Pflügers Arch (1982) 395:331–340

Pflügers Archiv
European Journal of Physiology

Instruments and Techniques

The Automated Analysis of Data from Single Ionic Channels

F. Sachs[1], J. Neil[1], and N. Barkakati[2]

[1] Department of Biophysical Sciences, SUNY, Buffalo, NY 14214, USA
[2] Department of Electrical Engineering, University of Maryland, College Park, MD, USA

Abstract. The development of single channel recordings has brought with it the need to analyse enormous amounts of data. The data analysis is time consuming and subject to observer biases since the events are random in time and are contaminated with uncorrelated noise. We have developed a heuristic pattern recognition program which identifies with high precision single channel currents and rejects contaminating noise. The program interactively provides for a variety of amplitude and duration measures. Analysis is flexible and rapid: a file containing over 10,000 events can be analysed in under 2 h.

Specific detection features include variable lowpass filtering, automatic baseline restoration, and adaptive amplitude thresholds. A record is analysed through duration histograms, binomial estimates of the number of active channels present, cross-correlation estimates between parameters, spectral analysis of events and background noise, and stationarity of mean channel current. The graphic output facilities can plot raw data (after filtering and baseline restoration) with the idealized signal superimposed or with detected events underlined. A batch processing facility has been included to allow processing of data during periods of low computer demand.

Key words: Ion channels – Single channels – Electrophysiology – Computer-patch clamp

Introduction

Since the development of the single channel recording techniques by Neher and Sakmann (1976), the method has been applied to many different preparations (Neher 1981; Hamill et al. 1981). The currents from single channels have the form of discrete jump-like events of random duration imposed upon an approximately Gaussian background as shown in Fig. 1. Since the relevant types of data are the current amplitude, the duration of channel "*open*" and "*closed*" times, and possible correlations between parameters, extracting the statistical properties of ionic channels from such records can be extremely time consuming. Many events have to be measured in order to gather statistically meaningful estimates of the transition probabilities between states.

We have worked for several years on automated analysis of this data, and over that period of time, the algorithms for analysis have undergone many cycles of complexity and simplification. It is our intent here to discuss what we have found as the most useful approaches and some of the hazards and pitfalls. There are many possible alternatives and undoubtedly more sophisticated approaches. The program has been designed to deal with stationary data and has been used exclusively for that purpose, but the modifications necessary to analyze channels with time dependent transition probabilities, such as single sodium channels (Patlak and Horn 1982), are minor. The analysis will be discussed using mostly nicotinic receptor-channel data as examples, since the bulk of our experience comes from that system. In what follows, we shall use the term "*channel*" interchangeably with "*unitary current*", to refer to the current that passes through a single channel. The term "*unitary current*" may be more precise, but is often clumsy in use.

The analysis program, titled IPROC, analyzes, off-line, the data time series for amplitude and duration information. Amplitude information is categorized in several ways. The most unbiased output is a total amplitude histogram, following baseline correction, of every data point in the record. This histogram contains information on the variations in channel amplitude and properties of fluctuations in open and closed states, possible dose response data (Sachs and Barkakati 1980), as well as information on the probability of activation. That is, the area under each peak in the amplitude histogram may be taken as proportional to the time spent in that state, or the probability of occurrence. The amplitude histogram can then be analyzed according to a binomial or Poisson distribution (see below).

A second amplitude histogram is generated from all those events which are judged by the program to be single channels. This is a highly biased amplitude histogram since only events occurring within definite amplitude limits are accepted, and weighting factors for the amplitude, such as the duration or steadiness of the amplitude, are lost in the compilation. Nonetheless, this histogram is useful for checking program operation, and searching for discrete populations of channels with different conductance.

A third amplitude histogram is constructed for events which are multiples of the unitary current amplitude (cf. Fig. 1B). Each data point which falls within an error band about a multiple of the unit event amplitude is entered into the histogram. This discrete level histogram is subjected to a binomial probability maximum likelihood estimate of the number of channels contributing to the record and the probability of each channel being open (cf. Korn et al. 1981; Patlak and Horn 1982).

Finally, as a rough measure of stationarity, the mean current of each record is calculated and plotted as function of record number.

Send offprint requests to F. Sachs at the above address

0031-6768/82/0395/0331/$02.00

Kinetic information is gathered in several forms. The simplest measures are the distribution of single channel open times and closed times. Currently, multiple amplitude events are not analyzed kinetically because of the ambiguity in rates introduced by not knowing which channel is undergoing a transition. Provided that the multiple events arise from independent channels, unambiguous equivalent information can be obtained from single events, although a bias may be introduced by the greater likelihood of multiple events occurring during long bursts (see discussion). Since all closed states are electrically equivalent, and there are usually several closed states, the opening rates must be scaled by the number of channels available to make the transition.

Closed periods may cover a wide range of time scales, from tens of microseconds to tens of seconds. In order to preserve the information with a wide dynamic range, three closed time duration histograms are constructed with bin-widths an order of magnitude apart.

A third type of duration accounting is concerned with activity that occurs in bursts, as shown in Fig. 1C. These bursts of activity are categorized by a class of closed times much shorter than the class of times between bursts, and generally represent repetitive activity of one channel (Neher and Steinbach 1978; Nelson and Sachs 1979, 1981). There is always some ambiguity in defining criteria for identifying these bursts, but in relatively sparse records, the ambiguity can be minimal (Colquhoun and Hawkes 1981). We categorize a burst as a series of events of unit amplitude preceded and followed by a closed time greater than some threshold. The burst parameters that we currently measure are the burst length, and the number of events per burst. The motivation for these measures can be seen from the simple sequential model for the nicotinic receptor-channel shown below, where A represents an agonist, Rc represents the channel in its closed conformation and Ro the channel in its open conformation, and rate constants as indicated

$$\mathrm{ARc} + \mathrm{A} \underset{k_d}{\overset{k_a}{\rightleftharpoons}} \mathrm{AARc} \underset{\alpha}{\overset{\beta}{\rightleftharpoons}} \mathrm{AARo}. \qquad (1)$$

Interpreted according to the above model for binding and channel opening, the mean intraburst closed time represents the inverse of $(\beta + k_d)$, the intraburst open time represents the inverse of α. The number of events per burst combined with the opening rate gives the dissociation constant, k_d (Colquhoun and Hawkes 1977), and with an assumption about the number of channels, k_a can be derived (Nelson and Schs 1982).

We have adapted the program to analyze, for example, amplitude-duration cross correlations for open and closed times (cf. Auerbach and Sachs, in preparation), power spectra of open and closed channels and bursts, and duration correlations between events.

