

Asynchronous Sampling via Time Encoding Strategies

Marek Hilton

November 22, 2024

Supervised by Prof. Pier Luigi Dragotti

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Electrical and
Electronic Engineering of Imperial College London and the Diploma of Imperial College London

Abstract

Time encoding is an alternative approach to sampling that encodes a signal in the timing of samples and not their amplitudes. This departure from conventional approaches has promised primarily to provide significant increases in efficiency by emulating some characteristics of the highly efficient biological neural networks that they are based on.

Time encoding generates inherently asynchronous samples and in this thesis I look at how such encodings can be used to more efficiently sample and reconstruct signals which themselves exhibit significant sporadic behaviour. By sporadic behaviour, I refer to signals that are comprised of periods of activity interspersed with periods of quiescence. Such signals are typically oversampled by existing time encoding methods which generate samples at a rate similar to conventional sampling techniques.

To this end, I focus largely on signals with finite rate of innovation as a model for sporadic signals. These include bursts of pulses and piecewise constant signals. I extend existing results regarding time encoding and reconstruction of signals with finite rate of innovation to novel sampling kernels, namely, hyperbolic kernels and the biologically inspired α -synaptic function.

Time encoding is inherently a non-uniform sampling technique and consequently, the majority of reconstruction algorithms, including our aforementioned contributions, proceed by applying standard non-uniform reconstruction techniques. In this sense, the signal driven timing of samples is not exploited to the fullest extent in standard approaches to time encoding.

Taking an alternative perspective, I present a new time encoding architecture called a time dilating integrate-and-fire encoder (TDIF) that uses dilation of time measurements to affect a form of filtering of an equivalent signal in *amplitude*-domain. I show that time encodings derived from such a sampler can be treated as uniform samples and thus, the standard techniques of uniform sampling can be applied. Since the local sampling rate of such a system is determined solely by the local power of the signal, it is ideally suited to sampling highly sporadic signals.

Finally I present some multi-encoder networks that complement the complement the reconstruction strategies presented in this thesis.

Acknowledgements

I would like to thank my supervisor Prof. Pier Luigi Dragotti for his diligent supervision over many years. I would also like to thank all my colleagues in the research group under his supervision and the wider communications and signal processing research group. I will certainly miss our many interesting conversations on a wide range of topics. In particular, I am grateful to Renke for our discussions on many ideas, philosophy, and the trials and tribulations of academic work.

To my family, thank you for supporting and encouraging me throughout. To Susanna, thank you for always being there for me and providing your invaluable counsel.

I would also like to thank my examiners Dr Chris Gilliam and Dr Wei Dai for reading this thesis, their insights, and interesting discussion on the topics written within.

Finally, I would like to thank the makers of Pandoc for freeing me from the shackles of LaTeX.

Statement of Originality

I declare that this thesis is my own work and that any content that is not my own is clearly stated as such and that appropriate references and attributions are given.

License

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Contents

1	Introduction and outline	11
1.1	Achievements and Innovative results	15
1.2	List of publications	16
2	Time encoding preliminaries	18
2.1	Architectures for sampling in amplitude	18
2.2	Reconstruction of bandlimited and shift-invariant signals	27
2.3	Reconstruction of signals with finite rate of innovation	32
2.4	Summary	44
3	Time encoding and recovery of signals with finite rate of innovation	46
3.1	The prototype IF-TEM	47
3.2	The α -synaptic function	48
3.3	Recovery guarantees with the α -synaptic kernel	55
3.4	Piecewise constant and polynomial signals by sampling the derivative	64
3.5	Higher order α -synaptic kernel	66
3.6	Signals with structure - recovery of power signals	68
3.7	Reconstruction with hyperbolic and rational kernels	73
3.8	Recovery from differential samples	85
3.9	Summary	88

