajectory of an aircraft, for example. Countless analog computers were used throughout the second world war to produce “smarter” weapons capable of predicting where a target would be some time in the future, to determine the optimal firing conditions to ensure with high probability that a shell would intercept it then. The ability to “program” an analog computer rapidly by changing amplifier gains and connectivity made it an indispensable tool for both weaponry and academic research. Mechanical analog computers had preceded electronic ones, and changing gears and linkages made reprogramming a cumbersome affair.


The desire for compactness and universality led to the development of general-purpose high gain blocks intended for use within a feedback loop, the op-amp idea that is “obvious” today. Perhaps the most famous op-amp of the vacuum tube era was the Philbrick K2-W, whose particularly elegant design served as the inspiration for transistorized counterparts throughout the 1950s and 1960s.

Early efforts at integrated circuit op-amps attempted to replicate on a one-to-one basis the connectivity of earlier discrete designs. These pioneering efforts led to impractical results, for the reasons we’ve discussed numerous times: IC technology offers a different set of strengths and weaknesses, and these different rules (should) result in different design approaches. Early IC op-amps, practical or not, were first implemented in bipolar technologies, and it is in those technologies that many now-familiar topologies debuted. The collection of op-amp schematics here traces that history in order to gain an appreciation for why IC op-amp topologies are what they are. Even though these are mainly bipolar examples, a great many of these topologies are perfectly amenable to implementation in CMOS form as well, with appropriate modifications (e.g., for biasing, etc.). At the very least, these examples may serve as a valuable source of inspiration for new designs.

The dates that follow the part numbers correspond to the first year of publication of a paper or data sheet, or the first year of product shipment.

1.1 The Fairchild μ A702 and 702A/702C (1964)

The first op-amp designed by Bob Widlar (pronounced “wide-lar”) is the 702. Even this early design (using only nine transistors!) is an example of the highly creative thinking that was to characterize all of his later efforts:

FIGURE 1. μ A702A

This op-amp consists of 2.5 voltage gain stages (the “0.5” is to be explained later). The first stage is a differential amplifier (Q1/Q2, with bias current source Q7, slaved to master current source Q8). As in most differential designs, there is the problem of how to convert to a single-ended output without sacrificing 50% of the gain. Here the young Widlar solves this problem with a circuit that foreshadows his later use of mirror loads. Think of Q3 and R3 (which together form a simple common-emitter amplifier) as a cheesy op-amp within this op-amp. To the extent that this combination does behave as a reasonable approximation to a high-gain element, the voltage at the base of Q3 moves much less than that of the collector. As a useful approximation, assume that the base voltage of Q3 doesn’t move at all. Then we may exploit the “virtual ground” idea that is so useful in ideal op-amp analysis.

Most of the signal current from the collector of Q1 flows through resistor R1. Suppose that this collector current increases by some amount, so that the drop across R1 also increases. Since the base voltage of Q3 is assumed not to move much, an increase in the voltage across R1 implies that the voltage at the top of R1 must increase. Now, differential symmetry says that the decrease in Q2’s collector current is the same as the increase in Q1’s collector current. The voltage drop across R2 consequently diminishes, a change which aids the increase in voltage at the top of R1. Hence, both halves of the differential pair contribute an increase in the signal that ultimately drives Q4. Thus, a differential-to-single ended conversion has taken place, with no loss in gain (ideally).

The second gain stage involves transistor Q4, which is connected as a pure, textbook common-emitter stage. The emitter is actually grounded, and the resistive load R4 is as simple as it gets.

As with all purely-NPN designs, this one suffers from the common problem of “DC-level compression,” where the output common-mode voltage range of each successive stage of a cascade moves ever closer to VCC (because the voltage at each collector needs to be higher than that of the corresponding base). Widlar solves this problem by using a clever level shifter which is the functional equivalent of a battery, placed in series with a signal. This “battery” is implemented as a resistor in series with a current source. According to Professor Ohm, we thus may expect a voltage drop $V = IR$. Here, Q9 is the current source, and R5 is the resistor. Together, they synthesize the level shifting battery that subtracts a voltage from the output of emitter follower Q5. Widlar being Widlar, however, the level shifter is not quite as simple as that: the battery voltage is not constant, for reasons we’ll explain momentarily.

Transistor Q6 is a final buffer (emitter follower) to provide reasonable output drive capability. Examination of that subcircuit, however, reveals that the emitter load R6 does not return to VEE, as one would expect of a textbook implementation. Rather, it returns to a resistive network in the level-shifting current source Q9. Looking closely at the loop formed by Q6, R6/R10/R11, and Q9, it is apparent that Widlar has synthesized a *positive* feedback loop. The loop transmission of this positive feedback loop is controlled to a value less than unity to avoid instability. Nonetheless, it’s large enough to get some voltage gain out of the level shifter (this is the “0.5” part). Widlar apparently felt motivated to use this unusual trick because the two resistively loaded gain stages provide a worst-case gain of under 1000 by themselves. The final stage boosts this by another factor of 2.5 or so, enough so that the overall gain is guaranteed to exceed 1000 (a magic marketing number) by a healthy margin under worst-case conditions.

To improve the stability of circuits built with the 702, a series RC network can be tied between the base of Q6 and ground. The lag network thus formed reduces the op-amp gain at high frequencies by creating a voltage divider with the level shifting resistor.

An alternative compensation method that improves, rather than degrades, high-frequency performance is also possible, by shunting R5 with a small capacitance. The resulting lead network pushes out the pole associated with any capacitance connected to the bottom end of R5. Useful bandwidths of 20-30MHz are obtainable from this all-NPN op-amp built in a very early IC technology, a remarkable achievement that was not to be matched for another decade.

1.2 The Fairchild μ A709 (1965)

Despite its innovations, the 702 was not a great commercial success. Its initial retail price of approximately \$300 limited sales to military and aerospace consumers. The relatively low gain, plus serious limitations on the input common-mode range (forced in part by the grounding of the emitters of Q3 and Q4), as well as limited output drive capability, further constrained the part’s appeal.

Widlar went back to the drawing board and developed the first analog functional block that could be called a certified “hit.” The 709 op-amp was wildly popular because it provided much larger open-loop gains (order of 60,000), and had an input common-mode range that included positive voltages.

The 709 clearly shares a great deal with its progenitor, the 702 (see schematic). The first stage is a resistively-loaded differential stage (Q1/Q2, biased by current source Q14 which, in turn, is slaved to Q15), with a differential-to-single ended converter that is a slightly more sophisticated version of that used in the 702 (here, Q3/Q5 and Q8 together comprise the “op-amp within the op-amp”). Emitter follower Q5 is biased to a low current without the use of large value resistors by making the voltage across R3 depend on the *difference* between two diode voltages (those of Q5 and Q7). This clever trick reduces the area of the chip significantly (remember: one Widlar-derived rule in IC design is “minimize passives by replacing them with transistors wherever possible”). This same trick is used to synthesize a low current in Q14 without requiring absurdly large resistor values.

An additional biasing technique extends the common-mode input range over that of the 702. Rather than simply grounding the emitters of the next stage (which prevents the inputs from going positive by any significant amount), Widlar uses a common-mode negative feedback loop to bias the front end of this op-amp. The voltage at the common emitter connection of Q6 and Q7 is reflective of the common-mode output voltage of the first stage. A rise in this common-mode voltage would increase the current through the Q15-Q14 mirror, which would tend to bring the common-mode voltage back down. Hence, the input common-mode voltage is allowed to vary without interfering directly with the bias of the next stage. The op-amp inputs are thus allowed to possess common-mode voltages within about 5V of the positive supply.

The second gain stage is a resistively-loaded common-emitter amplifier (Q6) which is preceded by emitter follower Q4. Emitter follower Q9 participates in a downward level shift, in conjunction with common-base lateral PNP transistor Q11. The 709 is thus the first commercial product in which a lateral PNP transistor makes an appearance. Widlar’s design accommodates the dreadful characteristics of these early devices, which include a highly variable β that is nominally two. Widlar’s design continues to function even if β is as low as 0.2! (Even so, low yields plagued the manufacture of this part.)

To achieve the high open-loop gains demanded by users of op-amps, this design has a third gain stage, comprising Q12 as a resistively-loaded common-emitter amplifier. The output of this third gain stage drives a textbook complementary emitter follower. Thanks to the grounded collector configuration, the PNP transistor Q13 can be implemented as a *vertical* PNP device, whose characteristics (β and f_T) are a much better match to those of an NPN than is a lateral PNP.

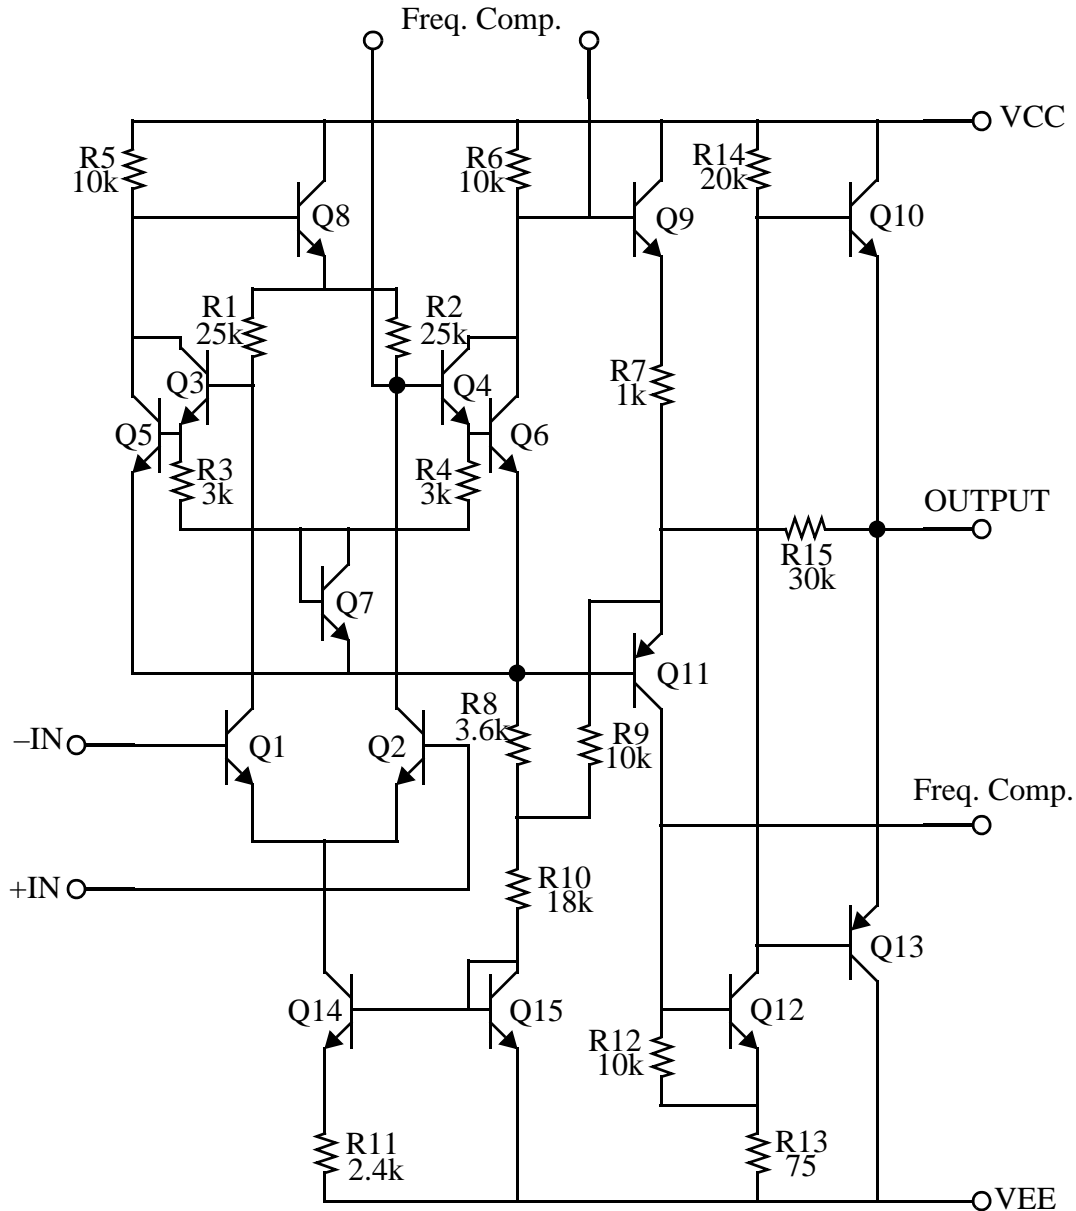
Unfortunately, a simple complementary follower possesses a “dead zone” in its input-output transfer characteristic. That is, the base voltage must be $\sim 0.7V$ above the desired output when the NPN is driving, or $\sim 0.7V$ below the output when the PNP is driving. There is thus a roughly 1.4V range of base voltages over which neither transistor is on. Widlar

employs negative feedback (through R15) to reduce the resulting distortion, although the diminished effectiveness of this loop at higher frequencies is definitely noticeable.

Using three gain stages produces both large open-loop gains and a stability challenge. To enable the user to improve phase margin, Widlar makes externally accessible essentially every high-impedance node in the amplifier. The user has the burden of designing, and connecting, a host of RC networks (many suggested by Fairchild in the 709's data sheet and applications notes) in an effort to achieve satisfactory stability, bandwidth and settling time. Sometimes, the desired results are even achieved.