Methods

A unitary current can be easily recognized by even the casual observer as a sharp step in the current-time trace, as shown in Fig. 1. For long steady events, with a good signal-to-noise ratio, and only a single population of identically sized unitary currents, the identification task is simple. With signal-to-

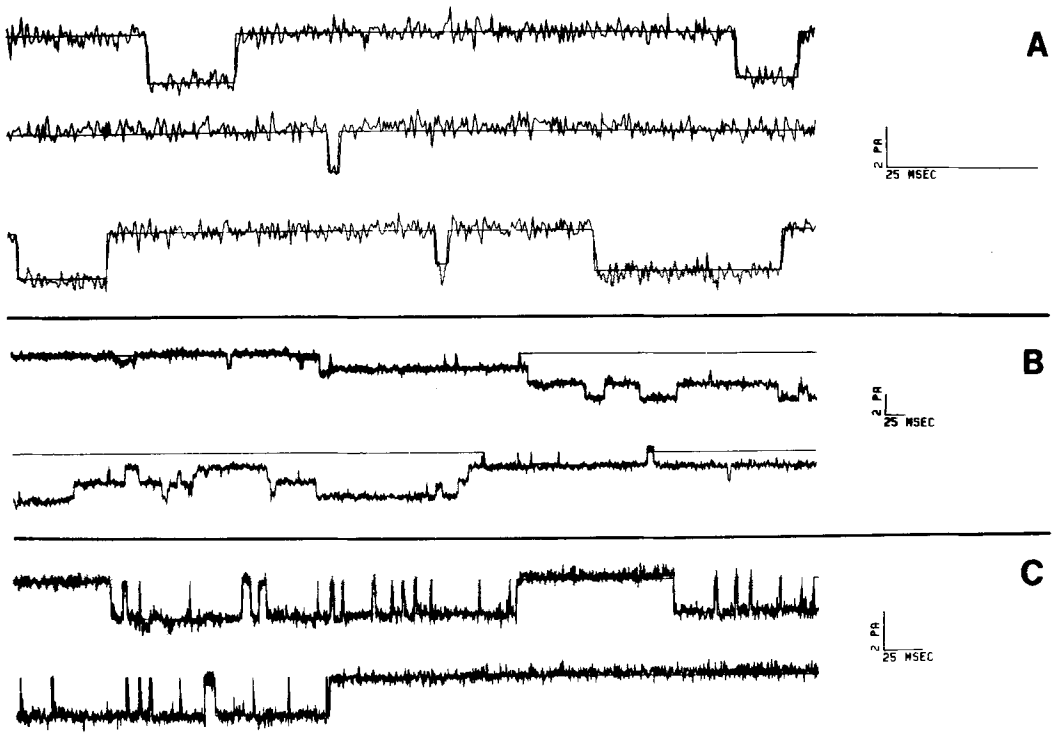

Fig. 1A–C. Different types of unitary current records recorded from chick myotubes in tissue culture. In all panels inward currents are shown downward. The program has identified valid unitary events by drawing straight lines through the events at the unitary amplitude. **(A)** A record of single events recorded in 100 μM carbachol, 18° C, − 80 mV, with a bandwidth of 2 kHz. No bursting is seen in these records. The mean number of events/burst over the whole experiment was 1.3 (DT142460. CA). **(B)** A record containing multiple events, illustrating the need to deal with multiple level events and events crossing record boundaries. This record has little baseline present. Much of the record has been ignored by the program because of the presence of multiple events. (50 μM carbachol at 10° C, − 50 mV, bandwidth 2 kHz, DT142930.CA). **(C)** A record of burst behavior recorded with 80 nM acetylcholine at 10° C. Membrane potential − 116 mV, bandwidth 2 kHz. The burst crosses the record boundary (DT130550.AC)

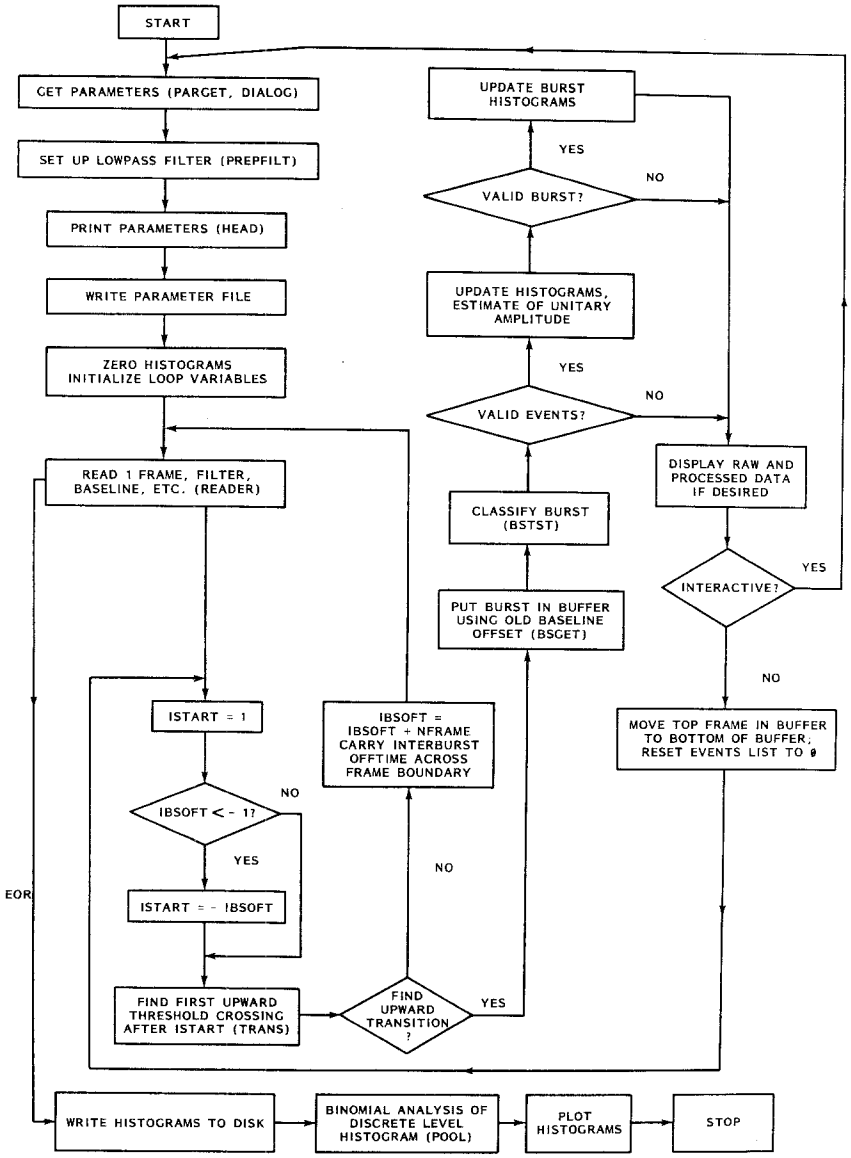

Fig. 2
Flow chart of the main program flow of IPROC. Important subroutine names are shown in parentheses. Some variable names are shown for brevity (Note, EOR stands for end of record). Refer to the text for more details

noise (p-p) levels less than five, and the need to analyze short events in the possible presence of multiple channel activity, the analysis is much more difficult. The algorithm currently employed is diagramatically illustrated in Fig. 2, and described in the following text.