4	Amplitude-domain and the time dilating integrate-and-fire encoder (TDIF)	92
4.1	Time encoding is a form of uniform sampling	93
4.2	Bandlimited and piecewise constant signals in amplitude-domain	98
4.3	Sampling kernel shaping via time measurement dilation	102
4.4	Interleaved time-dilating integrate-and-fire samplers	108
4.5	Piecewise-constant signals	112
4.6	Summary	120
5	Networks of time encoding machines	122
5.1	Inhibitory connections	123
5.2	Differential pairs	125
5.3	Amplitude domain synchronisation	128
5.4	Summary	130
6	Conclusion	132
6.1	Future directions	133
A	Innovation amplitudes & phase without least-squares	137
	References	139

List of Figures

2.1	A typical reference crossing encoder.	20
2.2	Example time encodings generated by a crossing encoder	21
2.3	A Lazar-style integrate-and-fire time encoding machine in the form of an asyn- chronous delta-sigma modulator (ADSM).	23
2.4	A typical integrate-and-fire time encoding machine	24
2.5	The intermediate signals of an IF-TEM	26
2.6	An example time encoding and reconstruction of bandlimited signal ($\Omega = 5 \text{ rad s}^{-1}$) using the iterative reconstruction algorithm of Eq. (2.9). Intermediate results for $k =$ $0, 2, 4$ are given.	29
2.7	Sporadic signal comprised of two of bursts of three Diracs	35
2.8	A filtered piecewise constant signal equivalent to the integral of Fig. 2.7.	37
2.9	Reproduction of three exponentials using an E-Spline which can reproduce the expo- nents $\alpha = -0.5, 0, 1$	39
2.10	Examples of non-uniform reproduction on the interval $[2.4, 3]$ from non-uniformly spaced E-splines that can reproduce exponents $\alpha = -0.5, 0, 1$	42
3.1	Two alternative depictions of the prototype integrate-and-fire time encoding machine.	47
3.2	The α -synaptic function with $a = 1$	49
3.3	Reproduction of the exponential e^{-t} using three shifts of the integral of the α -synaptic function, $\psi(t)$	53

3.4	An FRI signal which consists of pulses with an α -synaptic response ($\alpha = 1$). The Dirac magnitudes have been scaled by 0.5 in the plot for readability.	55
3.5	A two pulse input signal, $x(t)$, filtered with an α -synaptic kernel ($\alpha = 1$) and integrated to get $f(t)$. The corresponding quantity D is denoted with a red brace.	56
3.6	A degenerate case of detecting the level set.	58
3.7	Comparison of the constraint on ratio $x_{\min}/x_{\max}$ imposed for different numbers of Diracs	62
3.8	Selection of valid samples from an IF-TEM with α -synaptic activation	64
3.9	The derivative of the α -synaptic function $\alpha = 1$	65
3.10	Reproduction of two different exponentials with the integral of the same second order α -synaptic function, $t(e^{-\frac{3t}{2}} + e^{-3t})$. In the plot, $\frac{d}{dt}\psi(t) = \varphi(t)$	67
3.11	The hyperbolic secant and its derivative.	74
3.12	Example time encoding and reconstruction of a two Dirac burst using the hyperbolic secant kernel	78
3.13	Reconstruction of a piecewise constant signal from output samples generated by integrate-and-fire method. $\hat{x}(t)$ is the estimated signal.	82
3.14	The hyperbolic secant squared, $\text{sech}^2(x)$, and its integral $\tanh(x) + 1$	83
3.15	The spikes generated by multiple channels of differential samplers (aggregated) and the resulting reconstruction	89
4.1	Time domain signal, $h(t)$, and its integral, $f(t)$, plotted alongside the amplitude-domain signal, $g(c)$	94
4.2	Illustration of the time span property.	97
4.3	Sampling and reconstruction of a signal with a bandlimited and periodic representation in amplitude-domain.	100
4.4	Generalised models for conventional, time encoding, and time dilated sampling. . . .	103
4.5	An example of a time-dilated signal	107
4.6	Multi-channel sampler with 3 channels.	109
4.7	Integral and amplitude-domain signals after dilation.	110

4.8	A linear spline dilating function, $\varphi(x)$, supported on $[0, 0.3]$. and its periodic extension $\beta(x)$	111
4.9	An example of a non-negative piecewise constant time-domain signal and its resulting amplitude-domain signal	117
5.1	Types of inhibition.	124
5.2	A simple differential pair consisting of two IF-TEMs	126
5.3	The integrals of a differential pair with global reset against time.	126
5.4	Two ring synchronisers with single and triple feedback	128
5.5	Integrals of constituent IF-TEMs in ring synchronisers	129

List of Tables

1	Common abbreviations.	9
2	Common terminology.	9
3	Common mathematical notation.	10
3.1	Comparison of reconstruction from cumulative samples or differential samples	86
3.2	Comparison of FRI sampling kernels for time encoding	90

Common abbreviations, terminology, and mathematical symbols

Table 1: Common abbreviations.