The spectacular success of the 709 was associated with production demands high enough to cause rapid and steep price reductions (despite yields that were simply terrible for mysterious reasons). This op-amp, introduced at approximately \$70, was the first to break through the \$10 barrier (and, later, the \$5 barrier), guaranteeing extremely widespread use. It is also the first op-amp used (and destroyed) by the author. By 1969, op-amps were selling for around \$2.

The success of the 709 emboldened Widlar to request a significant enhancement in his compensation. When this request was denied by his boss, Charles Sporck, Widlar left Fairchild in 1966 to join a fledgling National Semiconductor. In an ironic twist, Sporck was to become president of National only a year later, becoming Widlar's boss once again. In his new role, Sporck had to accept Widlar's compensation package, one which permitted Bob to retire in 1970 (before his 30th birthday). Except for a brief period in the early 1980s when he worked with Linear Technology, Widlar continued to produce designs for National on a consulting basis for the rest of his life.

FIGURE 2. μ A709

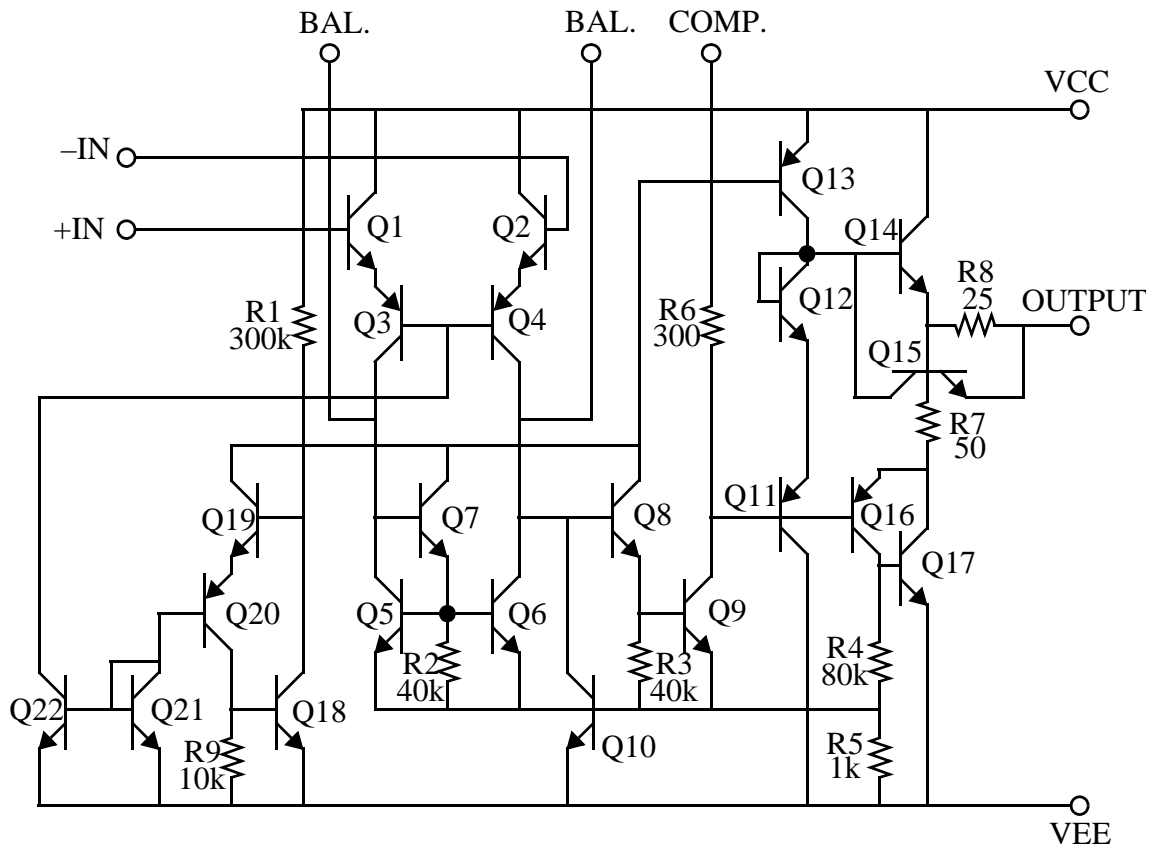
1.3 The LM101 (1967) and 101A (1968)¹

The first op-amp Widlar designed at National sought to repair several shortcomings of the 709. He wanted to improve further the input common-mode range, and make the part easier to use (requiring op-amp users to design compensators was asking a great deal, apparently). He also wanted to increase the open-loop gain to well over 100,000 under worst-

1. For these parts, National uses 1XX to denote military temperature range components, and 3XX for commercial temperature range parts. Thus an LM301 is characterized for operation from 0°C to 70°C.

case conditions. Finally, he wanted to protect the part against accidental output short-circuits, something that caused many 709s (like my first) to die prematurely.

FIGURE 3. LM301



To improve input common-mode range, Widlar had to develop a more satisfactory solution to the level shift problem. A battery-based shifter, after all, more than simply re-centers the common-mode level; it also necessarily subtracts from it. The traditional alternative to a battery is simply to alternate NPN- and PNP-based stages, where the downward common mode shift of the latter compensates for the upward shift of the former.

Such a level shift strategy, of course, presumes the availability of both NPN and PNP devices. However, inexpensive bipolar technologies are optimized for NPNs (because the mobility of electrons in silicon is roughly double that of holes). Semiconductor layers used to make NPNs are then simply reused to make PNP. Unfortunately the initial deficit of lower hole mobility is exacerbated by the highly suboptimal device design resulting from such a simple strategy. The consequent poor β and f_T make it a challenge to design with PNPs made this way.

The input stage is meant to provide the downward common-mode shifts that are normally associated with a PNP-based differential amplifier. However, direct use of PNP devices in a textbook differential amplifier configuration would result in unacceptably large input currents, owing to the poor β of lateral PNPs. Bandwidth would also suffer, from Miller

effect conspiring with poor f_T . Widlar's solution is remarkable: Connect the PNPs (Q3/Q4) in a differential common-*base* configuration (to eliminate Miller effect), and combine them with a pair of NPN-based emitter followers (Q1/Q2) to overcome the input-current problem. This connection does demand rather clever biasing to provide the base current for the PNPs, but it is a solvable problem, as we'll discuss shortly.

High gain, and conversion to single-ended mode, is enabled by mirror Q5 and Q6. Rather than tying Q5's collector directly to its base, as in a textbook mirror, the connection is performed through transistor Q7. This transistor is added simply to maintain symmetry, given that the base current of transistor Q8 tends to disrupt the equality of mirror currents.² By biasing Q7 to the same current as that of Q8, each mirror branch is disturbed by the same amount, eliminating any contribution to systematic offset.

The second gain stage is a current-source loaded, NPN-based common-emitter amplifier. Transistor Q8 is simply an emitter follower that drives Q9. The collectors of several transistors share a common connection with that of Q8 simply to save area (a single area-hungry n-epi layer can be shared, along with the base region of Q13). Transistor Q13 is a PNP current source which provides the load for this stage. This current is supplied to Q9 through two level-shifting devices, Q12 and Q11, which function to bias the complementary emitter follower (Q14/Q16/Q17) in a fashion that largely removes the dead zone of the 709.

Again because of the poor characteristics of PNPs, the output follower departs somewhat from standard textbook configurations: the pulldown device is a combination of an NPN and a PNP that mimics the basic polarities of a PNP. The overall effective β is the product of NPN and PNP β s, allowing the combination to possess the current drive characteristics of an NPN.

Widlar's engineering cleverness extends to the biasing circuitry as well. Transistors Q18 to Q22 control the biasing for the entire amplifier. The feedback loop formed by Q18/Q19/Q20 forces the voltage across R9 to equal about 0.7V (the base-emitter voltage of Q18). The current through Q19/Q20 is consequently 0.7V/10k Ω . This current is augmented by two 0.7V/40k Ω contributions by Q7 and Q8 to set the base current of PNP transistor Q13 equal to approximately 100 μ A. Assuming that Q11 and Q13 match reasonably well, their base currents are roughly equal and, consequently, the value of the current source load for Q9 is also about 100 μ A.

Setting the current through Q20 also sets the bias current for the input stage. The challenge here is that Widlar chooses to control the collector current of the input devices indirectly, by setting the base current of the lateral PNP devices. To avoid a consequent undesirable sensitivity on the value of β , the base current of PNP transistor Q20, whose collector current is controlled to a known and stable value, is measured and replicated by NPN mirror

2. In some mirrors, a follower is used to reduce current ratio errors due to base currents by reducing the error-inducing currents themselves. However, careful examination here reveals that Q7's base current is of the same magnitude as that of the mirror devices. Consequently, Q7's role is simply to provide a compensating error current.

Q21/Q22. The output of the mirror then drives the common base connection of the input lateral PNPs. To the extent that all PNP β s match, this clever circuit generates just the right β -dependent base current to produce a β -independent collector current.

A welcome refinement is a greatly improved ability to tolerate output short circuits to ground. Although marketing claims for the 709 included that its output stage could tolerate short circuits “for a time,” the 301 was the first that could tolerate shorts to ground indefinitely. This feat is accomplished by limiting the maximum output current to a sustainable value. Transistor Q15 is normally off, but if the op-amp attempts to source a current in excess of about 25mA, the drop across R8 causes Q15 to turn on. Doing so robs Q14 of base current, limiting further increases in output source current.

To protect the op-amp from sinking excessive current, it is not acceptable to pull up on the base of Q16 because Q9 itself can sink a substantial current, and the resulting fight between it and the protection device would inevitably resolve destructively. This problem does not arise in the sourcing direction because the protection device only has to overcome the limited drive of current source transistor Q13. To provide the desired protection, the sinking current is sampled through R5. When the voltage drop across R5 gets large enough to turn on Q10, current is stolen from the base of Q8, which ultimately limits the output sinking current to a safe value. The LM301 thus tolerates short circuits to ground for any length of time. However, it is important to observe that the op-amp is not necessarily protected against short circuits to the supplies.

Because of the current-source loading, two stages suffice to provide a nominal gain of over 500,000. The small number of stages simplifies frequency compensation, which is provided by connecting a suitable network between the collector of Q4 and the pin labeled “comp.” In most cases the compensator is as simple as a capacitor, to provide pole splitting Miller compensation.

One somewhat undesirable consequence of a constant collector current is a base current (which must be supplied by the external circuit) that varies with temperature, thanks to an increase in β with increasing temperature. In general, it is better for parameters to remain stable over temperature. A modified version of the 301, the 301A, achieves this stability by making the first stage bias currents more or less proportional to temperature. Transistors Q19-Q22 effectively form a thermometer to provide the desired behavior. The increasing collector current compensates the increase in β to yield a roughly constant input current. Furthermore, the transconductance of the input stage, which is proportional to $I_{bias}/(kT/q)$, also becomes much more temperature independent, stabilizing op-amp bandwidth.

A detailed analysis of the thermometer bias circuit is rather tedious, but a simplified version begins by neglecting the voltage drops across R4 and R9 (they’re small anyway). In that case, the voltage that appears across R1 is approximately

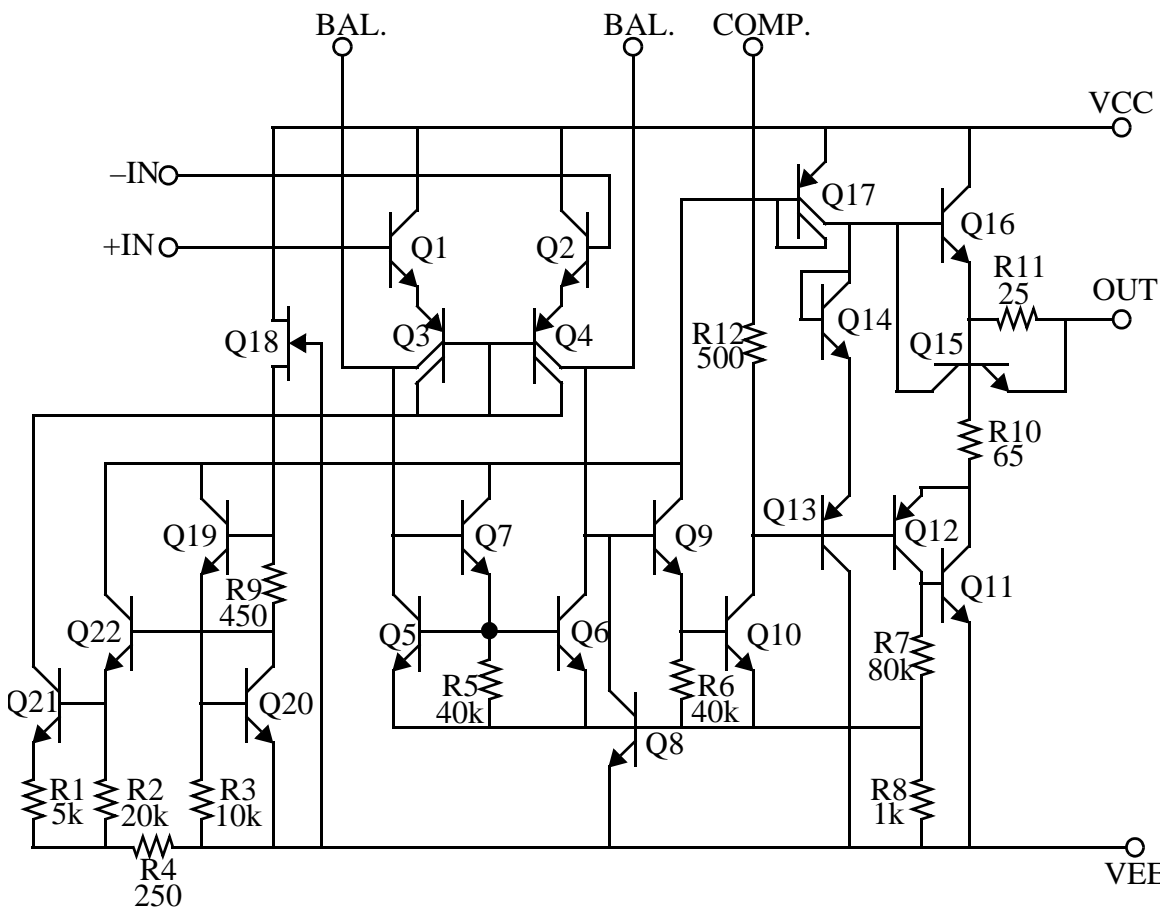
$$V_{R1} = (V_{BE19} + V_{BE20}) - (V_{BE21} + V_{BE22}) = \frac{kT}{q} \ln \left(\frac{I_{C19} I_{C20}}{I_{C21} I_{C22}} \right) . \quad (\text{EQ } 1)$$

That is, the voltage across R1 is the difference between a sum of diode voltages. Such a voltage is PTAT (proportional to absolute temperature), so the current through R1 would itself be PTAT if the resistance were stable over temperature.