Data is digitized from analog magnetic tape, stored in binary on a 5 MB hard disk, and read in 0.25 to 4 kword blocks. (In what follows "*kword*" indicates thousands of 16 bit computer words.) For consistency in programming all significant deflections are taken as positive. Thus, inward currents, normally recorded as negative going deflections are inverted prior to analysis.

In order to provide for variable filtering of the data, subroutine ILPF uses a finite impulse response (FIR) digital filter to reduce the bandwidth below the normal Nyquist limit applied to data conversion (Peled and Liu 1976). The FIR filter chosen for its simplicity of implementation was modelled after an ideal low pass filter (zero phase, flat frequency response to cutoff, and zero thereafter). For infinitely many terms, such a filter overshoots by 17%, but if the coefficients are Hamming tapered, and the number of coefficients is small, the overshoot can be reduced to a few percent (at the expense of the slope of the cutoff region of the filter). This filter offers the advantage of extremely steep cut-off in the frequency region just above the 3 dB cutoff. If a monotonic step response is desired, other filter functions such as a Gaussian or Bessel could be implemented. These filters, however, have a slow transition from the pass band into the stop band, and for the same cutoff frequency, pass more noise. No effort is made to decimate the data (remove correlated or redundant data points) following filtering, an operation that would increase processing speed proportionately. The filter routine is currently being expanded to permit frequency response corrections for the limited bandwidth of the head stage amplifier (Ostrem and Falconer 1981).

After filtering, the next step is to find the baseline using subroutine BASE (see Fig. 3). This operation has produced a surprising number of complications. Low pass filtering will not produce a reliable baseline, since the presence, duration and location of unitary currents within the record is variable, and filtering will give the mean rather than the zero current baseline. Our simplest and most reliable algorithm fits a

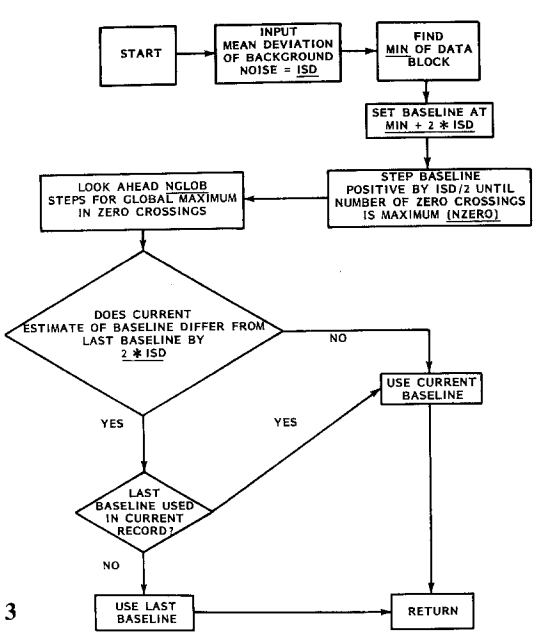

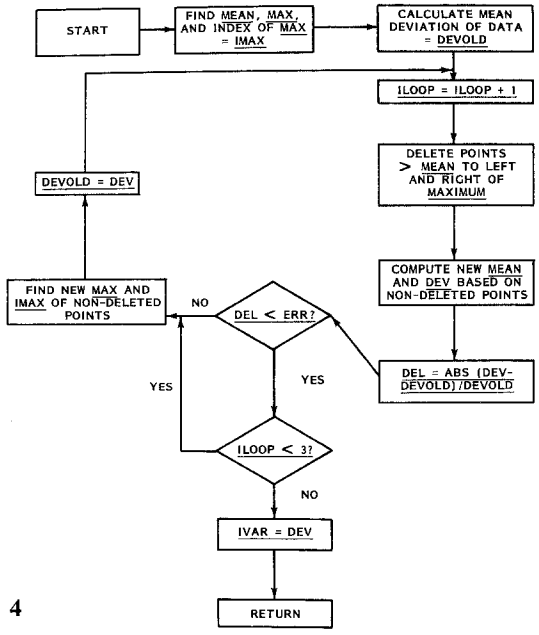

Fig. 3. Flow chart of subroutine BASE which finds the baseline of a block. The reader is refered to the text for further details

Fig. 4. Flow chart of subroutine IVAR which calculates the mean deviation of the background noise. Data is delivered to this routine following baseline correction. The mean deviation of the data and the mean is computed relative to the baseline. The maximum amplitude point is chosen, and all contiguous data points greater than the mean of the data are deleted to the right and left of the maximum. A new mean and mean deviation are computed, and the cycle repeats until the change in deviation is less than some error limit, ERR.

straight baseline which maximizes the number of zero crossings. The algorithm operates as follows. Given that the currents are positive going, the most negative point in the record is chosen and a straight line is passed through it. The line is stepped in the positive direction until a significant maximum in the number of zero crossings occurs. Significance is measured by an arbitrary set of standards which have proved satisfactory in most circumstances: (1) there must be at least ten crossings in a 1 kword record sampled at the Nyquist limit; (2) there are no other maxima within ten steps in the positive direction from any suspected maximum. (Selection of the appropriate step size is discussed below.) Once the baseline is found, it is subtracted from the data record, and the data is returned to the main program for further processing.

We have encountered two types of problem with the algorithm. With very long duration events, such as those shown in Fig. 1 B, i.e., those which occupy most of a record block, the program may locate the baseline at the channel top rather than at true zero. By taking long enough record lengths, this problem may be minimized. Long record lengths, however, lead to another problem. The baseline may not be flat, but sloping or even curved. We have handled the linearly sloping baseline by dividing the record in half, fitting a flat baseline to each half and passing a straight line through the midpoints of each subdivision. This strategy is generally successful, but the subdivision of the record causes problems for currents occupying most of a subdivision by again confusing channel tops for baselines. We have also used higher order spline fits for curved baselines, but these algorithms are even less stable although they can work very well on some data sets.

The problem of confusing the top of the channel with the baseline could be avoided by having a more global view of the data than a 4 kword record permits. This, however, entails baselines with more drift and curvature. Our experience is that the straight baseline, although not fitting as well as might be desired, rarely makes drastic errors, and evaluation of the baseline, record by record, rather than bringing in data from previous records, prevents errors from propagating. The record length should be chosen to be short enough so that the baseline slope errors are minimized, and there is enough "*closed time*" to clearly define the baseline (at least 10 % of the record should be closed time, using Nyquist sampling frequencies). We have introduced one simple trap to extend the effective record length. We test whether the current baseline is within, typically, two standard deviations of the background noise (see below) from the previous baseline. If the current baseline estimate is deviates by more than two standard deviations from the previous baseline, the previous baseline is used for the current record. This process is only carried on for a limited number of records to prevent propagating errors.