TEM	Time Encoding Machine
C-TEM	Reference Crossing Time Encoding Machine
IF-TEM	Integrate-and-fire Time Encoding Machine
ADSM	Asynchronous Delta Sigma Modulator
TDIF	Time Dilating Integrate-and-Fire
E-spline	Exponential spline
B-spline	Polynomial basis spline
FRI	Finite Rate of Innovation

Table 2: Common terminology.

Trigger mark	The threshold value of the integral $z(t)$ above which the TEM generates a spike
--------------	--

Inter-trigger timespan	The time difference between two spikes $\tau_n - \tau_{n-1}$
Dilation function	The input-output curve of the amplifier in a TDIF
Amplitude-domain	A signal comprised of time varying against amplitude (e.g. integral of IF-TEM)
Time scale	The input to the integrator of the timing branch of a TEM (see Fig. 4.4)
Spike triggered reset	The reset of the integral of an IF-TEM is returned to zero after each spike

Table 3: Common mathematical notation.

x_k	k^{th} amplitude innovation parameter
$x_{\min/\max}$	Minimum and maximum bounds on amplitude $ x_k $
t_k	k^{th} time innovation parameter, time of pulse or edge
$\Delta_{\min/\max}$	Minimum and maximum separation between input Diracs $t_k - t_{k-1}$
C	Threshold/Trigger mark
τ_n	n^{th} Spike time
z_n	Spike amplitude, typically $\pm C$
$\beta(c)$	Non-linear amplifier
θ_n	Inter-trigger timespan ($\theta_n = \tau_n - \tau_{n-1}$)
$u(t)$	Heaviside function
$\delta(t)$	Dirac delta function
$(f * g)(t)$	Convolution of functions f and g
$z(t)$	Reset integral of TEM
$f(t)$	Non-reset integral of TEM
$y(t)$	Filtered input signal $y(t) = (x * \varphi)(t)$
$x(t)$	Input to TEM
$\varphi(t)$	Filter response

Chapter 1

Introduction and outline

Sampling theory is a fundamental area of research in signal processing. It concerns how physical signals can be represented by discrete sets of measurements called samples, when these samples are sufficient to uniquely determine a signal they represent, and the algorithms by which the signal can be reconstructed from its samples. By signals, we typically refer to measurable continuous quantities that vary over a continuous domain such as time or space. For example, a signal may be the illumination varying across a 2-dimensional space of an image plane, or a 1-dimensional voltage varying over time.

The conventional and most intuitive approach to sampling such signals has been to record measurements at specific times or positions. For example, in the case of a time varying signal, samples would have the form of tuples, (τ_n, x_n) , where x_n is a measurement of the time varying quantity of the signal at a specific time τ_n . For brevity, from hereon in I will refer only to such time varying signals but the concept generalises to signals that vary over other domains. Numerous sampling theorems regarding how to reconstruct signals from such samples have been developed since the early twentieth century. This began with the work of Kotelnikov, Whittaker, and Shannon on sampling and reconstructing bandlimited signals with uniformly spaced samples [1,2]. Since then, sampling

theory has been extended to non-uniformly spaced samples and has been developed for a wider variety of signals. A comprehensive summary of these developments can be found in [3].

