Specialists in Fluorescence Imaging, Electrophysiology, and Microscopy

OptoPatch

Instruction Manual

www.cairn-research.co.uk

Free Phone: 08453301267 (UK only)

Tel: +44 (0) 1795590140
Fax: +44 (0) 1795594510

Contents

1 Introduction

1.1 Development

Designing the Optopatch has been by far the most challenging project that we have ever undertaken, but it has also been a tremendously interesting one. Patch clamping is itself a challenging technique, but since it is now so well established, those people who politely asked us during the design stage why we were bothering to get involved in it now, certainly had a point. To be honest, at first we rather agreed with them. And yet, the Optopatch came intobeing because people were specifically asking us for the improvements it incorporates. The design is all ours, and we're very proud of it, but we are also very grateful to the people who pointed out to us the problems that needed to be solved. There is a great deal of unique electronic circuitry in the Optopatch, allowing us to offer the three innovations. These are the optical current-passing technique (from which of course the product takes its name), the automatic hardware-based capacitance measurement facility, and greatly improved current-clamping performance. We also threw in a few other useful improvements, such as a wide-range high-frequency filter, a wide-range frequency generator, and extensive use of electronically controlled ("analogue") switches in the signal pathway in order to keep the path lengths shorter and the front panel wiring simpler.

It all started at a Physiological Society meeting in Leicester in the spring of 1993, when someone asked me if I had any better ideas about headstage design, in particular with regard to current passing, as he felt there could be room for improvement over both the standard high-value resistor technique and the more recently introduced integrating (capacitive) technique. At the time we had no plans in that area, and I wouldn't have given the matter any serious thought, but following the Society Dinner and its rather alcoholic aftermath, I was left with a more than usually severe hangover. To check that there had been no (further) lasting mental damage, I started to think about the problem, and at the time an optical technique seemed to be a most interesting alternative solution. It still seemed a good idea a few days later, and a few preliminary electronic experiments suggested that it could be made to work well in practice, but at the time we were particularly busy with other projects, so we had to wait for nearly eighteen months before we could proceed any further.

The problem was that a patch clamp headstage is literally the tip of the iceberg, so we knew that designing the rest of the electronics would be no small task. Although we weren't ready to go public on the headstage idea, we nevertheless sounded out the possibility of there being a Cairn patch clamp one day, and asked people what facilities they might like to see on it. As a result of that, our attention was drawn to the interest in following vesicular secretion from cells by membrane capacitance measurement, and this is what determined our design approach for the main electronics. We were impressed by the EPC-9 computer-controlled patch clamp, which can do this and other functions in software, but we also liked the idea of the hardware approach using a lock-in amplifier, which can make fast and sensitive capacitance measurements directly, but is not so easy to use. Faced with the choice of a "me-too" computerised approach, or the new challenge of an improved analogue system, we opted for the latter. What this means is that where the EPC-9 might implement a gain control using a digital-to-analogue converter (DAC) programmed by the computer, the Optopatch might implement it using an amplifier with a gain that is set by an analogue control voltage. This means that we can implement automatic control functions by deriving the appropriate control voltages internally, instead of digitising the relevant signals and then performing the control via programming of the DACs by the computer software. Although the comparative performance of DACs and gain-controlled amplifiers is asubject in itself, the resolution of the amplifiers we use is at least as good as a 16-bit DAC, and we think that either component is capable of giving satisfactory results. Over time we have also come to appreciate that basing the design on our own hardware platform has provided a much more stable development environment than software operating systems could have offered, so our (relatively few) design updates have been implemented solely to provide useful improvements.

Our approach tries to offer the best of both worlds, so most functions are set by conventional front panel controls. A computer is not essential for the Optopatch! However, we are well aware of just how useful it is to be able to record additional parameters such as a telegraph voltage to indicate the gain setting, so outputs representing the important switch and control settings have been provided. Although we did not set out to design a computer-controlled product, it was only a small extra step to make these signal lines bi-directional, so that a computer can set them as well as read them. Whether or not a computer is used to control the Optopatch, the internal analogue signal processing allows outputs such as direct linear measurements of membrane capacitance to be generated in real time. These can sent directly to any general-purpose data acquisition system, and there is no need to process them any further by special software. This approach is particularly useful when other data (such as fluorescence signals) are being captured at the same time. While of course we would love you to use our own software, you remain free to make your own choice. From the design point of view, our approach left us a fair amount of extra electronics to implement, but the individual functions were all simple enough. However, the overall schematic is quite impressive....

We thought all this would be enough to be getting on with, but towards the end of the design process, our attention was drawn to the fact that the current-clamp performance of other patch clamps can be noticeably less than ideal. The problem and its solution are described in detail in a later section, but we were made aware of the problem in time to reconfigure the interconnections between the headstage and the main electronics, in order to allow the easy implementation of its solution.

We have given considerable thought to making the construction of the Optopatch as easy as possible. The product's electronic complexity would justify us charging a premium price, but we wanted to price it competitively with all the other patch clamp amplifiers, in spite of its unique facilities. We have minimised the construction costs by placing most components on a single circuit board, although those specific to the optical headstage design are on a separate daughter board. This is to give us options such as to introduce headstages of more standard design, e.g. for passing higher currents in other applications, where the resistive method is as good as any other, and/or to compensate for even larger (i.e. nanofarad) membrane capacitances. All active devices on both boards are socketed, and all external connections are via pluggable connectors, allowing easy identification and repair of any fault. The use of analogue switches greatly simplifies the wiring to the front panel controls, as well as keeping the signal paths short. For example, the filter range switch is a single-pole rather than a four (or eight!) pole type, and instead of being in the analogue signal pathway, the switch connections need to take only digital control voltages. This arrangement gives improved performance and reliability at lower overall cost. Furthermore, the actions of some switches depend on the settings of others, and this is easily handled when only control voltages are involved, since standard logic circuitry can be used to produce the desired effects.

Towards the end of 1999 we introduced a revised version of the Optopatch, which incorporated a variety of improvements, mainly directed towards making the amplifier easier to use. In particular, in the RC compensation system we changed the variable conductance control to a variable resistance control. Conductance is arguably the more logical choice, but in practice it is far more convenient to be able to read the resistance directly. We also incorporated vastly improved metering facilities, which can display resistance and capacitance (amongst other things) as well as voltage and current, with additional indicators to display the appropriate units in all cases. The other significant change was to extend the potential range of the RC compensation system from +/-100mV to +/-200mV (the range without RC compensation remains at +/-1V, so the Optopatch can still be used for cyclic voltammetry measurements). Various other internal circuit refinements were also made then, but the performance improvements from these are relatively small, and are not necessarily noticeable in practice.

In 2000 we introduced the option of an eight-pole rather than a four-pole Bessel filter. The extra filter sections are on a separate daughter board, which can be retrofitted to all amplifiers made from that date onwards.

The other major development has been that of our (long-promised!) "gold" headstage. This uses a specialist field effect transistor (FET) in order to give improved noise performance. Its development was delayed by difficulties in finding higher-performance versions of some of the other headstage components (so that the full performance benefits of the new FET could be fully realised), and also by our move to our new and much larger premises in 2003. However, this headstage is now (summer 2004) at last going into production, but our regular headstage will also continue to be available, as its noise performance is still quite respectable and the gold version is necessarily a significantly more expensive product.

1.2 How to do Patch Clamp Experiments

This manual concentrates purely on the technical aspects of patch clamping, and in order to supply the user with the vitally important practical information, we can provide free of charge with the Optopatch a copy of "Microelectrode Techniques: The Plymouth Workshop Handbook" (2nd edition, 1994). This book includes several chapters on patch clamping, and much of the other information it contains is also relevant. A practical course such as the one at Plymouth (which takes place every September) forms an excellent introduction to the subject, and anyone who has a

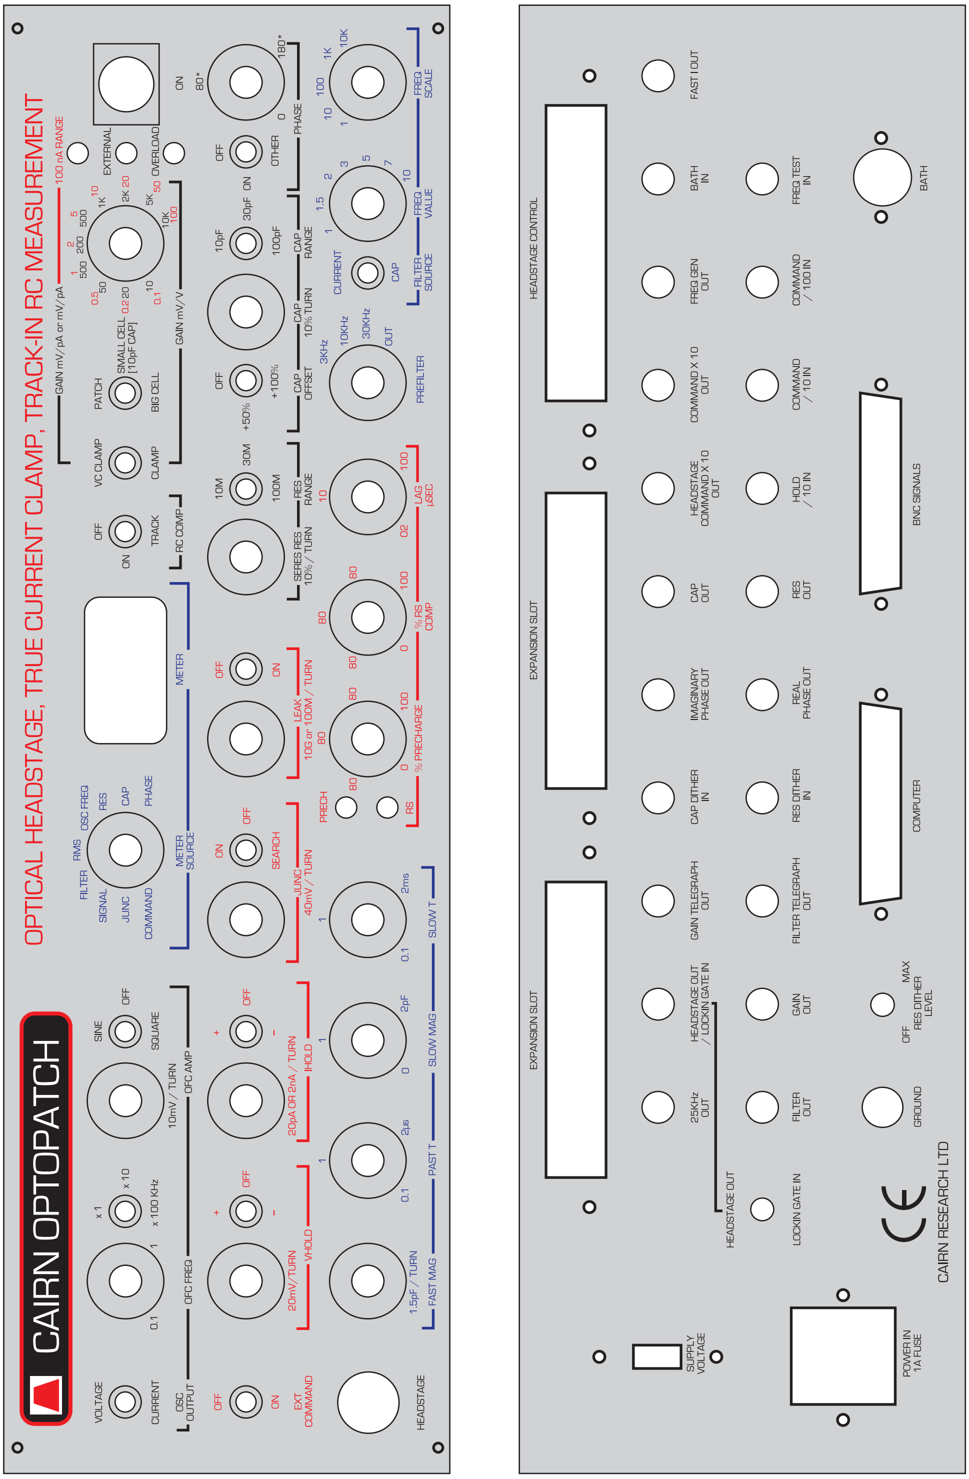

Figure 1 - The Optopatch Front and Rear Panels

chance to attend one is strongly recommended to do so. Newcomers to patch clamping would be well advised to read the appropriate parts of the book before delving too deeply into this manual.

The other book that may be appropriate to mention here is the second edition of "Single Channel Recording" (editors Sakmann and Neher, Plenum Press, 1995), which covers both the theoretical and practical aspects of patch clamping in very considerable detail. However, newcomers are recommended to read the Plymouth book first.

Both books also discuss the analysis of patch clamp recordings in great detail, so we shall not cover that topic here either. At the time of writing this section (October 1996) Cairn does not have any commercial software available for this purpose, although as a result of planned University collaborations, that situation could change in the future. Meanwhile, in addition to the commercial software available from other companies, we can also recommend the public-domain software produced by John Dempster at Strathclyde University (his full address is in the Plymouth book).

2 Front Panel

2.1 Top Row

This section provides a brief description of each of the front panel controls and rear panel sockets, and is primarily for reference. It will help experienced patch clampers to familiarise themselves with the specific features that the Optopatch provides, but other users may prefer to skip this section for the time being. We'll start with the front panel, which for reference is shown in Fig. 1, and describe the controls in a left to right sequence, beginning along the top row.

'OSC OUTPUT' Switch

In the upper (voltage) position, sine or square waves generated by the internal oscillator are added to the command potential. In the lower (current) position, they allow current pulses of up to 10nA to be applied to the headstage through the compensation capacitor instead, via an integrator to preserve the original waveform.

'OSC FREQ' Control

A single-turn control to vary the oscillator frequency over a 10:1 range. The control is not calibrated, but the actual frequency can be read precisely on the meter.

'OSC RANGE' Switch

The frequency ranges selectable on this switch are x1, x10 and x100, corresponding to 100Hz-1KHz, 1KHz-10KHz and 10KHz-100KHz.

'OSC AMP' Control

A ten-turn calibrated control for setting the oscillator output level. The maximum sine wave output is 100mV rms (242mV peak), or 10nA when the osc output switch is in the current position, and the maximum square wave output is +/-50mV, i.e. 100mV peak-to-peak.

'OSC AMP' Switch

This switch should be in the centre off position if the oscillator is not in use, otherwise it selects sine or square wave generation.

'METER SOURCE' Selector

The signals that can be read on the meter, together with the units that are also displayed, are as follows:

Command	The total command potential, but excluding any bath or junction potential corrections. Displayed units are mV, or pA or nA in current clamp.
Junc	Junction potential, set either by the front panel control, or automatically in search mode. Units are mV.
Signal	The headstage current output, or voltage output in current clamp mode. Units are pA or nA for current, or mV for voltage.
Filter	As above, but after gain and filtering. The output amplitude in volts is displayed.
Noise	The rms noise voltage on the Optopatch output after gain and filtering, giving a general purpose rms meter for noise or other measurements.
Osc freq	A precise readout, displayed in KHz.
Res	The effective resistance control setting, which takes the resistance range switch into account as well, plus any automatic adjustment made by the lock-in amplifier in tracking mode. Units are Mohms or Kohms as appropriate (but they are always Mohms when our standard headstage is used).
Cap	The effective capacitance control setting, which takes the capacitance range and offset switches into account as well, plus any automatic adjustment made by the lock-in amplifier in tracking mode. Units are always pF
Phase	The error signal generated by the phase-tracking circuit. The output range of +/10V covers a phase range of about +/-45 degrees. This is intended for information rather than precise measurement, so the display units are volts..

'METER'

The meter reads up to +/-19999 and is autoranging. The location of the decimal point is also adjusted as appropriate for the units that are currently being displayed.

'RC COMP' and Mode Select Switches

These three switches select the main operating modes of the Optopatch. They select whether the Optopatch is operating as a voltage or a current clamp (central switch), the current range (right-hand switch), and enable cell capacitance and series resistance (RC) compensation in whole-cell recording modes (left-hand switch). Some of the switch settings are interlocked with other controls on the Optopatch, to ensure correct operation.

For example, the track position on the RC comp switch, which selects automatic compensation and measurement of cell capacitance and series resistance, is active only if the phase switch is also on. Note also that in voltage clamp mode, conventional "series resistance" (RS) compensation, which attempts to compensate for steady-state errors in the cell membrane potential due to current flow through the electrode, is active regardless of the RC comp switch setting, but it can be disabled elsewhere (by the %RS compensation control) if RS compensation is not required. This arrangement has been adopted so that these two forms of compensation can be selected and controlled independently if preferred. In current clamp, the RC comp switch enables a bridge circuit to allow correction for the voltage drop across the electrode resistance due to the command current.

The central switch also has an I=0 position as well as voltage and current clamp settings. In this mode, the various command inputs, which select the command voltage in voltage clamp or command current in current clamp, are all inactive, so that the Optopatch GAIN

A control to vary the gain of the signal output (voltage or current) from the headstage, from 1 to 1,000 in a 1 2 5 sequence. The selected gain can be relayed to other equipment by a telegraph output on the rear panel, and it also takes the gain of the headstage into account in the voltage clamp modes.

The right-hand switch selects the operating current and cell capacitance ranges. In the patch position, the Optopatch operates as a true patch clamp amplifier with a full-scale current range of 1nA (100pA/V). The RC compensation and RS ("steady-state" series resistance) compensation facilities are inactive in this mode. The small cell and big cell positions differ primarily in the cell capacitance compensation ranges. In small cell mode, the cell capacitance range is up to 20pF, and the compensation signals are applied via the electrode capacitance compensation capacitor in the headstage. In big cell mode, the compensation range is up to 200pF, and the signals are provided via a larger (10pF) cell capacitance compensation capacitor, which is electronically connected to the input in this mode. The current range is 1nA (100pA/V) full scale in small cell mode and 100nA full-scale (10nA/V) in big cell mode. Please note these figures all apply to the standard optical headstage. Other headstages, with different current and capacitance ranges to these, can in principle also be supported.

'GAIN' Rotary Switch

The actual gain range of this switch is 1-500, but for convenience it has been calibrated in the appropriate units. In the patch and small cell voltage clamp modes, this corresponds to a range of 10-5,000mV/pA (black scale), and in the big cell mode it corresponds to a range of 0.1-50mV/pA (red scale). In current clamp the range is 10-5,000 mV/mV (black scale again). In earlier amplifiers the highest actual gain was 1,000, which was found to be unnecessarily high, so it was subsequently limited to 500, but even this is somwhat generous!

LED Indicators

The 100nA range LED is illuminated in the big cell voltage and current clamp modes.

The external LED is illuminated if a computer or other external device has full control of all the Optopatch switch functions. Please note however that switch functions can also be controlled individually as described later, and that in practice this form of control has proved to be more useful!

The overload LED will illuminate for about half a second if the output from the gain stage transiently exceeds +/-10V, or if high-frequency oscillations are detected on the command potential, either because of excessive series resistance compensation in voltage clamp mode or excessive electrode capacitance compensation in current clamp mode (in which case the appropriate facility is disabled for as long as the LED is iluminated). It will remain illuminated indefinitely if either condition persists. Note that the overload circuitry is fast enough to detect signal clipping caused by transient noise peaks, so it is a useful indicator of when (more) prefiltering is appropriate.

'ON'

If none of the other controls appear to be working, try pressing this one.

2.2 Middle Row

'EXT COMMAND' Switch

When this switch is on, signals applied to the command/10 and command /100 BNC sockets are included in the command potential.

'VHOLD' Control

A calibrated ten-turn potentiometer to set the holding potential in voltage clamp mode (it has no effect in current clamp mode). Full-scale is 200mV, and minimum is zero.

'VHOLD' Switch

A switch to select whether negative (the usual case!) or positive holding potentials are required. A centre off position is also provided.

'IHOLD' Control

A calibrated ten-turn potentiometer to set the holding current in current clamp mode (it has no effect in voltage clamp mode). Full scale is 20pA in patch and small cell modes, and 2nA in big cell mode. Minimum is zero.

'IHOLD ' Switch

A switch to select whether negative or positive currents are required. A centre off position is also provided.

'JUNC' Control

A calibrated ten-turn potentiometer to provide an additional holding potential adjustment from -200mV at its minimum setting to +200mV at its maximum setting, i.e. zero corresponds to five turns on this control.

'JUNC' Switch

A switch to disable the junction potential potentiometer if this facility is not required. It also selects search mode, in which a feedback loop slowly adjusts the voltage clamp command potential in order to keep the average current at zero, while not affecting fast currents. This is useful when searching for a good patch.

'LEAK' Control

A calibrated ten-turn potentiometer, active in voltage clamp mode, that subtracts a proportion of the command potential from the patch current, to allow purely resistive (leakage) currents to be removed from the current signal. It follows a resistance law, with a range of 1-101Gohm in patch mode or 100M-1.01Gohm in big cell mode.