In the process of finding the baseline, the step size for advancing the putative baseline could be increased if the background noise level were known. For the first record, where we do not yet have a measure of the background noise, we have conservatively taken the step size to be the minimum value, i.e. one bit (typically 2.5 fA). The step size would ideally be computed as a fraction of the standard deviation of the background noise. Unfortunately, if the channels have not yet been identified, there is no way to measure the baseline noise. An estimate of the standard deviation could be input by the operator, but we found that unnecessary since the algorithm used in subroutine IVAR produces good estimates of background noise.

Once the baseline has been found for the first record, the standard deviation of the background noise is computed by

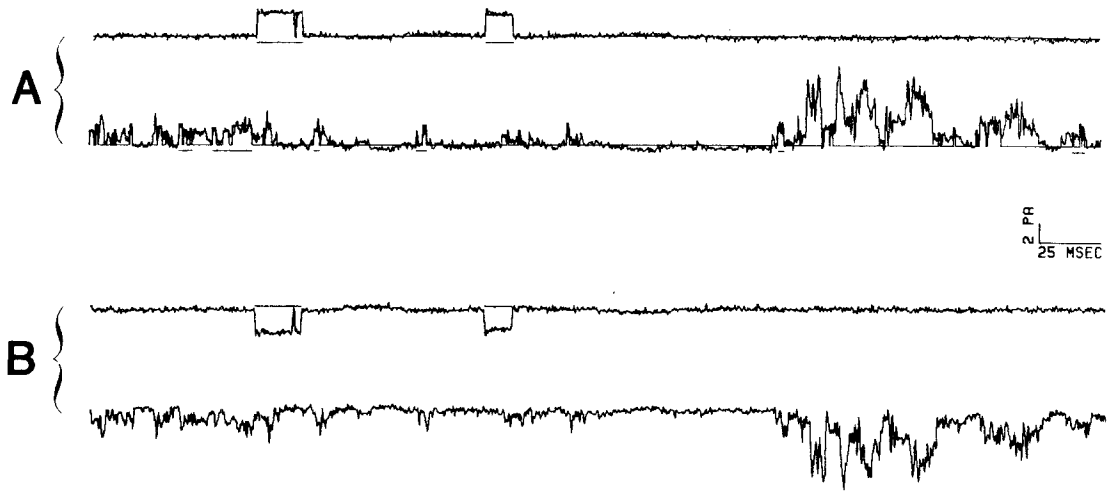

Fig. 5A and B. Comparison of algorithm performance in the presence of noise. (**A**) A record containing both single channel data (top line) and significant amounts of noise (lower line) was analysed by a simple threshold-crossing algorithm, with extensive false identification. In this figure, inward currents are shown upward going, baseline and channel detection are drawn in, and bursts are underlined. (**B**) Same record analysed by the burst variance algorithm. Amplitude thresholds were identical to those used by the threshold-crossing algorithm ($\pm$ 5 s.d.). Data (inward currents) are shown downward going, identified bursts are underlined, but baseline and channel detection are not plotted

subroutine IVAR. The essential algorithm is shown in Fig. 4. The routine seeks out the largest (most positive) events and deletes them from the record. The deletion is discontinued when deletion of an event results in a negligible change in the remaining record variance, as expected for deletions of stationary noise centered about the baseline. We have found it unnecessary to calculate the standard deviation for each record, since it varies little, and consequently we usually use the value obtained from the first record. The first record should be of relatively low duty cycle to reduce errors.

The next step is to detect valid transitions in the current. The main feature of a transition is that there is an abrupt change in amplitude. Since transitions between multiple, independent, channels yield events which are ambiguous in duration, we have chosen to ignore the kinetics of multiple level transitions. For transitions between zero and one event, we have found that a simple threshold crossing criterion is the most reliable. The threshold is set symmetrically at one half the unitary current amplitude, so that on the average, there is no net time skew from the low pass filtering. The initial estimate of unitary current amplitude is input by the operator. The subroutine named TRANS detects transitions. It searches a given record starting at the location of the last transition, and returns with the location and polarity of the next valid transition.

We have tried identification of unitary currents using 2 different algorithms. In our initial efforts, the amplitude of individual events was determined by a maximal zero crossing technique similar to that used for baseline correction. Channel acceptance was based only on the comparison of this amplitude to upper and lower adjustable thresholds symmetrically distributed about the estimated unitary current amplitude. This method offered ease of implemention, fast execution, and worked reliably for well characterized data such as shown in Fig. 1. Unfortunately, not all data from single channel preparations has such ideal characteristics. A common non-ideality can be seen on the second line of Fig. 5A. This artifactual signal, which we currently believe to be a membrane-glass seal breakdown, is often observed at high electrode potentials. To avoid accidentally accepting any of these events we developed a second algorithm which compares the raw data to a binary switching model.

Visual examination of Fig. 5 reveals that there are two essential differences between the desired single channel signal and the artifactual signal. First, all events in the desired signal are reasonable approximations to a rectangular pulse, which is not necessarily true for the artifactual signal. Second, and most important, the desired signal consists of a burst of equal amplitude events (the unitary current), while the events comprising the artifactual signal are of essentially random amplitude. These observations suggest that we could get better artifact rejection by analysis of the entire burst as a single entity.

This concept is implemented in subroutine BVAR, which determines the amplitude for a burst of single channel events and the rms (root mean square) deviation of the data from the model, which consists of a string of equal amplitude pulses. The amplitude of the events in a burst is obtained by computing the mean of those points above the detection threshold. The value returned as the rms deviation from the model is divided by the already computed baseline standard deviation. The entire burst of events is accepted if the scaled deviation is less than an adjustable threshold and the amplitude is within acceptable limits as discussed above. Due to finite amplifier settling times, the amplitude of a burst calculated from the mean may be significantly below the true amplitude depending on the channel open time. This amplitude, however, is only used for detection. The true channel amplitude is estimated as a mean of open channel currents measured after allowing for settling.

Once a potential burst has been identified, a running estimate of the channel amplitude is updated. The estimate is a weighted average of the amplitude of all unitary events, and since the threshold for detection is tied to the current amplitude, this provides an adaptive feature that simplifies program setup and increases reliability. The weighting used in calculating the best estimate of channel amplitude is the length of the burst divided by the standard deviation from the model. Long, steady amplitude events contribute more than short or unsteady events.