Resistor R9 is added to provide some compensation for the variation in Q18's current with supply voltage. As the supply voltage increases, the current through Q18 increases. The voltage at the base of Q22 would consequently increase, causing an increase in amplifier bias currents. Inserting R9 provides an additional voltage drop that reduces the base voltage of Q22, thus offsetting the increase in Q18's current.

Because the resistors used in the circuit all vary with temperature, the bias current behavior is not precisely as desired. However Widlar is able to exploit the differing temperature coefficients of the various resistor options to kludge together a network with acceptable characteristics. In particular R4 is added to compensate for the non-uniform temperature coefficient of the other resistors, leading to the desired final result: fairly constant input current.

FIGURE 4. LM301A



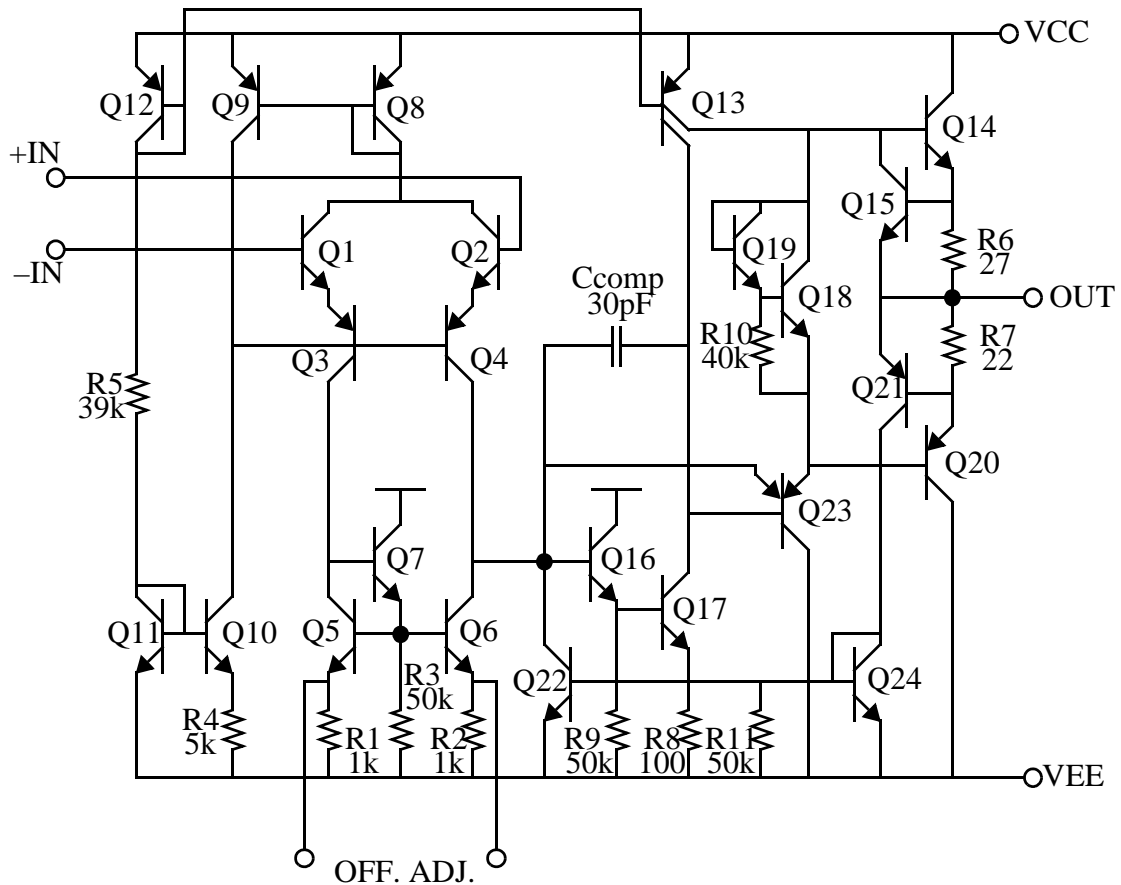
Another change, enabled by process improvements that produce higher PNP β , is simplified biasing of Q3/Q4 and Q17. The multiple collectors may be thought of as a shorthand notation for mirrors, where the current ratios are set by collector area ratios. Finally Q18,

drawn as a JFET, is a “pinch resistor,” in which the cross-sectional area of an n-epi resistor is reduced by a p-diffusion (the same as is used to make the base of an NPN) to increase the resistance for a given area. The p-diffusion “gate” is reverse-biased relative to the n-epi body, so that the depletion region further reduces (pinches) the cross-sectional area.

1.4 The μ A741 (1968)

Back at Fairchild, it was left to Dave Fullagar to carry on the work of designing op-amps. Fullagar rose to the challenge and developed the 741, the most popular op-amp of all time.

FIGURE 5. μ A741



Perhaps Fullagar might bristle at the suggestion that his design may have been inspired by those of Widlar, but it is certainly true that the basic architecture of the 741 is identical to that of the 301. The input stage is the same clever differential combination of NPN emitter follower and PNP common-base amplifier, while the second stage remains a current source loaded common-emitter amplifier. A textbook complementary emitter follower provides output drive that is limited to about the same maximum current as is the 741. The current limiting is also achieved in a similar way, with sourcing current limited by robbing base drive from the NPN follower, and sinking current limited by reaching all the way back to the output of the first stage.

Transistors Q18 and Q19 bias the output devices to reduce dead zone width, while an additional emitter on Q23 implements a variation on a Baker clamp that prevents Q17 from saturating. This consideration is important because bipolar devices do not turn off until injected minority carriers are removed. Preventing saturation reduces excess minority charge storage, and thus speeds overload recovery. Furthermore this circuit also limits the base current of Q23, providing additional protection against destruction.

The input stage uses feedback biasing to establish the base and collector currents. The $\sim 700\mu\text{A}$ master current is set by Q11, Q12 and R5. A Widlar mirror formed with Q10 and R4 reduces this master current to about $20\mu\text{A}$. A Wilson mirror is used to transfer this current to the input stage, while supplying the right PNP base current automatically.

Those minor variations aside, the chief appeal of the 741 is that it possesses a fixed internal compensation capacitor. Unlike the 301, the op-amp comes ready to use without requiring any external components. For simple resistive feedback connections, a simple Miller compensation capacitor is a good choice, so Fullagar's default is appropriate for those common situations. However, it should be clear that the optimum value of the compensating capacitor very much depends on the closed-loop gains one seeks. In particular, the compensation capacitor must be chosen large enough to handle the nominal worst case condition (100% feedback, as in a voltage follower). The closed-loop bandwidth will be unnecessarily reduced in other than unity gain follower connections. Despite a consequent near-guarantee of suboptimal performance for most applications, the ease of using the 741 has made it tremendously popular, proving Fullagar's assumption that engineers are basically lazy (I mean, very time-efficient).

Finally, as in the 301 and 301A, the 741 has provisions for nominal nulling of the offset. By connecting the main terminals of a potentiometer to the nodes marked "OFF. ADJ," and the wiper to VEE, a compensating imbalance in first-stage currents can be introduced. Unfortunately, zero offset and zero temperature drift rarely occur simultaneously.

After a subsequent tenure at Intersil, Fullagar went on to co-found Maxim Integrated Products, where he remains a vice president of R&D.

1.5 The LM110 (c. 1969-1970)

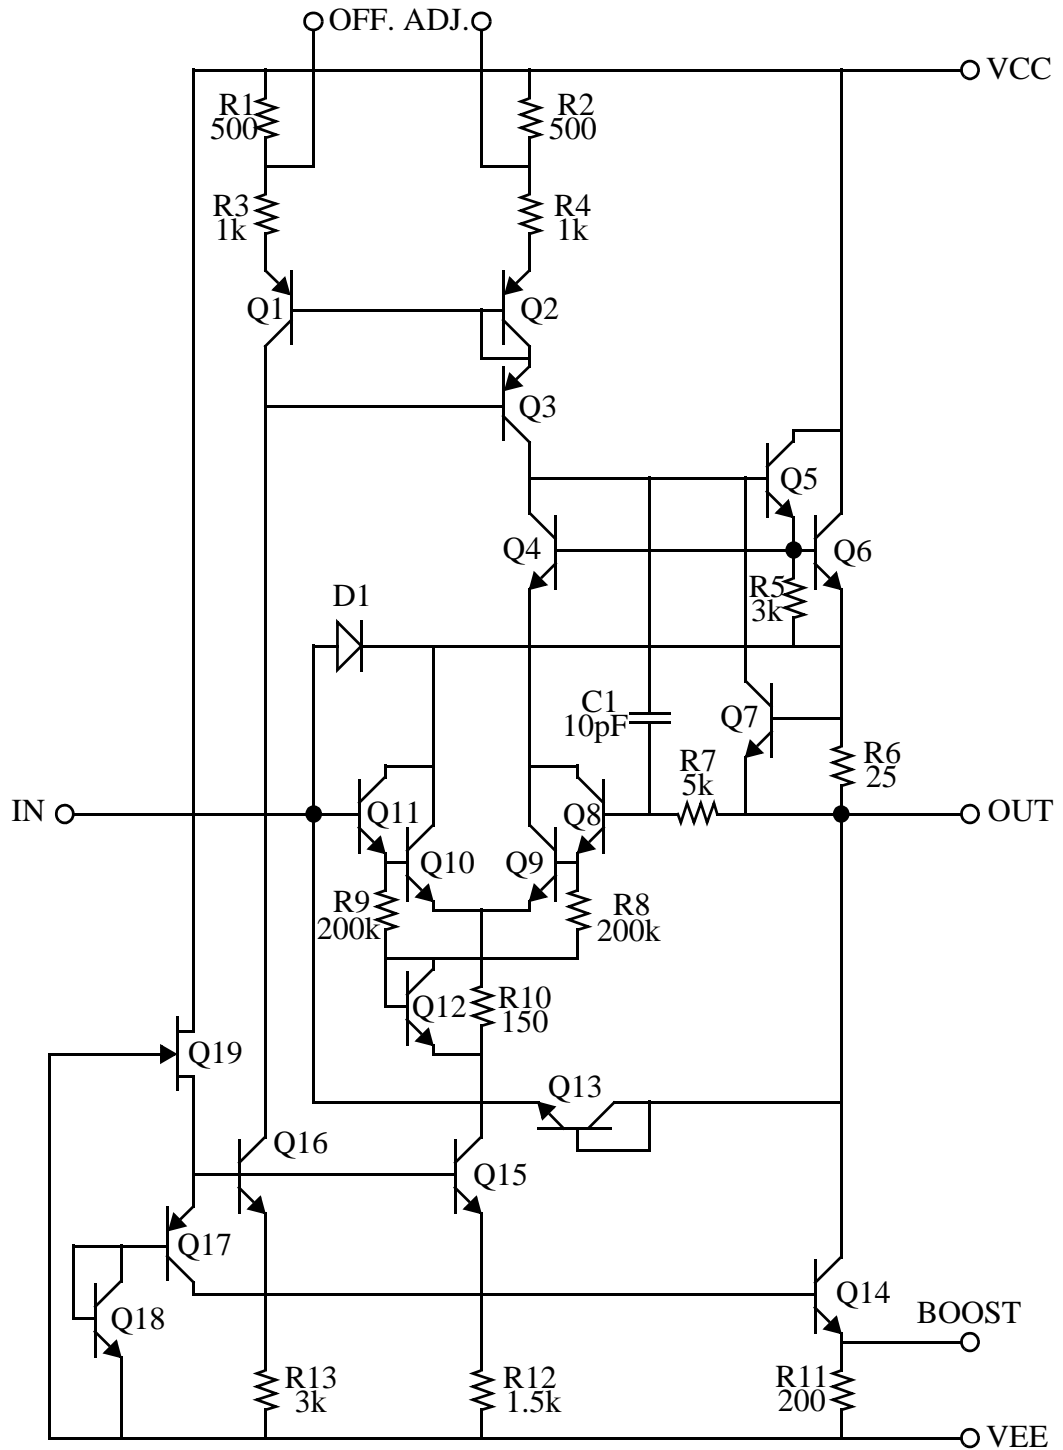
The circuits described so far are all general purpose op-amps. That is, they can be used in the standard catalog of op-amp applications, from a simple voltage follower to an integrator, to inverting and non-inverting amplifiers. The need to accommodate such a wide range of possibilities constrains design choices and necessarily involves numerous tradeoffs.

