Electric Circuit Analysis, KTH EI1120

N. Taylor

Topic 05: Opamps

1 Introduction

This is a whole Topic dedicated to a single new component, the *operational amplifier (opamp)* (sv: *operationsförstärkare*). For circuit analysis this is shown as a component with at least *three* terminals. We will use the following three-terminal symbol. The + and − inputs may sometimes be drawn the other way up.

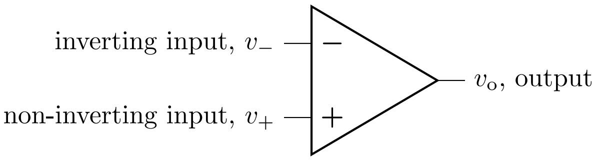

The properties of an *ideal opamp* are shown in the following diagram, using a model based around a voltage-controlled voltage source (VCVS).

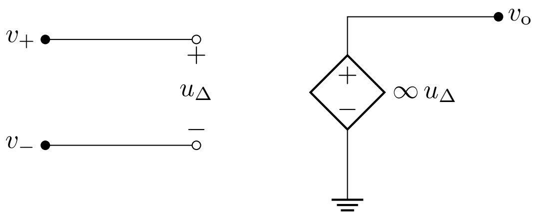

This model is a common way to think about the ideal opamp or to approximate it in a circuit simulation program (with ∞ replaced by a large value such as 10^9). The factor v_o/u_Δ, by which the output is bigger than the input, is called the *gain*.[1]

The sign of the output potential depends on which of the inputs has higher potential: if $v_+ > v_-$ then v_o shoots off to a huge positive value; if $v_+ < v_-$ then v_o shoots off to a huge negative value.

In practices and exams we will consider only ideal opamps ... just as we've assumed ideality for all our earlier components. However, to avoid reasoning about 'infinite gain' we might give some explanations based on opamps with nonidealities such as finite but high gain, and limited range of output; these cases can be easier to discuss in a practical way, and to relate to the lab work. The Extra Section at the end gives a lot more detail about physical implementation and nonidealities.

[1] The term *open-loop gain* is a more specific description, making clear that the gain applies to the opamp itself, from inputs to output. Later (Section 2.3) we will see that common opamp circuits add other components around the opamp, resulting in a closed loop of feedback from output to input. The complete system of the opamp and these other components can then be designed to have a *closed-loop gain* that could be small or even less than 1, depending on the component values.

We can summarise that an ideal opamp is:

> a *differential voltage amplifier*, meaning that it looks at the *difference* between the + and − inputs[2], not at the potentials relative to e.g. the opamp's power supply;
>
> with *infinite gain*, meaning the limit where the tiniest hint of a difference between the input potentials causes a huge change in output potential;
>
> and with *infinite input-resistance*, meaning that no current goes into the inputs, but they just *measure* the potentials v_- and v_+, as is suggested in the above model where u_Δ is across an open-circuit;
>
> and *zero output resistance* meaning that the output behaves like an ideal voltage source — its potential depends on the inputs, and the output then supplies any current that is then required by the circuit connected to it;
>
> and *unlimited range and speed*, meaning that there is no limit to the range of possible output and input voltages, nor to the speed at which the output can respond to the input.

Some of the ways in which practical opamps differ from the 'ideal' are given in the Extra section of this document. The most immediately obvious in the lab is that the output voltage has upper and lower limits.

1.1 What's the point?

Many interesting and complicated electronic devices are available that, like an opamp, are made up of several transistors and other components in a single package. So why do we give special attention to the opamp, by including it with much more fundamental types of component in a basic Linear Circuits course?

The reason is that opamps are very versatile, and therefore very widely used. By connecting different groups of simple components around an opamp, we can make an amplifier with a gain of e.g. 0.2, −274 or 10^4, or an analog adder or integrator or differentiator. Special selections of feedback even allow the opamp to behave as a close-to-ideal current source.

Although opamps may not seem directly related to electric power, they are used as amplifiers, integrators, comparators and so on, in applications such as instrumentation and control in power equipment other physical systems.

[2] The normal opamp does not, however, have a differential *output*; the output is a potential, i.e. relative to some fixed node such as a power supply ground.

2 Modes of Operation

We consider now three distinct behaviours of an opamp, defined by how its output potential responds to a voltage signal applied to one input. What defines the differences is how the connections *around* the opamp *cause the opamp output to affect the opamp inputs.*

[If you only care about the basics that get used in calculations in this course, you can skip to "Negative Feedback", which is commonly *assumed* to be the situation in circuit-analysis questions with opamps. The earlier parts are more for understanding of what opamps do, and of other functions that opamps perform.]

2.1 Comparator

When the output does not affect the inputs, the opamp is just amplifying the difference between the two input potentials that we choose to give it.

If these potentials are equal, the output will be zero (as long as $\infty \cdot 0$ doesn't get us scared ... just treat the gain as 'high but finite'). If the opamp has a very high gain, then even a small difference between the inputs will cause the output to become a very high potential (if $v_+ > v_-$) or very low [negative] potential (if $v_- > v_+$). For an 'ideal' opamp, this might sound rather dangerous, or at least not very useful.

A *real* opamp has very serious limits to its output, such as $\pm 10\,\mathrm{V}$, but *can* have a high gain such as 10^6. When a real opamp is used as a comparator, its output will therefore be practically always at its positive or negative limit; very seldom, with realistic signals, will the difference between the inputs be so close to zero that the output is *not* at one limit or the other.

The opamp in this case is working as a *comparator*. It compares the two inputs and reports a very clear answer — of a strong positive or negative potential — about which input has the higher potential.

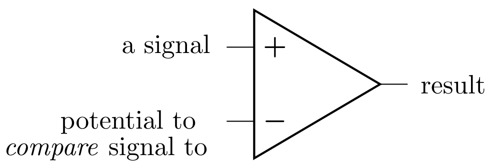

This comparator behaviour can be useful for turning analog signals to binary outputs like 'yes' or 'no'. For example, it could compare an analog signal to a reference, as a part of the decision-making about whether a particular 'bit' should be high or low when converting to a digital signal. Or it could turn on an alarm bell when a signal goes beyond a set level.