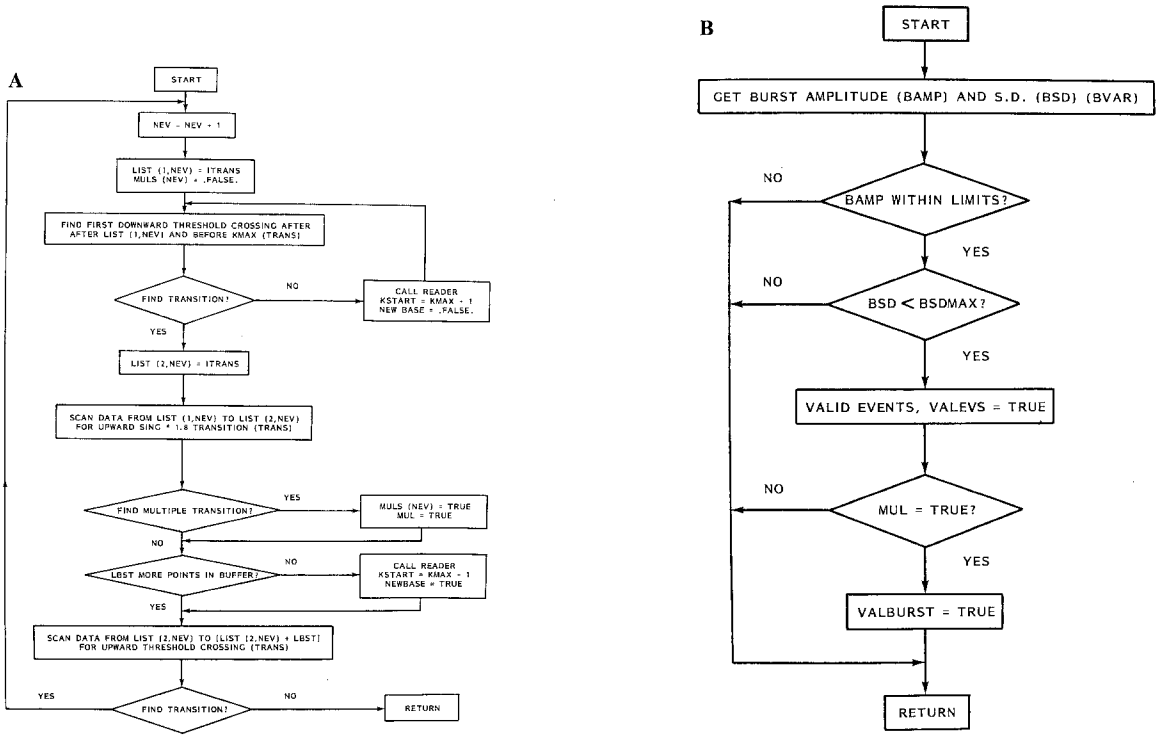

Fig. 6A and B. The analysis of burst behavior. (**A**) Flow chart of subroutine BSGET which is called at the first opening transition of a burst. The routine moves a segment of data containing a putative burst into an analysis buffer, and constructs a list of opening and closing transition times. NEV is the number of events in the apparent burst. LIST (1, NEV) is the location of the opening transition for event number NEV. LIST (2, NEV) holds the location of the closing transition. Each open period is checked for multiple activations, and a flag (MULS[NEV]) is set if the upper amplitude threshold is exceeded. Additionally, flag MUL indicates the presence of a multiple activation anywhere in the burst. LBST is the intraburst closed time threshold (in sample intervals). (**B**) Flow chart of subroutine BSTST. BSTST is called after the apparent burst has been stored in the analysis buffer by BSGET. Subroutine BVAR returns the mean amplitude of the open periods and the mean deviation of the burst from the model (see text). Acceptance is based upon comparison of these values to adjustable thresholds. The amplitude acceptance thresholds are continuously updated as a weighted mean of the amplitude of accepted events, with weight w = (burst length) * SD (baseline)/SD (burst) (SD = standard deviation)

The measurement of channel closed times is constrained to measure the times between events of unit, or larger, amplitude. Not uncommonly, a noise spike, or perhaps a channel located under the pipette rim will cross the amplitude threshold, but when tested by BVAR will have an amplitude below the lower cutoff for valid single channel events. These low amplitude events do not terminate a closed time. Closed times are defined as times between unitary or larger currents. Thus, a closed time can only be measured after a unitary or multiple event is terminated and validated. Closed times are entered into the three closed time histograms.

Bursts which cross record boundaries necessitate elaborate bookeeping. We have simplified the analysis by going from a record-oriented to an event-oriented system using a dynamically allocated buffer to hold each burst (subroutine BSGET). The flow of control in subroutine BSGET is shown in Fig. 6A. The buffer for burst storage is divided into an integral number of frames. When BSGET is called, the beginning of the putative burst is placed in the first frame of the buffer. BSGET searches from the opening transition to the end of the frame for a closed period longer than the threshold selected to terminate bursts. If a long enough closed time is not found, the next frame of the record is read, stored in the next frame of the buffer, and searched for a terminating closed period, and so on. During this process, the program constructs a list of opening and closing transition locations (measured in sample periods). Each open period is tested against an upper amplitude threshold to assure that it is not the result of multiple channel activations. Any interval which fails this test is flagged as a multiple event.

After acceptance of the string of unitary events by BVAR, the event string is validated by BSTST (Fig. 6B). Bursts are defined as a series of unit amplitude events separated by closed periods less than some threshold duration. This series must both begin and end with closed times greater than the burst defining threshold. The intrusion of multiple events invalidates the measurement of burst length, but will not interfere with the measurement of intraburst kinetics. Isolated single events are included in the accounting as bursts consisting of one event with a length equal to the open time.

An important part of the analysis is a constant comparison of the raw data and the recognition program output. As shown in Fig. 1, both the real and idealized data can be plotted superimposed, or alternatively, only identified bursts may be underlined (cf. Fig. 5B). Switches set in the interactive portion of the program permit the display of either all records, only those records containing unitary events, or none at all. The latter option is often used when detection criteria are remaining fixed, but classification criteria, such as histogram resolution or threshold need to be changed.

At the end of the complete data file, the mean and the total number of counts of all histograms are printed out in order to provide a quick summary of the data and provide starting values for subsequent curve fitting.

In order to estimate the number of channels active in the patch so that rate constants for independent channel opening

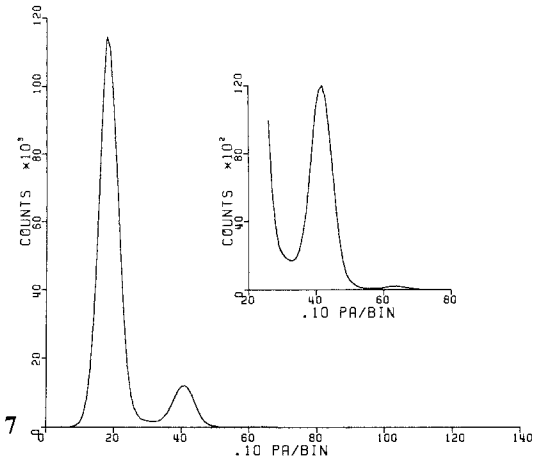

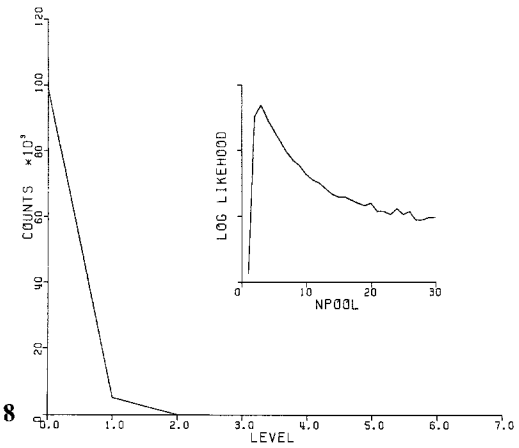

Fig. 7. Total amplitude histogram. Bin 20 has been used as the origin to permit baseline noise to be fully displayed. The three peaks shown correspond to 0, 1, and 2 channels open. Symmetry of the single channel peak at bin 40 indicates that few channels are in the sealing region. (Record DT311150.AC)