'LEAK' Switch

A switch to enable the leak subtraction if this facility is required.

'SERIES RES' Control

A calibrated ten-turn potentiometer, also referred to in this manual just as the res control, for compensation of the series resistance in the whole-cell voltage clamp recording modes. Note that the setting of this control is used by both the RC and RS compensation circuits in the whole-cell (big or small) voltage clamp modes, and by the bridge circuit in the whole-cell current clamp modes. Also note that the control range is actually from 5% to 105% of full-scale, rather than from 0% to 100%. The adjustment range on the ten-turn dial is therefore from 0.5 to 10.5 turns, so that the dial readings remain correct. An output voltage proportional to the control setting (10V at 100%) is available on the rear panel.

'RES RANGE' Switch

This switch allows selection of a full-scale resistance of 10Mohm, 30Mohm or 100Mohm, although the maximum settings are actually 5% higher as described above. For quantitative measurements the lowest usable range is recommended. Alternative headstages may actually offer different ranges to these, but in that case the correct values will be shown on the meter.

'% CAP OFFSET' Switch

A switch to shift the range of cell capacitances that can be compensated, by either 50% or 100% of the currently selected full scale value.

CAP' Control

A calibrated ten-turn potentiometer for compensation of the cell membrane capacitance in the whole-cell voltage clamp recording modes. When the cap range switch is set to its highest value, the full-scale capacitance is 10pF in the small cell mode and 100pF in the big cell mode. Capacitances of up to double these values can be compensated by using the cap offset facility. As with the series res control, the actual adjustment range is from 5% to 105% of full scale, and an output voltage proportional to the control setting (10V at 100%) is available on the rear panel.

'CAP RANGE' Switch

This switch allows selection of a full-scale cell capacitance of 10pF, 30pF or 100pF in the big cell voltage clamp mode. In the small cell mode, the full-scale capacitance is always 10pF, regardless of this switch, as the meter will show. The meter will also correctly show the capacitance values for alternative headstages, which may offer different capacitance ranges. In all ranges, the cap offset switch can subtract either 50% or 100% of the full-scale capacitance, to allow over-range capacitances to be measured. For example, the cap offset switch can be used to shift the 100pF range, which is normally 5-105pF, to either 55-155pF at 50% offset or 105-205pF at 100% offset. Again, the meter display always deals with this correctly.

'PHASE' Switch

A three-position switch to activate the control electronics for the lock-in amplifier when in either the on or dither positions. If the lock-in amplifier is not in use, this switch should be off, in order to prevent the unnecessary generation of the internal switching signals. It should also be off to enable precharging and series resistance compensation, the use of which are (in our opinion!) inappropriate for membrane capacitance measurement by a lock-in amplifier. To use the lock-in amplifier, the phase switch must be in either the on or dither positions, and the RC comp switch must be in either the on or the track positions. Each of these four switch combinations gives a different operating mode!

With the RC comp switch in the on position and the phase switch in the on position, conventional lock-in amplifier operation is obtained. In this mode, and when the lock-in amplifier is properly set up, the real and imaginary phase outputs are (at least for reasonably small changes) proportional to changes in series resistance and membrane capacitance respectively, but they will need to be calibrated somehow, e.g. by use of the dither facility described below.

With the RC comp switch in the track position and the phase switch in the on position, our exclusive "track-in mode" is in operation. In this mode, the series resistance and membrane capacitance settings are automatically adjusted to maintain perfect nulling of the (sinusoidal) current signal, and the real and imaginary outputs from the lock-in amplifier now represent the linear and calibrated adjustments that have been made, i.e. the actual changes in the resistance and capacitance.

When the phase switch is in the dither position (or when the RC comp switch is in the track position and the phase switch is on), the resistance and capacitance dither inputs on the rear panel become active. When the RC comp switch is in the on position (for conventional lock-in amplifier operation), the dither facility can be used to calibrate the real and imaginary phase outputs of the lock-in amplifier, by applying known signal levels to the resistance and capacitance dither inputs. However, please note that to use the resistance dither input, the res dither level control on the rear panel must be in its fully anticlockwise (off) position. A 1V input on either input give an effective resistance or capacitance change of 1% of the currently selected full scale for that control.

When phase switch is in the dither position and the RC comp switch is in the track position, then the automatic phase tracking facility is activated (to use the dither inputs in track mode without phase tracking being in operation as well, the phase switch should be left in the on position). This requires a low-frequency sinusioidal signal to be applied to the resistance dither input, and for convenience an internal 70Hz oscillator is provided for this purpose, with its level being set by the res dither control mentioned above. In this mode, the switching phase of the lock-in amplifier is automatically adjusted to minimise the effect of any resistance changes on the imaginary phase (capacitance) output of the lock-in amplifier. However, please note also that when the RC comp switch is in the on rather than the track position, in order to give conventional lock-in amplifier operaiton, this signal is useful for adjusting the phase control manually, so as to zero the 70Hz signal appearing on the imaginary phase output (i.e. the same as is done automatically in the phase tracking mode). This condition ensures that the imaginary phase output corresponds as closely as possible to the capacitance signal, so that the effects of changes in the series resistance are minimised.

'PHASE' Control

A single-turn control to set the switching phases of the control signals for the lock-in amplifier. The total phase shift provided by this control is adjustable from 0 to 180 degrees relative to the internal oscillator signal, and is independent of frequency. The phase relationships of the design are such that control settings of around 90 degrees should be aimed for (by selecting an appropriate oscillator frequency).

2.3 Bottom Row

'HEADSTAGE' Connector

Signals from this connector communicate with the main electronics by an interchangeable driver board, which gives increased flexibility for the possible support of alternative headstage designs for other applications.

'FAST MAG'

A ten-turn control for the compensation of up to 15pF of electrode capacitance. This is the most important of the electrode capacitance compensation controls, and the extra resolution provided by a ten-turn control is useful here.

'FAST T'

A single-turn control to adjust the response time of the fast mag control, in order to obtain the best transient response to a command potential step.

'SLOW MAG'

A single-turn control for the compensation of up to 2pF of additional electrode capacitance, which may be charged more slowly on account it being in series with part of the electrode resistance. Not to be confused with cell capacitance compensation.

'SLOW T'

A single-turn control to adjust the response time of the slow mag control, in order to obtain the best transient response to a command potential step.

'PRECHARGE' LED

This indicator illuminates to show when precharging is active. For that to occur, the %precharge control must be away from its fully anticlockwise position, the RC comp switch should be on, and the phase control switch should be off. The amplifier should also be in voltage clamp and either small cell or big cell mode.

'RS' LED

This indicator illuminates to show when series resistance compensation is active. The same conditions apply as for precharging, EXCEPT that RC compensation does not necessarily have to be on.

'% PRECHARGE'

A single-turn control to set the extent to which the command potential in the whole-cell voltage clamp modes is modified to charge the membrane capacitance more rapidly during a potential transient. The RC enable switch must be on, and the series res and (cell) cap controls must be appropriately set for it to operate correctly. Note that under these conditions the capacitive charging currents are passed via a capacitor in the headstage, so do not appear on the current output. Therefore this control may not seem to have much effect, even though it may actually be substantially reducing the time required to change the membrane potential. Use of this facility is recommended, although the near-maximum settings are likely to be too severe in practice. It may also be necessary to make a slight adjustment to the res and cap controls in order to preserve a flat current trace. If no precharging is required under conditions in which it is available, this control should be turned fully anticlockwise until a click is felt.

'%RS COMPENSATION'

A single-turn control to set the extent to which the command potential in the whole-cell voltage clamp modes is corrected for the effect of "steady-state" (i.e. ionic rather than capacitive) currents. The neutralisation of these currents by the voltage clamp causes a clamping potential error in consequence of the resulting potential difference across the series resistance. Larger settings of this control can cause instability, so use with caution. However, our oscillation killer circuitry makes this control much safer to use than on other patch clamps, since it automatically disables RS compensation before the instability becomes great enough to drive the command potential to lethal amplitudes. This control is active whether or not the RC enable switch is on, and the series res control should be appropriately set for it to operate correctly. If RS compensation is used in the absence of RC compensation, then the capacitive charging currents in response to step changes in membrane potential (now directly observable in the absence of RC compensation) will be shortened but will be of correspondingly greater amplitude. If RS compensation is not required, this control should be turned fully anticlockwise until a click is felt.

'LAG'

A single-turn control to slow the response of the RS compensation. This can permit larger settings of the %RS compensation control, which overall gives better compensation of more slowly-varying currents.

'PREFILTER'

The gain stage of the Optopatch can amplify signals enormously if required, but the wide bandwidth of the preceding stages (at least 100KHz for both current ranges) means that a fair amount of high-frequency noise may be present. This limits the amount of gain that can be applied, since the noise increases the peak amplitude of the signal. More gain can therefore be applied if the signal is filtered first, but steep-cut filters, as needed for patch clamping, give best results when presented with relatively high signal levels, since their signal-to-noise ratios are generally not as good as that of a well-designed gain stage. The best solution to this dilemma is therefore to have a simple prefilter (in our case two-pole) before the gain stage, with the main filter after it. As well as an OUT position, filter frequencies of 30KHz, 10KHz or 3KHz can be selected.

'FILTER SOURCE'

A switch to divert the input from the output Bessel filter, so that it can process the capacitance output signal from the lock-in amplifier, instead of the amplified current signal (or voltage signal in the current clamp modes).

'FREQ VALUE' Selector

A seven-position rotary switch to set the cut-off frequency of the four-pole Bessel filter, allowing the frequency set on the freq range selector to be multiplied by values between one and ten.

'FREQ RANGE' Selector

A five-position rotary switch for the output Bessel filter, to give base frequencies of 1Hz, 10Hz, 100Hz, 1KHz and 10KHz, which can be multiplied up to tenfold by the freq value selector. The lowest ranges are likely to be of assistance in making high-resolution membrane capacitance measurements. (Note that an internal accessory board can be fitted inside the Optopatch, to convert the standard four-pole filter to an eight-pole type, as preferred by some users for single-channel recordings.) The selected frequency (range and value) can be relayed to other equipment by a telegraph output on the rear panel.

3 Rear Panel

3.1 Expansion Slots

There are three expansion slots along the top of the rear panel. One of these is occupied by the headstage driver electronics, and the other two are general-purpose. They are intended for current or future optional accessories such as the eight-pole Bessel filter upgrade.

The reason for putting the headstage driver electronics on a separate board was to allow the basic Optopatch design to accommodate other types of headstage for different applications, e.g. a standard resistive feedback design for higher current ranges, although it has to be said that we have yet to find a need for this! It also facilitates calibration adjustments, as four presets on the driver circuit board (described in more detail elsewhere) are directly accessible from the rear panel. The adjustments are for headstage voltage offset, current-passing offset, and bandwidth in the patch and (big) cell current ranges. These presets should be adjusted to match the headstage in use, while calibrating the four optical sensitivity presets on the headstage at the same time. The driver subpanel also has switch for enabling the reference input in the headstage. In addition, there is a 9-pin D connector, which provides a variety of outputs that may be useful for calibration.

3.2 BNC Sockets and Associated Switches

The following BNC inputs and outputs are available. Please refer to the appropriate section of the manual for more detailed information than that given here.

'FAST I OUT'

This is the equalised signal from the current-to-voltage converter circuit, which we recommend for any applications that exploit the unusually wide (at least 100KHz) bandwidth of the Optopatch.

'VBATH'

The bath potential input (unity gain). This and other inputs can be left floating if not used, as all inputs have a relatively low input impedance of 10K ohms.

'FREQ TEST IN'

This input drives an internal integrator that converts the input voltage to current pulses that are applied directly to the input via the electrode capacitance compensation capacitor in the headstage (this is referred to as a freqtest input, since one use is to set or verify the frequency response of the system). A 10V voltage step will give a 10nA current step in both the 1nA and 100nA current ranges, so the signal level should be chosen appropriately! The capacitative coupling means that this input is suitable for AC signals only, but the full-scale current amplitude extends down to 200Hz. This facility is very useful, but it is normally used in conjunction with the internal oscillator, and an external input will rarely if ever be needed.

'FREQGEN OUT'

The output from the internal oscillator, at 1V full scale, i.e. ten times higher than the internal signal level.

'COMMAND/100 IN'

Signals here are attenuated 100 fold, and are added to the command signal (command potential in voltage clamp or command current in current clamp) if the ext command switch is on.

'COMMANDX10 OUT'

This output is at x10 gain, and does not contain any bath and junction potential offsets.

'COMMAND/10 IN'

Signals here are attenuated tenfold, and are added to the command signal if the ext command switch is on.

'HEADSTAGE COMMANDX10 OUT'

This output is also at x10 gain, and includes bath and junction potential offsets in voltage clamp mode, plus the effect of any precharging that may have been applied. As its name suggests, this is the command signal that is actually sent to the headstage.

'HOLD/10 IN'

An input that is attenuated tenfold, and which is permanently active in both voltage clamp and current clamp modes. Any signals provided here are added to any that are provided by the front panel holding potential and holding current controls.

'CAP OUT'

This output represents the setting of the cap control, at 10V full scale. When the RC enable switch is in its track position, the output from the automatic compensation circuitry is added in as well, so this signal normally represents the TOTAL membrane capacitance. Since the automatic component can also be as much as + or -5V, there is a theoretical possibility that the output voltage limit of about 14V will be reached, although the automatic compensation circuit will still work correctly under these conditions, and the automatic component can still be read from the imaginary phase output (see below). Note that correct interpretation of the output voltage requires knowledge of the settings of the capacitance range and capacitance offset switches, and also depends on whether the Optopatch is in small cell or big cell mode, but these parameters can be read and/or set via the computer connector.

'RES OUT'

This output represents the setting of the series res control, at 10V full scale. The same remarks apply (where appropriate) as for the cap output.

‘IMAGINARY PHASE OUT’

An output that is a measure of changes in membrane capacitance, provided by the lock-in amplifier. When the RC comp switch is on but not in the track position, it is affected directly by the gain setting, and at unity gain a 1V rms sine wave gives a DC output of 1V. When the RC comp switch is in the track position, an output of 1V represents a change in capacitance of 10% of the full-scale value of the cap control. The output voltage is independent of the gain setting in this mode, but the gain setting affects the response time of the feedback loop. Higher gain quickens the response, up to a limit determined by system stability considerations, but response times below 10msec can easily be achieved under most conditions. Thus in track mode, one can either read the capacitance changes from this output (preferred for high-resolution measurements) or the total capacitance from the cap output.

‘REAL PHASE OUT’

The corresponding output for changes in series resistance. The same remarks apply as for the imaginary phase output.

‘CAP DITHER IN’

An input that is active when the phase switch is in the dither position (or in the on position when the RC comp switch is in the track position). A 1V input changes the effective setting of the cap control by 1% of full scale. Note that, to simulate an INCREASE in membrane capacitance, a NEGATIVE input is needed to reduce the control setting by that amount. The permissible signal range is +/-10V.

‘RES DITHER IN’

This input is basically the counterpart of the cap dither input, and operates in the same way, but it is also associated with the automatic phase compensation system. When the phase switch is in the dither position AND the RC comp switch is in its track position then a low-frequency sinusoidal dither signal applied here will cause the switching phase of the lock-in amplifier to be adjusted so as to minimise the level of this signal on the cap output. This signal can be generated by an internal 70Hz oscillator, and its amplitude can be varied by the res dither level control on the rear panel. When the oscillator is in use, the 70Hz output signal is available on this socket. To allow an external input to be used instead, the res dither level control should be rotated fully anticlockwise to operate the built-in “off” switch.

‘GAIN TELEGRAPH OUT’

The output voltage here represents the overall gain setting. An output of 2.5V represents a sensitivity of 1mV/pA in voltage clamp mode and 1mV/mV (unity gain) in current clamp mode (although the minimum gain in current clamp mode is actually ten, since the membrane voltage is already amplified by this amount before being sent to the gain control rotary switch). Each gain increment, in a 1-2-5-10 sequence on the gain selector, increases the output by 0.5V, and the current range of the headstage is taken into account in computing the output in voltage clamp mode. For our standard headstage, this gives an output voltage range of 1-5V in big cell voltage clamp (100nA current range); 4-8V in patch and small cell voltage clamp (1nA current range); and 4-8V in current clamp (both current ranges);

'FILTER TELEGRAPH OUT'

This output represents the setting of the filter frequency control, at 10V full scale. It consists of one component of 0, 2, 4, 6 or 8V, corresponding to a frequency range of 1Hz, 10Hz, 100Hz, 1KHz or 10KHz, to which is added another component of between 0.2 and 2V, proportional to the frequency multiplier of x1 to x10 set by the frequency value control. For example, a frequency of 3KHz gives an output of 6.6V.

'HEADSTAGE OUT/LOCKIN GATE IN'

The current output (or voltage output in current clamp modes, which is actually ten times the membrane potential) from the headstage. Note that in voltage clamp mode, the signal will also have been subjected to prefiltering and leak subtraction (if these facilities are in use). However, if an unsubtracted and unfiltered current signal is also required while these facilities are in use, it is always available from the FAST I OUT socket.

If required, this socket can be switched to provide a gating control input to the lock-in amplifier. In that case, a logic high level (nominally +5V) will switch off the lock-in amplifier and disconnect the frequency generator from the command potential. This is to allow command potentials to be applied, e.g. to stimulate secretory processes, while minimising the generation of artefactual resistance and capacitance signals. In tracking mode, the real and imaginary phase outputs maintain their previous values while the lock-in amplifier is gated off, and they will rapidly assume their new values when the gating signal is removed.

'GAIN OUT'

The headstage output signal as described above, but after amplification according to the gain control setting.

'25KHz OUT'

The same as the gain signal, but after three-pole Bessel filtering at a fixed frequency of 25KHz.

'FILTER OUT'

The output signal after four-pole (or eight-pole with internal accessory board) Bessel filtering at the chosen frequency.

3.3 Other Connectors

In addition to the 21 BNC connectors, there are two 37-way D connectors, a 5 pin circular connector (180 degree locking DIN type) and a 4mm grounding socket.

The circular connector is an alternative bath (or other reference) potential input, and it also has +15V and -15V power supply connections, to allow a powered preamplifier to be used here. The input resistance is 10K ohms, going to a virtual ground. The gain can be reduced by including an appropriate resistance in series, e.g. 90 K ohms if a preamplifier with a gain of x10 is used. Any signal applied here is added to any signal applied at the BNC input. However, patch clamps supplied from July 2000 onwards have incorporated a differential high-impedance input in the headstage, which effectively makes the original function of this connector redundant. We are retaining it because it could be useful for a variety of other purposes in the future.

The first of the two 37 way connectors is a socket which is a duplicate of all the BNC sockets except the bath potential input and the fast current output. The upper row of 19 pins carries the remaining signals, alternating in a bottom row/top row BNC socket sequence, and going in the same direction along the panel, i.e. from right to left we have freqtest in, freqgen out, vcommand/100 in, vcommandx10 out, etc.. This corresponds to a pin sequence of 1,2,3,4,5 etc., using the numbering convention for this connector. All 18 pins in the shorter row (20-37 inclusive) are connected to signal ground.

The other 37 way connector is to allow a computer (or other external equipment) to read and/or set the major controls of the Optopatch. Please note that this connector is a plug, which means that the pin numbering is in the opposite direction to that of the socket. The sequence of connections given below is the sequence in which they appear in a ribbon cable when an IDC (insulation displacement, i.e. the type that presses directly onto the cable) D connector is used, and numbering from right to left when viewed from the rear panel, since the various functions are grouped in this way. The D connector pin numbering, which is visible on the connector itself, is shown in brackets for reference.