A comparator circuit with this output limit is an approximation of the [Sign function], but with the sign function's ± 1 range scaled to $\pm v_{\mathrm{o,max}}$. The following plot shows how the output could look when a high-gain opamp with output voltage limits of about ± 1 is comparing two signals that vary with time. The axis labels with 'a.u.' denote 'arbitrary units'. No time-delay of the opamp output is considered. The parameter called slew-rate describes how quickly a real opamp's output can change.

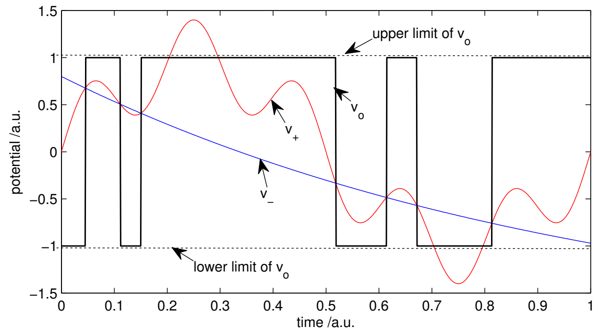

One common way to use a comparator is with a fixed reference, such as the voltage U in the following diagram.

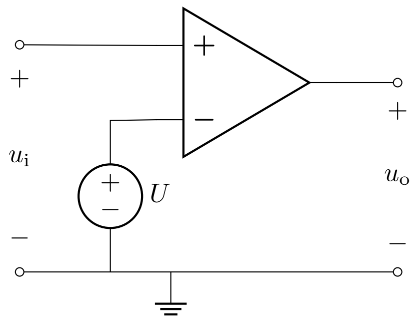

This is the principle for deciding which coloured lamps will be turned on in the cable-finder (kabelsökare) that has for many years been the Elkretsanalys lab task near the end of the course.

2.2 Positive Feedback

If the circuit around an opamp includes a connection that can make the potential of the non-inverting input rise or fall when the opamp's output respectively rises or falls, then the opamp has *positive feedback*. A crude example is the following.

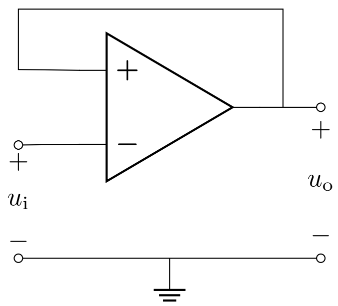

This has some similarity to a comparator, but the tendency to go all the way to the upper or lower output limit is made even stronger. For example, as soon as v_- becomes at all lower than v_+, the opamp's output will start to increase, which due to the feedback will make v_+ increase, and so the difference $v_+ - v_-$ becomes even greater.

For a real opamp with positive feedback, any difference between the input u_i and output u_o will cause the output to move quickly to one of its limits. Because feedback will hold the non-inverting input v_+ to this output voltage limit, the only way to make the output change is to make the input go *beyond* this limit. For example, if the output is at its upper limit, such as 1 V, the only way to change it is to make the input become $>$ 1 V, in which case the output will move quickly to its lower limit such as -1 V. Then it will not change again unless the input goes below -1 V!

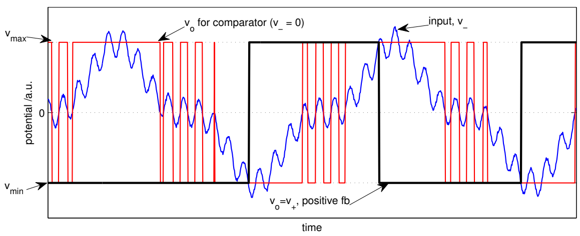

This behaviour can be described as a comparator with *hysteresis*, which means that when it has made a choice it needs an extra push to change its mind back again. In the above example, the blue curve shows an input signal, the red curve shows the output of a comparator that compares this to zero, and the thick black curve shows the case where the signal goes into the inverting input while the non-inverting input gets full positive feedback from the output.

The comparator responds to the wiggly input with lots of changing between its output limits. Positive feedback avoids the changes except at extremes of input. By putting a voltage divider between the output and input, the circuit could be made so that only, say, 1% of the output potential appeared at the non-inverting input. Then the circuit would be quite similar to a comparator that compares the inverting input to ground, except that there would be a small hyteresis effect that would stop the comparator responding to very small changes in the input.

2.3 Negative Feedback

It is the situation of *negative feedback* that tames a wild comparator or positive feedback circuit, to a well-controlled linear amplifier with moderate gain.

The trick is just the opposite of positive feedback: one must ensure that the potential of the *inverting* input will rise or fall when the opamp's output respectively rises or falls.

In this way, if v_+ becomes higher than v_-, so the output starts changing towards a higher potential, then the feedback will cause the potential v_- to increase as well. If the opamp's open-loop gain is high, this change will only stop when $v_- \simeq v_+$. The negative feedback has thus ensured that the output will adjust to whatever value is needed to force v_- to stay equal to v_+.

3 Examples with negative feedback

The following are some simple, classic circuits built around opamps with negative feedback. The condition $v_- = v_+$ will be assumed, due to the negative feedback.

(With real opamps it is wise to check that the output voltage would not exceed the limits imposed by maximum voltage, maximum speed of change of voltage, and maximum current. If the output voltage becomes limited by these nonidealities, then the opamp will not manage to enforce the condition $v_- = v_+$.)

3.1 Buffer amplifier (voltage follower)

This is the simplest case of negative feedback, where the inverting input is held to the output potential. The output u_o will therefore adjust to follow the input u_i.

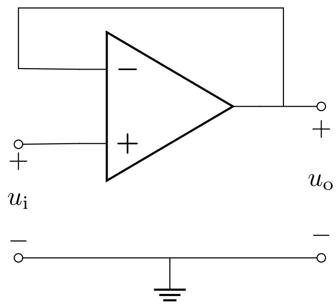

What could be better about this circuit than about a simple piece of wire, which also gives $u_\mathrm{o} = u_\mathrm{i}$?