Fig. 8. Discrete level histogram and the log of the binomial likelihood function. We observed the overlap of two unitary events in the record. The likelihood function peaks with a pool size of three. The noise in the likelihood function at large pool sizes is due to single precision numerical errors, and points out the futility of estimating large pool sizes with a low probability of occurrence

can be scaled properly, the discrete level histogram is subjected to a maximum likelihood analysis assuming a binomial distribution of independent channel activity (Korn et al. 1981). Subroutine POOL calculates the likelihood that the number of observations at each level came from a binomial distribution of N channels each having a probability p of being open at any one time. The likelihood function, L, is given by,

$$L = \prod_{i=0}^{N} P(i, N, p)^{x_i}$$

where $P(i, N, p)$ represents the binomial probability of observing level i with N channels in the pool, each with a probability p of being open, and x_i represents the number of observations at level i. The logarithm of L is a sum of terms rather than a product and somewhat simpler to manage. L is analytically maximized with respect to p by setting to zero the partial derivative of L with respect to p (see Korn et al. 1981). Thus, $p = \langle i \rangle / N$. The likelihood function is maximized with respect to N by calculating the $\log(L)$ for each N from the observed maximum level up to 80 channels and searching for the maximum. The parameters N and p are printed out and the log of the likelihood of function is plotted to further judge validity of the maximum (cf. Fig. 8).

IPROC has been built with a number of convenience features that simplify setup and repetitive, or batch processing. For setting up the run time parameters, an interactive dialog (DIALOG) presents default parameters for potential modification. The output is generally directed to a Tektronix graphics terminal to display the results of identification. Following the presentation of each data record, the operator has the option of altering any of the parameters, typically bandwidth, estimates of the single channel amplitude, or error limits. Following each adjustment, the parameters are written out to the parameter file for that data record. When satisfactory parameters have been fixed, the file may be rewound and analyzed, or the program terminated for later batch processing. In batch mode, the data file is specified following the call to IPROC, and the existing parameter file is used for setting run time parameters. A local edit mode switch can be set after the file name to permit editing an existing parameters file. This feature is useful for making minor modifications in a preexisting parameter file.

Results

We have found that with data having a signal-to-noise ratio of greater than five to one and a steady baseline, the reliability of the program is virtually as good as a trained observer. The program is constructed in a conservative manner so that false positive identifications are rare. That is, virtually no multiple events enter the histograms, but some single channel events are missed, usually due to excessively short duration or low amplitude. The usual sort of tradeoffs of time resolution and reliability operate here. If very short events are to be identified (i.e., those events of only one point duration), then the number of false positive triggering events will increase. Similarly, excessive low pass filtering will increase the loss of short events. The program has not been optimized for running speed, although integer operations have been used wherever possible to increase speed and reduce storage requirements. The current analysis time is about three seconds per 2 kword block of data, although the actual runtime may be limited by the plotting output speed. Typical output histograms of the runtime analysis are shown in Figs. 7–10.

Figure 7 shows a typical total amplitude histogram. In the example shown, the width of the peak corresponding to single channels, approximately bin 40, has about the same width as the background noise peak at bin 20 (note, the origin of the histogram has been shifted to bin 20 to encompass the positive and negative components of the background noise). Some amplitude histograms show pronounced skewing toward the low amplitudes indicative of currents arising from channels within the sealing region between the membrane and the pipette (Sachs and Barkakati 1980) or speed limitations. We have seen cases, notably with suberyldicholine at very positive membrane potentials, where the channels appear to be

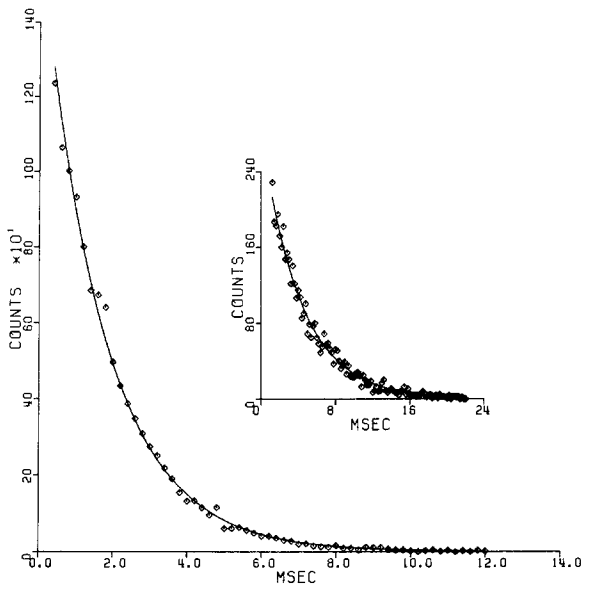

Fig. 9. Open times and burst length for a 70 pS potassium selective channel in tissue cultured chick muscle. Data was curve fit with a single exponential function (solid lines) by a non-linear regression program. (Data courtesy A. Auerbach.) Fit parameters were for the open times: 12,800 ± 300 counts and 1.68 ± 0.040 ms and for the burst length 5,830 ± 190 counts and 4.21 ± 0.15 ms. Error limits are for 90% confidence based on "*support plane*" estimates which takes into account cross correlations between parameters (Marquardt 1964)

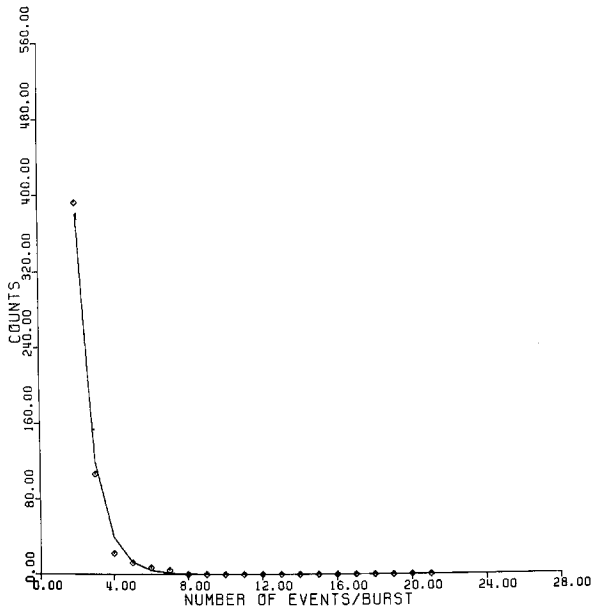

Fig. 10. Histogram of the number of events per burst, with a burst defined as a series of unitary events separated by closed times less that 5 ms. The smooth curve represents a weighted nonlinear regression to a geometric probability law. $P(n) = P^{n-1}(1 - P)$. For these data, P is given by 0.67. (100 μM carbachol, 10° C)

flickering at frequencies beyond the bandwidth of the system, so that there is a wide variation in recorded channel amplitude.