1	(19)	+5V out
2	(37)	0V out
3	(18)	Osc frequency out at 1V/KHz (or 0.1V/KHz on the 10-100KHz range)
4	(36)	Logic high out on the 10-100KHz oscillator range
5	(17)	5V out (normal), 0V or >5V, according to type of headstage in use
6	(35)	Auto component of phase control voltage (for reading only)
7	(16)	Auto component of res control voltage (can either read or control)
8	(34)	Auto component of cap control voltage (can either read or control)
9	(15)	Filter variable control voltage (can either read or control)
10	(33)	Lock-in amplifier phase voltage (can either read or control)
11	(14)	Multiply gain x2 input if logic high
12	(32)	Multiply gain x5 if high
13	(13)	Multiply gain x10 if high
14	(31)	Multiply gain x10 if high (the effects of these four inputs are cumulative)
15	(12)	Set 10Hz filter range input if high (also overrides any higher range setting)
16	(30)	Set 100Hz filter range if high (ditto for this and the other ranges)
17	(11)	Set 1KHz filter range if high
18	(29)	Set 10KHz filter range if high (none of these high selects 100KHz)

19	(10)	Selects computer control of the following switches if high, otherwise reads
20	(28)	Search mode (on junction potential switch)
21	(9)	Vclamp
22	(27)	Iclamp (neither of these two gives I=0)
23	(8)	Big cell
24	(26)	Patch (neither of these two gives small cell)
25	(7)	RC enable off
26	(25)	RC auto (neither of these gives RC on)
27	(6)	10pF
28	(24)	100pF (neither of these gives 30pF)
29	(5)	Cap offset off
30	(23)	Cap +100% (neither of these gives cap +50%)
31	(4)	Resistance x1
32	(22)	Resistance x0.1 (neither of these gives resistance x0.3)
33	(3)	Phase off
34	(21)	Dither (neither of these gives phase on)
35	(2)	ground
36	(20)	ground
37	(1)	ground

The four signals on IDC connection numbers 7-10 inclusive as defined above (i.e. DIN connector numbers 16, 34, 17 and 33) are effectively bi-directional. They provide the signals as described, via a source impedance of about 10K ohms, and if they are to be used as outputs, they should be connected to inputs of 1 Megohm or higher input impedance. In this regard, note that buffered (100 ohm source impedance) outputs of three of these signals, i.e. the series res, cap and filter frequency control settings, are available on BNC sockets and on the other D connector, so the use of those outputs for monitoring is generally to be preferred. The phase control voltage, which is not available elsewhere, is 0 to 10V for a 0 degree to 180 degree phase shift. (Although 5V gives an exact 90 degree phase shift, the overall voltage relationship is not precisely linear, but it is predictable and therefore capable of linearisation by software.) Alternatively, if these signals are connected to low impedance voltage SOURCES, then the source voltages will determine these signal levels instead. This allows external control of lock-in amplifier phase, filter variable frequency, and the auto components of the res and cap control settings. In connection with the following paragraphs, please note that overall computer control does not have to be asserted in order to drive these functions from external sources.

We had originally intended an all-or-nothing approach to computer control of the switch operations, in which computer control of all of them is asserted by placing a logic high level on IDC connection 19 (DIN connector pin 10). However, in practice there has been much more interest in controlling just some of the switch functions individually, and the design does also permit this. The only restriction is that when overall computer control is not being asserted, the associated front panel switch for that function has to be kept in a specific position so that it does not interact with the external control signals. In all cases the front panel switch is a three position one, and the switch needs to be kept in the central position. To summarise, the following switch functions can be controlled individually, by asserting logic high levels on the appropriate connector pins as listed above.

Junction potential (search mode only), if panel switch is in the off position

Voltage/current clamp, if panel switch is in the I=0 position

Patch/cell, if panel switch is in the smallcell position

RC comp, if panel switch is in the on position

Cap range, if panel switch is in the 30pF position

Cap offset, if panel switch is in the +50% position

Res range, if panel switch is in the 30M position

Phase, if panel switch is in the on position

Please note that there is an additional restriction concerning holding potentials and currents, which arose because we provided independent front panel controls for voltage and current (as opposed to a single control) as a design upgrade rather than as part of the original design. The effect of this is that when the voltage/current clamp switch is controlled externally (with the front panel switch in the central I=0 position), the front panel holding potential and holding current controls both remain switched off. However, the required potential or current can of course be provided from an external source via the external hold and/or command/10 and command/100 BNC inputs, and arguably this may make more sense anyway when any form of external control is being used.

Finally, it is also possible to control the output Bessel filter externally without asserting overall computer control, by keeping the freq scale rotary switch in the 10K position and selecting the appropriate lower range if required by putting a logic high level onto the corresponding connector pin as specified above. The setting of the freq value control can be overriden by applying an analogue voltage of 0-10V on the filter variable control voltage pin of the connector. External control of the gain is currently not possible unless overall computer control is asserted, but in principle the internal logic could be reprogrammed to allow this, so please ask about this if you may be interested. Also note that absolute gains of up to 1,000 can be selected externally (by driving all four gain control inputs high), whereas the highest gain that can be set by the front panel rotary switch corresponds to an absolute gain of 500. We originally provided a maximum absolute gain of 1,000, but in practice it proved to be unnecessarily high, so this position was removed from the switch, but for historical reasons it can still be selected externally.

4 Circuit Operation

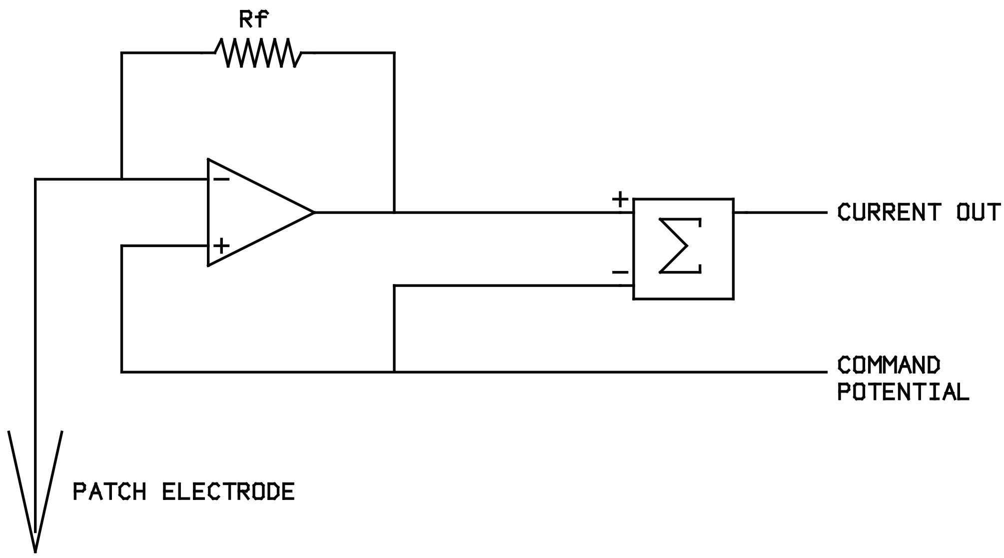

Figure 2 - Resistive Feedback Headstage.

4.1 The Optical Headstage

The headstage design is proprietary, and we're still not ready to say too much about its operation just yet, but we would like to publish full technical details in an appropriate journal in due course. We were (at last!) granted a patent for it in August 2000. However, we do explain the basic operating principles here.

The signal-to-noise problem with patch clamping is not just one of accurately passing very small currents - although standard feedback arrangements and components with very low leakage currents can look after that - but also of accurately measuring them. In the standard patch clamp headstage arrangement shown in Fig. 2, in which a feedback resistor is used to pass current, any difference between the command potential and the electrode potential causes a current to be passed through the resistor in such a direction as to oppose this difference. The current through the feedback resistor causes a potential difference to appear across it, which can be monitored by subtracting the command potential from the amplifier output. Unfortunately, resistors with values in the range normally used in electronic circuits, i.e. up to a few Megohms, generate levels of thermal noise (due to the random movement of electrons within them) that are large compared with the voltages that represent the patch currents. The noise voltage actually increases as the resistance increases, but only as the square root, whereas the voltages that represent the patch currents increase linearly with the resistance. The overall situation thus improves only with the square root of the resistance, and this favours the use of very high-value feedback resistors, of ten gigohms or more, but the effects of stray capacitance then become very important. Even minute amounts will significantly attenuate the high-frequency response, and practical circuits must be designed carefully so as to control and equalise these effects. Part of the problem is that some of the capacitance tends to be distributed along the resistance, which means that neither it nor the resistor itself can be accurately modelled as single components. Instead, one is dealing with some sort of network, which significantly complicates the equalisation, and it may have adverse noise implications too. However, the equalisation circuitry is not shown in Fig. 2, because we're dealing with an ideal resistor in this model.

Although resistive headstages can be made to work well, the alternative method of using a capacitor to pass the patch currents has become a viable alternative. This method is illustrated in Fig 3. Capacitors don't generate thermal noise, so in principle capacitive (also called integrating) headstages are somewhat quieter than resistive ones, but other problems must be solved in order to achieve satisfactory overall performance. First, a steady current will cause the capacitor to charge, so some method of discharging it periodically must be included. Second, the output voltage is now the time integral of the current, and a differentiation must be performed in order to recover the actual current waveform. Such circuits necessarily have very high gain at high frequencies, so they must have both low noise and low susceptibility to interference pickup in order to realise the potential benefits of the capacitive technique. Finally, the performance of the capacitor itself must be exemplary. In view of the very high AC gain levels any less than ideal capacitor behaviour, such as non-linearity with voltage or microphonic pickup, become increasingly important. Although these problems have been overcome, we didn't relish the thought of going down that road ourselves.

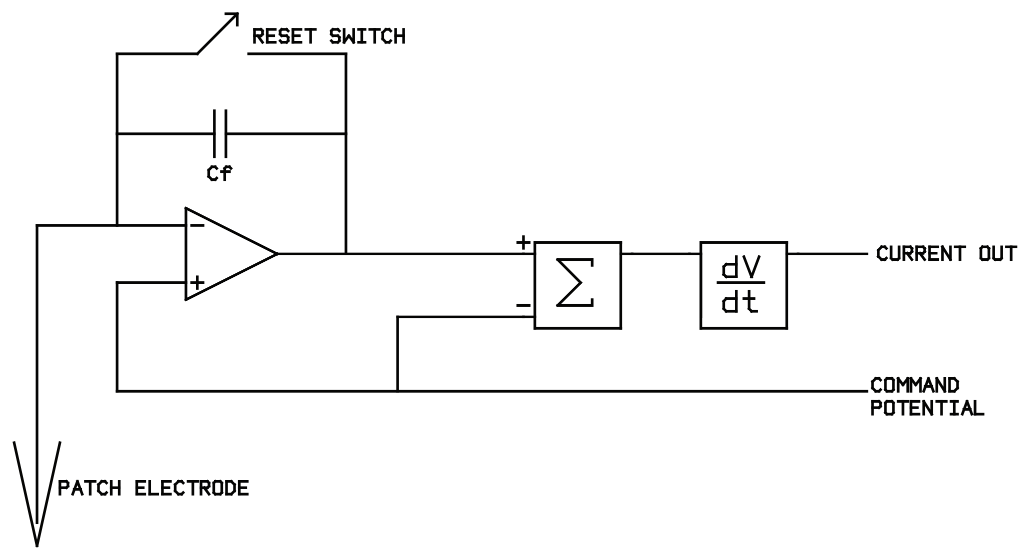

Figure 3 - Capacitive Feedback Headstage.

This is where the optical technique comes in. We liked the idea not just because it was new and different - although that was certainly attractive - but because it seemed such a clean solution to the problem. This is not to say that the optical technique doesn't pose problems of its own, but rather that the nature of those problems doesn't threaten the main performance requirements of high bandwidth and low noise. The main practical problems to solve with the optical approach are those of calibration and linearity, which can be tackled successfully in ways that don't compromise the performance of the system in respect of noise and bandwidth.

In the Optopatch, currents are passed by shining light onto miniature photodiodes which are connected between the differential inputs of the headstage amplifier, as shown in Fig. 4. This method gives the advantage of steady current-passing that the resistive method provides, but in this case there is no resistance to contribute any (additional) thermal noise, so the theoretical performance is the same as the capacitive circuit. The light is produced by passing linearised control currents through light-emitting diodes, which are also incorporated within the headstage in order to achieve a compact and totally enclosed optical design, although the drive electronics are on the daughter board in the main enclosure in order to minimise the size of the headstage. By using different illumination pathways of different overall efficiency, we have been able to implement both the current-passing ranges of +/-1nA and +/-100nA full scale by the optical technique. Although the optical method doesn't offer any significant advantage over the resistive method in the higher current-passing range, it simplifies the system design to have both ranges operating in the same way, but we should emphasise again that the overall system design can in principle support other types of headstage as well.

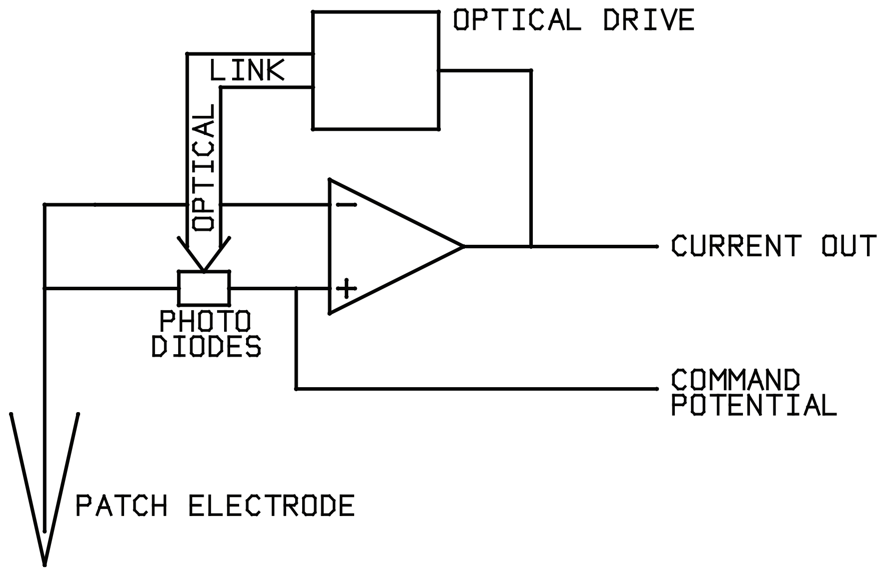

Figure 4 - Optical Feedback Headstage.

In the circuit of Fig. 4, any difference between the command and electrode potentials generates an error voltage at the output of the headstage amplifier. This provides the input to the optical current-passing stage, causing a photodiode current that tends to restore the input potential difference, and the amplifier output voltage is a direct measure of the current. The operation of the circuit is thus similar to a resistive headstage, except that the command potential does not appear on the amplifier output, and so does not need to be subtracted from it.

4.2 Membrane Capacitance Measurement

Actually there's not too much we can say here, because our objective of making measurements of membrane capacitance simple, accurate, fast and sensitive (and of course in real time too), determined the design of much of the main electronics, so there are relatively few aspects that we can describe separately. Instead, most information will be given in connection with the detailed description of the system operation. However, we can explain the general approach in this section.

Basically, we have provided facilities for the accurate measurement of membrane capacitance (and also of the series resistance, which needs to be determined at the same time), by extending the operation of the compensation controls that patch clamps traditionally provide for removal of the membrane capacitance charging currents from the recorded current signals. In our system, these controls are adjusted to (approximately) their correct settings, and the system can then be switched into an automatic mode. In this mode, a built-in lock-in amplifier, which measures the currents produced by modulation of the command voltage by the built-in frequency generator, generates error voltages that maintain the control settings at their correct values. These voltages are direct linear measurements of changes in membrane capacitance and series resistance, so no software or other signal-processing is required in order to follow changes in either parameter.

The lock-in amplifier technique is very powerful, and we'll discuss it in detail in a later section, but for now we just need to note that it involves gain-switching in synchrony with the command potential frequency, and for best operation it is important that the phase of the switching is set correctly relative to that of the command frequency. The correct phase depends on the membrane capacitance and series resistance, so it may change during an experiment. We have therefore provided an additional facility, using a low-frequency modulating signal, that automatically maintains the correct phase at all times.

These modifications give lock-in amplifiers the user-friendliness of software methods for membrane capacitance measurement. The sensitivity of these amplifiers can be extraordinarily high, and although in principle anything can be emulated in software, the amount of data and processing power required to emulate their performance would be very considerable. In some respects it is difficult to say how unique our system is, because the software writers might be able to argue that their algorithms are doing the same thing. However, this is difficult for us to assess, because it is not always clear what the algorithms actually are! Nevertheless, we believe we are the first to design an integrated solution to the problem in hardware, and the way we have gone about it most certainly is novel. We don't want to make extravagant claims about the performance of our system, but its results on model systems do give us some cause for optimism. Emulation's can be great, but on the other hand you can't beat the real thing.... Details of its operation and performance have now been published (Johnson, Thomas and Kros, 2002), and we can provide a pdf of the paper to anyone for personal use on request.

4.3 Current Clamping Improvements

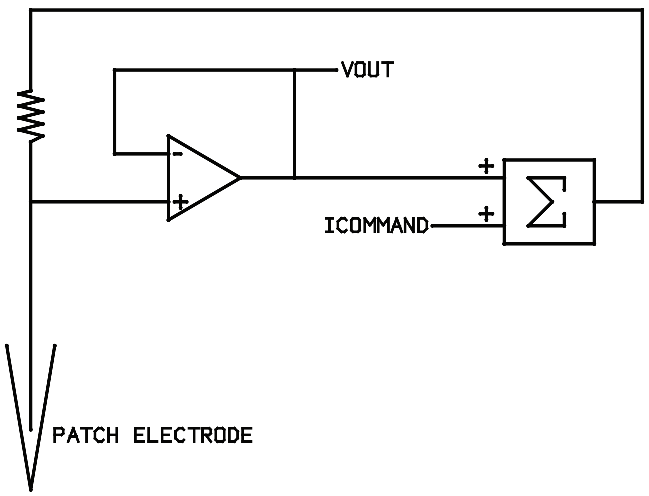

Figure 5 - Ideal Current Clamp.

Another fundamental improvement that we have been able to incorporate into the Optopatch is the ability to perform TRUE current clamping. We had been too busy concentrating on the other innovations to give current clamping much thought until Dr. David Ogden (MRC Mill Hill, London) kindly pointed out to us in September 1995 that the standard patch clamp current clamping circuit has performance limitations that are often detectable in practice. Dr. Ian Forsythe (Leicester, UK) independently mentioned the same problem to us a few weeks later, which further encouraged us to investigate it, and the following paragraphs explain both the problem and its solution. It turns out that the optical feedback makes implementation of the solution even cleaner than it would be otherwise, but in order to emphasise the generality of the solution, the discussion assumes that a high-value resistor is used to pass the current.

First of all, it is easy to design a true current-clamping circuit, and Fig. 5 shows all that is needed. The headstage amplifier is connected as a voltage follower, and the current is provided by a voltage source connected to a resistor. The load resistance into which the current is passed is the resistance of the electrode in series with that of the cell, and for the purposes of this discussion we shall assume that the electrode resistance is sufficiently low compared with the cell resistance that it can be ignored (in practice that may well not be true, but the problem can be dealt with in other ways, as described in the section on resistance and capacitance compensation, so for this discussion we can indeed ignore it). To approximate to a true current source, the value of the current-passing resistor should be high compared with the load resistor, so that the voltage at the amplifier input remains low compared with the current-passing drive voltage, since the current flow is actually given by the difference between these two voltages. With the optical current-passing method, this condition is met automatically, but when a resistor is used, the equivalent effect can be obtained by adding the input voltage to the current-passing command voltage, as Fig. 5 shows. So why don't other patch clamps work this way?

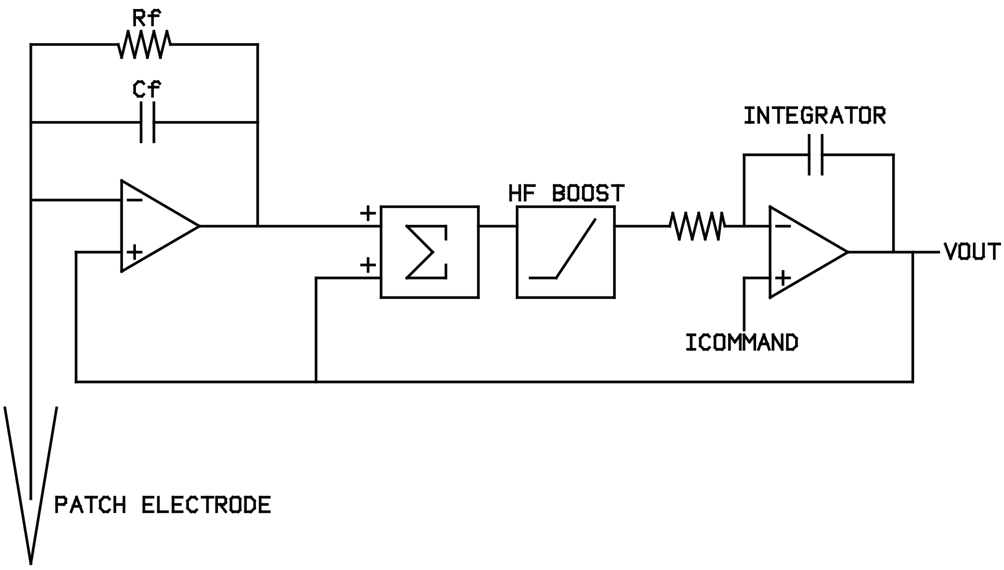

Figure 6 - Conventional Current Clamp

The answer is one of convenience, since it is probably fair to say that current-clamping was added to the standard patch clamp design as something of an afterthought. The problem here is that the electrode in Fig. 5 is connected to the noninverting input of the headstage amplifier, whereas for voltage clamping it needs to be connected to the inverting input as shown in Figs. 2-4. To prevent the need to reconfigure the input amplifier when switching between voltage clamp and current clamp, patch clamp amplifiers conventionally perform current clamping with the circuit shown in Fig. 6 (we show here a resistive headstage for generality, but the other headstage types are equivalent in this respect). It consists of the basic voltage clamp circuit of Fig. 2, but this time including the necessary frequency compensation stage as well, and with the output feeding back to the voltage command input via an integrator. A feedback capacitance, Cf, is shown in parallel with the feedback resistance, Rf, not just because some capacitance here is unavoidable, but because it is usually deliberately increased to give a well-defined high-frequency rolloff, that can then be compensated more precisely by the subsequent high-frequency boost stage as shown (see Sigworth, 1995, for a full theoretical treatment). As far as the frequency response is concerned, we can therefore notionally remove Cf and the boost circuit, and regard the integrator as being connected directly to the headstage. Just as in voltage clamp mode, the headstage output - after subtraction of the command voltage - is the voltage across Rf, which is a direct measure of the current through it. Any difference between this voltage (current) and the noninverting input of the integrator, to which the current-passing command voltage is applied, is fed back as an error signal to the command voltage input of the Figure 5 - Ideal Current Clamp.headstage. The effect of this circuit is therefore always to maintain the electrode at a voltage such that the required current is passed through the feedback resistor, and hence into the electrode.