These tradeoffs can be relaxed considerably if the specific application is known a priori. As a specific example, suppose the op-amp is to be used only in a voltage follower connection. It may not be necessary to seek extremely high open-loop gains in such a case, affording a reduction in the number of stages, a consequent reduction in the number of poles and a simpler compensation problem. Bandwidth may then improve, and the result-

ing circuit may exhibit characteristics that are far superior to what conventional op-amps could provide.

The LM110, Widlar's second effort at a dedicated voltage follower after the very similar LM102 (designed in 1968, which evidently was a busy year for him), demonstrates that significant performance gains are indeed possible.

FIGURE 6. LM110



As one might expect, the LM110 is based on an op-amp, but with the feedback connection already provided internally. The op-amp at the heart of the LM110 is a single stage current-source loaded common emitter differential amplifier. The differential pair, Q8-Q11 consists of Darlington-connected devices to reduce base current (which, as always, must be supplied by the external source that drives the amplifier input). The actual input transis-

tors, Q8 and Q11, are provided with some bleed current through R8/R9 and Q12 (the three elements forming a Widlar mirror with Q9/Q10). This bleed current serves two valuable functions. One is to assure that the collector currents of Q8 and Q11 are not totally dependent on the base current of Q9 and Q10. This consideration is important because the base-emitter voltages, and hence the op-amp offset, is sensitive to collector-current mismatch. The other benefit of providing some bleed current is that the f_T of the input devices is boosted, facilitating the attainment of high stable bandwidths.

The process technology used in the 110 is somewhat modified to provide two types of NPN transistors. One is “normal,” and the other possesses exceedingly thin base regions (for these, a second emitter diffusion follows the first, resulting in greater encroachment of the emitter into the base region, shrinking the latter). As a result, NPN transistor β is boosted significantly (typical values are in the astounding range of several thousand). Regrettably, a side effect of thinning the base is a gross reduction in the breakdown voltage of the device. At relatively low collector voltages (e.g., 2V), the depletion layer can extend all the way across the base, resulting in a condition known as punchthrough. In this condition, the device current is no longer controlled by the base, and the transistor is easily damaged if the external circuit does not limit current to safe values.

To build a practical circuit with such devices requires extraordinary care. Widlar exercises this care by arranging cleverly for the thin-base (“superbeta”) transistors (Q8-Q11) never to see more than a diode drop from collector to emitter. For example, Widlar abandons any attempt at a sophisticated differential-to-single ended conversion (e.g., with a mirror load), instead choosing simply to throw away half of the gain. The collectors of transistors Q10 and Q11 would thus be tied off to VCC in a simple textbook implementation. However, doing so would result in the application of excessive voltage across those devices, precluding the use of superbeta devices. Recognize, however, that these input transistors merely need a source of DC current, and that this source need not be at VCC potential. Widlar simply connects the collectors to the output follower Q6. Since a follower should follow, the base voltage of Q11 should generally be close to the voltage at the emitter of Q6, resulting in a zero collector-base voltage for Q11 (and a single diode drop for the collector-base voltage of Q10). That is, the voltage from collector to base is bootstrapped to zero. If, on a transient or overload basis, this bootstrapping fails, Widlar provides diodes D1 (drawn as a diode) and Q13 (drawn as a transistor, in keeping with National’s own published schematics), which serve to guarantee that the input and output never differ by more than about 0.7V.

Transistors Q8 and Q9 need to be loaded by a current source (for high gain), so their collectors cannot be tied to the output, unfortunately. Widlar solves this problem by inserting Q4 in series with the current source load (provided by Wilson mirror Q1 through Q3). By tying the base of Q4 to that of output follower Q6, the emitter of Q4 is at about the same potential as the output. Thus, the nominal collector-base voltage of Q8 is also bootstrapped to zero, and we see that the four superbeta input devices are indeed protected against excessive voltage, independent of input voltage.

To protect against excessive load current, transistor Q7 samples the sourced current as the voltage dropped across R6. When the current exceeds about 25mA, the base drive current for Q5 is prevented from increasing any further by diversion through Q7.

Explicit protection in the sinking direction is not necessary because the output stage is a simple current-source loaded double follower (Q5 and Q6). The current source transistor Q14 provides all of the pull-down current (the latter may be adjusted by connecting an external resistor across R11), and thus inherently limits the sinking ability of the 110.

The bias current sources are arranged in a slightly unconventional manner to provide protection against other pathological conditions. Transistor Q19, which should be regarded as simply a high-value resistor (or, alternatively, a current source), acts in concert with master Q19 and mirror slaves Q14, Q15 and Q16 to bias the whole circuit. Transistor Q17 is normally operated in saturation, so that the base potential of Q15 and Q16 is a diode drop above the voltage across Q18. If Q14 should pull down so far that it saturates, Q17 drops out of saturation, leaving largely unperturbed the voltage at the bases of Q15 and Q16. If Q14's base were tied to that of Q15 and Q16, saturation of Q14 would result in the pulling down of those base voltages, upsetting the biasing for the entire amplifier. Recovery from that condition could be rather slow, degrading the large-signal dynamics considerably.

Finally, the amplifier is stabilized with a simple Miller network, consisting of R7 and C1.

Thanks to the use of a single all-NPN stage, this follower exhibits 20MHz bandwidth and a 30V/ μ s slew rate. This bandwidth is achieved along with a 1nA input current (two orders of magnitude lower than that of a 741), as a result of using superbeta devices.

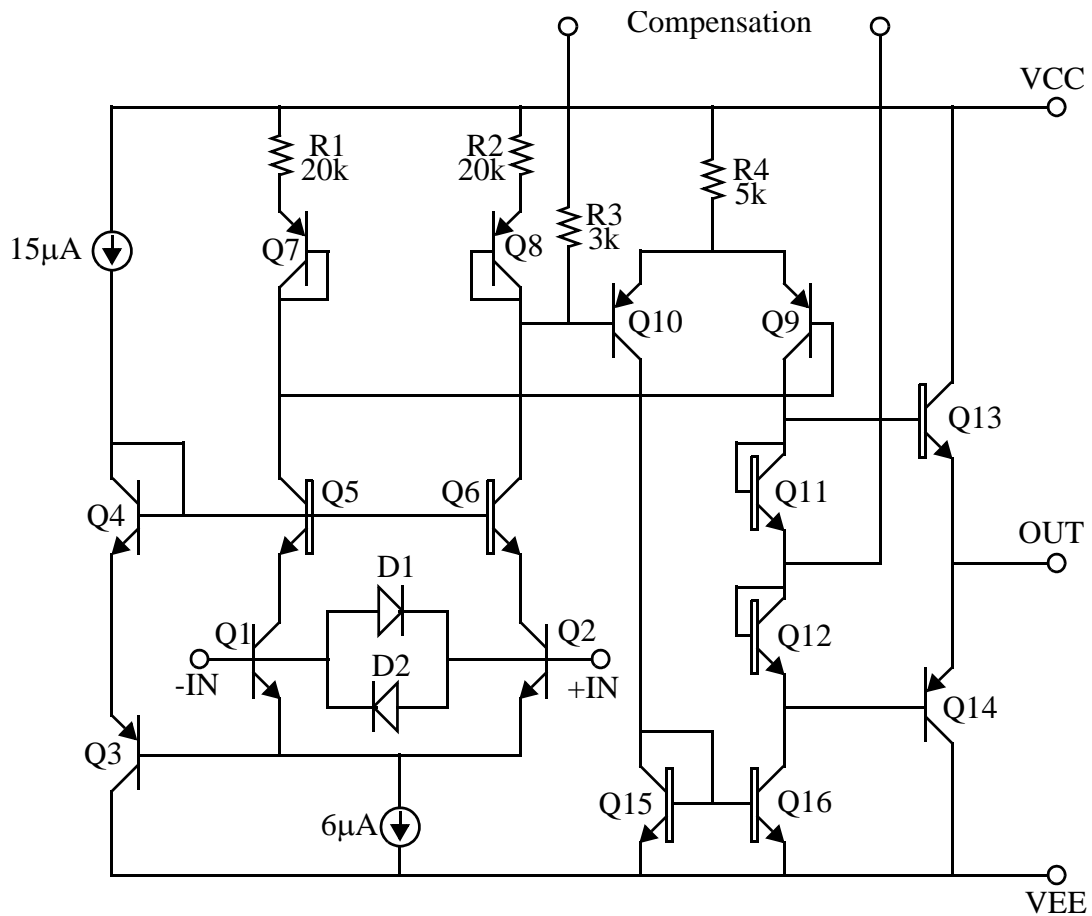
1.6 The LM108 (1969)

The zero nominal input-output voltage difference of a follower makes it relatively easy to use low breakdown voltage devices. The difficulty factor rises considerably, however, when the goal is to use superbeta transistors in a general purpose op-amp. The LM108, another Widlar creation, manages to pull off this minor miracle by expanding the use of the bootstrapping tricks found in the LM110. As a result, input currents remain about 1nA.

Shown in Figure 7 is a simplified schematic that preserves the essential features of this op-amp. Since the NPN transistors come in two flavors (ordinary and superbeta), the standard wide-base devices are drawn with the double lines.

This two-stage amplifier consists of a resistively-loaded differential NPN amplifier, followed by a mirror-loaded differential PNP stage. Superbeta input devices Q1 and Q2 are protected by wide-base bootstrapping transistors Q5 and Q6. The common emitter potential of Q1/Q2 is shifted upward a diode drop by PNP emitter follower Q3, and another diode drop by diode Q4. The nominal collector-emitter potential of the superbeta devices is thus maintained at a diode drop, independent of the common-mode potential of the inputs. Furthermore, protection against excessive differential voltages is provided by diodes D1 and D2.

FIGURE 7. LM108 (simplified)



1.7 The OP-07 (1975)

Although it is certainly true that Bob Widlar gets credit for a disproportionately large share of the analog IC innovations of the 1960s and 1970s, it would be terribly unfair to convey the impression that no one else was contributing to the development of the art. The OP-07 from Precision Monolithics (now a part of Analog Devices) shows that others were hard at work as well.

This three stage op-amp introduces two valuable ideas. One is the use of active base current cancellation, rather than reliance on superbeta input devices (and their attendant non-standard device processing). The other is post-fabrication trimming to reduce offsets by more than an order of magnitude over conventional approaches.

Current at the input terminals is reduced by well over an order of magnitude by measuring it, and then supplying it internally. Here, transistors Q3 and Q4 are not conventional cascodes at all. Rather, they are dummy devices whose sole purpose is to allow the measurement of base current. To the extent that the base currents of the “cascoding” transistors match those of the main input transistors (Q1/Q2), then the mirrors Q5/Q7 and Q8/Q6 will supply to the bases of Q1 and Q2 precisely the right amount of current. The external

sources driving the op-amp input terminals only have to supply the current resulting from incomplete cancellation.

The bias voltage for the dummy devices is made to track the input common-mode voltage by the use of the three series-connected diodes and a current source which drives the emitters of Q5 to Q8. Resistors R3/R4, acting with clamping diodes, protect the input stage against excessive differential voltages. This protection is needed because large enough differential voltages can cause base-emitter breakdown which, in turn, can cause shifts in β (as well as increase device noise at low frequencies). Since the bias-current cancellation depends on β matching, such shifts are much less tolerable here than in most cases.

The resistively loaded first stage contains numerous series-connected resistive segments, each having a reverse-biased junction in parallel with it. At wafer test time, the offset of the amplifier is measured, and an algorithm decides which resistive segment(s) should be shorted out to minimize the offset. Then, a large current is passed through the corresponding reverse-biased junction, causing the aluminum metallization to spike through the junction and short out the resistor in question. Although it may seem that this technique (called “zener-zapping” by its inventor, George Erdi who, by the way, had also worked at Fairchild for a time) couldn’t possibly be reliable, it allows the routine and robust attainment of sub-100 μ V offsets.