Figure 8 shows the discrete level histogram and the log of the likelihood function. The maximum observed level was 2, and the likelihood function peaked with pool size of 3 and a probability of 0.019 per channel of being open. Most likelihood estimates have been in the range of 1 to 7 channels. The meaning of this number is not immediately clear since all measurements of channel density with labelled bungarotoxin give densities of hundreds to thousands/square micron (Sytkowski et al. 1973; Elson 1979). With pipette openings of a square micron or so, there should be many more than 3 channels present. Several possible explanations are that the densities observed with labelled toxin do not reveal micron level inhomogeneities, that all toxin binding sites are not acitve channels, or that a rapid component of desensitization removes channels from the observed pool.

Data on transition rates is summarized in histograms of the open and closed times, and the burst lengths. These histograms are generally multiple exponentials, although some channels appear to produce single exponentials. One of these exceptionally simple distributions is found in a potassium selective channel (ca. 70 pS) in cultured chick muscle (data courtesy A. Auerbach). Figure 9 shows the distribution of open times and burst lengths (inset) for a 2 megaword record containing about 5,000 bursts and 10,000 open times. The data in each case have been fit with a single exponential function using a nonlinear regression program weighted with Poisson statistics for counting processes and corrected for finite binwidths. For the open times (outer panel), the time constant was 1.6 ± 0.040 ms and for the burst length, the time constant was 4.21 ± 0.15 ms. The error limits stated correspond to 90% support plane confidence limits, which estimate the confidence in a parameter regardless of the values of the others (Marquardt 1964). The details of the kinetics are not important here, but the parameters are quoted to emphasize the program's ability to analyse large amounts of data and thereby increase resolution.

Finally, Fig. 10 shows the number of events per burst for a 100 μM carbachol activated nicotinic channel with a regression to a geometric probability law (Colquhoun and Hawkes 1977). The geometric probability function comes from considering the sequential scheme of Eq. (1). When a channel closes it may either reopen or the agonist may dissociate. A burst with n openings comes about by not dissociating for $n - 1$ times and undergoing dissociation the last time. Thus, $p(n) = p^{n-1}(1 - p)$, where $1 - p$ is the probability of dissociation.

Discussion

The automated nature of the analysis permits the accumulation of large numbers of events, thus increasing the sensitivity for testing of kinetic models. In addition the automated system lowers the labor cost for such time consuming characterizations as cross correlations, conditional probability estimates and power spectra of selected portions of the record (such as open and closed channels). For example, we have made cross correlations of amplitude and duration of flickers that occur during bursts and find that some apparent closing transitions go to a non-zero conductance level (Auerbach and Sachs, in preparation).

Given the fact that the data are random with regards to amplitude, duration and duty cycle, we found it useful to let the program be somewhat adaptive in its identification procedures. This reduces the need for operator intervention. The main adaptive feature is the adjustment of the threshold for event identification based upon a weighted average of all successfully detected unitary events. The procedure has

proven to be robust and reliable. The program could be made more adaptive by using multiple pass procedures, such as one pass to examine the amplitude histogram for peaks and setting the thresholds accordingly, but in the interests of speed and simplicity, we have found that operator input of a few starting parameters is simple and straightforward.

Although we have tried a variety of complex identification algorithms, we keep coming back to the simplest versions, since the probability of serious errors is minimized. Given the wide variability, particularly in amplitude and duty cycle, complex algorithms which may be successful on certain data sets are unreliable on others. The simplest algorithms remain the most reliable and least in need of adaptation to particular data sets.

We have examined some of the dynamic limitations of the analysis scheme, particularly with regard to detecting short closed times and the effect of missing short closed times on the interpretation of open times. The Bessel response filters used for anti-aliasing and the digital filter in the program, provide an undistorted record of event widths if the events pass a little beyond half of the unitary current amplitude, which is why we have chosen to use that point as the threshold for detection. Since the program measures the event duration at half maximum, we examined the full width at half maximum (FWHM) of pulses sent through both the four pole hardware anti-aliasing filter and the software FIR filter. The FWHM is essentially independent of bandwidth until the product of bandwidth (BW) and pulse duration (Tp) is less than 0.4. When BW * Tp = 0.31 the output pulse just reaches half amplitude and the FWHM = 0. Thus, in the absence of noise, the system will miss all events shorter than 0.31 of the inverse analog bandwidth. Most of the longer events will be correctly measured for duration. In a noise-free record the probability of detection is a steep function of the duration-bandwidth product, varying from essentially one with a product greater than one, to 0.5 with a product of 0.33, to 0 at a product of 0.31. (There is a probability even in the noise free system since the sampling clock is uncorrelated with the data and a sample may or may not be taken during the time that the pulse is beyond threshold.) Thus, the shortest events will obviously not be digitized with full resolution, but those longer than about 0.3 of the inverse bandwidth will be recorded and placed in the proper bin of the histogram. The presence of noise will reduce the steepness of the detection curve so that events which are too short to be detected in a noise free system may yet be recorded, and similarly, some events which are large enough to be recorded in a noise free system will be missed because of opposing noise. As a rough guideline, however, we expect to detect all events longer than 0.5 of the inverse bandwidth and miss all events shorter than 0.3 of the inverse bandwidth.

Missing short closed times will tend to make the open times appear longer, and decrease the number of events per burst. We can estimate the magnitude of the effect by considering the distribution of open times when the system is unable to detect events shorter than some dead time T_d (F. Sachs and R. S. Spangler, unpublished). Assuming an exponential distribution of true open and closed times with respective time constants T_o and T_c, and absolute dead time T_d and assuming, for simplicity, that $T_c \ll T_o$, one can show that the open times will be distributed exponentially, but with a time constant given by:

$$T_o\,(\text{observed}) = T_o \times \exp(T_d/T_c).$$

Suppose we take T_d to be 0.3 of the analog bandwidth, which for the sake of illustration, we take to be 1 kHz, $T_c = 1$ ms, then T_o (observed) $= 1.8\,T_o$. If we use a more conservative estimate of the dead time, i.e., 0.5 of the inverse bandwidth, then T_o (observed) $= 2.7\,T_o$. We have observed mean closed times in the range of 0.2–2 ms. To reduce the error in the open time (also the number of events/burst) to 20% with a T_c of 0.3 ms, requires a deadtime of less than 0.05 ms, or a bandwidth of between 6 and 10 kHz depending upon the dead time assumptions. This bandwidth is difficult to achieve with good signal to noise ratios. Readers are referred to Hamill et al. (1981), for a discussion of head stage circuit limitations.