An integrator is used in order to introduce significant low-pass filtering into the overall feedback loop. The filtering is necessary because the bandwidth of the signal from the headstage, even after equalisation by the boost circuit, is likely to be no more than 100KHz at best, and in practice it may be considerably less. To ensure stability of the overall feedback loop, its bandwidth needs to be constrained to a still lower frequency, of perhaps only a few KHz, whereas general amplifier circuits can easily have unity-gain bandwidths in the MHz range. In particular, this applies to the input stage in Fig. 5, where the feedback is a direct connection, since the amplifier is just a voltage follower in this circuit. There is only the bandwidth of the input amplifier itself to consider in this case, and to push that well into the MHz range is a trivial matter.

Unfortunately, the low bandwidth of the Fig. 6 circuit causes a problem here, for the following reason. At frequencies beyond the open-loop bandwidth, there is no longer any feedback, so the amplifier input impedance falls to the value given by the parallel impedance of Rf and Cf (since the output of the input amplifier stage then acts as a signal ground). The effect of feedback around the circuit is to raise this impedance proportionately with the amount of feedback applied, and very substantial amounts of feedback would be required over the entire operating frequency range in order for the circuit to approach the performance of that in Fig. 5. As the frequency falls, the gain of the integrator rises, but the gain of the boost circuit falls, so the open-loop gain - and hence the amount of feedback - actually does not begin to increase until the corner frequency of the boost circuit is reached, i.e. the frequency at which the impedance of Cf equals that of Rf. What this actually means is that lowpass corner frequency of the integrator must be less than the combined bandwidth of the headstage and the high-frequency boost circuit if there is to be any feedback at all at frequencies above the corner frequency of the boost circuit!

There is also yet another problem to consider, which concerns the headstage's input capacitance, and in practice this problem may also be serious. It may not be immediately obvious that there is a problem, since the input capacitance neutralisation circuit (not shown in Fig. 6) continues to compensate correctly for capacitances between the input and ground, so the system's performance remains the same in this respect as for a step in the command voltage under voltage clamp. However, there is another component of the input capacitance, which appears between the inverting and non inverting inputs of the input amplifier. This capacitance does not load the input to a significant extent in voltage clamp mode, because local feedback, mainly provided by Cf at high frequencies, acts quickly to reduce the differential input voltage towards zero in response to a command voltage step, so it is effectively completely compensated. However, in current clamp, a step change in the voltage at the electrode will initially appear entirely as a differential input voltage, which will be reduced much more slowly by the current feedback loop, with a time constant determined by the integrator. Under these circumstances, the current to charge the differential input capacitance (which could easily be on the order of 10pF in a typical circuit), will initially be drawn from the electrode. The fact that the action of the integrator then discharges the capacitance again actually makes matters even worse, because this means that the electrode effectively has all the charge returned to it later on!

The combination of these effects is most certainly noticeable in practice. Consider for example, the case in which passing a steady current in whole-cell mode causes the cell to fire action potentials. Under these conditions both the retention of high input impedance up to a relatively high frequency and the preservation of a good transient performance are particularly important, and what can happen in practice if these conditions are not met is, in the words to us of one current-clamper, "...as if the patch clamp is trying to voltage clamp the action potentials."

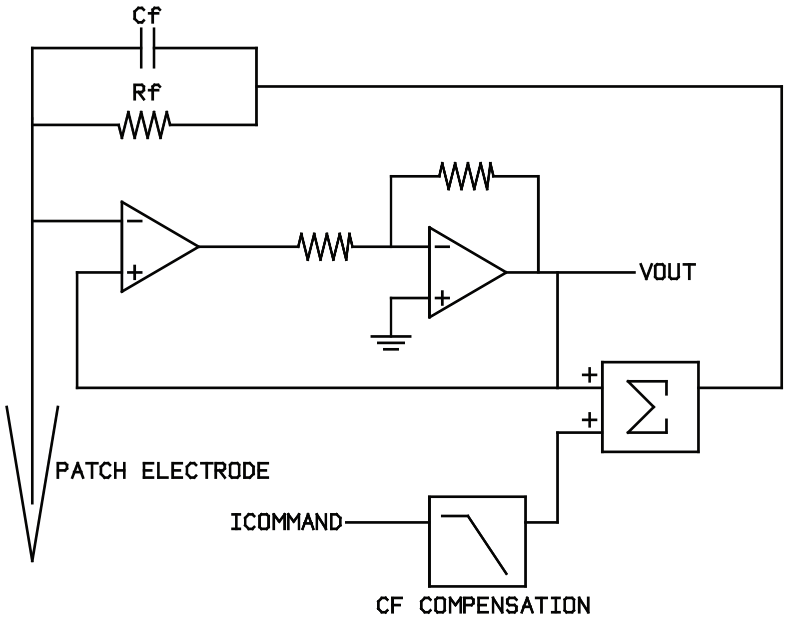

Figure 7 - True Current Clamp

Obviously the current-passing performance of a particular patch clamp amplifier will depend on the actual component values, which will in any case also depend on the current range selected, and without having our own version of Fig. 6 to analyse, it does not seem appropriate to quantify the problems any further here. The only general point we can note is that when a headstage is operating in the capacitive (integrating) mode, where Cf may be higher than for the resistive mode, it might require boost compensation up to a higher frequency for comparable performance, but on the other hand the frequency response is likely to be easier to correct in this mode, so in practice there may not be much if any difference. Since we are using another current-passing method altogether in the Optopatch, we prefer to be neutral on this subject!

A much enhanced general implementation of current clamping is shown in Fig. 7. Swapping the connections to the inverting and non inverting inputs of the headstage amplifier when switching between modes would clearly be inconvenient, but we can obtain the same effect by inverting its output instead. Feeding the inverted output back to the command input of the headstage as shown effectively reverses the polarity of the inputs, making the headstage amplifier into a voltage follower. Although the feedback has to be applied via another amplifier stage, the signals here are all of low impedance, so it is easy to ensure wide closed-loop bandwidth, resulting in a performance that is to all intents and purposes identical to a conventional voltage follower. We can now pass current directly though what was Rf, driven by an amplifier which corrects for the input voltage in the same way as in Fig. 5.

Two possible practical problems intrude. First, the capacitor Cf no longer serves any useful purpose, so ideally it should be removed, but in practice this may involve additional switching in the headstage. However, in practice it could stay connected across Rf, and is effects could be compensated for by a low-pass filter (of characteristics that are complementary to those of the high-frequency boost circuit required for equalisation in voltage clamp mode) in the current clamp command input. Such filters can be made to have a much more ideal frequency response, giving much more precise correction, and in any case the correction is outside the main feedback loop, so does not have any influence on its performance. However, it would still be better for Cf to be relatively small, since it still lowers the uncorrected impedance of the current source. The second possible problem is that in the conventional headstage design of Fig. 6, the subtraction of the command voltage from the current output may take place in the headstage itself, which is not appropriate for the current clamp circuit of Fig. 7. The fact that we were made aware of the current-clamping problem during the design stage of the Optopatch made it much easier for us to cure it than might have been the case otherwise.

The headstage design in the Optopatch, as shown in Fig. 4, effectively consists of just the input amplifier stage, in which the connection between its output and the current-passing circuitry (equivalent to Rf and Cf in Fig. 3) is made indirectly by interconnections that go back to the main electronics, plus part of the optical drive system. This allows the circuit to be reconfigured between the standard voltage clamp and true current clamp modes without the need to perform any switching within the headstage itself. Although the configuration of the optical headstage made it particularly easy to implement true current clamping, equivalent arrangements could nevertheless be made for the other types of headstage as well.

Since this section was first written, Magistretti et al (1996) published an article in TINS (see bibliography), which also discussed this problem, and it provides useful further reading material on this subject. A more detailed report has now also been published (Magistretti et al., 1998).

5 Differential Operation - Advantages and Pitfalls

Patch clamps supplied from July 2000 onwards incorporate a high-impedance differential input in the headstage. We were encouraged to do this because of the grounding problems that arise when two or more patch clamps (of whatever make) are used together, as they often are for whole cell applications, e.g. for making simultaneous pre- and postsynaptic recordings. The difficulty is that when the bath electrode is properly grounded as far as one patch clamp is concerned, the other patch clamps see a signal here (mainly at the power line frequency and its harmonics), because their grounds are actually all at slightly different potentials. The Optopatch can now get around this problem by measuring signals differentially with respect to the bath potential, so the bath can now be taken to whatever ground connection is most convenient for other reasons (note that it still needs to be grounded to something!).

However, differential operation is NOT recommended for single-channel recordings! The problem here is that the reference amplifier generates additional noise. Even if it were as quiet as the main amplifier, the overall amplifier noise would be increased by a factor of $\sqrt{2}$, and since it is actually a relatively standard integrated circuit its noise contribution will be significantly higher than that. Fortunately, differential operation is unlikely to be appropriate for single-channel recordings in any case, since even in a multiamplifier setup one is unlikely to attempt this on more than one amplifier at a time. The differential reference has therefore been made switchable, and the switch can be found on the driver board subpanel on the rear panel of the Optopatch. To ensure lowest noise for single-channel recordings, the reference for that amplifier should be switched off, and the bath electrode should be connected to the ground pin of the headstage. It is recommended that the reference input should also be connected to the ground pin rather than left floating. Any other amplifiers should then have their references switched on and with their inputs connected to this same ground pin (not their own).

Note that the much lower source impedance of the whole cell recording environment (because the amplifier "sees" the resistance and capacitance of the entire membrane rather than of just a small patch) means that the current noise from the source is much higher than for the patch recording environment, so in practice it will significantly exceed the additional current noise from the reference amplifier. This may sound counterintuitive, but it follows from the same theoretical consideration that favours the use of the highest possible feedback resistance in a resistive headstage. It is therefore unlikely that using the reference input in whole cell mode will significantly affect the overall noise performance. Of course, one is unlikely to notice that the whole cell configuration has a fundamentally higher current noise, because the currents that are being measured are so much higher, and they also have an additional noise component resulting from the opening and closing of the individual ion channels that give rise to them.

6 Detailed Description of Operation

6.1 General Facilities

This section describes the basic operating modes of the Optopatch, plus a few other features that don't naturally belong to any of the subsequent sections. First, we'll consider the basic operating modes. There are two main sets of modes, selection within each being made by a three-position toggle switch, and they determine whether the Optopatch is operating as a voltage clamp or as a current clamp, and whether the range of currents that can be passed is appropriate for true patch or for whole-cell recordings. They also affect the operation of some of the other controls in what should be an intuitive manner, and those effects will be noted briefly here for reference, although the sections describing those controls should be consulted for further information.

6.2 Mode Control Switches

The most important selection is between the voltage clamp and current clamp modes. This selector switch also has an intermediate position, labelled I=0, which is identical to current clamp except that the normal command inputs do not cause any current to be passed. In voltage clamp, the Optopatch passes a current so as to keep the electrode potential equal to the command potential, and the current is measured. In current clamp, the Optopatch passes a current that is proportional to the command voltage, and the electrode potential is measured. In I=0, the electrode potential is measured in the absence of any applied current. This mode is very useful for setting up, or for using the Optopatch as a conventional microelectrode amplifier.

The other selection is the current-passing range, and the three modes here are called "patch", "small cell" and "big cell". In patch mode, the current sensitivity is high, giving a maximum of 1nA, which corresponds to a 10V output. This situation is equivalent to a 10 gigohm feedback resistor from the signal amplitude point of view. Such currents are large enough to voltage-clamp some whole cells, as well as patches, so we have provided a "small cell" mode, in which the additional facilities needed for whole-cell recordings come into operation, but which can retain the same high current sensitivity. "Big cell" mode corresponds to normal cell mode on other patch clamps, and in this mode currents of up to 100nA can be passed, which is equivalent to a resistive headstage with a 100 Megohm feedback resistor. Since the optical current-passing system works well over a very wide current range, we did not feel it necessary to include any current ranges intermediate between these two.

In both cell modes, the circuitry for membrane capacitance and series resistance compensation - and measurement - comes into operation. The resistance compensation range is from 0.5Mohm to 100Mohm in both cell modes, but the capacitance range increases from 10pF (or 20pF when an offset is selected) in small cell mode to 100pF (200pF) in big cell mode.

Membrane capacitance compensation is active only in voltage clamp mode, but series resistance compensation also occurs in current clamp mode, although it is implemented in a different way. The operation of those controls in voltage clamp mode is quite complicated, so we won't go into that subject here. However, the situation for series resistance compensation is much simpler in current clamp mode, so we can describe it briefly here as well. Referring back to the circuit diagrams that showed the various current clamp circuits, one can see that in all cases the applied current passes through the electrode resistance. Therefore there is a potential difference between the tip of the electrode (which is where we want to be making the measurements) and the headstage input (which is where we are actually making them). However, it is possible to correct for this effect by subtracting an appropriate proportion of the command current input from the voltage output, using what is often, but not entirely correctly, described as a bridge circuit. Again, the Plymouth book describes the use of this type of facility in detail. Note that the series resistance ALSO causes errors in voltage clamp mode, but they are not so easy to correct. This will be discussed in detail later.

6.3 Leak Subtraction

At this point we can also usefully discuss the leakage current control, which performs a complementary action in voltage clamp mode, and works in a similar way. The problem in this case is that there may be leakage pathways that effectively appear between the headstage and ground. In true patch recording mode, an imperfect seal of the patch electrode against the cell membrane results in a leakage current pathway through the seal to the external solution, which is at ground potential or effectively so. A similar leakage pathway can also be present in whole cell recording mode, of course, but in this case there may be a membrane leakage resistance as well, which appears in parallel with the patch seal leakage. Unlike the effects of series resistance, such leakage pathways will not affect the accuracy of the voltage clamping (except insofar as they may worsen the errors due to the series resistance if the leakage current is high). The primary effect of a leakage pathway is that a component of the total current will flow through it, in proportion to the potential difference across it, so this can be corrected by subtracting an appropriate proportion of the command voltage from the current signal.

Although the leakage pathways are also present in current clamp mode, trying to compensate for them in this case serves no useful purpose, as their only effect is to take a constant proportion of the current. This effect is easily corrected if required by applying a proportionately larger current. In contrast, leak subtraction of the current record in voltage clamp mode IS useful, because it changes the overall form of the record by affecting just one component of it. It is therefore most appropriate that leak subtraction is provided in voltage clamp mode only, as we have done here.

The leak subtraction control is a calibrated ten-turn potentiometer, calibrated in units of resistance. However, since there has to be some lower limiting value of leakage resistance which can be compensated, this control has a minimum setting of one-tenth of a turn, giving an operating range of 0.1-10.1 turns. The resistance range of the control is 10Gohm per turn or 100Mohm per turn for the 1nA (patch and small cell) and 100nA (big cell) current ranges respectively. The associated on/off switch allows this control to be inactivated if required.

6.4 Noise and Other Metered Measurements

The Optopatch includes a digital meter, with a selector switch to choose any of nine different parameters. The meter is autoranging, and one of eight LED indicators illuminates to display the appropriate units. It has an extra digit compared with standard digital meters, displaying up to 19999, with the decimal point placed according to the gain range and the displayed units. Much as we would like to describe the internal circuitry which does all this (because it is quite complicated), there is actually no point, as the meter just shows you want to know without any fuss. For example, just compare working out what capacitance value corresponds to 4.7 turns on the cap control at 30pF full scale with 50% offset, with reading 29.10pF directly off the meter, and you will soon get the idea.

Most of the parameters that the meter can measure are described elsewhere in the manual, and the description of the front panel also provides a complete summary. However, the RMS position is not described in detail elsewhere, so it is covered here. This allows true rms measurement of the final output signal after gain and filtering so it can be used to measure any signal. The signal's low frequency cut off is produced by a single pole filter at 100Hz and the high frequency cut off is set by the output Bessel filter.

Signals are sent to an integrated circuit that performs an rms (root mean square) calculation of the signal levels. In retrospect, we rather wish we'd designed our own circuit to do this, as the single dedicated IC we use is by far the most expensive component in the entire Optopatch design, but at least it does the job.

AC signal levels are normally quoted in rms units, as the equivalent power, i.e. when an rms voltage is multiplied by an rms current, is then the same as for DC signals of the same values. If an AC signal waveform is predictable, e.g. a sine wave, then it is possible to measure its average or peak value after rectification to give a DC voltage, and then to apply an appropriate conversion factor to give the rms value. This is how most AC meters are calibrated. However, the main reason for our providing this facility is for measurement of noise, where true rms measurement is essential, hence our use of a dedicated IC to do a proper rms calculation. Note that the squaring action inherent in rms measurement also performs a rectification, so rms values are always positive.

Since the rms measurement facility is a general-purpose one, the meter displays rms signal levels directly as volts. For noise measurement in patch mode, the following conversion factor may therefore be useful. The signal from the headstage is 10mV/pA, so if it is measured at an absolute gain of 100 (i.e. 1K/mV in patch or small cell mode), 1V on the meter corresponds to 1pA of noise.

6.5 Command Voltage Generation

The Optopatch provides the standard range of controls for generating the command voltage (or command current when in current clamp mode), plus an oscillator for generating sine, square and triangular waves over a wide frequency range. Although not normally provided on patch clamp amplifiers, the sine wave generation facility is particularly useful for cell membrane capacitance measurement, and it will be discussed further in that section of the manual.

The controls provided are as follows. First of all, the cell holding potential in voltage clamp mode can be set anywhere between -200mV and +200mV on a ten-turn potentiometer. A switch is provided to select either positive or negative operation. (Note that there is an independent control for setting the holding current in current clamp mode, as described below). If the ext command switch is on, then step or other potential waveforms, applied externally via the command/10 and/or command/100 input sockets on the rear panel, are superimposed on the holding potential (i.e. all three signals are summed). As their names suggest, these inputs are attenuated by a factor of 10 and 100 respectively, so for example a signal of +500mV at the commandx10 input will generate a step in the command potential of +50mV. If superimposed on a holding potential of say -70mV, then the overall command potential will change from -70mV to -20mV. However, it may be found more convenient to include the holding potential in the external command potential, so a centre off position is provided on the holding potential polarity switch. The centre position is labelled ext, since it actually switches to another external input (hold/10 in), which is also subjected to a tenfold attenuation. This input can either be used to accept an external holding potential, with modulating potentials continuing to be applied via the command/10 and/or command/100 inputs (if the ext command switch is on), or the two potentials could be combined together and applied on any one of these inputs. The possible advantage of the external hold input for this purpose is that the holding potential switch allows the user to switch between an internally generated holding potential, which may be useful for setting up, and an externally generated potential that provides the entire command potential. The hold/10, command/10 and command/100 inputs have a relatively low input impedance of 10Kohms. This means that it is not necessary to ground any input if it is not used, and if greater attenuation is preferred, this can easily be achieved by an appropriate series resistor, e.g. 90Kohms of series resistance would increase the attenuation by a further factor of ten.

The other possible component of the command potential is the output from the frequency generator, which is variable from 100Hz to 100KHz in three switched ranges. The variable frequency control is not itself calibrated, but the square wave frequency can be read precisely on the meter when the selector switch is in the osc position. In order to read the frequency, square waves need to be generated as well. These are converted into constant-length pulses, which are then filtered to give a voltage directly proportional to frequency. This works very well, but there is always a risk of interference to very low-level signals whenever large rapidly-changing signals such as these are present inside the same piece of equipment. In case this is ever a problem - although it may well not be - we have provided a jumper on the main circuit board, which allows sine waves to be generated without the accompanying square waves. The only drawback of this arrangement is that sine wave frequencies can no longer be read on the meter. In any case, square waves are never generated when the osc switch is off, which would be the usual condition when very small signals, i.e. single-channel currents, are being measured. The meter displays the frequency directly in KHz. The signal amplitude is set by another ten-turn potentiometer, and the signal is always symmetrical about a zero central potential. When the osc selector switch is in the sine position, low-distortion sine wave signals of up to 100mV rms (i.e. 242mV peak-to-peak) are added to the command potential, and in the square position, signals of up to + and - 50mV (i.e. 100mV peak-to-peak) are added. These outputs are also available at a tenfold higher level on the freqgen output on the rear panel.