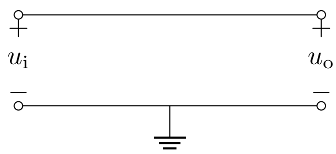

The possible advantage of the opamp buffer, compared to the wire, is that the input of the opamp takes ideally no current from the circuit that it is connected to, and the output of the opamp can supply any current that an external circuit tries to take from it. This is clearly very different from a piece of wire, for which all the current into the output must come from the thing that is connected to its input.

The buffer therefore allows a very delicate circuit (one with high Thevenin resistance, where the voltage will change significantly when even a small current is taken from it) to have its voltage measured, and the opamp can then supply a similar voltage to things that require some current, such as instruments for measuring the voltage, or long cables that require significant charging when changing their voltage.

3.2 Non-inverting amplifier

The *non-inverting amplifier* uses the input signal as the reference, and has feedback that passes through a voltage divider on the output.

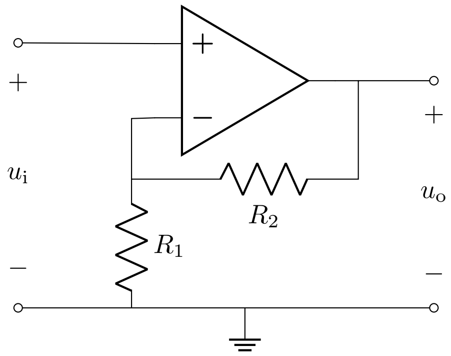

It is known that no current goes into the opamp inputs, so R_1 and R_2 can be considered as series-connected. voltage division can therefore be used to state that $v_- = \frac{R_1}{R_1+R_2} u_o$. Putting this together with the assumption that $v_- = v_+ = u_i$, the non-inverting amplifier's gain is found to be

$$\frac{u_o}{u_i} = \frac{R_1 + R_2}{R_1} = 1 + \frac{R_2}{R_1}.$$

Assuming that both resistors have finite positive values, we see that this gain must be greater than 1.

One distinctive feature of the non-inverting amplifier is that the input to the amplifier circuit is connected only to an input of the opamp. With an ideal opamp, this would mean that the input resistance of the whole amplifier circuit (seen at the terminals marked u_i) is infinite. In practice, a non-inverting amplifier, using an opamp with high input resistance (see the table in Section 7.7), allows sensitive measurements of voltage without drawing significant current from the circuit being measured.

3.3 Inverting amplifier

The inverting amplifier is a very common configuration. In its simplest form, an input resistor and feedback resistor both meet at the inverting input.

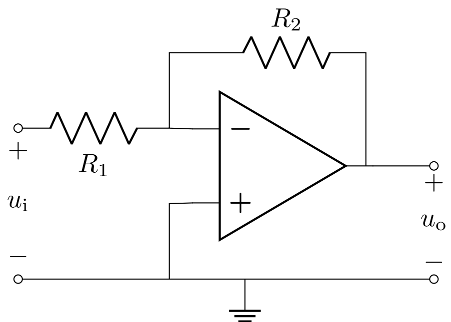

The non-inverting input is connected to ground potential, so the opamp's output becomes whatever value is needed to force the inverting input to zero potential too.

This zero potential at the inverting input, caused by feedback, is called a *virtual ground*. It has the ground-node's potential, but it is *not* part of the ground node: a separate KCL can be written for this node.

Taking KCL at the inverting input,

$$\frac{u_i - 0}{R_1} + \frac{u_o - 0}{R_2} = 0,$$

which gives the relation

$$\frac{u_o}{u_i} = \frac{-R_2}{R_1}.$$

One way of thinking of this is as a voltage divider between the output and input potentials, where the middle point is known to be at zero. The inverting amplifier gives, as its name suggests, a negative output for a positive input. The value of its gain can be anything less than zero, so the output can have a bigger, smaller or similar magnitude relative to the input.

3.4 Current input

All the preceding cases have considered voltages or potentials as input and output quantities. Another way of connecting an opamp is to get an output voltage proportional to an input *current*.

The inverting amplifier could be thought of as doing this, except that the input current is driven by an input potential that forces it through R_1. A circuit designed for measuring current should ideally have very low input resistance (consider how an ideal ammeter is like a short-circuit). This can be achieved by removing the input resistor on the inverting amplifier, leaving just a feedback resistor R.

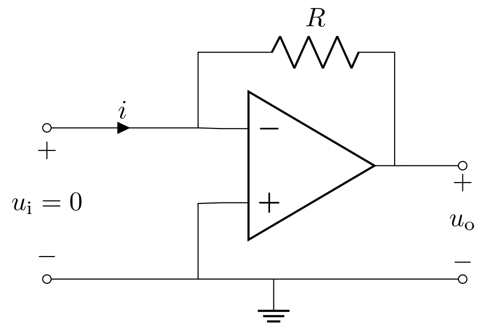

Then the input current i comes into a constant 0 V (virtual ground) and the output voltage is $u_o = -Ri$.

3.5 Adder

An inverting amplifier with multiple inputs provides a way of adding the values of the inputs. The feedback holds the inverting input to zero to match the non-inverting input: it is a virtual ground. The current into each of several input resistors is therefore proportional to the potential applied to the input end of each, such as v_1 and v_2 in the following diagram.

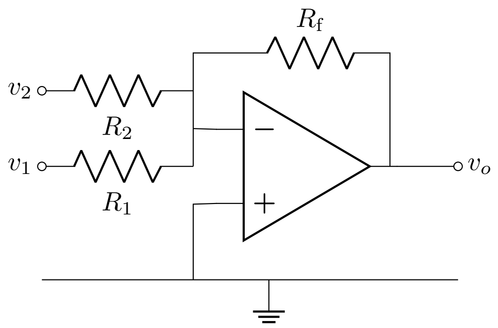

The output voltage has to allow all this current to flow through the feedback resistor R_f. Thus, $v_\mathrm{o} = -R_\mathrm{f}\left(\frac{v_1}{R_1} + \frac{v_2}{R_2}\right)$. By making $R_1 = R_2 = R_\mathrm{f}$ the output potential is equal to the sum of the input potentials; varied resistances permit weighted sums and amplification.