A second type of bias in durations comes from eliminating multiple events from the burst duration accounting. The error may estimated as follows.

Multiple events arise from the overlap of bursts. Let the mean burst duration be called T_B, and $k_B = 1/T_B$. Now, if the probability of observing double events is small, the occurrence may be treated as a Poisson process with mean rate,

$$k_2 = P_2/T_B$$

where P_2 is the probability of having two or more channels open at the same time. The probability of no multiple events for a time t is given by,

$$P(\text{no multiples in } t) = \exp(-k_2 t).$$

The probability density for a burst of duration t containing no multiple events is then (assuming an exponential distribution of burst lengths),

$$pdf(t) = k_B \exp(-k_B t) * \exp(-k_2 t).$$

The mean duration of the analysed bursts, that is, those of duration t and containing no multiple events, is given by

$$\langle t_B \rangle = \int_0^\infty t\,pdf(t)\,dt = T_B/(1+P_2)^2.$$

For $P_2 = 0.01$ (cf. Fig. 8), $\langle t_B \rangle = 0.98\,T_B$. The error is small, and if necessary, can be used to recursively correct the data. The same type of error applies to the measurement of open times in general.

The program, as it currently exists, requires a large amount of data storage, since there is no data compression. There are three possible approaches to dealing with storage requirements, and possibly increasing speed. A brute force approach is to simply increase storage capability. Digital magnetic tape can provide efficient storage of sequential data so that a 90 Mbyte tape can store up to 30 min of data digitized at 20 kHz. This method preserves all the data, although possibly requiring substantial time for analysis. We are adapting the program to deal with data stored this way in a two pass system. In one pass, events are identified and their locations are stored in a list on disk. A second pass using a related program looks only at the identified events and performs the requisite analysis on those segments. Thus, data may be filtered to a low bandwidth for reliable identification of channels, and the post-processor can reread those events with a high (possibly software corrected) bandwidth. Since many channels appear to have multiple open states (cf. Auerbach and Sachs, in preparation), understanding channel kinetics requires measuring amplitudes as well as open-closed transition probabilities.

A second scheme to reduce digital storage requirements is to process the data on-line either from analog tape or from direct recording (i.e. sodium channel records). The program

as it currently exists can process on-line data at about 250 samples/s. A few modifications to the code would be necessary to allow multitasking of data conversion and analysis. Additionally, further effort could be made to speed up the most time consuming parts of the program, the software filter and the baseline correction.

A third scheme to reduce storage requirements would delete irrelevant data prior to storage. An abbreviated baseline and threshold crossing routine (hardware or software) could trigger storage of only record blocks with active events. This method obviously has biases, but can produce enormous space savings in low duty cycle records.

With regard to generalizations, the main limitations of the program are the necessity for events to be unidirectional and moderately uniform in amplitude. If several species of channel with different conductances are present, such as junctional and extrajunctional nicotinic receptors (Hamill et al. 1981), each channel type can be separately identified and cataloged in multiple passes by setting narrow limits on the amplitude acceptance criteria. The program could be modified to keep track of several classes of unitary amplitudes in the same pass. To analyze data containing bipolar currents, the duty cycle would have to be relatively low so that the baseline could be reliably determined from a global maximum in the number of zero crossings of a presumptive baseline.

Acknowledgements. This work was supported by grant NS13194. The authors would like to thank Dr. Deborah Nelson, Mr. Howard Jachter, Mr. Victor Demjanenko and M. N. Prakash for advice and assistance in developing this program. The programs are written in FORTRAN V and were run on a Data General S-130 computer. Copies of the source programs (50 pp) are available on letterhead request with postage for 14 oz. plus a $3.00 copying fee. Alternately, floppy or cartridge disk copies will be supplied in Data General format upon receipt of a formatted disk with proper mailing containers and postage.

References

Colquhoun D, Hawkes AG (1981) On the stochastic properties of single ion channels. Proc Roy Soc B211:205–235

Colquhoun D, Hawkes AG (1977) Relaxation and fluctuations of membrane currents that flow through drug-operated channels. Proc Roy Soc 199:231–262

Elson HF (1979) A possible modulation of acetylcholine receptors of embryonic chick muscle cells by α-bungarotoxin. J Supramol Struct 10:39–50

Hamill OP, Marty A, Neher E, Sakmann B, Sigworth F (1981) Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches. Pflügers Arch 391 (2):85–100

Korn H, Triller A, Mallet A, Faber DS (1981) Fluctuating responses at a central synapse: *n* of binomial fit predicts number of stained presynaptic boutons. Science 213:898–901

Marquardt DW (1964) Least squares estimation of nonlinear parameters. IBM Share Library 3094:15–20

Neher E (1981) Unit conductance studies in biological membranes. In: Baker PF (ed) Techniques in cellular physiology. Elsevier/North Holland, Amsterdam, pp P121/1–P121/16

Neher E, Sakmann B (1976) Single channel currents recorded from membrane of denervated frog muscle fibers. Nature 260:799–802

Neher E, Steinbach JH (1978) Local anesthetics transiently block currents through single acetylcholine receptor channels. J Physiol (Lond) 277:153–176

Nelson D, Sachs F (1979) Single ionic channel observed in tissue cultured muscle. Nature 282:861–863

Nelson D, Sachs F (1981) Ethanol decreases dissociation of agonist from nicotinic channels. Biophys J 33:121a

Nelson D, Sachs F (1982) Agonist and channel kinetics of the nicotinic acetylcholine receptor. Biophys J 37(2):321a

Ostrem J, Falconer D (1981) A differential operator technique for restoring degraded signals and images. IEEE trans. Pattern Anal Mach Intell PAMI-3(3):278–284

Patlak J, Horn R (1982) The effect of *n*-bromoacetamide on single sodium channel currents in excised membrane patches. J Gen Physiol 79(3):333–351

Peled A, Liu B (1974) Digital signal processing. Wiley and Sons, New York

Sachs F, Barkakati N (1980) The amplitude distribution of single channel currents can yield dose response information. Biophys Soc Abstr 2425:174

Sytkowski AJ, Vogel Z, Nirenberg M (1973) Development of acetylcholine receptor clusters on cultured muscle cells. Proc Nat Acad 70:270–274

Received November 23, 1981/Accepted August 19, 1982

Note in Proof

The program is now available on magnetic tape, 1,600–3,200 bpi, in Data General dump format. The reference (Auerbach and Sachs, in preparation) is in press: Auerbach A, Sachs F: Flickering to a subconductance state of the nicotinic ion channel. Biophysical J