The high bandwidth of the Optopatch means that a fast command voltage step produces an initial current transient which cannot be completely removed by adjustment of the electrode capacitance controls. Their operation is described in detail later on, but for now we can note that the best settings of these controls leave a biphasic transient, lasting about 5usec (on a model cell), which can be reduced to a negligible size by appropriate filtering. This is a general characteristic of patch clamp amplifiers, but since the Optopatch can have a bandwidth of up to 200KHz, the transient can be somewhat larger than for other amplifiers with lower bandwidth. A simple and effective way of reducing the transient is to slow the rise of the command potential, and therefore our command potential circuit includes a low-pass filter with a time constant of 1usec, corresponding to a corner frequency of 160KHz (purists please note- this time constant does not appear in the series resistance compensation pathway). Although short, this time constant gives a significant reduction, and of course the command signals can be further filtered externally if required. Command signals can also be applied without any filtering, by passing them as currents as described next, which we regard as a preferable method in any case. Finally, we should note that the unfiltered frequency response of the Optopatch can be "tamed" so that is more similar to other patch clamps if preferred. Jumper links on the headstage driver board allow the unfiltered bandwidth to be reduced from 200KHz to 100KHz independently for the 1nA and 100nA current ranges. This approximately halves the remaining transient.

The Optopatch also includes an integrator, for passing command CURRENTS via a small capacitor in the headstage (in both voltage and current clamp mode) rather than voltages, which can be driven either internally from the sine or square wave signals, and/or externally from a BNC input on the rear panel. Circuits of this type produce symmetrical voltage ramps when driven by square waves, so they are sometimes also known as ramp generators. The frequency generator output is connected to the internal input of the integrator instead of to the command potential circuit when the sine/square switch is in the current position, but the external BNC input is always active, and the integrator will add the two inputs together if both are present. The external input is effectively attenuated by a factor of ten, so a 1V signal here is equivalent to a 100mV signal from the internal oscillator. The output of the integrator is connected to the input of the patch clamp via a small capacitor, allowing AC currents to be passed. This is the same component that is used for compensation for the input and electrode capacitance, as described in the next section, and the integrator signal is simply added to the compensation signal, so the provision of this facility does not require any additional complexity in the headstage.

The current passed through the capacitor depends on the rate of change of the voltage across it, which means that the capacitor re-differentiates the integrator output to give a current waveform of the same shape as the voltage input to the integrator. There is of course a scaling factor to be calculated, which depends on the time constant of the integrator (i.e. the rate at which its output rises for a given voltage input) and the value of input capacitor. At the nominal full-scale input, i.e.

100mV for the internal frequency generator or 1V for the external input, the integrator output changes at 10V per millisecond, and this is coupled to the input via the electrode capacitance compensation capacitor in the headstage, to give a full-scale current of 10nA, which is conveniently intermediate between the patch clamp's two current ranges of 1nA and 100nA full scale, thereby making it useful in either. The relatively high integrator gain causes its output to saturate for square wave frequencies below 200Hz (where the peak-to-peak output is 25V), but it is still possible to pass half the full-scale current at 100Hz, and for sine waves the situation is slightly better, as the clip frequency for full-scale output is about 150Hz in that case. The integrator includes a DC servo circuit to hold the average output at zero, but it has no significant effect on its signal performance over the frequency range that is of interest here.

This facility has a variety of uses. In current clamp it provides an alternative method of passing (AC) currents, but its greatest use is in voltage clamping, where it allows the bandwidth and gain of the recording system to be verified (these parameters can be modified by other controls, which will be described later). Unlike the command potential generation circuitry, there is no low-pass filtering, and the inherent bandwidth of an integrator/capacitor combination is very wide, which makes it a fundamentally better way of applying high-frequency signals in any case. Even in voltage clamp mode, the capacitor still passes a current into the input, and in order for the input voltage to remain constant, the voltage clamp's feedback loop must correspondingly reduce the feedback current, hence this facility is sometimes referred to as a "speed test", since complementary currents will be observed on the current output, and their time course will depend on the bandwidth of the current-measuring system. In contrast, the time course of a current step caused by a change in command potential also depends on the time required for the potential change to occur at the preparation, so this is not a proper measure of the recording bandwidth.

The advantage of using a capacitor to pass current is that essentially instantaneous current steps can be generated by applying voltage ramps, whereas a step voltage across a current-passing resistor is likely to cause an initially larger current transient, caused by the presence of some parallel capacitance across the resistor. Therefore the capacitor method provides a reliable speed reference for the recording system. For the same reason, it may also be a somewhat better way of producing fast current steps in current clamp mode, although in the Optopatch the two methods should be more nearly equivalent than in other designs.

When the patch electrode is in the bath, a potential of zero should ideally be recorded in I=0 mode and no current should flow in Vclamp mode, but in practice, electrode and bath potentials will result in a residual voltage. Although this could be taken into account by the applied potentials, it is clearly preferable to have some means of independent correction, which is provided by the junc control and switch. This is a centre-zero control that can provide up to + or - 200mV of additional voltage offset when the switch is in the on position. The switch also provides a "search" facility for use in voltage clamp mode. The search facility automatically adjusts the junction potential offset so as to drive the electrode current towards zero, but with a relatively long time constant, so that transient currents are not significantly attenuated. As its name suggests, this facility allows patches to be searched for at high current gain, while preventing the relatively large currents that would otherwise flow through the relatively low electrode resistance in response to only small potential differences (once a patch has been formed, the resistance of course increases dramatically, so this is no longer a problem). The applied junction potential offsets that are automatically generated in search mode can be read on the meter when the selector switch is in the junc position (when the switch is in the on position, the actual junction potential set by the control is read). When a patch has been found, the meter reading can be used as an additional guide to selection of the appropriate control setting before the switch is moved to the on position.

Although under most conditions the experimental bath will be connected directly to ground via a low-resistance bath electrode, there are also circumstances which may require external potentials to be imposed here, or in which the zero potential of the bath cannot be guaranteed. A bath potential input is therefore provided on the rear panel. However, please note that the provision of an external reference input on the headstages of amplifiers supplied from July 2000 onwards has made this facility more or less obsolete, so it is retained and now described primarily for historical reasons. If any signal that is present in the bath is also applied to the bath potential input, then the patch clamp command potentials will be generated with reference to this signal rather than ground, so that the operation of the patch clamp will not be affected by it. Once again, the input is of low impedance (10Kohms), so there is no need to ground it if it is not used. It is a unity-gain input, so signals should be applied at their original level. Because of the low input impedance, it is not suitable for direct connection to a reference electrode, but in any case it is greatly preferable to use such electrodes with a buffer amplifier sited close to them, rather than with a long unbuffered connection to the recording equipment (provision for connecting and powering such an amplifier is made by a five-way circular connector on the rear panel, but that function has now been taken over by the reference input that is now provided in the headstage). However, there is one useful exception. If the bath electrode is of low impedance (so as to be able to support a reasonable length of signal connection without coupling mains hum or interference into the bath, and to drive the bath effectively), then it can be connected directly to the bath potential input, in which case it is earthed by the 10Kohms input resistor (or an even lower resistor could be connected in parallel between this input and ground). External signals can then be imposed on the bath by applying them at this point, where they will drive both the bath electrode and the bath potential input.

The total command potential thus consists of some combination of an internal or external holding potential, an external command potential, an internal oscillator signal, an internal junction offset potential, and an external bath potential. (The command potential that is actually applied to the electrode may be further modified by the compensation facilities described later.) A command x10 output is available on the rear panel, i.e. at a tenfold higher level than the command potential applied to the cell. This output does not, however, contain the bath or junction potentials, since they are effectively not seen by the preparation either. The same signal combination is displayed on the meter at unity gain when the meter selector switch is in the command position.

In current clamp mode, the command potential instead becomes a command current. However, junction and bath potentials are no longer summed into this signal, and they are instead added to the headstage output, which now represents the electrode potential (actually amplified tenfold), in order to act in the same way as in voltage clamp mode. A full-scale current of 1nA or 100nA requires a voltage input of 10V at the input to the current-passing circuit, and the external holding and command inputs supply this signal at unity gain, so the external /10 command sensitivities are 100pA/V and 10nA/V for the two current ranges. To retain the same relation between internal and external sensitivities as in the voltage clamp modes, these equate to 10pA/V and 1nA/V for the holding and frequency generator potentials.

Sometimes it may be desirable to switch from voltage clamp to current clamp (or vice versa) while maintaining the cell at some specified potential. This means changing from one value of command potential to another value of command current, at the same instant that the system is switched from voltage clamp to current clamp. Although that can in principle be done by a computerised control system, we decided to provide separate controls for the holding potential in voltage clamp mode and for the holding current in current clamp mode. In current clamp the holding current range can be varied between + and – 200pA full scale in patch and small cell modes, and between + and - 2nA in big cell mode, by a dedicated ten-turn control and polarity switch, which again has a centre off position.

6.6 Signal Gain and Filtering

Comprehensive facilities for amplifying and filtering the current signals (or voltage signals in current clamp mode) are included. The absolute gain is switchable from x1 to x500 in steps following a x1 x2 x5 sequence, and filtering can be applied both before and after the gain stage. The incoming current signals have a value of 10V at full scale on the two ranges of 1nA and 100nA, giving sensitivities of 100pA/V and 10nA/V respectively at minimum gain. The incoming voltage signals in the current clamp modes have already been amplified by ten. We doubt whether the highest gains will be of practical value for patch clamping, but we included them because the capacity was there and because they might be useful in other applications.

It can be particularly useful to have a record of the gain setting, so we have included the popular facility of an output voltage that can be recorded by the data-capture system to provide the gain information. We provide similar voltage values to those used on other commercial equipment, and the telegraph output also takes the gain of the headstage into account. The output at a gain of 1mv/pA in voltage clamp mode is 2.5V, rising in 0.5V steps for each gain increment, so it varies between 4V and 9.0V over the gain range of 1 to 500 in patch mode (10mV/pA to 5000mV/pA), and between 1V and 6.0V (0.1mV/pA to 50mV/pA) in big cell mode. The telegraph output in the small cell mode varies according to the current range selected for this mode on the rear panel switch, which can be either of these. In current clamp, a voltage output of 1mV/mV (i.e. unity gain) would also give 2.5V, but since the voltage is already at x10 gain before being sent to the gain selector, the total gain varies between 10 and 5,000, giving a telegraph voltage range of 4V to 9.0V.

It is also worthwhile if not essential to include adjustable low-pass filtering, since patch currents are inevitably noisy, and the noise depends on the signal bandwidth. Filtering is often useful in the whole-cell modes too. The Optopatch provides two low-pass filters, to remove noise that is at higher frequencies than those that are of interest for current measurement. There is a good reason for providing two filters rather than one, which we'll discuss further below. Unfortunately, filters do not have an ideal (step) frequency response, so their use entails some compromises. In the simplest, or single-pole, type, which just consists of a single capacitor and resistor, the gain halves with each doubling of frequency above a given "corner" or "cutoff" frequency, but steeper responses can be obtained by cascading several such sections together, the steepness increasing geometrically with the number of poles, i.e. the gain of a two-pole filter decreases by a factor of four for each doubling of frequency, that of a three-pole filter by eight, and so on. In filters having more than one pole, the shape of the response around the corner frequency can be modified by feedback to give a more or less sharp transition between a flat response at lower frequencies and a smoothly falling response at higher frequencies. A sharp transition may seem better, but such filters distort the waveforms of transient signals to a greater extent, and we use a Bessel characteristic for both filters in order to achieve the best compromise.

Noise comes from a variety of sources. The simplest source to consider is one that generates noise equally at all frequencies; for example, thermal noise generated by a resistor has this characteristic. Noise sources, being random, add in an rms rather than a linear manner. This has many implications, amongst which is that for a source that generates the same amount of noise at all frequencies (i.e. white noise in the usual jargon), the total noise increases with the square root of the bandwidth rather than linearly. However, for reasons which are well described by Sigworth (1995) in the single channels recording book, the internal noise generated by a patch clamp is NOT independent of frequency. Instead there is a noise component which increases linearly with frequency, and which tends to be dominant at frequencies above a few KHz, so the dependence of the rms noise on the signal bandwidth is much greater. This tends to favour use of multipole filtering. Four-pole filtering is often chosen as the best compromise between performance and complexity, whereas purists tend to favour eight-pole filtering, but there is general agreement that there are no practical advantages in having more than eight poles. We have compromised by providing four-pole filtering in our main filter, which is optionally extendable to eight poles by an internal accessory board.

Changing the cut-off frequency of a four-pole filter requires four components (usually resistors) to be switched, and for a larger change the other four frequency-determining components (i.e. the capacitors) may have to be switched as well, which makes for a messy operation in practical terms. We therefore opted for an alternative method, which also has the advantage of allowing the cut-off frequency to be continuously variable if required (via the computer connector). We had already devised a variable filter circuit for our fluorescence equipment, using transconductance amplifiers - the output of which depend on a control current as well as the incoming signal - to simulate variable resistors, and this turned out to be an ideal application for it. The other possible technique is the switched-capacitor filter, and some very smart integrated circuits of this type are available, but this method requires a digital control frequency, and we were therefore concerned about the possibility of digital interference reaching the sensitive parts of the patch clamp electronics. Such filters also tend to have a more restricted frequency range.

Although the transconductance technique requires a larger number of devices than other methods, it gives excellent performance without causing any interference risks. We use a switched control voltage to vary the frequency over a tenfold range, and we then switch the four capacitors to change the range by further factors of ten. The optional accessory board, to extend the filtering to eight poles, works in exactly the same way. As mentioned in the introduction, the use of analogue switches here, as elsewhere in the Optopatch, allows us to perform multi-pole switching of the analogue signals by only single-pole switching of the digital control voltage, which feeds all the analogue switches together. It also removes the need to send any of the signals all the way out to the front panel and back, which can be a problem, especially at higher frequencies, because of the unavoidably greater stray capacitance associated with the wiring. Purists will be aware that extending a filter from four to eight poles requires the characteristics of the existing four-pole section to be changed, and we accomplish this by moving some jumper links on the main circuit board when we fit the accessory board. The performance of the accessory board in early 2000 has fully matched our expectations, so we strongly encourage its use.

From the foregoing it should be clear that multi-pole filters, however they are implemented, tend to be quite complicated. Compared with a simple amplifier stage, they tend to be more susceptible to various forms of interference pickup, and they also tend to generate more noise of their own. Therefore, in a general-purpose gain and filtering circuit, it is preferable to put the gain stage first, so that the filter acts on the amplified signal. However, if there is a lot of high-frequency noise on the signal, this configuration introduces another problem. The gain stage also amplifies the noise, causing the amplifier to saturate at a possibly much lower gain than if the noise were not present. Siting the filter first would have avoided this particular problem. We therefore provide an additional filter prior to the gain stage. This filter can be much simpler than the output filter, so there are accordingly fewer performance compromises, and accordingly we provide a two-pole Bessel filter with corner frequencies of 3KHz, 10KHz and 30KHz, plus an "out" position. It is referred to as the prefilter in view of its location in the signal pathway, and we recommend that it is set to a frequency that includes the range of interest, which can then be defined more precisely by the main filter that follows the gain stage.

The base frequencies of the output filter are 1Hz, 10Hz, 100Hz, 1KHz and 10KHz, set by the freq range selector, and they can be varied over the range x1 to x10 by the freq value selector. We had originally intended to make the freq value selection by a continuously variable control, but consultations with experienced patch-clampers suggested that it would be clearer to provide a selection of fixed values instead. However, we have provided a wider choice of values than is usual (x1, x1.5, x2, x3, x5 x7 and x10), and the possibility of providing a variable control voltage via the computer connector, when the system is under computer control, has been retained.

The lowest frequency ranges are not appropriate for patch clamping, but they may be appropriate for other applications. In particular, they may be useful for filtering the signals generated by the capacitance measurement facility, and a switch is provided to connect the capacitance signals directly to the filter. To do this, the switch should be in the "cap" position, when the filter takes its input from the imaginary phase output, otherwise it should be set to the "current" position.

In addition to the voltage output representing the gain, several signal outputs are available on the rear panel. These are:

(1) The current output from the headstage, at the minimum gains specified above, including the two-pole prefilter. It is available in voltage clamp mode only.

(2) The current output from the headstage, at the minimum gains specified above, including the two-pole prefilter and also the effects of any leakage current compensation in voltage clamp mode. It is available in both voltage and current clamp modes.

(3) The same signal, after amplification by the gain stage.

(4) Another output from the gain stage, but filtered with a three-pole Bessel response at 25KHz, which can be useful for general monitoring purposes.

(5) The output from the four- or eight-pole Bessel filter.

6.7 Capacitance and Resistance Compensation

Because of the very high source resistances associated with patch clamping, even very small amounts of capacitance associated with the cell and with the recording equipment can significantly affect the recorded current signals. Furthermore, in the whole-cell recording mode, the capacitance of the cell membrane and the resistance of the electrode together act to slow the time course of changes in the cell membrane potential in response to changes in the command potential, and the electrode resistance can also cause the cell membrane potential to differ from the command potential as a result of the voltage drop across it due to the clamp currents. However, all these problems can be compensated for to a greater or lesser extent, as will now be described. Most of this section deals with compensation for standard patch (voltage) clamping, in both the true patch and the whole-cell configurations, but the differences for the current clamping situation will be discussed subsequently.

For dealing with the capacitance, it is important to distinguish carefully between "fast" capacitance compensation, which compensates for the capacitance between the amplifier input and ground, and "slow" capacitance compensation, which compensates for the capacitance between the electrode tip and ground (or in practice the experimental bath, but this is also at ground potential, or effectively so). In whole-cell recording mode, the slow capacitance can be relatively much more significant, since it represents the capacitance of the cell membrane, which may be tens of picofarads or even higher, whereas uncompensated fast capacitances are typically only a few picofarads at worst. In electronic terms, the difference between these capacitances is that the fast capacitance is connected directly to the amplifier input, whereas the slow capacitance has the resistance of the electrode in series with it. Fig. 8 shows the equivalent circuit, and we once again show a standard resistive headstage for the purposes of this illustration. Without compensation, the current-passing resistor (or equivalent component) in the headstage would have to charge both these capacitances in response to a change in the command potential.

This is undesirable because the process takes a finite time and because the charging currents appear directly on the current output from the patch clamp. In both cases, the capacitance is compensated by providing an appropriately set amount of charge to the input via a small capacitor in the headstage, in response to any change in the command potential, with the settings being adjusted to minimise the residual signals appearing on the current output. The circuitry for the two forms of capacitance compensation has to be distinct, because the fast capacitance can be charged (almost) immediately, whereas the slow capacitance has to be charged via the electrode resistance. The simpler case of fast capacitance compensation will be considered first.

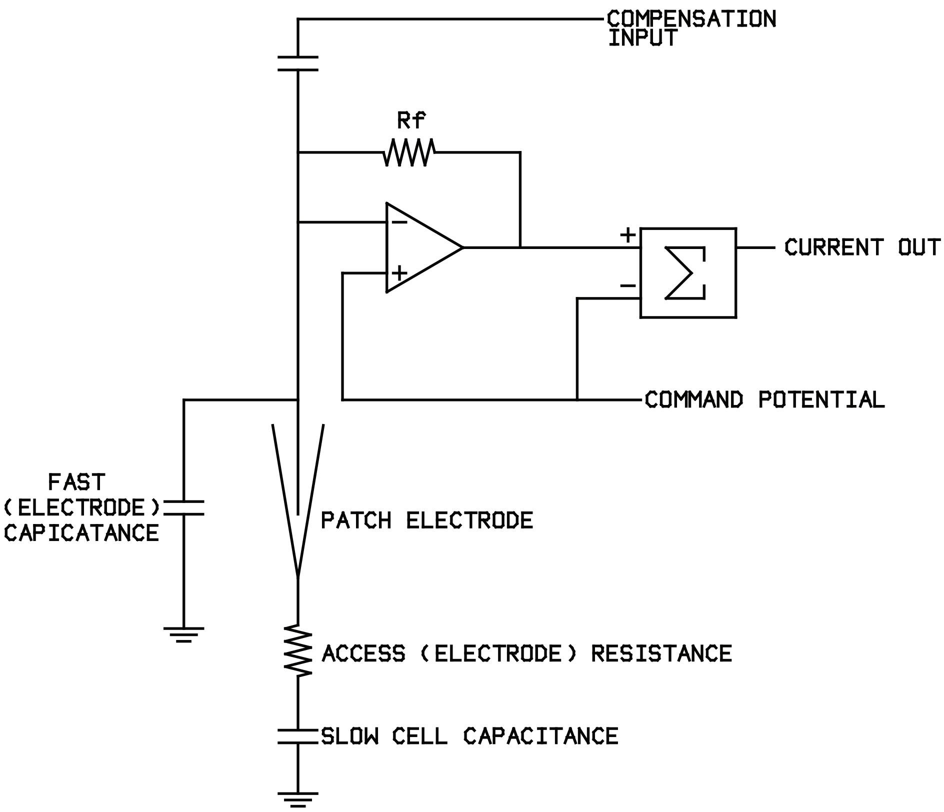

Figure 8 - Capacitance Compensation Circuit

The fast capacitance consists of the input capacitance of the amplifier, plus the capacitance between the electrode and the bath. In true patch recording mode, there is a very high resistance seal immediately on the other side of the electrode resistance, and although there is also a capacitance associated with this seal, its value is so small that it can be neglected under normal recording conditions. In general, we therefore only need to consider the fast capacitance in true patch mode, although for completeness we should note that a few brave volunteers are prepared to attempt to study the very small patch capacitance as well. For simplicity and clarity, we shall in future refer to the fast capacitance as the electrode capacitance, and at the risk of offending those who dare to do otherwise, we shall assume that this is the only capacitance that needs to considered under true patch recording conditions. With ideal electrodes and ideal amplifiers, this capacitance could in theory be compensated totally, to allow recordings of arbitrarily high bandwidth (although noise considerations would still favour there being no capacitance to compensate in the first place) but in practice there are limitations, because, for example, the electrode does not behave precisely as a single capacitance and a single resistance to ground. Therefore in practice, it is not possible to achieve either a truly instantaneous voltage step at the electrode or a current that is completely free of a capacitive transient as a result of the voltage step, but the approximation can be improved by modifying the time course of the compensating signal. In particular, a certain amount of high-frequency filtering of this signal is usually helpful, and under some circumstances - particularly when the electrode resistance is relatively high - it can also be useful to supply a small additional compensation component that has been filtered at a somewhat lower frequency.