To a very great degree, sampling theory has focused on this modality where a quantity is measured at specific points in time that are independent of the signal. This is even apparent from the terminology that we use for these processes. The measurement of the quantity itself is typically referred to as quantisation whereas the process of quantising at particular points in time is often separately referred to as sampling. The approach also maps quite naturally to practical implementation in synchronous circuits driven by digital clocks. For example, in the uniform sampling context, samples are taken periodically at multiples of the clock period requiring no explicit measurement of the time of a sample.¹ This is why, while we talk about quantisation of the amplitude measurement, x_n , we typically refer to clock jitter and not time quantisation of the time measurement, τ_n .

An alternative approach to the problem is to sample such that the timing of the samples and not the amplitude contains the information used to reconstruct the signal. Put another way, the timing of the samples are signal-driven whilst the amplitudes of the samples are in some way independent of the signal or otherwise implicit. Since the signal is encoded in the timing of samples, this is referred to as time encoding.

The time encoding approach is largely inspired by the integrate-and-fire models that describe the operation of biological neurons and consequently the most common form of time encoding machine (TEM) is the integrate-and-fire time encoding machine that mimics the behaviour of these neurons. In the most general terms, a neuron accumulates (integrates) an electrical stimulus (input signal) in the form of a potential difference across its cell membrane. When this membrane potential reaches a threshold value, the neuron releases it in the form of a pulse of electrical current or spike called an action potential. These are then conducted via electrochemical processes through the neuron to neighbouring neurons via synapses [4]. An oft used analogy is that of tipping bucket rain gauges that accumulate rain water, intermittently dumping their content each time the weight becomes sufficient to tip the bucket. The production of action potentials is sporadic and occurs at signal

¹Similarly, in a spatial domain, the positions of sensors are fixed and might be measured a priori.

driven times when the neuron’s membrane potential reaches a fixed amplitude. As such, action potentials can be viewed as a form of time encoding.

There are several potential applications of time encoding. First and foremost is in the design of sensing systems that are simpler or more power efficient. The biological neural networks that inspired integrate-and-fire time encoding exhibit extraordinary efficiency and it seems reasonable to assume that some of this efficiency is derived from the use of action potentials to propagate information.² Potential efficiency gains arise in two forms. Firstly, the actual hardware required to implement an integrate-and-fire system may be simpler than conventional sampling systems that must use complicated quantisers and sampling mechanisms. The simplest time encoding machine can consist of little more than a few operational amplifiers and a clock. Secondly, for certain types of sparse or sporadic signals, the integrate-and-fire time encoding machine, due to its event driven nature, can be setup to generate a more optimal and finite number of samples.

Time encoding approaches have been demonstrated in practice most notably with the development of dynamic vision sensors (DVS) [5]. These cameras do not operate with the conventional frame based approach. Instead of each pixel being scanned and quantised periodically, pixels generate samples asynchronously when their inputs change rapidly and the timing of the sample is recorded, not the amplitude. In this sense, the DVS represents a form of time encoding machine where the derivative of the signal drives the generation of samples.

In a conventional camera, large numbers of samples are redundant. For example, the input to many pixels changes only minimally from frame to frame, particularly for pixels situated in the background to the subject of the image. Despite this, a conventional frame-based camera will repeatedly sample these pixels regardless. The DVS solves this issue since areas with little to no activity produce no response from their respective pixels. This drastic reduction of the number samples has obvious benefits for power efficiency. Furthermore, coupled with the asynchronous nature of sample production, the DVS is also well suited to high speed low latency photography. A comprehensive survey of the DVS and various algorithms used for reconstruction from DVS data can be found in

²Though it is a matter of academic debate how important specifically the timing of action potentials is in information processing in biological neural networks.

[6]. The majority of approaches to reconstructing images from DVS data revolve around statistical or learning techniques. Time encoding theory based approaches are notable in their absence, perhaps due to the relative immaturity of time encoding theory.

When considering the theory of time encoding, two natural questions arise. Firstly, how and under what conditions can the input signal or stimulus be reconstructed from the timing of these spikes and secondly, what benefits can be gained. The former question has been answered for a number of different signal classes including bandlimited signals, signals in shift-invariant space, and some signals with finite rate of innovation. A review of these results, the time encoding machine architectures they pertain to and, the benefits they bring is the subject of Chapter 2. In passing, I also defined what I mean by signals with finite rate of innovation and how the concept can be used to give a more precise definition of sporadic signals.