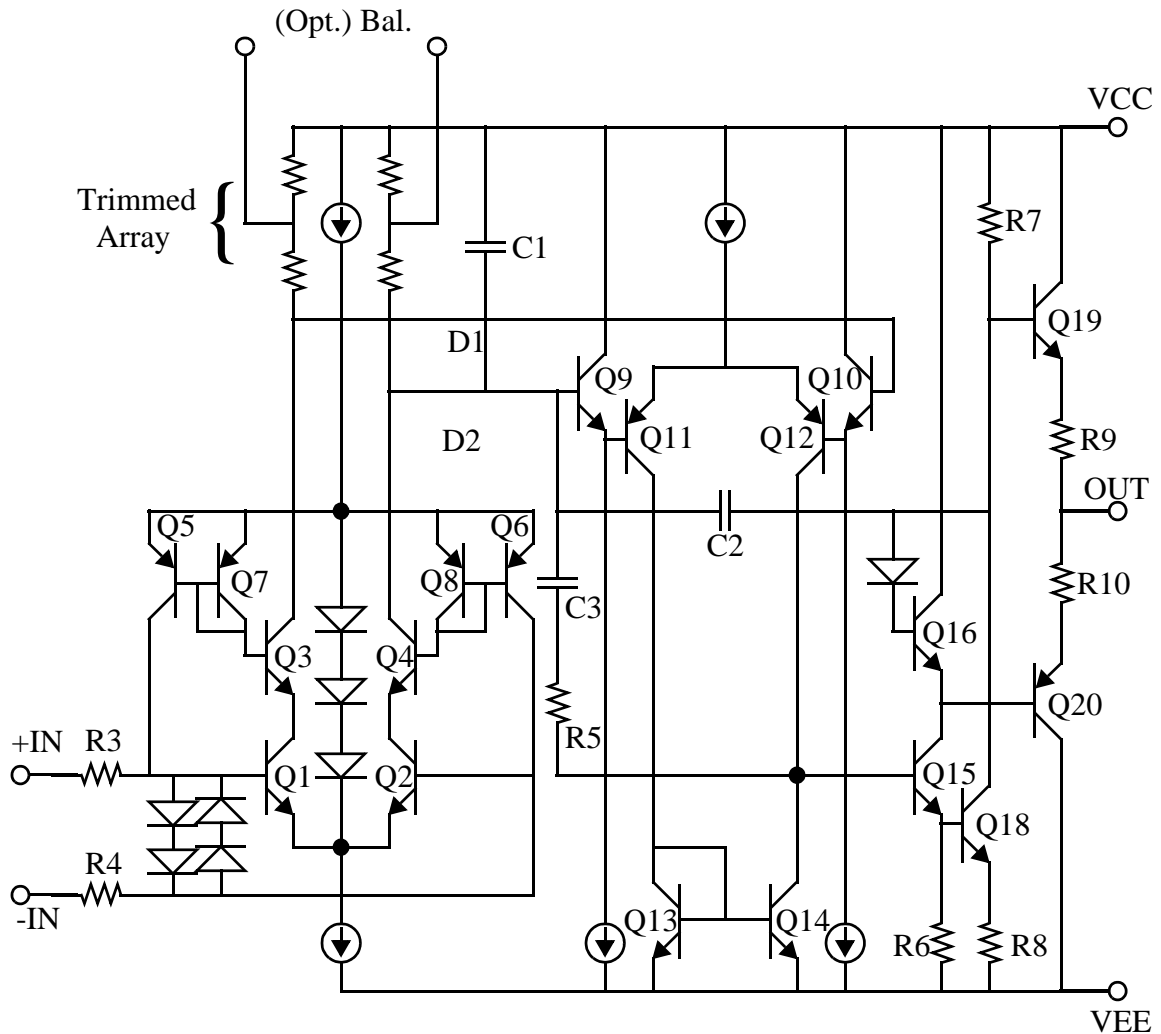
The second gain stage is a follower-driven PNP differential amplifier. Conversion to a single-ended output is performed the usual way, with an NPN mirror. That single output drives another resistively-loaded common emitter stage (Q18/R7). A standard complementary emitter follower (Q19/Q20) completes the op-amp. As in the 741, a substrate PNP is used for the output pulldown device, with its drive obtained through a two-diode drop level shifter implemented with a diode in series with the base of Q16.

To shape the gain-phase behavior of this op-amp, feedforward around the PNP stage is provided at high frequencies through R5 and C3. Miller compensation is provided around the last two stages with capacitor C2; this compensator generates the dominant pole for the op-amp. Finally, capacitor C1 rolls off half of the input stage’s contribution at high frequencies, implementing a gentle lag compensator (with an α of two).

Overall, the three stages confer a gain in excess of a million. In combination with the very small offsets, this amplifier is particularly well suited to high-precision applications. At the same time the dynamic performance is sufficiently similar to that of a 741 that the OP-07 may also be called a general purpose amplifier.

George Erdi left PMI in 1981 to co-found Linear Technology.

FIGURE 8. OP-07 (simplified)

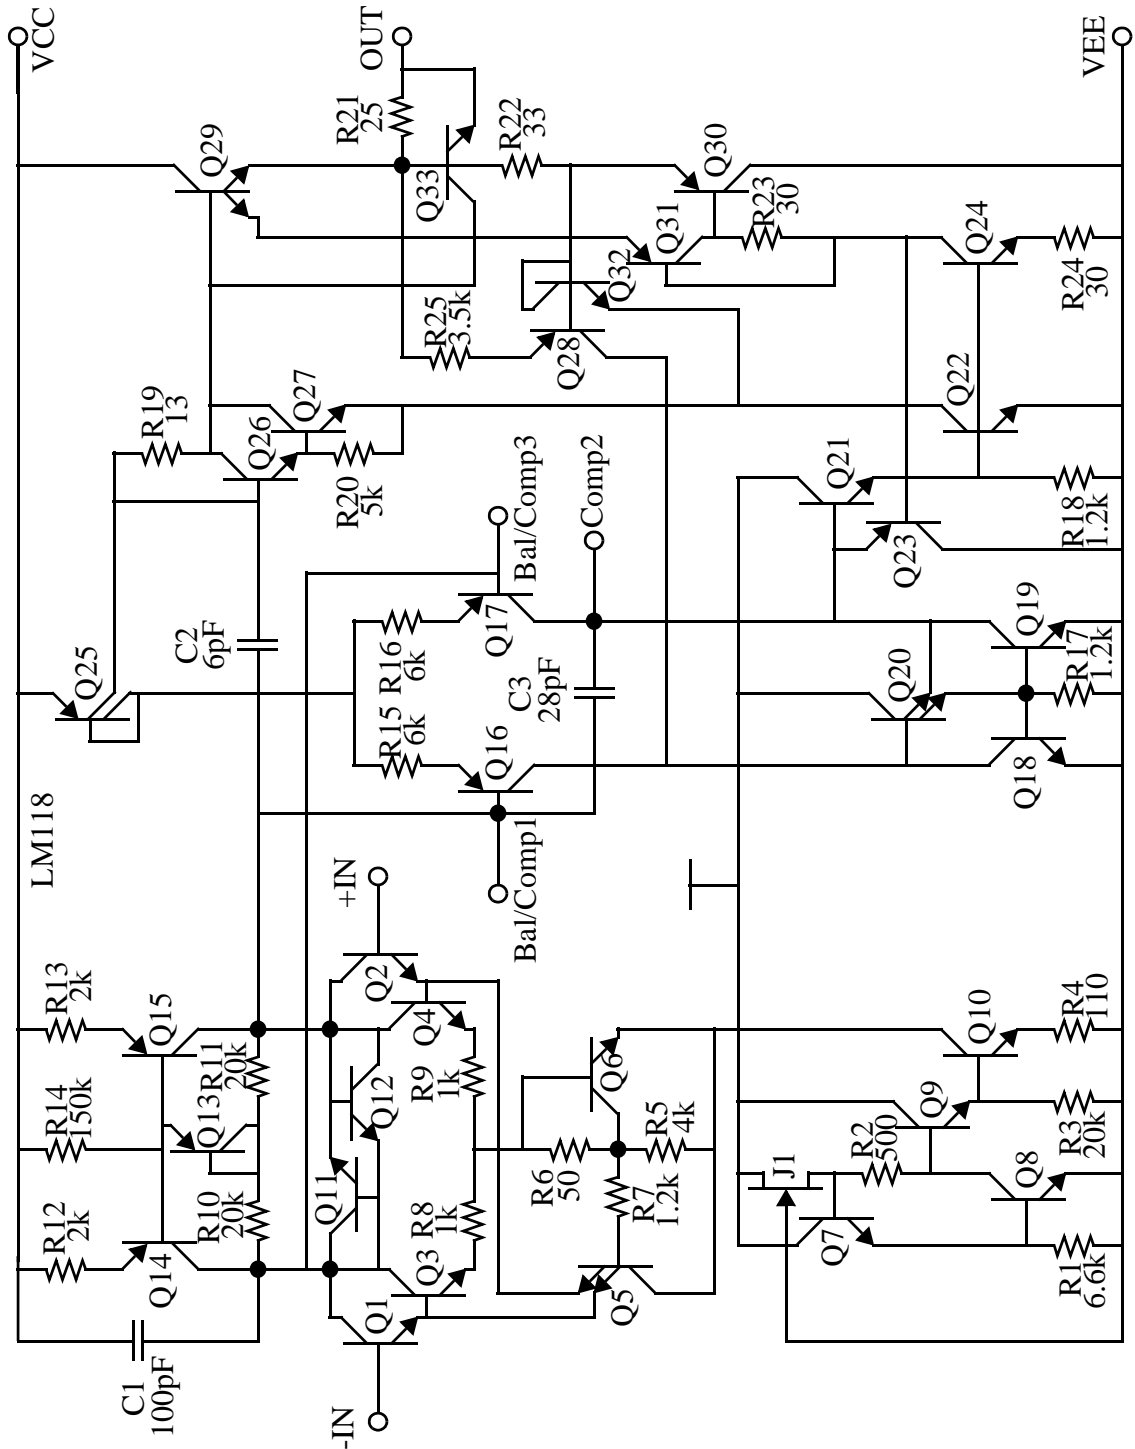


1.8 LM118 (1971)

While the OP-07 addresses the needs of the high-precision world, engineers are also congenial speed freaks. Improvements in process technology help, of course, but often the greatest speed gains are the result of clever topological choices. The LM318, designed by Linear Technology co-founder Bob Dobkin, proves this assertion. Built in a process technology not much different from that used to make the LM301A, the LM318 achieves 15MHz unity gain bandwidths and 50V/ μ s slew rates. That is, the bandwidth is improved by an order of magnitude, and the slew rate by two orders of magnitude.

Although the schematic may look somewhat intimidating at first, we may readily discern several important features. First is that the input stage is a differential NPN Darlington amplifier (Q1 through Q4). The second stage is a PNP differential amplifier (Q16/Q17), with a current source load for high gain and conversion to single ended output. There is a third gain stage (NPN Q22), loaded with a current source (Q25) for high gain.

FIGURE 9. LM118



The first stage is biased with a cell that is extremely similar to that used in the LM301A. Current source J1, as well as Q7 through Q10, generate a PTAT current through Q10. The bleed current for Q1 and Q2 is derived from Q5, which is operated with the roles of emitter and collector reversed. Although many device characteristics degrade with such a

reversal, the goal here is to minimize the common mode capacitance at the emitters of Q1 and Q2. Since the emitter area of Q5 can be made much smaller than that of the collector tub, the capacitance is correspondingly much smaller. The poor current gain and f_T of Q5 in the inverted mode are totally irrelevant here.

In any amplifier intended for high speed operation, it is important to consider both small- and large-signal performance. In particular, rapid recovery from overload conditions is highly desirable. Here, diodes Q11 and Q12 act as clamps on the output of the first stage to prevent more than a diode drop of differential swing.

The first stage is loaded with a current source (Q14/Q15) as far as common mode signals are concerned, but with resistors (R10/R11) for differential mode signals. This arrangement allows one to decouple bias point choices from gain considerations. Diode Q13 is another clamp which prevents the common mode output level from drifting too far from the common base voltage of Q14/Q15.

The second stage is a slow PNP stage (the PNPs are the same terrible lateral PNPs of the 301A), and is present only to provide sufficient DC gain. Feedforward capacitor C3 shorts out this stage at high frequencies so that the overall amplifier degenerates to two NPN stages below crossover. Unfortunately, use of a simple capacitor also guarantees that C3 provides positive feedback around the second stage. If that minor loop is itself unstable, nothing in the exterior main loop can fix it. Regrettably, a substantial fraction of LM318s display minor loop instability, providing a low-level output even with the inputs grounded. This problem could be avoided by placing an emitter follower in series with C3, to force signals to flow forward only. Conventional pole-splitting Miller compensation is provided between the output of the third stage and the output of the first stage through C2.

The differential-to-single ended converting mirror that loads the second stage contains a special transistor, Q20. Ignoring the emitter that is connected to the collector of Q19, it is clear that one function of Q20 is to provide a base current on the collector of Q19 that offsets the base current of Q21 on the collector of Q19 (note that R17 and R18 are equal in value). The second emitter of Q20 is normally reverse biased relative to the base of Q20. However, if Q19 ever approaches saturation, that second emitter begins to conduct, clamping the collector of Q19 at a potential roughly equal to Q19's base voltage. This clamping prevents saturation and a consequent slow return to normal operation.

Protection of the output stage from overloads is performed in ways similar to what we've seen for the 301A and the 741. In the sourcing direction, R21 and Q33 limit the output current once again to about 25mA. Further current increases are prevented by stealing base drive away from Q29. In the sinking direction, Q28 measures the current through R22. Limiting is accomplished by reaching back to the collector of Q16 and stealing its current away through Q28.

Again to speed recovery from overloads, Q23 acts as a clamp to prevent saturation of Q24. If Q24 approaches saturation, Q23 turns on, stealing base drive from Q21, thus preventing increases in the base drive of Q24.