The Optopatch provides the four controls necessary to perform this compensation, i.e. a true fast capacitance amplitude compensation with variable filter time constant, for compensation of up to 15pF of electrode capacitance, and a smaller amount of additional compensation (up to 2pF) at a longer but also variable filter time constant. The best way to understand the operation of these controls is to use them, while taking care to remember that overcompensation would normally result in oscillation (i.e. near-certain cell death) when the Optopatch is in current clamp mode, so the two amplitude controls should always be adjusted upwards from zero when setting the compensation for the first time. However, the Optopatch incorporates a built-in oscillation detector, which will disable the compensation before the oscillation builds up sufficiently to kill the cell, as described below. It should also be remembered that any change in the length of the electrode immersed in the bath (which can occur if the level of solution in the bath changes, as well as if the electrode is moved) will change the electrode capacitance and hence the control settings for optimum compensation. Apart from that, the operation of these controls should be both reliable and stable, and fast capacitance compensation can and should be used routinely in all experiments.

A potential problem with fast capacitance compensation is that in current clamp mode, overcompensation causes the amplifier to become unstable, and the resulting high-frequency oscillation is an excellent way to kill a cell. This is a general problem of the fast capacitance compensation technique, rather than being specific to the Optopatch, but we can now offer a specific solution to it. In July 2000 we incorporated a simple accessory circuit which detects a series of fast voltage transients in current clamp mode, of the sort that occurs when an oscillation is building up. This circuit then disconnects the compensation signal for about half a second, which it indicates by illuminating the overload LED, and will keep disconnecting it if the problem persists. The voltage threshold for operation is set by PRE2 on the headstage driver board, accessible after removal of the Optopatch top panel. This feature is normally supplied enabled, and with PRE2 set to maximum sensitivity, but it can be disabled by turning PRE2 fully anticlockwise and/or by changing the setting of jumper J3 on this board. When this facility was tested at the September 2000 Plymouth course, we were delighted to find that even the small (and normally very fragile in this respect!) granule cells in brain slices were protected by it, so we promptly extended it to operate for series resistance compensation as well (see below).

In the whole-cell configuration of the patch clamp, in which the membrane patch at the electrode tip has been destroyed so that the patch electrode is now effectively intracellular, the situation is more complicated. In the following discussion, it will be assumed that the electrode capacitance has been effectively completely compensated as described above, before the patch has been destroyed. This is an important assumption, since (especially with small cells), uncompensated changes in the electrode capacitance may significantly interfere with the signals observed in whole-cell mode.

In this configuration, the capacitance of the cell is connected to the electrode via the electrode resistance. These components together form a significant time constant, and although it is possible to compensate reasonably well for their effects, the situation is much harder to handle than the fast capacitance. Typical electrode resistance values are in the megohm range or greater, and the capacitance of the cell can easily be 10pF or more, giving time constants of tens or perhaps even hundreds of microseconds. Furthermore, the resistance of the interconnection between the electrode and the cell interior is generally somewhat higher than the resistance of the electrode before it was patched onto the cell, and it can also vary during the course of an experiment, and in common with usual practice we shall use the term "series resistance" to describe this parameter.

Just as for the electrode capacitance, it is possible to compensate for the cell capacitance, so that the charging currents are provided via a capacitor in the headstage, but in this case the compensating signal must be matched to the membrane charging time constant, which means that the compensation circuitry must have adjustments for both the cell membrane capacitance and the series resistance. Where the cell capacitance is relatively small, i.e. below 20pF, the compensation signal can be supplied via the same capacitor that is used to compensate for the electrode capacitance (as shown in Fig. 8), simply by adding the two signals together, and this is how the Optopatch operates in the small cell mode. In this mode the DC current-passing capacity is the same as in patch mode, i.e. equivalent to a feedback resistance of 10Gohm (although in the Optopatch the resistance value is of course infinite from the noise point of view), giving a maximum current of 1nA. It differs from true patch mode in that the cap, res and the related controls described below are all disabled in patch mode.

In the big cell mode, the equivalent value of the feedback resistor is 100Mohm, giving maximum DC currents of 100nA. The cell capacitance compensation is now applied separately through a tenfold larger capacitor, which is automatically connected to the headstage in order to increase the cell capacitance compensation range to 200pF. Cell capacitances below 20pF can also be compensated in this mode, so there is a fair amount of overlap. The main use of the small cell mode is to allow capacitance measurements to be made at higher current gain. Note that capacitances above 10pF in small cell mode or above 100pF in big cell mode need to be measured using the capacitance offset facility, which switches in capacitance offsets of either 50% or 100% of full scale in both cell modes, allowing cell capacitances of up to 20pF and 200pF to be compensated. However, since the drive circuitry for the compensation capacitor is required to produce relatively larger voltage excursions under these circumstances, the maximum permissible command potential excursions will be correspondingly reduced, as described below. For measuring even larger capacitances, an alternative headstage with a larger cell compensation capacitor (100pF instead of 10pF) is available, capable of compensating cell capacitances up to 2nF. This could also be useful for other applications such as lipid bilayer experiments.

At full-scale cell capacitance, without offsets, the capacitance drive circuit can support command potential excursions within the range of -200mV to +200mV. As the compensated capacitance increases beyond 10pF in small cell mode or 100pF in big cell mode, the command potential excursion rage is progressively reduced, to +/-100mV for 20pF and 100pF capacitances respectively. However, this is still likely to be sufficient for the vast majority of applications. Note that the command potential range in the absence of membrane capacitance compensation is greater still, at +/-1V, to allow the Optopatch to be used for cyclic voltammetry experiments, where this type of compensation is inappropiate.

For measuring smaller membrane capacitances, the full-scale capacitance set by the cap control can be reduced to either 30pF or 10pF in big cell mode (the range is usually fixed at 10pF in small cell mode, although a 3pF option can be provided on request), which assists in the use of the control, which is a ten-turn potentiometer, and its associated electrical output (described below) for making accurate measurements. The effects of 50% and 100% capacitance offsets are similarly scaled down so that they still work correctly, although the command voltage range restriction outlined above still applies. However, this is of little if any practical significance, because under these circumstances it is probably easier to switch to a higher capacitance range instead.

The series res control is also a ten-turn linear potentiometer, and it covers the same ranges in both the small cell and big cell modes. Full-scale resistances of 10Mohm, 30Mohm or 100Mohm can be selected on the res range switch. Note that the minimum setting for both the res and the cap controls corresponds to 5% of full scale rather than zero, so the minimum mechanical position read on the dial is one-fifth of a turn. The maximum dial setting is therefore 10.5 turns, corresponding to 105% of full scale. The dial readings are therefore precise under all conditions, and can be relied upon for making quantitative measurements, although in practice it will be easier to read the values directly from the digital meter. Output voltages proportional to the settings of both controls are available on the res out and cap out BNC sockets at 10V full scale. If the Optopatch is in track mode (see below), these outputs, and also the displayed meter values in the res and cap positions, will include the additional components introduced by the automatic resistance and capacitance compensation circuit, which is described in the next section.

One effect of having such a wide range of resistance and capacitance control settings is that certain combinations are not going to be very realistic. For example, the controls theoretically compensate for a series resistance as low as 0.5Mohm in combination with a membrane capacitance as low as 0.5pF. This would give a charging time constant of only 0.25usec, which would require incredibly fast electronics to achieve, and in practice there is no need to compensate for such a short charging time constant anyway. Our RC compensation circuit has therefore been designed to deal with all practical resistance and capacitance combinations, i.e. for charging time constants down to a few microseconds, and this includes combinations where one or other parameter is at or near its minimum value. At the same time, we have taken care to ensure that setting both of them at or near their minimum values does not cause any stability or other problems.

The series resistance and cell capacitance compensation facilities can be disabled in both cell modes if not required (with the exception of conventional series resistance compensation in voltage clamp mode - see later), by setting the RC comp switch to its off position, otherwise this switch should be on. It is actually a three-position switch, and the third position, labelled track, enables the automatic capacitance tracking facility. Use of that facility requires a sinusoidal command voltage, and it is intended for high-resolution measurement of the cell capacitance (and series resistance) rather than for their compensation during command voltage steps. The compensation adjustments should therefore be carried out by the methods described below with the switch in the on position.

When the res and cap controls are correctly set, the amplitude and time course of the current through the compensation capacitor precisely match those required to charge the membrane capacitance, so no signal is observed at the current output of the patch clamp. This is the best criterion to use for obtaining the optimum control settings, although it is also possible to use a sinusoidal rather than a square command voltage waveform, which will be discussed further in the description of the automatic capacitance tracking facility. However, even though the current trace will now be settling more quickly than before, the cell membrane potential is still changing with the time constant given by the series resistance and cell capacitance, rather than instantly following the command voltage. Furthermore, the series resistance also causes steady-state errors in the potential to which the cell is clamped, as a result of the voltage drop across it due to the clamp current. Fortunately, it is possible to correct for both these problems to a reasonable extent. However, not only is this is a relatively complicated topic to discuss, but the other patch clamp manufacturers have also provided different selections of controls for the carrying out the necessary adjustments, and particularly with earlier designs only a partial implementation may have been provided in any case.

At the risk of causing yet further confusion, we have provided our own name for one of the controls, i.e. that of "precharging" for compensation of the capacitive charging current, since we think it most clearly describes what the control actually does, and we then use "series resistance compensation" to describe compensation of the voltage drop across the series resistance due to the clamp (i.e. ion channel) currents. In other commercial designs, these parameters may be adjusted in tandem by a single control, or there may be controls that perform apparently equivalent functions to ours, but under different names such as "prediction" or "correction", so there is no uniqueness about the operation of this part of the Optopatch (although its implementation is different in order to support the operation of our capacitance-tracking circuitry, described below), and in any case, this topic has been covered in depth by Sigworth (1995).

The correction for the time constant for changing the membrane potential is best described with reference to voltage steps, although it works correctly for all waveforms. The principle is to charge the cell capacitance more rapidly by adding a transient component at the start of the command voltage step. For example, if the effect of the transient is initially to double the size of the step, then the cell capacitance will initially be charged twice as quickly. The additional component must be transient, so that it decays away as the membrane capacitance charges, but if it has the appropriate time constant, then its effect will be to double the charging current at all times, thereby preserving the single-exponential form of the membrane potential change, but halving its time constant. In practice, much larger transients can be generated, and speed improvements approaching an order of magnitude can be achieved. The correction sounds complicated, but in fact the appropriate waveform can be generated automatically by the patch clamp electronics, since the res and cap control settings, determined by minimising the transients on the current ouFigure 8 - Capacitance Compensation Circuit

tput as described above, provide the information that is needed to derive it. The amount of precharging is set by a variable control, which give zero precharging at its minimum setting, and nominally perfect precharging, i.e. a very large and very short pulse, at its maximum setting. {This is the basis of an alternative method, termed "supercharging" (Armstrong and Chow, 1987), which can be used in patch clamps that do not have the precharging facility, where an initial voltage spike of appropriate area (amplitude multiplied by time) is provided on the command pulse itself.} In practice, non-ideal behaviour of the internal electronics sets a limit to the accuracy of the precharging waveform that can be achieved as the maximum setting is approached, but the real limit is set by the validity of the approximation of the cell's equivalent circuit to a single capacitor connected to a single resistor.

In general, precharging is extremely effective, and its use is recommended. Although there is no need to be concerned with what the precharging is actually doing to the command potential, interested users can nevertheless monitor both the original command potential (command out) and its precharged derivative (headstage command out) via the rear panel outputs provided for this purpose. This can be reassuring because the effects of precharging may not otherwise be so immediately apparent, since the capacitive charging current is not seen directly, but the time constant of the precharging transient on a command potential serves as a useful indication of the time constant of the membrane potential change.

Precharging is appropriate only under certain conditions, so to avoid possible conflicts it is not always available. An LED indicator shows when precharging is in operation. The precharging control must be away from its fully anticlockwise position, and the Optopatch must be in small cell or big cell voltage clamp modes, with the lock-in amplifier switched off.

In contrast, correction for the voltage drop across the series resistance due to ionic currents, which is normally referred to as "series resistance compensation", is less reliable, although a reasonable improvement can be achieved in many cases. The problem with this form of compensation is that it involves positive feedback, which inevitably reduces the stability of the system. Specifically, the voltage drop across the series resistance is proportional to the current, and it can in principle be corrected for by using the current signal to derive a compensatory increase in the command voltage. Once again the appropriate scaling factor can be generated directly by the patch clamp electronics, from the res control setting made previously, and the fraction of the series resistance that is to be compensated is then set by the % RS comp control. However, there are problems that limit the practical effectiveness of the technique. First, there is a finite response time from when a change in current is detected to when the appropriate change in command voltage can be made, and since this feedback loop includes the cell and electrode parameters as well as the patch clamp electronics, its behaviour is neither perfectly ideal nor perfectly predictable. Second, overcompensation will result in complete instability, since the effective series resistance will then be negative. This condition can occur if the actual series resistance is less than the equivalent value set on the res control, even if the % RS comp control is below its maximum setting, and to make matters still worse in practice, the series resistance may change both with the current and the time during a voltage command pulse. This is generally a greater problem during the sustained ionic currents that occur in response to the command pulse than during the transient currents that occur during precharging, and in any case series resistance changes during precharging will cause only a temporary error in the cell membrane potential compared to the command voltage, which does not directly threaten the stability of the system. Therefore, the fraction of the series resistance that can be safely compensated in practice is generally less than the fractional correction that can be applied for precharging, which is why we have chosen to provide independent controls for setting how much of each of the two compensations is applied.

The series resistance may also limit the current that can be passed into the preparation. In the worst compensatable case, the series resistance may be 100Mohm, but the maximum command voltage at the headstage is nominally only 1V (in practice about 1.2V), so the maximum current passable through such an electrode is only 10nA. In practice there is no point in providing any larger voltage, because this recording situation is already hopeless!

Another problem with series resistance compensation is that even when only a fraction of the series resistance is compensated, the process is inherently regenerative. Without series resistance compensation, an increase in current will reduce the cell potential relative to the command potential, which will mean that the actual value of the current is less than what it would have been if there were no series resistance. Although an error, it actually has a stabilising effect on the system. However, even only partial series resistance compensation will cause the command potential to increase, which will cause a further increase in the current, which in turn will cause a further (albeit smaller) increase in command potential, and so on. This type of regenerative effect tends to result in damped oscillations of both the voltage and current in response to a sudden change in either of these parameters. Feedback loops tend to be more stable when dominated by a single time delay, and in practice it is helpful to provide such a delay in the series resistance compensation circuit by low-pass filtering the current error signal that is fed back to the command potential. In common with usual practice we call this the "lag" control, and by increasing the lag it is often possible to compensate for more of the series resistance. However, there is a trade off in that when a larger lag setting is used, series resistance compensation is initially less effective in correcting the command potential following a sudden current change, so the best settings of the compensation and lag controls are inevitably a compromise. Although the actual cell membrane potential cannot be measured, the headstage command output on the back panel can be used to monitor the effect of the series resistance compensation (as well as that of any precharging) on the command potential, which gives a useful guide to the voltage errors that the series resistance compensation is trying to correct.

We have arranged for series resistance compensation always to be active in the small cell and big cell voltage clamp modes, i.e. it is unaffected by the setting of the RC comp switch. However, it is otherwise available only under the same conditions as for precharging, and an LED illuminates to show when series resistance compensation is in effect. If series resistance compensation is not required, it can be disabled by turning the %RS comp control fully anticlockwise. The advantage of this arrangement is that it allows the two forms of compensation to be selected entirely independently of each other. Note that if RC compensation is off, then capacitive charging currents in response to command potential steps will now need to be provided by the headstage itself, so they will directly appear on the current output. Under these circumstances, series resistance compensation can make the charging currents correspondingly faster and shorter, rather like precharging, but in our opinion RC compensation with precharging is a better way of doing this.

We mentioned in the description of electrode capacitance compensation that our oscillation-killer facility did indeed protect cells very effectively from amplifier instability caused by overcompensation in current clamp mode. We have therefore extended its operation to include RS compensation. In fact, if it detects a command potential oscillation, it disables electrode capacitance compensation, RS compensation and (for good measure, although it rarely causes problems) precharging as well for about half a second. If oscillation recurs when these functions are re-enabled, it promptly switches them off again, and it will continue to do so until the control that is causing the problem is readjusted. We had originally intended to make the oscillation detection circuit user switchable, but in practice we could find no advantage in ever switching it off. However, it can be disabled if required, by appropriately changing the position of the link on the "stabiliser" jumper on the headstage driver board. The sensitivity of the circuit can be adjusted by PRE2 on this board, although in practice this is best left at maximum, i.e. fully clockwise.

The situation for current clamping is different in some but not all respects. For current clamping, as already described, the patch clamp must have a high input impedance to function correctly. Amongst other things, this means that its input capacitance, and that of the electrode, must be compensated exactly as for voltage clamping. Once this is done, the Optopatch will act as a pure current source, and unlike voltage clamping, there is no need to make any compensation for the cell capacitance or series resistance in the whole-cell mode. However, the current passed by the patch clamp must pass through the series resistance, and although this does not affect the accuracy of the current clamping, it does mean that the potential recorded by the electrode differs from that at the cell by the voltage drop across the series resistance. This is the similar to the series resistance problem with voltage clamping, but there is an important difference. As explained above, series resistance compensation in voltage clamp mode occurs within the main feedback loop and tends to destabilise it, but the errors in current clamp are OUTSIDE the feedback loop, making them much easier to compensate. All we need to do is to subtract an appropriate proportion of the current clamp command signal from the recorded potential, in order to compensate for the IxR drop across the series resistance. In practice this is easy, and is in fact a relatively direct alternative way of measuring the series resistance. If there is zero series resistance, and the cell membrane resistance is high, the potential recorded in response to a step in current will be integrated by the cell capacitance to form a ramp. The effect of a finite series resistance will be to cause a voltage step at the start of the ramp, so to correct for it one merely needs to subtract enough of the current command signal to cancel the step. To do this in either cell mode, the RC comp switch should be in the on position, and the adjustment should be made using the series res control, which as before provides a direct measure of the series resistance. The cap, %precharging, %RS compensation and lag controls are all inactive in current clamp mode.

To summarise, compensation for the electrode capacitance is relatively easy and straightforward, and compensation for the effects of the cell capacitance can also be achieved with reasonable accuracy, whereas compensation for the effects of sustained currents through the series resistance under voltage clamp is generally less satisfactory. In that case the best approach - to the extent that the cells allow, of course - to ensure that the series resistance is as low as possible in the first place. However, under many circumstances the cells may not be generating significant ionic currents, so errors from this source are not a problem. Instead the cell can just be modelled as a membrane capacitance connected to the series resistance, with the possible addition of a much higher resistance across the membrane capacitance to represent the relatively small ionic currents. This situation allows extremely precise measurement of the membrane capacitance, using the techniques described in the next section.