In this thesis I look primarily at how time encoding, a form of asynchronous sampling, can be used to sample signals that exhibit sporadic behaviour with greater efficiency than conventional or existing time encoding strategies. To this end, in Chapter 3 I consider time encoding and reconstruction of signals with finite rate of innovation using novel sampling kernels. In particular, I look at the exponential reproduction technique for the α -synaptic function and the benefits this confers over existing works using the exponential spline. I also consider hyperbolic kernels such as the hyperbolic secant and reconstruction using a rational approximation approach.

The state of the art and my contributions in Chapter 3 typically use time-domain reconstruction algorithms. In Chapter 4 I take a different approach to time encoding that achieves reconstruction in the ‘amplitude-domain’ and thus more comprehensively exploits the duality of conventional sampling and time encoding. I do this by considering time encoding as a uniform sampling of an equivalent amplitude-domain signal. Building on this perspective I present a new time encoding architecture called a time-dilating integrate-and-fire encoder (TDIF) that manipulates time measurements during sampling. I show that such an architecture allows standard theorems of uniform sampling to be used to reconstruct a broad range of signals from their time encodings. I show that in the case

of piecewise constant signals, the new architecture produces an efficient number of samples while still allowing for perfect recovery.

In Chapter 5, I look at how networks of time encoding machines can be combined to derive some desirable sampling behaviours at the hardware level. This includes sampling networks that eliminate non-idealities through differential sampling and networks that conditionally sample based on the power of the input signal as well networks that amplitude-domain synchronisation of TEMs.

Finally, I conclude with some remarks in Chapter 6.

1.1 Achievements and Innovative results

This thesis contains several innovative results that can be summarised as follows:

- *Reconstruction algorithms from time encodings for novel signal models.* In particular, I focused on finite rate of innovation signals since they are an appropriate model for signals that exhibit sporadic behaviour. Examples of such signals are reconstruction algorithms for bursts of pulses interspersed with quiescence as well as bursts of edge transitions (of a piecewise constant signal) interspersed with quiescence. I develop reconstruction algorithms for hyperbolic functions as well as bursts of pulses with α -synaptic shape. In the case of the α -synaptic kernel the results are extended to higher kernels as well as filtered piecewise constant signals. The reconstruction approach for hyperbolic kernels is also extended to recovery of power signals as well as recovery from differential samples.
- *A novel time measurement dilating integrate-and-fire architecture for time encoding (TDIF)* is developed that leverages the integrate-and-fire aspect of the sampling process. By viewing the integrate-and-fire sampling procedure as a level set sampling or ‘amplitude’-domain sampling, standard techniques of uniform sampling can be used for signals that are bandlimited in amplitude-domain. I present a novel *time dilating* integrate-and-fire architecture and demonstrate that it generates uniform samples in amplitude-domain with arbitrary kernels. Thus sam-

pling theorems for uniform samples in time-domain can be applied with these modified samples in amplitude-domain.

- *Several reconstruction theorems for the TDIF.* The TDIF architecture allows standard uniform sampling in the amplitude-domain. I explicitly derive theorems for signals that are shift-invariant in amplitude-domain, and non-negative piecewise constant signals in time-domain (a signal with finite rate of innovation). In fact, the latter is a demonstration of the fact that piecewise constant signals are also piecewise constant signals in amplitude-domain
- *Multi-TEM networks.* Multiple TEMs can be arranged in topologies that provide desirable behaviours in sampling and sampling unrelated contexts. I present two. The first is a network called a differential pair designed to generate differential samples that can be more resilient to non-idealities in the constituent TEMs. The second are a class of networks called ring synchronisers that achieve the synchronisation of integrate-and-fire TEMs in amplitude-domain. This is an analogue to clock synchronisation in time-domain.