A danger inherent in any complementary output follower is the potential of thermal runaway. A temperature rise increases collector current, which increases the dissipation, which raises the temperature, etc. If the gain around this positive electro-thermal loop exceeds unity, temperatures will increase until destruction occurs. Here, the base of Q30 is biased through diode Q31 to minimize crossover distortion. However, more than a simple level shift is provided: If the base current of Q30 increases, the drop across R23 increases, reducing the forward bias on Q30, and thus preventing thermal runaway.

1.9 HA-2539 (c. 1982)

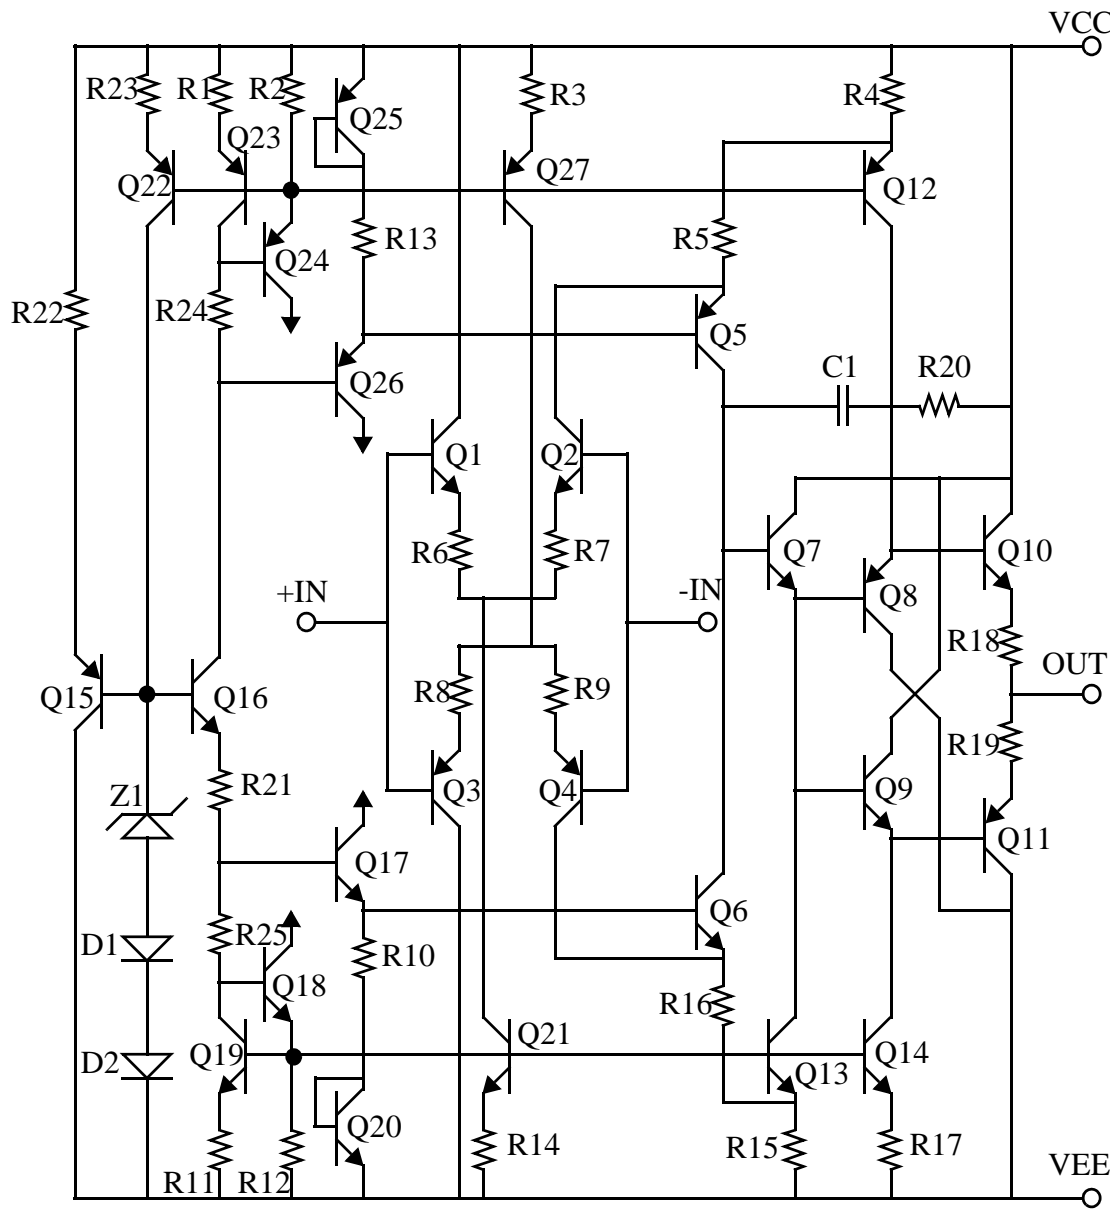
Many of the topologies we've seen are the result of attempts to work around the lack of good PNPs in an inexpensive process technology. Clearly, an alternative is simply to pay more for a complementary bipolar process technology to permit simpler circuit designs. Harris Semiconductor takes just such an approach in many of their designs, of which the HA-2539 high-speed op-amp is a typical example.

The circuitry on the left third of the schematic is just there to provide bias currents for the amplifier. The op-amp itself is a pair of folded-cascode amplifiers in parallel. One such amplifier is formed from differential pair Q1/Q2 driving PNP common-base Q5. The other is PNP differential pair Q3/Q4 driving common-base Q6. Each of these amplifiers throws away half of the gain; no effort is made to recover the lost gain.

The contributions of these two amplifiers are combined at the base of emitter follower Q7. A simple lag compensator (C1 and R20) is connected to this node. Finally two more follower stages complete the amplifier.

The simplicity of this op-amp permits exceedingly high bandwidths (600MHz in closed-loop gains of 10), and very high slew rates (600V/ μ s). However, the internal compensation is insufficient to assure stability in follower connections, and the single gain stage only provides a DC gain of about 15,000.

FIGURE 10. HA-2539



1.10 The LH0033

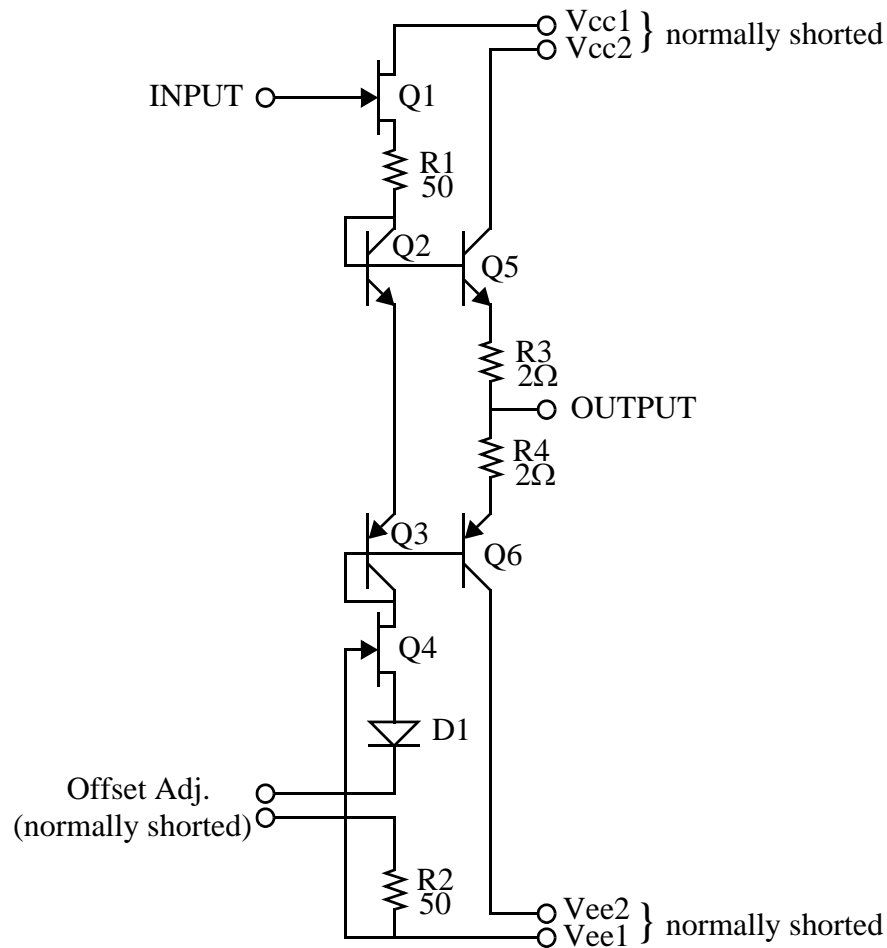
The designs we've studied up to this point are all intended for operation in closed-loop applications, and are also implemented as monolithic integrated circuits. Although a hybrid assembly rather than an integrated circuit, the LH0033 from National is worth examining because its performance underscores the attributes and tradeoffs associated with an open-loop buffer implementation.

This device provides a slew rate of $1500\text{V}/\mu\text{s}$ and bandwidths on the order of 100MHz . As seen from the schematic, the circuit is extremely simple: A JFET source follower drives a complementary emitter follower. Level shifting diodes Q2 and Q3 remove the dead zone

of the output stage, and the Q4/D1 combination produces a bias current for Q1 that results in a source voltage that is above the gate voltage by one diode drop, augmented by the IR drop across R2. Since R1 and R2 are of the same value, the voltage at the source of Q1 is just high enough so that, after subtraction of the drop across R1 and the diode drop from base to emitter of Q2, the output voltage is nominally the same as the input voltage.

This open-loop matching is not as precise as what one can expect of closed-loop designs (offsets are as bad as 25mV, as opposed to 1-5mV for typical monolithic implementations), but the high slew rate and high small-signal bandwidth more than compensate for many applications.

FIGURE 11. LH0033 “fast” buffer

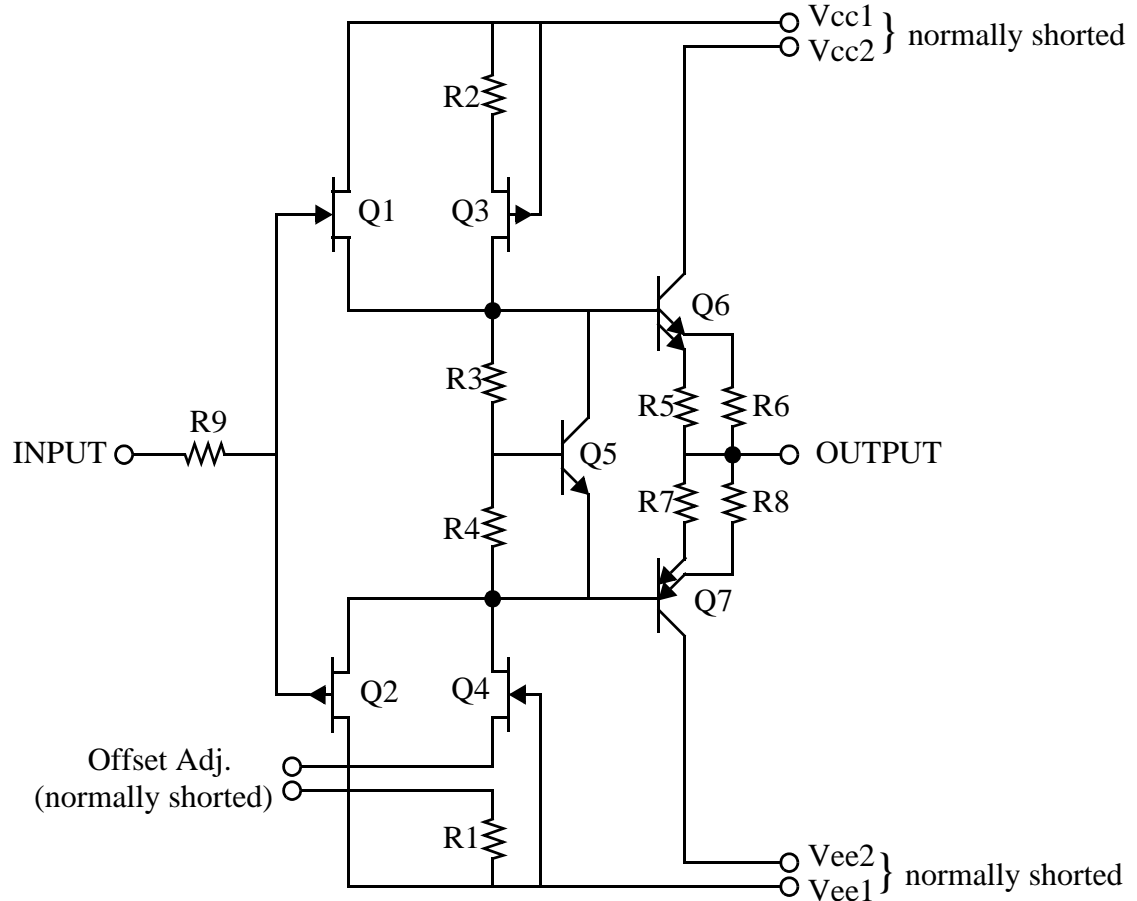


1.11 LH0063

Another hybrid buffer, the LH0063 is actually described in the data sheet as “damn fast.” Lest you think that this description is some obscure IEEE term, note that the slew rate is 6000V/ μ s. Clearly, the data sheet does not exaggerate: that’s damn fast! (Think of it this way: the slew rate is 6 volts per nanosecond!)