3.6 Integrator, differentiator, logarithm!

We're not really allowed to discuss this until the next Topic! But probably you're slightly familiar with capacitors and inductors and possibly diodes; so this might be of interest.

Imagine replacing one of the two resistors in an inverting amplifier with a capacitor. In one case, the output would be the time-integral of the voltage at the input … in the other case it would be the time-derivative. Using an inductor instead of a capacitor, to replace one resistor, would give the opposite effect. You are left to think about *which* resistors would need to be replaced.

In this way, differential equations as well as algebraic equations can be implemented as opamp circuits.

We could go further, to point out that the forward current in a diode has an exponential dependence on the voltage; using a diode in the feedback (e.g. replacing R_2 in the inverting amplifier) results in a logarithm function being implemented from input to output! That is sometimes useful when measuring a wide range of input signals.

4 Input and output resistance

In practical applications of amplifier circuits, the behaviour of the input and output are important. It might matter to you whether the amplifier would require an input current of 1 µA or 50 mA when you try to drive it with a 0.1 V signal: that current will affect the circuit that is producing the signal.

When a simple non-inverting amplifier is made with an opamp, the circuit's input goes into the opamp's non-inverting input only. This, as we know, would have infinite resistance (zero current) in the ideal case. In a real case, the input resistance may vary by orders of magnitude, depending on the design of the opamp, but it will be high compared to many electronic circuits on which we want to make measurements.

A simple inverting amplifier instead starts with a resistor R connecting from the input to a "virtual ground" point at the inverting-input. The current drawn by the circuit's input at potential v_i is therefore v_i/R, so this circuit is said to have an *input resistance* of R.

The output of an ideal opamp is like a voltage source, so it has zero resistance. That is a way of saying it's like a Thevenin source with zero resistance: any change in the current in or out will cause no change in the voltage. Real opamps have some fall of the output voltage when they more current is drawn from them; on the other hand, the feedback will attempt to force the actual output voltage to be the 'correct' value, which helps make the apparent output resistance lower, when used with negative feedback. A circuit containing an opamp may have other components between the opamp's output and the whole circuit's output, for example for protection of the opamp against excessive currents. The output of an opamp-based amplifier circuit may therefore have a significant *output resistance*.

5 Analysis of Opamp Circuits

This section is the most important for the work we do with opamps in a Linear Circuits course. It is all about circuits *where ideal opamps are assumed, and negative feedback is provided by the external circuit.*

As was seen in Section 2.3, this situation of negative feedback results in the opamp automatically adjusting its output to the value that is needed so that $v_- = v_+$.

Important! remember that the output gets its output current from a hidden place — so currents *can* come in/out of the opamp output, even though this appears to contradict KCL. The current at the opamp's output must return to the circuit's ground node. So, **do not** assume zero current in the opamp output, and do not try to use KCL on the 3-terminals of the opamp.

The dependent-source model of the opamp (Section 1) makes this clear by showing a ground-node hidden inside the opamp symbol; this is the same node as the circuit's ground node. A more detailed opamp symbol would show at least two 'power supply' connections on the opamp, through which the currents in or out of the opamp output are obtained from a voltage-source somewhere else in the circuit. This makes the diagram messy and more complex so we usually avoid it.

5.1 For simplifications-based analysis

One way to analyse a circuit that contains opamps is a step-by-step method, with perhaps a mixture of simplification, KVL, KCL, source-transformation, superposition, etc. Compared to a systematic application of nodal analysis, this might help us more to feel how the circuit behaves, and to avoid getting lots of equations to handle at once. On the other hand, we might sometimes get stuck by not being able to see what step can be made next.

The key assumptions used for analysis are that:

> The inputs are equal: $v_- = v_+$, because the opamp is ideal and has negative feedback.
>
> No current goes into the opamp inputs.
>
> The opamp's output potential will become whatever is needed in order to push v_- to match v_+; the output can supply any current that is needed to achieve this and to supply any loads connected to the output. Define an unknown potential if it's not already defined.

Adding the above assumptions to our previous set of KVL, KCL, Ohm's law, and various tricks such as voltage and current division, source transformation, two-terminal equivalents and superposition, a lot of opamp analysis can be done. It is useful to get familiar with the basic inverting and non-inverting amplifier circuits, as these often appear as sub-parts within a bigger circuit containing opamps.

5.2 For nodal analysis

Instead of working out a good strategy for step-by-step solution of the circuit, we can use a systematic approach with nodal analysis, requiring less thought about where to start. (A combination of simplifications and nodal analysis is often a good choice.)

Only a few rules need to be added to what we already know about nodal analysis, in order to include an opamp in our analysis:

> Define the opamp output's potential: e.g. v_o.
>
> If opamp output current is needed (for extended nodal analysis) then define it, e.g. i_o with an arrow. In the supernode approach, this is not necessary: the output would be seen as being part of the 'ground supernode' where KCL is not needed.
>
> (The inputs have no current, so they do not affect the KCL equations for the nodes where they are connected.)
>
> The assumption of negative feedback on an ideal opamp gives $v_- = v_+$. Thus, the potentials of the nodes connected to the opamp inputs can be set to be equal: this gives a further equation.

The output is like the output of any voltage-source that has its other side connected to ground; it can supply any current it is asked to supply. So we cannot assume this current to be zero.

If the opamp symbol is replaced with the dependent-source model shown in Section 1, and a *finite* gain such as k is used for the VCVS, then our existing knowledge of nodal analysis would provide a solution. The output potential is simply related to the differential input, as $v_\mathrm{o} = k\,(v_+ - v_-)$.

But if the opamp is assumed to be ideal, the infinite gain appears to cause a problem! We have the assumption that $v_+ = v_-$, so the above approach says that $v_o = \infty \cdot 0$. It is then not possible to use simple equations and numerical solutions to calculate the output voltage v_o from the gain and the controlling variable, in the way we would do with a normal dependent voltage source in nodal analysis. That makes us feel that we have lost some information: one of the equations in the nodal analysis has become useless. As the number of variables has stayed the same, the circuit won't be solvable. However, the assumption of infinite gain has also told us that $v_- = v_+$. This equation can be adding to the nodal analysis equations, to make the circuit solvable.