7 Membrane Capacitance Measurement

The features described in the previous section allow not only for the cell membrane capacitance to be compensated, but also allow its accurate measurement, since the cap control and the associated meter display are directly calibrated in capacitance units. However, there is also considerable interest in measuring time-dependent changes in membrane capacitance at the very highest possible resolution. The major application for such measurements is the study of cellular secretory processes, so we shall use this as the example in the following discussion. During vesicular excretion of substances from the cell, the contents are released from the cell by the fusion of the membrane of the vesicle with that of the cell, thereby increasing its area and thus also its capacitance. Fusion of an individual vesicle to the cell membrane may increase the cell capacitance by only a few femtofarads (i.e. thousandths of a picofarad) or less, but massive secretory events involving the fusion of very many vesicles, such as the de-granulation of mast cells, may result in at least a doubling of the membrane capacitance.

Various methods have been described for the accurate measurement of changes in the cell capacitance, but in the past none has proved ideal, although we modestly hope that the Optopatch will resolve the difficulty of which one to choose. The following paragraphs describe our approach in some detail, which has now also been published (Johnson, Thomas and Kros, 2002). It seems generally accepted that the most sensitive method of measuring small changes in capacitance is by using the lock-in amplifier technique described by Neher and Marty (1982), and Lindau and Neher (1988), but in its basic form this technique is not so suitable for measuring larger changes, as both the sensitivity and linearity are lost as the capacitance shifts away from its starting value. Larger changes have therefore generally been measured by a variety of software methods, such as by analysing the time course of the current in response to step potential changes (e.g. Lindau and Neher, 1988). For users of commercial patch clamps, the software methods also have the advantage that no other special equipment is required, whereas the lock-in amplifier technique requires special hardware. Although this hardware does not have to be integrated into the patch clamp in order to apply the technique, it certainly makes sense to do so, and that strategy has allowed us to develop it further in the ways that we shall shortly describe. An intermediate technique has been to use software to simulate a lock-in amplifier (Joshi and Fernandez, 1988; Fidler and Fernandez, 1989). The software of the EPC-9 patch clamp provides an equivalent facility, although from published information (Sigworth, Affolter and Neher, 1995) it appears that a somewhat different algorithm is used. Although software simulations can at best only approach the performance of the equivalent hardware in respect of both resolution and speed of operation, they can be made to work effectively over a wider capacitance range (Fidler and Fernandez, 1989), because it is easier to keep them optimally adjusted. The capacitance measurement facilities on the Optopatch have been designed to combine the performance advantages of dedicated hardware with the wider measurement range and greater ease of use of the software-based approaches, by using novel additional circuitry that allows the lock-in amplifier always to remain perfectly balanced. The following paragraphs describe the basic operating principle of the lock-in amplifier technique and of our enhancements to it.

The equivalent circuit of the recording configuration contains only three components, i.e. series resistance, membrane resistance and membrane capacitance, but a full theoretical analysis of it for the purposes of capacitance measurement is actually far from straightforward. The article by Gillis (1995) covers the subject in detail. It is quite impressive that the analysis of such a simple circuit requires so much mathematics! This and the other papers cited above should be referred to for more detailed information on the application of the lock-in amplifier technique to this type of circuit, but the following discussion should provide sufficient background to explain the operation and use of the facilities provided on the Optopatch.

We'll first consider the simple case in which the membrane resistance is infinite, so that the only circuit components are the series resistance, Rs, and the membrane capacitance, Cm. The lock-in amplifier technique can be applied to this system as follows. The cell is voltage clamped at a suitable holding potential, usually close to its normal resting potential, and a sinusoidal command voltage of typically a few tens of millivolts peak-to-peak at a frequency of about 1KHz is superimposed on it. It is important that the voltage excursions do not cause significant ion channel activation, so that the membrane resistance remains high compared with the series resistance. The sinusoidal command voltage is effectively applied across Rs and Cm in series, and the patch clamp supplies and records the resulting sinusoidal current. If the current pathway were purely resistive, the current would be exactly in phase with the voltage, and its amplitude would be a reciprocal measure of the resistance. On the other hand, if the current pathway consisted only of Cm, the phase of the current flow would be shifted by 90 degrees relative to the command voltage (the phase shift represents a differentiation, which arises because the current through the capacitor depends on the rate of rise of the voltage across it).

Strictly speaking, we should speak of capacitors drawing in and releasing current, because there is no actual current through them as there is through a resistor, but for AC signals capacitors behave as if there is, so it is a useful shorthand. The "resistance" of a capacitor falls with increasing frequency, so the more general term of impedance is used in this case. Unfortunately there is an inconsistency in the normally used measurement units, because the impedance of a capacitor is lower if the capacitance is greater, so from the current-passing point of view a large capacitance resembles a small resistance. However, the capacitance units make more sense when dealing with currents, and for consistency it is therefore sometimes preferable to quantify resistances in their reciprocal units of conductance (the corresponding term to include capacitance as well is admittance). However, since resistance seems to be a much better understood concept than conductance, we have now standardised on using resistance rather than conductance for all the Optopatch controls.

We are now ready to consider the case of a capacitor and a resistor in series. The phases of the currents through them relative to the applied voltage must be the same as each other, so the voltages across the two components must be at 90 degrees to each other in order to maintain the phase relationships described above. This relationship is best illustrated by vectors. The vector sum of the two voltages at right angles is equal to the command voltage, and this gives us sufficient information to calculate their magnitude and phase, as Fig. 9 graphically shows. The phase angle of the voltage across the resistor must also be the phase angle of the current through both components, and the voltage across it, i.e. the length of the vector, gives the amplitude. Conversely, this means that if we measure the magnitude and phase of the current, we can calculate both Rs and Cm. Clearly, a large change in either Rs or Cm will affect the voltages across both of them, and therefore both the magnitude and phase of the current, but let us consider instead the effect of a very small change in either, in which case the effect on the other can be ignored. The geometrically equivalent approximation in Fig. 9 is to say that the length of the other voltage vector remains the same, which we can accommodate by allowing the angle between the two voltage vectors to deviate very slightly from 90 degrees. We can see by inspection of Fig. 9 that if Rs decreases slightly, the magnitude of its voltage vector will increase slightly, but its phase remains essentially unchanged. There will therefore be a change only in the magnitude of the current. On the other hand, if Cm increases slightly, the magnitude of the Rs vector, and hence of the current, will remain unchanged, but the effect of changing the length of the Cm voltage vector is to swing the Rs voltage vector around slightly. In this case, therefore, the effect is to change the phase of the current. The effect of a small phase change is equivalent to adding or subtracting a small component at a 90 degree phase angle, as vector addition of the type shown in Fig. 9 will immediately demonstrate.

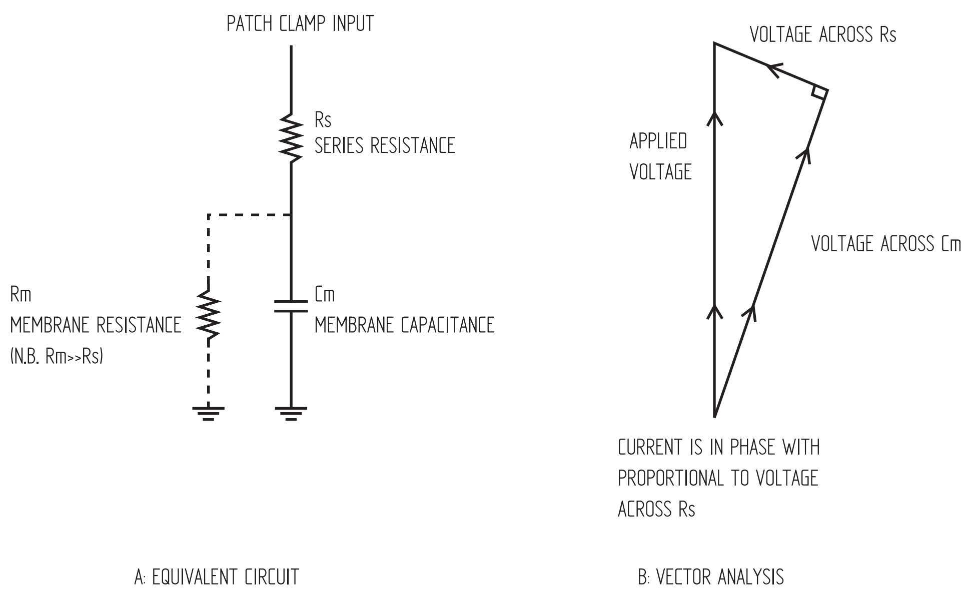

Figure 9 – Model System for Membrane Capacitance Measurements

This is all well and good, but the current waveform would need to be measured with considerable precision in order to quantify and discriminate between these small changes. however, the lock-in amplifier method allows us to play a trick that makes the task much easier. We've already seen that the current that charges the membrane capacitance in response to a step change in the command potential can be supplied by an appropriate signal via a capacitor, instead of by the current-passing resistor (or equivalent components) in the headstage. Such currents therefore do not appear on the current output signal from the headstage. Exactly the same situation can apply to sinusoidal changes in the command potential, and the required res and cap control settings are the same. In both cases, when these controls are optimally adjusted, all the current will be passed by the capacitor, and the current output signal from the headstage will remain at zero.

If either Rs or Cm change by a small amount after adjustment of these controls, a small sinusoidal current will appear. This current will be in phase with the current supplied by the compensation capacitor if Rs changes, and at 90 degrees to it if Cm changes. The advantage of this recording situation is that the changes can now be measured directly as small signals, instead of being calculated from the difference between two large signals. However, Rs can change significantly during an experiment, so we have to be able to distinguish between changes in Rs and changes in Cm. This is the basic reason for using a lock-in amplifier, as it provides the required discrimination.

A lock-in amplifier can do this because it works as follows. It has two independent signal-processing paths, each of which can be switched electronically to have a gain of either plus one or minus one, and the two switches are driven by square waves that are of the same frequency as the input waveform, but of different and variable phase. The two outputs are then filtered to give DC signals. The operation of such a system is illustrated in Fig. 10. From this it can be seen that if the gain switching coincide with the zero-crossing points of the input signal, the signal is full-wave rectified to give a maximum DC output. In this example the output is positive, but it would be negative if either the input or the switching waveform were inverted. Conversely, if the gain switching coincide with the maximum and minimum values of the input signal, the DC output is zero. Gain switching at other points give intermediate outputs. In a practical circuit, the phases of the gain switching are variable with respect to the input signal, but a constant phase difference is maintained between the two signal paths.

As a technical point, we should note that switching the gain with a square-wave signal also detects signals which are odd harmonics of the square wave frequency, since a square wave is itself such a harmonic series. This can impair the noise performance of the system, since noise at the harmonic frequencies will also be detected. Our lock-in amplifier includes a low-pass input filter to reduce the amplitude of any harmonic signals to insignificant levels, and the cut-off frequency of this filter varies automatically according to the operating frequency of the lock-in amplifier. A side effect of this filter is to shift the phase of the fundamental frequency by about 45 degrees, which turns out to be quite useful as will be discussed later.

By sending the current output signal to a lock-in amplifier, we can therefore make the required discrimination between changes in series resistance and changes in membrane capacitance, IF we set the phase of the gain switching so that one output represents purely the capacitance signal and the other represents purely the resistance signal. Since we've already seen from Fig. 9 that the required phase varies according to the values of these parameters, it also needs to be adjusted in the course of an experiment. Furthermore, the relation between a given small change in either parameter and the corresponding output from the lock-in amplifier is linear only for small changes, and the slopes also depend on the starting values (hence the possibility for so much mathematics). The various software approaches can be an attractive alternative, because however they work, they can hide any such complexity from the user, but the advantage of the lock-in amplifier method is that it is very sensitive. Not only has most of the current signal been balanced out before it is sent to the lock-in circuit, but also arbitrarily small signals at the gain-switching frequency can be measured if the outputs are sufficiently filtered, whereas other signals, including noise, are rejected. This makes it worthwhile to put up with the other problems.

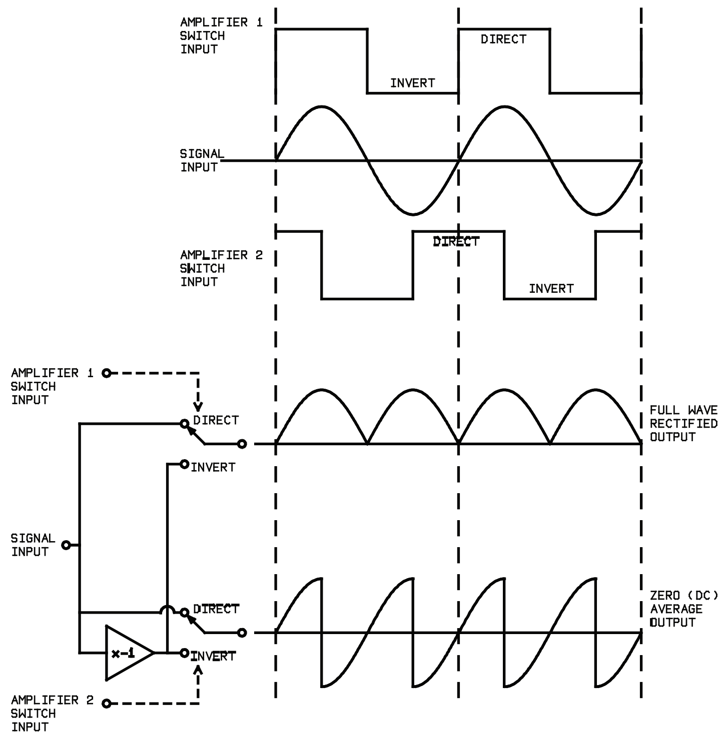

Figure 10 - Operating Principle of the Lock-In Amplifier.

One useful facility to assist the calibration can be to provide a capacitance "dither" switch or signal input, which changes the effective cap control setting by a precise small amount, thereby superimposing a calibrated marker pulse on the capacitance output signal. In addition, it is a useful way of determining the correct phase, since it will then affect the capacitance output only. An alternative but equivalent method of phase determination is to dither the effective res control setting instead, and to check that it only affects the resistance, which is less useful for calibration purposes, but has the advantage of minimising the disturbance to the capacitance signal when the phase is already nearly correct (needless to say, we're going to make use of this idea). The Optopatch provides both res and cap dither signal inputs, calibrated such that a 1V input corresponds to a dither of 1% of full-scale of each control, and there is also an internal 70Hz oscillator which can be used to dither the effective res control setting. In order to enable these two inputs, the phase switch should be in the dither position (see the previous description of the front panel controls for further information).

While discussing the correct phase to use, we should note that Joshi and Fernandez (1988) have pointed out that the presence of a finite membrane resistance changes the situation slightly from the ideal one described above. The effect is that the phases of the currents that are measured in response to small changes in Rs and Cm are no longer precisely at 90 degrees to each other. This is because some of the current through Rs now flows through the membrane resistance instead of the membrane capacitance, but as long as the proportion remains reasonably small, it does not significantly affect the capacitance measurement. From an accuracy point of view, the main effect is to introduce an error in the Rs estimation, which is of no practical consequence. However, in order to minimise the effects of changes in Rs on the Cm signal, the ideal approach is to set the phase by dithering Rs itself and adjusting for the minimum effect on the capacitance measurement, rather than by dithering the res or cap controls as described above. This is a nice theoretical point, but in practice the difference between the phase settings determined by the two techniques is likely to be fairly small, and Gillis (1995) is of the opinion that their method of dithering Rs, by switching in a 1 megohm resistor in series with the electrode, is prone to other errors. In any case, the Optopatch incorporates a rather different method of dealing with the problem.

The possible need for frequent readjustment of the res, cap and phase controls with the lock-in amplifier technique is clearly a disadvantage. It would be much nicer if these adjustments could be made continuously and automatically, so we have provided this as an optional facility in the Optopatch, which we call track mode. The idea is conceptually very simple. We use gain-controlled amplifiers in the automatic section of the RC compensation circuit, and we use the two outputs from the lock-in amplifier circuit to derive control voltages for these amplifiers. The outputs from these amplifiers are added to those from the main part of the RC compensation circuit (which operates in the conventional way), causing them to change exactly as if the res and cap control settings had been changed. The effect of this is to reduce the outputs from the lock-in amplifier, and by providing sufficient gain in the feedback loop, these outputs can be made to remain effectively at zero. Thus the circuit is always balanced, and even better, the resulting changes in the res and cap control voltages generated by the two feedback loops are always directly related to the membrane capacitance and series resistance changes! In the Optopatch, changes in both Cm and Rs are thus available as linear and calibrated signals, and they can also be read on the meter if required. The response time is also fast, i.e. in the millisecond range in our circuit, and our performance measurements confirm that the high resolution of the lock-in amplifier technique is preserved. We therefore suspect that users will always prefer to use the system in this way, although conventional operation of the lock-in amplifier also remains possible for comparative or other purposes.

To have the circuit simultaneously under the control of two independent feedback loops may sound impossible, but it is easy to prove that there is no significant interaction between them, by disabling one or other of them (although we could see no point in including this possibility in the final design). The reason is that a sine wave of any phase can be regarded as the sum of two independent components, of appropriate amplitude and at 90 degrees to each other. Therefore the lock-in amplifier with feedback can independently recognise and remove the resistive and the capacitive components of the residual current signal, since it actually sees the two components as two independent signals.

We also need to discuss the effect of the gain selector on the resistance and capacitance measurement facility. The lock-in amplifier is driven from the output of the gain stage, so when used conventionally, the real and imaginary outputs will both be amplified by an amount according to the gain setting. Before taking the gain setting into account, a 1V rms signal on the headstage current output will give a 1V DC output when the phase is optimised for detection of that signal, and the gain rotary switch control thus allows up to 500-fold amplification. In track mode, the gain control is within the two feedback loops, so it does not affect the resistance and capacitance outputs in this case, but it DOES affect the amount of feedback (it varies the open-loop gain), which in turn affects both the response time and the stability of the system. As the gain is increased, the response time of the system to a step change in capacitance (or resistance) is reduced, although just as in any other feedback system the circuit will become unstable if too much gain is applied. See Johnson, Thomas and Kros (2002) for further information including experimental data to illustrate this effect. The overall loop gain is affected by the actual values of series resistance and cell capacitance, and although these are very unlikely to change during an experiment by an amount sufficient to significantly affect the stability, the system is capable of operating over at least two orders of magnitude of membrane capacitance (from as high as 200pF down to 1pF or possibly even less). Such a wide range could not be optimally handled with a single fixed gain! For guidance though, we suggest a gain of around 100 for the 100nA full-scale current range (which corresponds to the 1K mV/mV position on the gain control), or just 1 (i.e. the 10mV/mV position) for the 1nA full-scale current range.

As a general rule, we recommend use of the 100nA current range for tracking mode, because the minimum gain on the 1nA range may be too high under some conditions (but we welcome user reports!). Note that on the standard Optopatch, only a single capacitance range of 10pF full scale (increasable to 15 or 20pF with the +50% and +100% switches) is available, which is the same as the lowest range in the big cell mode. The main advantage of the small cell mode is thus the ability to measure smaller currents more accurately, rather than to measure smaller capacitances. However, it is possible to modify the Optopatch to provide a 3pF full scale range (selected by the 30pF switch position in small cell mode), and we may make this a standard feature if it proves to be useful. Measuring such small capacitances requires the use of much higher command potential frequencies, typically on the order of tens of kilohertz, but the Optopatch circuitry is capable of handling this. We believe that the main practical difficulty in measuring very small membrane capacitances is likely to be the increasing interference from changes in the electrode capacitance (e.g. due to tiny changes in fluid level in the bath) as this is going to become relatively more important as the membrane capacitance is reduced, but we nevertheless suspect that the Optopatch will prove to be a rather suitable amplifier for making such measurements. Again, user feedback will be most welcome.

For the circuit to work properly, it is also important for the phase to be set appropriately, so we have included some additional optional facilities to cope with this. If the phase is badly wrong, the feedback pathways will become positive rather than negative, so the circuit will not stabilise, but as expected from feedback theory, phase errors even approaching of 90 degrees can be tolerated. As with the standard lock-in amplifier circuit, the effect of an incorrect phase setting will be that an error signal on the real phase (resistance) output, caused by a sudden change in Rs, will also be accompanied by an error signal on the imaginary phase (capacitance) output, but in this case it will do so only transiently. The reason for that is that even if the phase is somewhat incorrect, the double feedback loop will nevertheless succeed in reducing both outputs to zero, but during the settling period, the capacitance feedback loop will see a small signal to which it will react inappropriately until the resistance feedback loop, which still sees most of the signal, has removed it. Slow changes in Rs will not show up in the imaginary (capacitance) output at all, even if the phase is not optimally set, and we suspect this will prove to be a nice way of minimising any inaccuracies caused by the phase-shifting effect of a finite membrane resistance as pointed out in the Joshi and Fernandez (1988) paper referred to above.