This speed is achieved through a complementary version of the LH0033. By providing both pull-down and pull-up devices, the buffer can support those remarkable slew rates and provide small-signal bandwidths in excess of 200MHz.

FIGURE 12. LH0063 “damn fast” buffer



The output dead zone is here removed with a clever circuit. The combination of R3, R4 and Q5 constitutes a “ V_{BE} multiplier.” To understand how this circuit works, note that the voltage across R4 is a diode drop. If we neglect base current, the voltage across R3 must be a factor $R3/R4$ times that across R4. Hence, n here equals $R3/R4$, and one may synthesize a voltage between the bases of the two output transistors that is any multiple of a diode drop. By selecting this multiple close to, but somewhat less than, two diode drops, the output stage can be biased to the verge of conduction (to remove the dead zone), but still avoid the possibility of thermal runaway (from too much forward bias).

This V_{BE} multiplier actually made its debut in the first version of the 741 op-amp. However, it can exhibit stability problems of its own, since it is actually a feedback system itself.

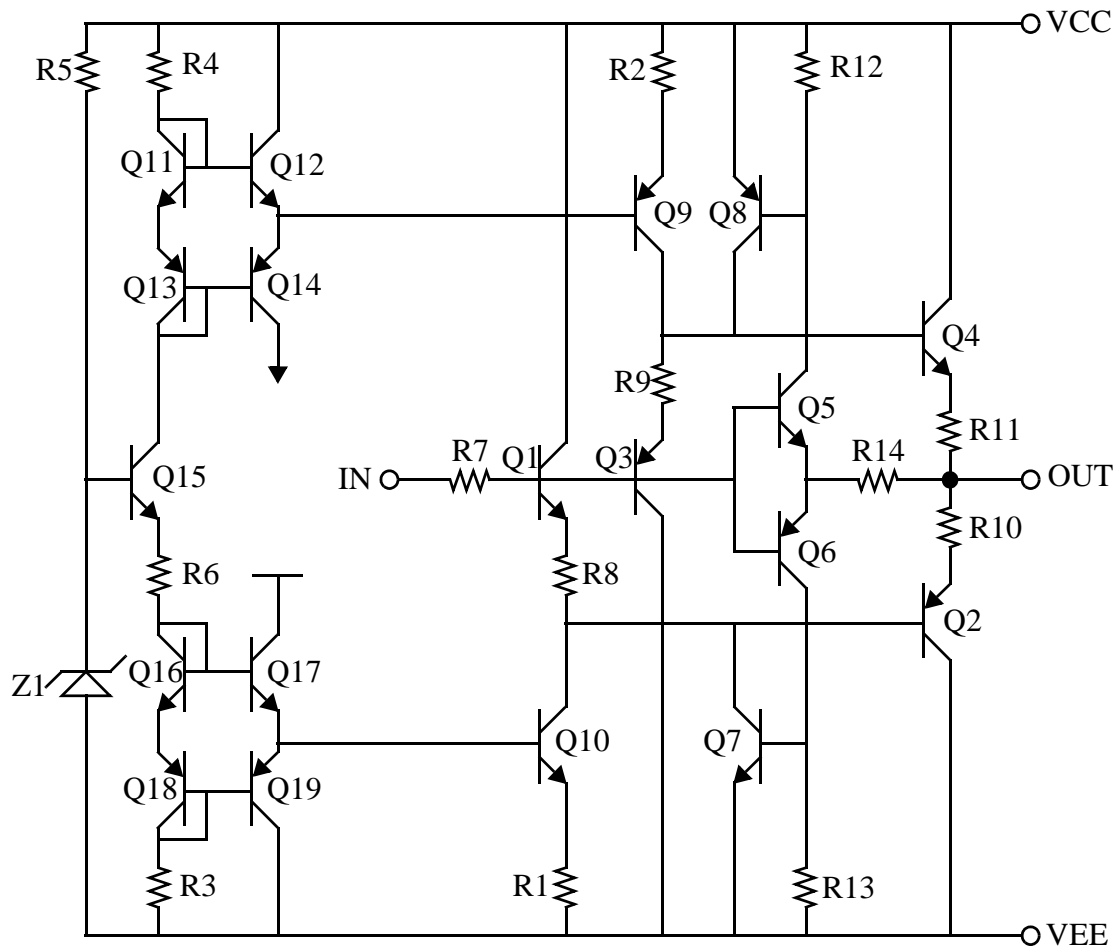
1.12 The HA-5033 (c. 1982)

Another approach to buffer design is to employ feedback on a transient basis. That is, the amplifier normally acts as an open-loop design (for speed that is unconstrained by stability concerns), but if input-output errors grow out of bounds, a feedback loop is temporarily activated to fix those errors.

The Harris HA-5033 is again made with the same complementary bipolar process used to make the HA-2539 op-amp described earlier. As in that device, there are two paralleled amplifiers here. One consists of Q1 and Q2 (whose level shifts nominally cancel), and the other of Q3 and Q4. Current sources Q10 and Q9 bias Q1 and Q3 respectively.

If the input-output differential exceeds a diode drop (indicating poor following), either Q5 or Q6 turns on. Assume for the moment that it is Q5 that turns on, as the result of an input that exceeds the output. As Q5 conducts, it induces a drop across R12 which, in turn, exponentially turns on Q8. Turning on Q8 provides additional current above that provided by Q9, enhancing the pulling up of the base of Q4 to reduce rapidly the initial error.

FIGURE 13. HA-5033



We see, then, that the HA-5033 contains the electronic equivalent of a turbocharger, where the part consumes a relatively small amount of power under most conditions, but is capable of temporary boosts in performance as needed.

1.13 BUF-03 (1979)

George Erdi, of OP-07 fame, also designed an open-loop buffer whose performance rivals that of many closed-loop designs, thanks to careful attention to, and mitigation of, error sources, in combination with a small internal feedback loop.

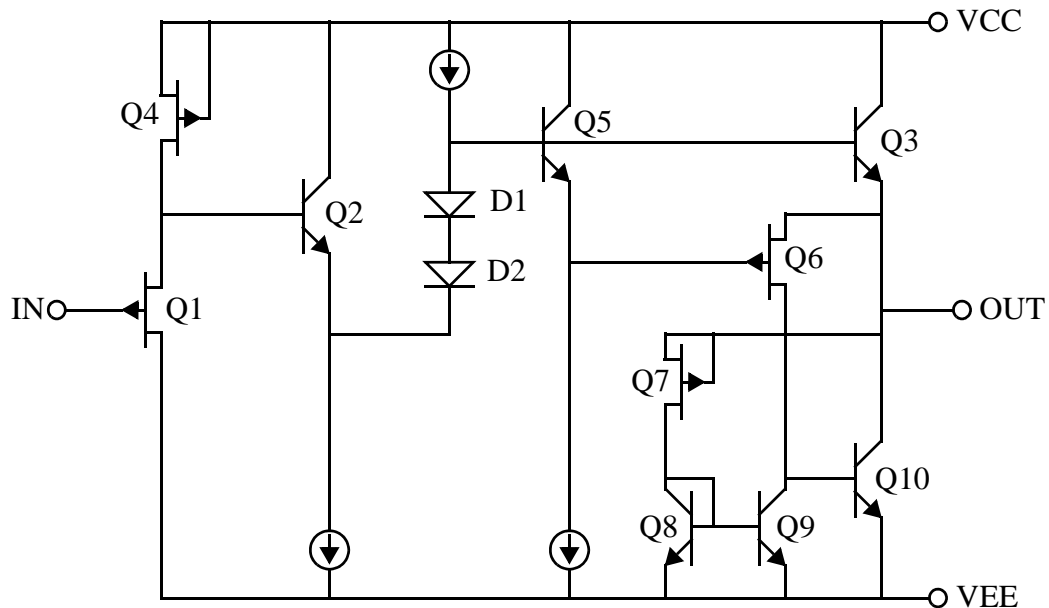
As can be seen in the (highly) simplified schematic, the design uses P-channel JFETs to reduce input current, and also to act as substitutes for PNPs. Clearly, providing JFETs requires some process modifications. In this case, the improvements in performance justified the modest increase in process cost.

Transistor Q4 is biased with zero gate-source voltage, and thus biases Q1 to zero gate-source voltage as well (remember: JFETs are generally depletion-mode devices, meaning that they continue to conduct current even at zero gate-source voltage). The base voltage of Q2 is consequently the same as the input voltage. Emitter follower Q2 shifts this signal down one diode drop, then D1/D2 shifts things back up two diode drops. The voltage at the emitters of the followers Q5 and Q3 is consequently the same as the input voltage.

One deficiency of a purely open-loop design is that increases in load current inevitably cause output voltage drops (it is impractical to create zero output impedance followers). To reduce load-dependent errors, Erdi cleverly builds an internal reference follower (Q5) whose output voltage is not load dependent. The output voltage of this reference follower is compared to the output voltage of the overall buffer through JFET Q6. Note that the base current of Q10 is the difference of the current sourced by Q6 and that sunk by Q9, and that the current through Q9 (Q8) is the current of P-channel JFET Q7 with its gate tied to its source. If the gate-source voltage of Q6 is zero (indicating zero error), no net base current exists to drive Q10. However, if the emitter voltage of Q3 rises above that of Q5 (e.g., because of load current changes), the drain current of Q6 increases, supplying base drive to Q10, thereby pulling the output back down. This local feedback provides two important benefits. One is the reduction in error, and the other is that it allows the use of an NPN device for pulling down as well as for pulling up. Hence, no PNPs of any kind are needed in the output buffer at all.

Thanks to this architecture, the buffer exhibits a 60MHz small-signal bandwidth and a slew rate of 300V/ μ s. Use of trimming (“zener zapping” once again) reduces nominal offsets to a quite respectable 2mV.

FIGURE 14. BUF-03 (simplified)



1.14 The LF155/156/157

The BUF-03 shows that P-channel JFETs can enable the elimination of PNPs without resorting to the expense of a fully complementary bipolar process. This idea is carried forward in the design of a complete op-amp, National Semiconductor's LF155 family.

This two stage design employs a current-source loaded P-channel JFET differential amplifier (J1/J2). Depletion-mode transistors conduct current at zero gate-source potential, and this characteristic is exploited in current sources J10/J11. Conversion to single-ended signals is deferred to the second stage, achieved simply by throwing away half of the gain.

The purely differential nature of the first stage produces a problem: how does one control the common-mode output voltage? That is, how do we guarantee, under all conditions, that the tail current equals *precisely* (not just to an engineering approximation) the sum of the currents in J10 and J11? The answer is: measure, and control, the consequences of a common-mode drift. Here, note that the voltage at the common emitter point of Q7 and Q8 is indicative of the common mode output voltage of the first stage. If this voltage rises, Q12 conducts more current, robbing some tail current away from the input pair, allowing J10/J11 to pull the common mode voltage back down. This feedback loop also greatly improves common-mode rejection.

Overall bias is provided by current-source J4, which sets the current through the other collector segments of Q1, as well as the current through the master parts of mirrors, D4 and D5.