6 Examples

Let's try the following example, by several methods. All node potentials have been marked, but are not 'known' quantities: as usual, just the component values are known.[3] We will assume that just the output potential v_5 is sought. Some of the other potentials might be calculated in the process of finding this.

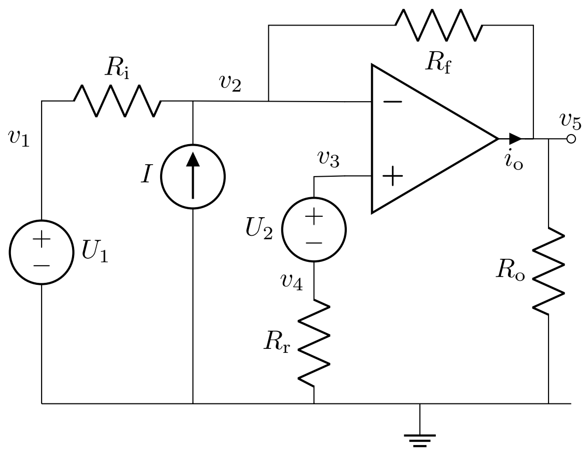

6.1 Simplifications, step-by-step

The current in R_r must be zero, as the only route for this current is through an (ideal) opamp's input. By Ohm's law, there can then be no voltage across this resistor. Thus, $v_4 = 0$ and consequently $v_3 = U_2$.

[3] Notice that the source U_1 or I could instead have been written as just a terminal with a marked quantity such as u_i or i_i, to show that the circuit is designed to be connected to receive input from another circuit (not shown) that provides a particular voltage or current to be amplified by our opamp circuit. A way to model this is just to represent that input voltage or current with a source, as we've done here.

By the usual opamp assumption, the negative feedback holds the two inputs to equal potentials: hence $v_2 = v_3 = U_2$.

By KCL in the node marked v_2,

$$\frac{U_2 - U_1}{R_\mathrm{i}} - I + \frac{U_2 - v_5}{R_\mathrm{f}} = 0,$$

in which the output voltage v_5 (the only unknown) is found as

$$v_5 = R_\mathrm{f} \left(\left(\frac{1}{R_\mathrm{i}} + \frac{1}{R_\mathrm{f}} \right) U_2 - \frac{1}{R_\mathrm{i}} U_1 - I \right).$$

6.2 Nodal analysis: write the equations

Let's show, first, extended nodal analysis with no simplifications. This would be good for programming a computer to solve everything (including currents in the voltage sources and opamp output) using simple rules but many equations.

First, KCL at all nodes except the ground node. This is just the usual procedure in extended nodal analysis. We have to define the unknown currents in voltage sources: let's define i_α and i_β into the +-terminals of sources U_1 and U_2 respectively, and i_o out of the opamp as shown in the above diagram.

$$\begin{aligned}
\text{KCL(1):} \quad 0 &= i_\alpha + \frac{v_1 - v_2}{R_\mathrm{i}} \\
\text{KCL(2):} \quad 0 &= \frac{v_2 - v_1}{R_\mathrm{i}} - I + \frac{v_2 - v_5}{R_\mathrm{f}} \\
\text{KCL(3):} \quad 0 &= i_\beta \\
\text{KCL(4):} \quad 0 &= -i_\beta + \frac{v_4}{R_\mathrm{r}} \\
\text{KCL(5):} \quad 0 &= \frac{v_5}{R_\mathrm{o}} + \frac{v_5 - v_2}{R_\mathrm{f}} - i_\mathrm{o}
\end{aligned}$$

Then, each voltage source brings its own relation between voltages,

$$\begin{aligned}
\text{VSRC(1):} \quad v_1 &= U_1 \\
\text{VSRC(2):} \quad v_3 - v_4 &= U_2,
\end{aligned}$$

and the opamp is slightly different in that it relates the node potentials connected to its inputs,

$$\text{OPAMP:} \quad v_3 = v_2$$

There are no dependent sources with controlling variables that need to be defined, so we're now finished with writing the equations: there are 8 equations, and 8 unknowns (5 node-potentials and 3 currents). Good luck with solving the above, for v_5.

6.3 Nodal analysis and simplification

We *could* solve those equations, but it's usually nicer to use simplifications, particularly when we only actually want to find a small proportion of the unknowns – just v_5 in this case.

So we can instead use the supernode approach to avoid considering the currents in voltage sources (including the opamp). We can also use some simplifications before writing the equations.

By the earlier step-by-step reasoning, $v_3 = U_2$, and thus $v_2 = U_2$. The components U_2 and R_f, and the opamp's non-inverting input, are a separate branch in the circuit, not affected by any other component. So all we need to know about them is the conclusion $v_2 = U_2$.

The resistor R_o does not affect the solution for v_5. One way of seeing this is that it is in parallel with a voltage source (the opamp output), and so does not affect the voltage. Another way of reasoning is that the opamp has to adjust v_5 in order to get the correct feedback to force v_2 equal to v_3: if we connect more load to the opamp, it just has to supply that load with whatever current the load demands in order to hold the necessary potential v_5. Thus, ignore R_o.

The opamp output is seen as a voltage source with its other terminal connecting to ground: by the supernode treatment, it is therefore part of the ground supernode, and needs no KCL. Thus, the current i_o is not needed in the equations when we solve for potentials.

Components U_1 and R_i can be treated as a single branch, when writing KCL at v_2. Another way of seeing this is that v_1 forms part of a supernode with the ground node, so v_1 is known and no KCL is needed at v_1.

So, we're left with just the node v_2, at which we are supposed to write KCL:

$$\frac{v_2 - v_1}{R_\mathrm{i}} - I + \frac{v_2 - v_5}{R_\mathrm{f}} = 0.$$

By substituting the known values of v_2 and v_1,

$$\frac{U_2 - U_1}{R_\mathrm{i}} - I + \frac{U_2 - v_5}{R_\mathrm{f}} = 0,$$

which is just what we had and solved in the first (step-by-step) approach.

This example may make nodal analysis (whether by the extended method or with simplifications) look more work than the first approach that didn't try to be systematic. But sometimes a circuit isn't easy to handle in little pieces, and a systematic method becomes better even when working by hand.