Nevertheless, we could not resist the temptation to include the option of automatic phase tracking as well. In principle, we could compute the required phase shift from the command voltage frequency plus the initial values and subsequent automatic changes in the res and cap control settings, but this approach would be complicated and prone to a variety of errors. We therefore prefer the alternative approach of determining the required phase directly, since this will always work correctly. The method is to modulate the res control voltage with a low-frequency sine wave (we use 70Hz), to simulate changes in this parameter. We then send the capacitance output signal from the lock-in amplifier to the input of another lock-in amplifier, which detects any 70Hz signal at that point. If there is any 70Hz signal on the capacitance output, then the phase of the main lock-in amplifier must be incorrect, and another feedback loop adjusts the phase (which is also a voltage-controlled parameter) in order to eliminate it. This lock-in amplifier needs to have only one switching element rather than two, and its switching phase is set so as to maximise its sensitivity to the 70Hz signal, but since this phase relationship doesn't change, that can be a fixed part of the circuit design. We chose 70Hz because it is a high frequency compared with the bandwidth over which the capacitance signals are normally measured, so any residual is easy to filter out, but it is a low frequency compared with that of the command voltage, so it is clearly detectable on the capacitance output. Both this and the resistance output are three-pole Bessel filtered with a cut-off frequency of 100Hz, so the 70Hz signal is relatively little affected by these filters. In practice, further filtering of the capacitance signal is advisable in any case, so that will easily remove any residual 70Hz component from it. Under our test bench conditions, the residual 70Hz output before filtering was only about 1mV, whereas the full-scale output is 10V, so the intrusion of the 70Hz signal really should be negligible.

The response time of the phase tracking circuit has deliberately been made relatively long, for several reasons. First, this minimises any possible interaction with the automatic res and cap feedback loops, which, as already explained, do not require a perfect phase setting in order to function. Therefore, on the time scale seen by the 70Hz loop, the automatic res and cap settings are always correct. Second, a longer integration period within the 70Hz loop allows a smaller 70Hz signal amplitude to be used (we suspect that the level of the 70Hz signal could be further reduced if there was any point in doing so). Third, relatively substantial changes in Ra and/or Cm are required to cause a significant phase shift, so the required phase is most unlikely to shift at all rapidly. Our circuit has a response time of a few seconds, which we expect will be fast enough to allow sufficiently accurate tracking under any reasonable experimental conditions.

The extent of the automatic phase adjustment can be displayed on the meter in the vphase position. Phase adjustments of up to +/- 45 degrees can be made, corresponding to displayed meter voltages of up to +/-10V. The relation is not entirely linear, so the meter display is just shown as volts rather than degrees, but quantitative phase information is of no particular interest anyway. The important point is that if the automatic phase adjustment is seen to be approaching its positive or negative limits, the phase control can be adjusted manually to bring it back within range.

The complete circuit is thus rather complicated, but the individual components are not particularly expensive, and the overall cost of this method is much less than that of the computer that would be required to do all this in software. Furthermore, it allows fast, sensitive, accurate and calibrated measurement of membrane capacitance in real time, which we do not believe any other method can offer. In general, the software approach to solving any problem is powerful, important, and occasionally wonderful (we even love some software ourselves), but it would take an enormous amount of digital processing to match what our analogue circuits are doing here!

The best way of describing the controls for capacitance measurement is to describe the typical experimental sequence. Of course, we recommend that experience should be gained on an appropriate model cell before attempting any real experiments, as there is quite a lot to learn here, but in the following description we'll assume that everything is for real. First of all, the cell is patched, the patch is blown to give whole-cell voltage clamp recording, and the frequency generator is set to give a sine wave of appropriate amplitude and frequency. Note that the internal oscillator MUST be used, as it generates two signals with 90 degree phase differences in order to drive the lock-in amplifier and phase control electronics, and an external signal would be unable to do this. For small capacitances, below 10pF, the small cell mode can be used, but depending on the series resistance, a frequency somewhat higher than the 1KHz recommended above may give better results. As a guide, we recommend a frequency such that the impedance of the membrane capacitance at that frequency equals the series resistance. Trial and error will show what is best, and the lock-in amplifier system will work at frequencies well in excess of 10KHz, so there are no limitations there!

Neither precharging nor series resistance compensation are appropriate for capacitance measurement, so both these controls are inactivated when the lock-in amplifier is on. To begin with we'll deal with use of the lock-in amplifier in its conventional, i.e. non-tracking mode. The holding and oscillator potentials should be set to traverse the membrane potential over a range in which no significant ionic currents are flowing, so that the membrane resistance remains high. This consideration limits the oscillator amplitude to probably no more than 20mV rms or so, whereas sensitivity considerations encourage use of the largest possible oscillator amplitude, so – again by trial and error – the best compromise needs to be found. To activate the lock-in amplifier, the RC comp and phase switches should both be on, and then the res and cap controls should be adjusted to cancel the sinusoidal output current. The lock-in amplifier takes the current signal after it has gone through the prefilter and gain stage, but before the four-pole variable filter. Therefore the gain control also determines the gain of the real and imaginary lock-in outputs, and in normal (non-tracking) mode the gain setting is just according to user convenience.

We are now ready to set the phase control. This control covers a total phase range of 180 degrees, to give a reasonable safety margin over the range required in practice, and unlike simple phase-shifting circuits, the phase shift it produces is independent of frequency. For best discrimination between resistance and capacitance changes, the sine wave frequency should be such as to give about a 45 degree phase shift in the current relative to the voltage (Fig. 9 shows the reason for this). As explained in the description of the lock-in amplifier, our design includes a tracking low-pass filter that introduces an additional phase shift of approximately 45 degrees, also regardless of frequency, giving about 90 degrees altogether. This very conveniently means that an initial phase setting of about one-half of full-scale should be aimed for.

There are at least two ways of checking the phase adjustment. One method is to de-tune the res and cap control settings manually, and to observe the effects on the real and imaginary phase outputs. The phase is correct when the adjustment of either control affects only the appropriate output. Note that, for example, to simulate the effect of an INCREASE in membrane capacitance, for which a positive output change is clearly appropriate, the setting of the cap control must be correspondingly REDUCED, so don't let this confuse you! The other method is to use the dither facilities, particularly the built-in 70Hz oscillator to dither the res signal (its level is variable by the control on the rear panel). The oscillator is activated by setting the phase switch to the dither position, and the phase control should then be adjusted to minimise the 70Hz signal on the imaginary phase output of the lock-in amplifier.

In this mode, the calibration of the real and imaginary phase outputs in terms of resistance and capacitance changes depends on a number of factors, including the sinewave amplitude, gain control setting, and a rather complicated dependence on the actual resistance and capacitance values. Rather than even trying to calculate the overall relation, the calibration should instead be performed by introducing known resistance and capacitance changes and measuring the consequent changes in the two outputs. Fortunately, such changes can be simulated by adjusting the res and cap controls as described above, either by physically changing their settings or applying appropriate dither inputs when the phase switch is moved to the dither position.

Having got this far, we can now explore the tracking mode of the lock-in amplifier, and to take things one step at a time, we shall assume that the phase switch is back in the on position rather than in the dither position. However, before moving the RC comp switch to the track position, the gain control setting should be checked. For safety, until you are thoroughly familiar with the system, we recommend that minimum gain is selected. In any case this is almost certainly essential if the 1nA gain range is being used in the small cell mode, as the current gain is 100 times higher in this case. We therefore recommend that you learn to use the tracking mode in big cell mode, and then perhaps try the small cell mode later on if it may be appropriate.

If the res, cap and phase controls were (more or less) correctly set, the real and imaginary phase outputs should remain at or near zero, but their function has nevertheless changed. They now carry the error voltages generated by the lock-in amplifier to maintain zero sinusoidal output current. These voltages have the same gain relationship as the res and cap output voltages, i.e. 10V for a full-scale change, although the actual output voltage range is limited to 50% of full scale, i.e. +/-5V. Note that the outputs can change in either direction, so negative values may also be generated, provided of course that the total control voltage remains positive. The fact that these outputs represent capacitance (and resistance) changes, and therefore start from zero, allow them to be amplified by other equipment if required, without ending up with large DC signals. If large DC signals do develop as a result of more substantial changes in series resistance or cell capacitance, they can easily be removed by readjusting the res or cap controls, but the development of large signals in no way affects the accuracy of the system, so long as the phase remains correct. On the other hand, these difference voltages are also added to the res and cap outputs, both on the meter and on the rear panel sockets, allowing the total capacitance and resistance to be monitored if preferred (this addition occurs ONLY in tracking mode, because only then do the real and imaginary outputs represent calibrated resistance and capacitance signals).

There is no longer any need to use the resistance and capacitance dither inputs for calibration purposes in tracking mode, but they can still be used, e.g. to measure the response time to step changes, as we'll describe below. However, there is a very important difference concerning use of the dither inputs in tracking mode. In this mode, unlike the conventional lock-in amplifier mode, it is NOT necessary to move the phase switch to the dither position. In fact, the dither position in tracking mode activates the automatic phase tracking facility, described below. Therefore, to apply resistance and capacitance dither signals in tracking mode without activating the phase tracking system as well, the phase switch should be left in the on position.

We can now investigate the effect of the gain control in tracking mode. Since the gain control is now within a feedback loop, it no longer affects the magnitude of the real and imaginary phase outputs for given resistance and capacitance changes. However, in this configuration, more gain now means more feedback, which in turn means faster and more response to resistance and capacitance changes. Just as for any other feedback amplifier, too much gain results in instability and oscillation, and under most recording conditions the system (in big cell mode) becomes unstable at gains above 200 (2KmV/mV position on the gain control). In practice, gains of up to 50-100 (500-1KmV/mV)seem safe enough! The easiest way of investigating the effect of the gain control is to apply a low-frequency square wave to the capacitance dither input in tracking mode (see previous paragraph!) and then observe the imaginary phase output. Note that both the real and imaginary phase outputs are low-pass filtered at 100Hz, but if you want to observe an unfiltered imaginary phase output, this can be done by moving the filter source switch to the cap position, setting the output filter to a reasonably high frequency, say 1KHz or above, and then observing the output there. Increasing the gain reduces the response time for a step capacitance dither change, but the system remains well-behaved, i.e. no overshoots, until stability is lost completely, which typically occurs at gains of 200 or higher.

We can also investigate another useful feature of the tracking mode, which we hadn't appreciated until we'd observed it for ourselves. When a lock-in amplifier is used conventionally, it is very important that the switching phase is set as accurately as possible. Otherwise, a resistance change will also cause some change in the imaginary phase (capacitance) output and vice versa. That also occurs in tracking mode, but only TRANSIENTLY in this case. The reason is that in tracking mode, both the res and cap control settings are continuously readjusted to preserve a flat current trace. There is only ONE setting of the res and cap controls which will achieve this, and so long as the switching phase of the lock-in amplifier is at least roughly correct, so that the system is stable, the tracking mode will find and maintain these settings. So, what happens when the switching phase isn't exactly correct is that to begin with, an error signal IS generated in the incorrect output of the lock-in amplifier, so it will begin to adjust the setting of the incorrect control. However, a larger error signal is generated in the correct output, and the consequent readjustment of the correct control will reduce the error signal sent to the incorrect control as well. The higher the gain setting, the more rapidly will the incorrect transient signal be removed, although its peak amplitude will depend on the switching phase error, so there remain some advantages to getting the switching phase exactly right, which brings us to the final option.

The final option is to select automatic phase tracking as well, by setting the phase switch to the dither position. An internal oscillator is used to dither the res control voltage at 70Hz as previously described, and another lock-in amplifier measures its amplitude on the capacitance signal output. The 70Hz lock-in amplifier needs to have only one output, which will be zero when the phase is set correctly, otherwise there will be a positive or negative output according to the direction of the phase error. The output is added to the voltage set on the phase potentiometer, to produce the total control voltage for the phase shifter. The effect is to shift the phase so as to minimise the 70Hz signal on the capacitance output. Since the required phase is unlikely to shift very far or very fast, we have limited the response speed and phase adjustment range of this control loop, in order to prevent any risk of it interacting in any way with the rest of the circuit. In order to indicate the extent of any automatic phase adjustment that is made in this way, the error output voltage can be displayed on the meter by selecting the "phase" position on the meter switch. The total phase adjustment range is nominally plus and minus 45 degrees, which corresponds to +10V and -10V on the meter. The relation between the control voltage and the phase shift is reasonably linear but not precisely so, and this facility is not intended for precise phase shift measurement (which is unlikely to be of much interest). Instead, it is intended to show whether the circuit is nearing either end of its adjustment range. If it is, the manual phase control should slowly be adjusted to reduce the automatically applied component back towards zero - or maybe better, the operating frequency should be changed to produce the same effect, especially if the adjustment would otherwise bring the phase control significantly away from its central position.

In general, only the capacitance signal will be of major interest during experiments, so the resistance signal can often be ignored, although it is advisable to pay some attention to it in case substantial changes are occurring. Even so, accurate recording of it is unlikely to be warranted. On the other hand, accurate recording of the capacitance signal certainly is warranted. Our measurements with model systems confirm that capacitance changes well below 0.1% of full scale can be resolved (the results we achieved were actually so good that we are reluctant to quote them, because it seems unlikely that real cells would perform so well, and we don't want to mislead). The actual capacitance resolution achievable depends on a number of factors, but in order to increase it (at the extent of poorer time resolution), additional filtering can be applied. We have therefore provided a switch to send the capacitance output signal to the output Bessel filter, and its operating range extends down to the relatively low frequencies, i.e. about 1-100Hz, that are appropriate for this application.

Just in time for the July 2000 update of this manual, we received the very nice data shown in Figure 11, from Stuart Johnson and Corne Kros at Bristol University. This record is from a mouse hair cell, stimulated with a 3KHz sinusoidal command potential of 13mV rms. This was briefly interrupted by a depolarising step, while the lockin amplifier was gated off. The capacitance signal was 8-pole Bessel filtered at 250Hz, and phase tracking was also in operation. This result nicely confirms our predictions concerning the sensitivity of the system, and it also shows just how well the lockin amplifier gating facility can work in practice, so we are delighted to be able to include it here. As previously quoted several times already (!), we then got together for some joint experiments (Johnson, Thomas and Kros, 2002), which nicely illustrate all the varioius points made above.

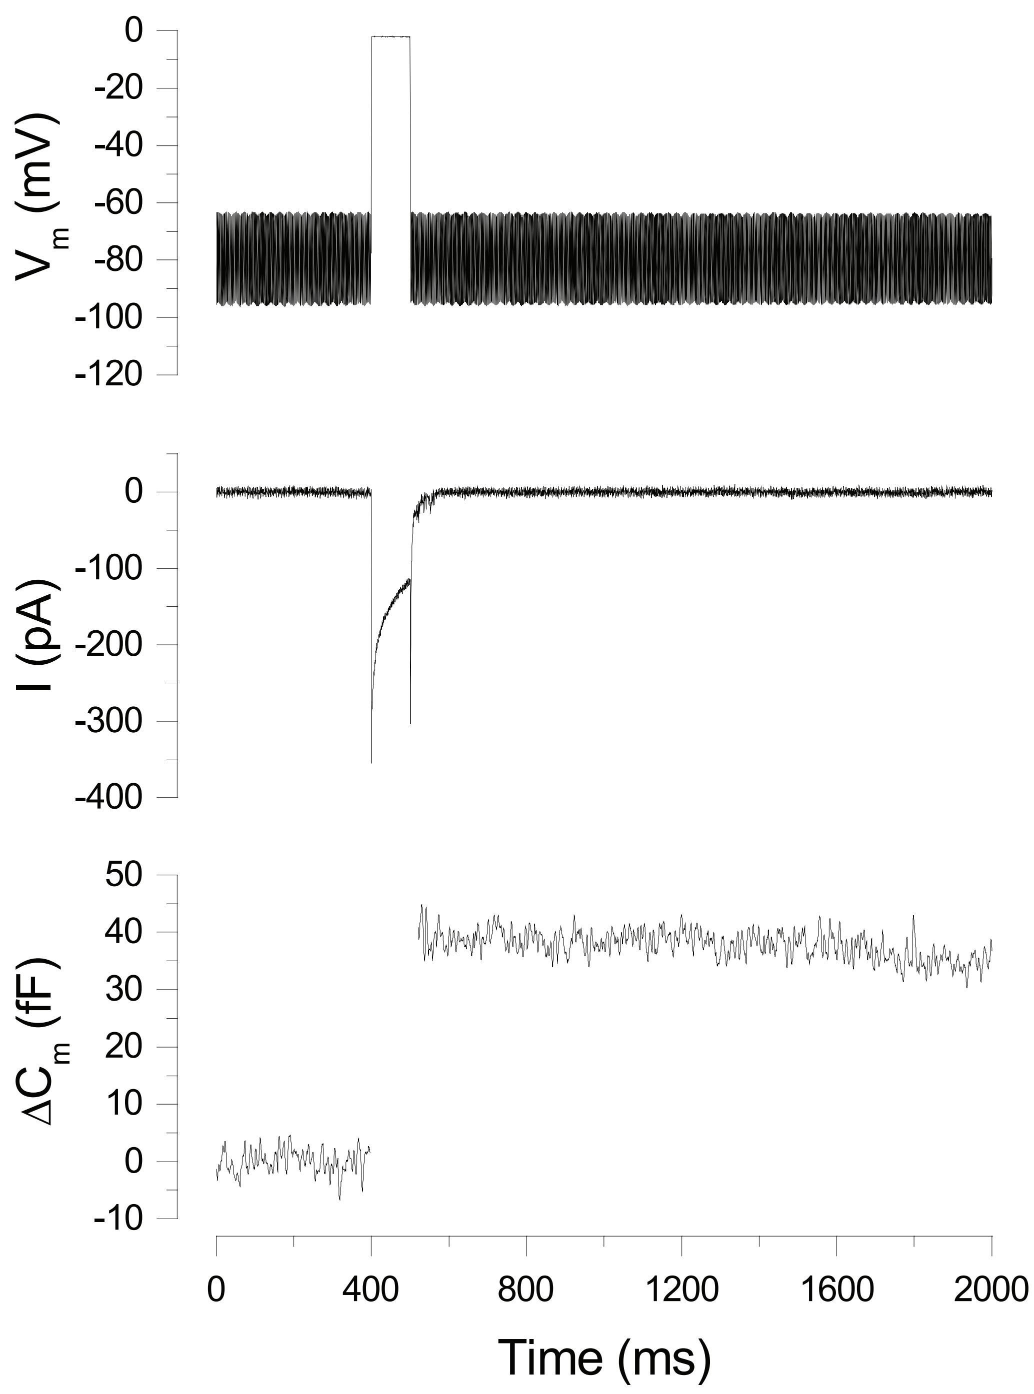

Figure 11 - An example of membrane capacitance measurement with the Optopatch

8 References

The Plymouth book supplied with the Optopatch is sufficiently useful that it may well have gone missing, so we give the full reference here in case you need to find another copy.
"Microelectrode Techniques: The Plymouth Workshop Handbook", editor David Ogden, 2nd edition (1994). Company of Biologists. ISBN 0 948601 49 3

The other general reference is:
"Single-Channel Recording", editors Bert Sakmann and Erwin Neher, 2nd edition (1995). Plenum Press. ISBN 0 306 44870 X

Individual references cited in this manual are as follows:

Armstrong, C.M., and Chow, R.H. (1987). Supercharging: A method for improving patch-clamp performance. Biophys. J., 52, 133-136.

Fidler, N., and Fernandez, J.M. (1989). Phase tracking: an improved phase detection technique for cell membrane capacitance measurements. Biophys. J., 56, 1153-1162

Gillis, K.D. (1995). Techniques for membrane capacitance measurements. In "Single-Channel Recording", pp 155-198.

Johnson, S.L., Thomas, M.V., and Kros, C.J. (2002). Membrane capacitance measurement using patch clamp with integrated self-balancing lock-in amplifier. Pfluegers Arch. 443, 653-663.

Joshi, C., and Fernandez, J.M. (1988). Capacitance measurements. An analysis of the phase detector technique used to study exocytosis and endocytosis. Biophys. J., 53, 885-892.

Lindau, M., and Neher, E. (1988). Patch-clamp techniques for time-resolved capacitance measurements in single cells. Pfluegers Arch. 411, 137-146.

Magistretti, J., Mantegazza, M., Guatteo, E. and Wanke, E. (1996). Action potentials recorded with patch-clamp amplifiers: are they genuine? TINS, 19, 530-534.

Magistretti, J., Mantegazza, M., deCurtis, M. and Wanke, E. (1998). Modalities of distortion of physiological voltage signals by patch-clamp amplifiers: A modelling study. Biophys. J., 74, 831-842.

Neher, E., and Marty, A. (1982). Discrete changes of cell membrane capacitance observed under conditions of enhanced secretion in bovine adrenal chromaffin cells. Proc. Natl. Acad. Sci. USA, 79, 6712-6716.

Sigworth, F.J. (1995). Electronic design of the patch clamp. In "Single-Channel Recording", pp 95-127.

Sigworth, F.J., Affolter, H. and Neher, E. (1995). Design of the EPC-9, a computer-controlled patch-clamp amplifier. 2. Software. J. Neurosci. Methods, 56, 203-215