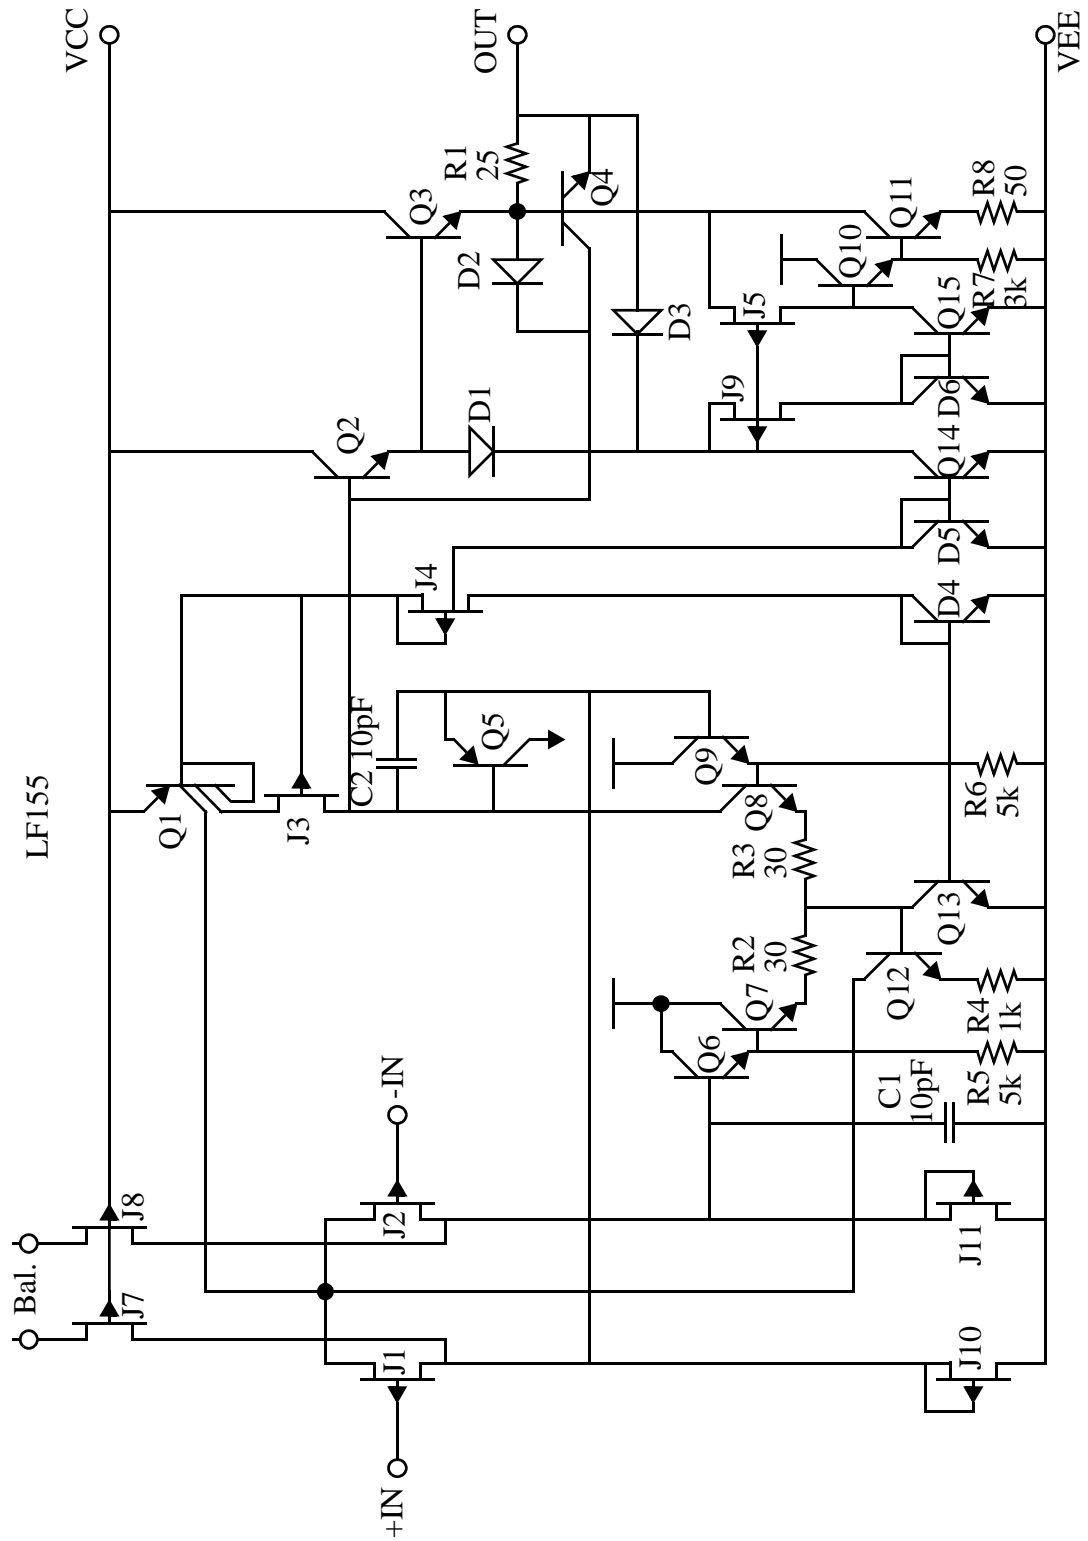
Current-source loaded Darlington differential pair Q6/Q6-Q8/Q9 constitute the second stage. Cascoding device J3 boosts the output impedance of Q1 to increase gain, while C2 provides the traditional pole-splitting Miller compensation. This capacitance ranges from

10pF for the LF155, to 2pF for the LF157. Device Q5, which is normally reverse-biased, provides protection against saturation of Q8 as we've seen in several other designs. As in all of those other cases, avoiding saturation speeds up recovery from overload conditions.

Buffer Q2 further drives output follower Q3. As in the BUF-03, the output stage here involves a local negative feedback loop, in which the output of a reference internal follower (Q2/D1) is compared with the actual amplifier output, and the drive for pulldown device Q10/Q11 generated appropriately. High accuracy is not the goal here, since global feedback around the amplifier takes care of that problem. Rather, the goal is simply to generate the correct drive for the NPN pulldown transistors.

Thanks to the elimination of troublesome PNPs, this op-amp provides $5\text{V}/\mu\text{s}$ slew rate and 2.5MHz bandwidth for the LF155, and $50\text{V}/\mu\text{s}$ and 20MHz bandwidth for the LF157. The former is stable in the follower configuration, while the latter is stable only for closed-loop gains in excess of five.

FIGURE 15. LF155

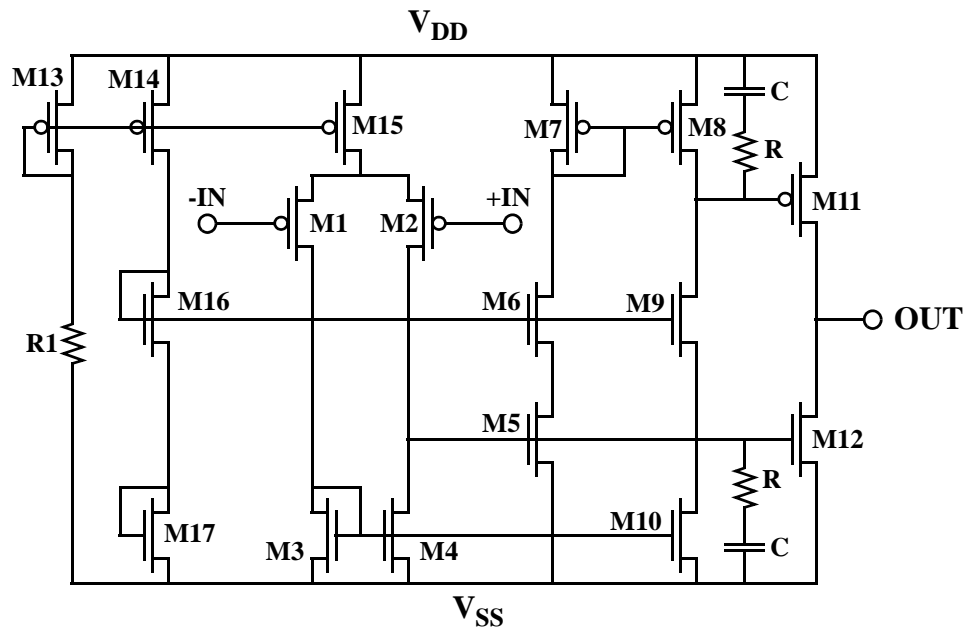


1.15 CMOS op-amp example

Starting in the early to mid-1970s, analog applications for CMOS technology gained increasing attention, initially at universities. Ultimately, that research made a transition into commercial products, as CMOS matured as a technology, and as engineers understood better how to recast analog design in a form amenable to practical realization in CMOS form. In particular, it took a fair amount of time for most engineers to understand that CMOS op-amps are most frequently used as subcircuits within larger systems, meaning that the load to be driven is usually well known. The design can then be tailored to that known load (indeed, the design can often exploit that load as part of frequency compensation).

To use CMOS in analog circuits requires mitigation of several problems, including low g_m/I , poor matching, and low output resistance. The relatively low gain per stage has inspired numerous solutions, including the gain-boosted cascode (“supercascode”) connection. Another alternative is simply to cascade more gain stages, but at the cost of having to solve a more difficult compensation problem. As an example of how one might proceed along these lines nonetheless, consider the following CMOS op-amp:

FIGURE 16. CMOS op-amp example³



Here, M13-M15 establish the bias currents for the amplifier, while the stack of M16 and M17 provide a bias voltage for common-gate cascode devices M6 and M9.

The overall topology is somewhat unusual in that it *sums* a cascade of three gain stages (for high gain), with a cascade of two gain stages (for stability, as will be explained

3. T. H. Lee and C. Cole, MIT, 1985.

shortly). The three gain stages include the first differential pair (whose main output is found at the drain of M4), followed by current source-loaded common-source amplifier M8 (driven through a voltage inverter formed by M5-M7), and finally followed by common-source amplifier M11. We may assume that there are three important poles associated with these three gain stages so, without compensation, this op-amp would be nearly useless in any feedback context.

In addition to the three-gain path described in the foregoing paragraph, this op-amp provides additional gain paths between input and output. For example, note that the input differential stage has a low-gain output (at the drain of diode-connected M3). This output drives a common-source amplifier (M10) which, in turn, drives one more common-source amplifier M11. Because there are two gain stages of significance, we may assume that there are two corresponding poles of significance. Note that the contribution of this two-gain stage cascade sums with that of the three-gain stage path.

There is yet another two-gain stage amplifier buried within this op-amp. Note that the main output of the input differential pair also drives the gate of M12. Finally, the output current of M12 sums with that of M11. In total, then, there are two two-stage amplifiers summed with one three-stage amplifier. The reason that this observation is significant is that a two-pole amplifier is easier to stabilize than a three-stage one. By summing a two-stage and three-stage amplifier, we produce an overall response that is dominated at high frequencies by the two-stage path (because the three-stage amplifier has rolled off by then). Stated another way, we enjoy the benefits of a three-stage amplifier's high gain at low frequencies, and a two-stage amplifier's simpler dynamics at high frequencies. It's a straightforward graphical exercise to see that the summation of three-pole and two-pole transfer functions produces a zero at a frequency where the two gains are equal.

Now, although a two-stage amplifier has simpler dynamics, they're still not simple enough. To guarantee sufficient phase margin, we still have to do additional work. In this case, lag compensators are sufficient. Here, series RC networks load the gates of the output transistors. For best high-frequency power supply rejection, each network is returned to the supply voltage that maximizes the filtering of noise coupled from the supply to the gate-source port of the corresponding transistor. For example, the PMOS device (M11) has a network tied across its gate-to-source. If the network were returned to V_{SS} instead, noise on V_{DD} would directly modulate M11's gate voltage, producing a noisy output current. Thus, not all incremental grounds are equal, because true grounds don't exist in nature.

Thanks to the cascade of three stages, this design can provide one or two orders of magnitude higher gain than conventional two-stage designs. Thanks to the summation with a lower-order amplifier buried within it, the stability properties are dominated by the more easily manipulated dynamics of two-stage amplifiers.

As a closing observation, the stability of a feedback system using this op-amp depends on the load capacitance (e.g.). Thus, this architecture is not a general-purpose one, for it exploits the particular circumstances of many CMOS analog ICs.

1.16 Current Feedback Amplifiers

The op-amps we've presented to this point all share certain characteristics. Among these are a very high input impedance (both differential- and common-mode), implying voltage-mode operation. As a consequence, capacitance is an extremely important limitation on bandwidth. Op-amps of this type have dominated the amplifier world since the very first vacuum tube implementations, and generally exhibit constant gain-bandwidth products. As a result, an increase in closed-loop gain must be paid for by a decrease in closed-loop bandwidth. Furthermore, the need to charge up node capacitances with limited currents produces the well-known slew limitations that vex conventional op-amps.

If high impedance is not strictly required (indeed, if it is purposefully abandoned), one may evade the tyranny of constant gain-bandwidth and, to a remarkable extent, limited slew rate. One class of amplifiers that exploits this idea is the so-called current feedback amplifier, a simplified schematic of which is provided in Figure 17. Although the basic topology actually traces back to 1930s-era vacuum tube circuits (including HP's famous Model 200 oscillator), a widespread appreciation of the attributes (and limitations) of this topology had to await the 1980s, when companies such as Comlinear (now a part of National) and élantec popularized the circuit.

As can be seen, the noninverting input is high impedance, but the inverting one is not. In fact, the inverting input is actually the *output* of an emitter follower, whose input in turn is the noninverting terminal of the amplifier. Any differential input voltage causes a current flow (and an exponentially sensitive one at that). This current flows through the inverting input, and its mirrored version causes a voltage to be developed at the high impedance node (collectors of Q6/Q9). Thus, we may say that the output voltage is proportional to the current that flows through the inverting input terminal. This observation is the key to understanding how to use and analyze circuits using these types of amplifiers.

The voltage developed at the high impedance node is buffered by a two-stage complementary emitter follower built out of Q11-Q14. Capacitor C1 is a compensation capacitor that works by creating a (hopefully) dominant pole at the high impedance node.

Practical closed-loop configurations with such amplifiers look very similar to their voltage-mode op-amp counterparts. For example, inverting amplifiers are formed with a feedback resistor R_f , grounded noninverting input, and a series resistor R_i between the input source and the inverting terminal. Since the output voltage is proportional to the current through the minus terminal, we can characterize the proportionality by an impedance Z_{21} . After a little work, we may therefore write

$$\frac{V_o}{V_i} = \frac{-R_f/R_i}{1 + R_f/Z_{21}}, \quad (\text{EQ 2})$$

where we have idealized the input emitter followers as perfect (zero input current, zero output impedance). As Z_{21} approaches infinity (the analogous condition to infinite gain in a conventional op-amp), the closed-loop gain approaches the same answer as for ordinary op-amp inverting amplifiers. Note that the input resistor does not appear in the denomina-

tor (which controls the poles). Thus, only the feedback resistor controls the dynamics, and one may adjust the closed-loop gain over wide limits without seriously affecting overall bandwidth.

Notable also is the absence of a biasing current source for the input stage. Unlike a conventional differential stage, the output current of the input stage is not actually bounded (ideally speaking; of course there are real bounds in practice); there is therefore no slew limit. That is, there is essentially no difference between large and small signal behavior (again, ideally speaking). Current-mode op-amps have become widely used in applications such as video line drivers because of these desirable characteristics.

FIGURE 17. Current-feedback amplifier (simplified)