You should do plenty of practice from the exercises and the past exams, with some attention to systematic methods as well as to simplifications.

7 — Extra —

7.1 Physical implementation

An opamp is a special combination of several parts, each of which is quite complicated and nonlinear. In the early operational amplifiers of the 1940s and 1950s, thermionic valves (electron tubes: elektronrör) were used as the building blocks for opamps. In a modern opamps there are typically around 15 transistors.[4] The following diagram shows the design of a very well known old opamp, the "741". Each of the three-terminal devices in the diagram is a transistor.

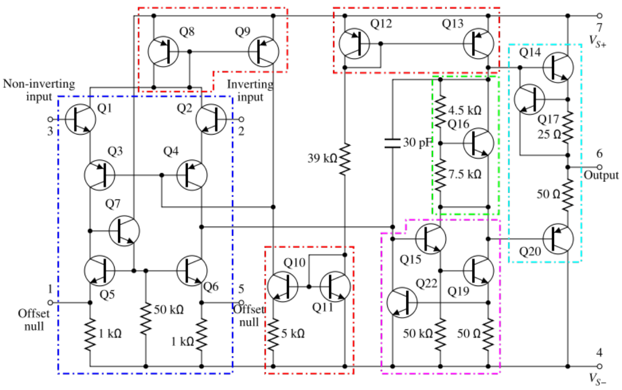

Source: Wikipedia. Drawing by D. Braun, 2007.

A physical opamp comes as a single package with a few terminals connecting to it. These will be at least the two inputs, one output, and two power-supply connections.

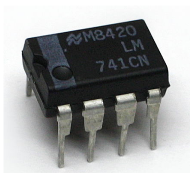

There may also be terminals for adjusting or compensating some feature of the opamp's behaviour, such as the gain or input offset. A single package may contain more than one opamp. For more detail about the history and constructional details of opamps, see the Wikipedia page.

7.2 Some history of feedback amplifiers

Feedback has been a feature of engineering for much longer than its relevance to electrical engineering. However, it was the big growth of electrical and electronic engineering in the first half of the 1900s that resulted in great interest in how to analyse and design electronic systems such as amplifiers, and more general physical systems with electrical control. An important step was the realisation that an amplifier with high gain but bad linearity (not a constant ratio of output/input) could be used with negative feedback to make it into a nicely behaved amplifier with lower gain. One of many tales about this idea is [Black's Box].

7.3 Where's the energy from?

This very sensible question came up from discussions the first time that we tried Lab1 with opamp circuits. The ± 15 V supply to the opamp board was obviously in some way important to making the opamp work, and yet we don't normally draw any supply-connections in our diagrams.

It is sort of true to say that an opamp "needs the power supply to give energy to the dependent source inside the opamp". But this doesn't help us understand the mechanism — what *is* the dependent source, physically, and *how* does it get supplied with energy?

A better explanation would be that the opamp uses transistors to *connect* the output to either the positive or negative supply rail, depending on whether the output voltage needs to be more positive or more negative. Transistors can do this in a gradual way, by varying their apparent resistance.

In this way, the complicated circuit inside the opamp results in the difference between the input voltages controlling the transistors, which control currents in another part of the circuit, consisting of the power supply (voltage source) and opamp output.

At least two terminals on a real opamp chip are supplied from a voltage source. For tidiness, these are often not shown in the opamp symbol; sometimes they are, in which case they can be vertical lines from the top and bottom, as shown below.

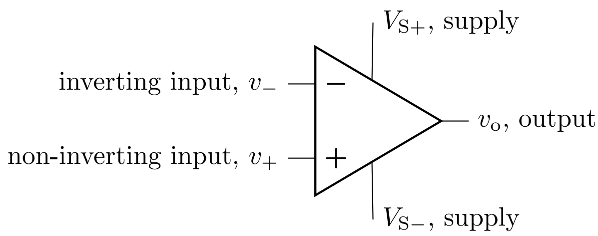

Current out of the opamp output terminal is coming into the opamp from the higher-potential supply terminal, V_{S+}, and current into the opamp output is passing out through the lower-potential one, V_{S-}. The circuit outside the opamp must, of course, have another connection to the voltage source, so that current through the opamp output can complete its journey to the source that it came from.

When one wants the output to be able to go positive or negative (with respect to a chosen 'ground' in the circuit), the inputs need to be from voltage sources at positive and negative potentials. If it is acceptable to

[4] Transistors: various types of nonlinear device with at least three terminals. A single transistor is often modelled as a controlled source, when one considers just small changes in the currents and voltages around a constant value. The control input is for example the current into one terminal, and the current into another terminal is controlled by this.

have an output only in a range of positive or negative values, then the supply to the opamp can be from a ground and a single potential.

The use of supply rails is clear even *inside* the 741 opamp, whose diagram was shown in Section 7.1. The lines (nodes) at the top and bottom provide the supply for the several branches of transistors in this circuit.

7.4 Supply rails

In electronic circuits it is common to have many chips (integrated circuits, like the opamp) and other branches that need a power supply at a particular voltage.

For simplicity one prefers to have a single voltage that everything can use, but in practice it is sometimes necessary to have several voltages for different components. Different voltages can be provided either by power supply hardware that behaves as several different voltage sources, or by having converters within the circuit to convert locally to a different voltage needed by a particular component.

A simple converter is the 'linear' converter, which provides a *lower*, dc voltage by having a transistor that behaves as a controlled resistor: when the output voltage becomes too low, the 'resistance' is decreased to let more current come through. This, of course, generates heat: to convert from 5 V to 3.3 V requires 1.7 V to be 'dropped' across the transistor, which means that 1.7 W will be turned to heat in the transistor for every 1 A supplied on the output. More modern converters are 'switch-mode', using transistors as switches that are either very low or high resistance, to store energy from a source into an inductor or capacitor, then move this to other inductors or capacitors on the converter's output.

In a modern desktop computer there is a *power supply unit* (sv: *nätaggregat*) as a separate box containing voltage sources for several voltages needed by different parts in the computer. It has output wires at 12 V, 5 V and 3.3 V, −12 V, and also at 0 V as the ground or 'return' conductor for all of these sources. These potentials are distributed by conductors on the computer's mainboard, to devices with various power requirements. See more about computer power rails on Wikipedia: [ATX-PSU].

7.5 Inidealities

A real opamp differs in several ways from the so-called ideal case. You do not need to know anything about this for our course; it is sufficient to do calculations on ideal opamps. However, anyone who is interested in using real opamps should have a good idea of how a real opamp differs from the ideal.

Output limits We have been discussing an output that "shoots off to a huge value" as soon as there is a tiny difference between the inputs. A real opamp's output is bounded by the power supply available to the opamp; usually it can't come closer than a fraction of a volt to the potentials of the opamp's power supply. So instead of shooting off "to infinity", a real opamp's output will go up or down to a value that depends on the power supply connections and the opamp design; this might be for example a range of 0.2 V to 2.8 V or −14 V to 14 V. When it has hit its limit, it can go no further even if the input value changes further. This is shown in the following figure, assuming a very small opamp gain of 10^3, power supply connections at ±5 V, and a margin of 0.6 V between the output limits and supply rails.

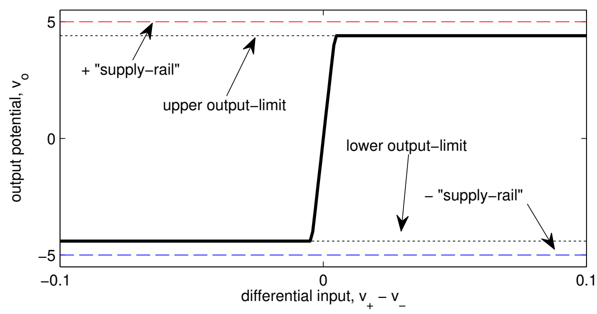

This change in behaviour is a nonlinearity: the opamp's *gain*, if defined as the slope of the v_o versus $v_+ - v_-$ relation, has suddenly changed from a large value to zero, since any further change in the input cannot make the output change any further.

Slew rate The opamp's output cannot change instantaneously. When a very quick change in potential is made on an input, the rate of change of output potential is called the slew-rate. This is quoted in units such as V/µs. Bearing in mind that some opamps will be used with frequencies such as 1 MHz and more, or for measuring quick changes, this sort of limit can be significant. Thinking about this limited rate can help us to understand negative feedback circuits: when the inputs have a different voltage, the output *starts going up or down* but cannot change instantaneously. Then, when the negative feedback has caused the inverting input to have the same potential as the non-inverting input, the change can stop. (What happens in reality will depend on the internals of the opamp and the external details of the feedback circuit: at high frequencies the feedback will stop working properly; but at least this idea of limited rate helps us understand why an opamp is able to work stably in negative feedback instead of its output oscillating up and down between big extremes as we might believe if thinking of a instantaneous response from input to output.)

Finite gain The opamp's gain can easily be around 10^6 for dc, but gains vary by several orders of magnitude (tiopotenser) depending on the design of the opamp. In many cases we can say that $10^6 \simeq \infty$, 'fpp' (for practical purposes!). The validity of modelling an opamp as ideal, instead of having to look up the real specifications and do longer calculations using these, depends on the application: if you want an audio amplifier with a gain of 5, then a gain of 10^6 is practically ideal, but if you are making a sensitive measurement instrument you might want to include the finite gain, input bias current, etc, in your calculations. Opamp gains are designed to fall when the signal (input voltage) comes up towards higher frequencies. A common specification is the *gain-bandwidth product*; this is quoted as a constant, implying that increases in frequency will cause proportionate decreases in gain. A reason for *desiring* reduced gain at higher frequency is that there otherwise can come a point where the frequency is high enough that there is a large *phase delay* in the opamp, causing the output to be delayed so much from the input that an ac signal would look 'upside down' (180 degrees displaced), in which case a negative feedback circuit would look like positive feedback and the circuit would be unstable.

Gain-Bandwidth product Not only is the gain finite, but it also gets lower as the frequency increases. In this part of the course we have only looked at dc circuits, so we don't consider time and frequency. Towards the upper part of an opamp's range of operating frequency, its gain falls as quickly as the frequency rises, so their product is approximately constant. This value is often cited as a specification; it is the frequency at which the gain becomes < 1.

Finite input resistance The inputs do allow some small current to flow, so they have not got infinite input resistance. The input resistance depends a lot on the type of transistors used, which in turn depends on what properties the opamp is optimised for.

Input bias current The inputs may also have some bias current, which means a fairly constant (independent of potential) current that the opamp's internal circuits try to send out of or into the input terminals. This can be a trouble when the circuits connected to the opamp input have very high resistance, or are purely capacitive: the bias current can then signficantly change the intended output voltage, and cause voltages on capacitors to keep changing over time.

Input offset current In some designs of circuit, the effect of an opamp's input bias currents can be largely removed from the circuit's output if the bias currents at both inputs are the same. The input offset is the difference between the bias currents of the two inputs; it is designed to be zero, but will have a different (and variable) value for different opamps.

Input offset voltage This describes the situation where the output changes between its positive and negative values *not* when the inputs are at the same potential. It is as if the opamp's differential gain were $G(v_+ - v_- - v_{\text{offs}})$. The difference between v_+ and v_- that is needed in order to make the output go through zero is called the input offset voltage.

Common mode input The ideal opamp's output depends only on the difference between the input potentials. A real opamp's output may depend a little on the absolute values of the inputs, not just their differential value. A potential that is on both inputs is called a common-mode potential. The common-mode rejection ratio describes how much more strongly the differential mode $v_+ - v_-$ is amplified on the output, compared to the common mode $\frac{v_+ + v_-}{2}$. The numbers are usually big, such as a million times … so they are commonly expressed in the logarithmic scale of *decibels*, which we come to when discussing filters in the ac section of the course.

7.6 Hysteresis and Control

This is mainly a continuation of the comparator and positive feedback.

Try turning the temperature setting of an oven (ugn) up and down around the temperature that the oven currently has. The indicator lamp (and heater) should turn on when the temperature setting is about the same as the oven's temperature and you're turning it up. Then it should turn off when you turn it down … but does it turn on and off at very close to the same temperature? The answer is usually 'no': there is a gap. When it has turned on, it is reluctant to turn off immediately for just a small change in the difference between the 'setpoint' and the measured value.

This is hysteresis in the thermostat, which is probably a simple mechanical device. Some degree of hysteresis is often desirable, to avoid the way that small changes of inputs to a pure comparator could cause the output to wiggle quickly between its extremes. For example, a thermostat controlling a room heating system that is based on pumped hot-water would be a nuisance if it kept turning a pump and fuel-burning heater on and off depending on little changes in its temperature when people move air around by their activities in the room. On the other hand, an electric room heater can be turned on and off easily, without much adverse effect.[5] The thermostats in these simple heaters sometimes have negative feedback by including a little extra heater wire *inside* the thermostat. This causes the thermostat to turn on and off quite quickly, e.g. every few seconds or tens of seconds, but the proportion of the time for which the thermostat turns the heater on is dependent on the room temperature and the thermostat setting. The result is a more steady room temperature, compared to an oven-like

[5] One effect of heaters turning on and off can be the lights changing slightly in brightness; that's due to the Thevenin equivalent of the supply source having a non-negligible 'resistance'. For the types of quite small individual heaters and phase-to-phase connections common in Sweden, this is seldom significantly noticeable.

thermostat where the measured temperature has to become significantly above or below the setpoint in order to cause the output to turn off or on.

The foregoing had no direct connection to normal *applications* of opamps, but it gives other examples of systems where feedback and hysteresis are present. These are part of the very broad (and interesting, and much more complicated than the above) subject of *control engineering* (sv: *reglerteknik*). The principles are found all around us, in mechanical, electrical, financial, social, etc., situations.

7.7 Some actual opamp specifications

Opamps are designed with a wide range of specifications, for a wide range of applications. Some need to perform well at high frequencies; others may be desired for low frequencies and even dc conditions but with extremely high sensitivity to small currents. The important parameters include gain, bandwidth (frequency-range), slew-rate, input resistance and capacitance, input bias currents, input offset voltage, input voltage range, output voltage range, power consumption from the supply, and (of course) price!

In the "ideal opamp" of circuit theory, we consider gain and slew-rate being infinite, input bias currents and offset voltage being zero, input resistance infinite, output resistance zero, and no limits to the output or input potential.

In real circuits, some of these properties may be less important than more practical properties. The cost of an opamp is clearly an important feature; it will often not be worth improving some parameter if this increases the price. Power consumption is important, especially when in a battery-powered circuit; the power taken from the source by an opamp is more than what it supplies from its output, as some power is used in its internal components. Non-zero output resistance can even be useful, in order to limit the short-circuit current. Other indealities such as finite gain might be quite unimportant in practice, as long as the gain is greater than some design threshold. For low-voltage circuits, such as those supplied from a 3 V or 1.5 V source, it is important that the output should be able to come close to the power-supply voltage: an opamp design that only lets the output come within 0.7 V of the input would result in the output voltage range being only about a half of the supply range.

Applications of opamps vary greatly. When manufacturing an opamp, some of the parameters are in conflict if they're to be moved towards the apparently ideal value: one cannot expect very high gain, slew-rate and input resistance, as well as low bias currents and offset voltage and power consumption, and all at a very low price! This can be seen as the reason for a wide range of opamp types, optimised for particular applications; the optimisation may make an opamp particularly *bad* for another application. One key feature in the design is the type of transistor used in each part of the opamp. In the lab task we have used an opamp based on JFET transistors, giving a very high input resistance.

The following links are to data-sheets for a few commercially available opamp chips. There are lots of others: the only reasons for choosing these ones are that they were easily found when looking for examples of opamps optimised for several different applications (power, sensitivity, cost etc.).

The [LM741] opamp is a very common general-pupose model, that is used in many products and old hobby designs.

The [LF356] is what we have used in the lab.

The [MAX477] is for high-frequency applications.

The [LT1970] opamp is designed for 'power', i.e. large output current; it has a 500 mA output.

The [LMC6001] is a 'precision' opamp, able to handle very small currents, in the fA range.

The following table compares some parameters of four of the above opamps.

Some of these parameters are quite tightly specified, such as the permitted limits of supply voltages and input voltages; on the other hand, one can normally operate well below the limit.

Other parameters are not tightly specified, and may depend strongly on temperature or on variations between batches and units of the opamp. For example, the 741 specifies input bias current of 80 nA as a normal value at the standard temperature of 25 °C, but up to 1500 nA when considering the worst case of operating temperature and manufacture of the chip.

Several parameters depend strongly upon the supply voltage that is used. It is typical that the input voltages should not go outside the range of the actual supply voltage, so the permitted input voltage range varies depending on the supply voltage.

The maximum output current depends on the supply voltage, as there is significant resistance, of tens of ohms, near the output of the opamp. The current used by the opamp itself depends on the supply voltage. Gain can depend significantly on the output current and on whether this current is going out of ('sourcing') or coming into ('sinking') the opamp output.

Opamp	LM741	LMC6001	MAX477	LT1970
Specialism	general, cheap	low current [in]	high speed	high current [out]
Supply voltage max	44 V	15.5 V	12 V	36 V
Input voltage range	26 V or supply	same as supply	0.3 V beyond supply	36 V
Output voltage range	$\sim$1 V of supply	$\sim$0.3 V of supply	$\sim$1.5 V of supply	$\sim$ supply
Gain (A_v)	2×10^5	3×10^5	0.2×10^5	1×10^5
Gain-Bandwidth (GBWP)	1 MHz	1.3 MHz	300 MHz	3.6 MHz
Slew rate (SR)	25 V/µs	1.5 V/µs	1100 V/µs	1.6 V/µs
Output current (shorted)	25 mA	25 mA	150 mA	500 mA
Input resistance (R_in)	2 MΩ	1 TΩ	1 MΩ	0.5 MΩ
Input bias current	80 nA	10 fA	1000 nA	160 nA
Input offset current	20 nA	5 fA	200 nA	100 nA
Input offset voltage	1 mV	1 mV	0.5 mV	0.2 mV