1.2 List of publications

1. M. Hilton, R. Alexandru and P. L. Dragotti, “*Time Encoding Using the Hyperbolic Secant Kernel*,” 2020 28th European Signal Processing Conference (EUSIPCO), Amsterdam, Netherlands, 2021, pp. 2304-2308, doi: 10.23919/Eusipco47968.2020.9287806.
2. M. Hilton, R. Alexandru and P. L. Dragotti, “*Guaranteed Reconstruction from Integrate-and-Fire Neurons with Alpha Synaptic Activation*,” ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Toronto, ON, Canada, 2021, pp. 5474-5478, doi: 10.1109/ICASSP39728.2021.9414759.
3. M. Hilton and P. L. Dragotti, “*Sparse Asynchronous Samples from Networks of TEMs for Reconstruction of Classes of Non-Bandlimited Signals*,” ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Rhodes Island, Greece, 2023, pp. 1-5, doi: 10.1109/ICASSP49357.2023.10095620.

4. M. Hilton and P. L. Dragotti, “*Playing with Time: Dilation and Encoding*,” under review for Applied and Computational Harmonic Analysis, Elsevier, 2024.

Chapter 2

Time encoding preliminaries

2.1 Architectures for sampling in amplitude

Sampling is the process by which an observable signal is discretised and converted into a countable set of digital numbers. The signal being observed can have a multidimensional varying element, but for the purposes of this work, I will focus on the case where signals consist of a single dependent quantity that varies against an independent property. The independent quantity must have an order, such as in the case of time and space. I will focus on signals that vary over time. Voltage across an electrical component or the pressure impinging on a microphone are examples of such time varying signals.

The objective of sampling theory is to determine under what conditions the original signal can be resynthesised from this set of samples. The most prevalent forms of sampling are where measurements of the observable signal are taken at predetermined points in time. Most commonly these are uniformly spaced in time which gives rise to the celebrated Shannon-Whittaker sampling theorem that gives a way to sample and reconstruct bandlimited signals. Sampling of this kind, where times are predetermined and magnitudes of the samples are signal driven, has been extended to

non-uniform sampling, where samples are non-uniformly spaced in time [7,8], and to larger classes of signals such as signals in shift-invariant spaces [9], and sparse signals [10–12].

In general, the samples of a one-dimensional time-varying signal can be thought of as a set of tuples $\{(\tau_n, x_n) \mid n \in [1, \dots, N]\}$ where N is the number of samples. In the context of conventional sampling strategies where timing of samples are not signal dependent, the times of the samples, τ_n , are assumed to be predetermined whereas the ‘magnitude’ of the samples, x_n , are assumed to be driven by the time-varying quantity of the signal, for example, the voltage on a wire. Because of this, the samples of a conventional sampler can further be considered to be just a sequence of measurements $\{x_n \mid n \in [1, \dots, N]\}$ since the measurements of time are implicit in the index n . For example, in uniform sampling, $\tau_n = nT_s$ where T_s is the known sampling period. Even in the non-uniform sampling context, the timing of the samples might be known a priori or are signal independent. If these timings are considered to be an ordered set then the sample times are similarly implicit in the index of the sample.

In a time encoding machine (TEM) the encoding of the signal into samples is intended to be precisely the opposite. The samples should be a set of measurements in time $\{\tau_n \mid n \in [1, \dots, N]\}$ where the measurements of the ‘magnitude’ at that time are implicit in the time samples. Much like in the conventional uniform sampling case where the time was implicit due to knowledge of the sampling period, the implicit magnitudes in a time encoding are given when we have prior knowledge of the parameters of the time encoding sampler. For example, in an integrate-and-fire encoder we have prior knowledge of the threshold, or in a crossing encoder we have knowledge of the reference signal. These aforementioned existing architectures which are the physical systems by which samples of this form can be realised is the subject of this section. The methods through which signals can be reconstructed from time encodings generated by these architectures are then discussed in Section 2.2.




Figure 2.1: A typical reference crossing encoder.

2.1.1 Crossing time encoders

The simplest and most intuitive form of time encoding machine is the reference crossing encoder illustrated in Fig. 2.1. It is a simple device consisting of a known reference signal, $r(t)$, typically a sinusoidal oscillation, ramp, or constant signal, and an input signal $x(t)$. The two signals are compared and when they cross (are equal) an output sample is generated. We can conceptualise this output either as a spike with infinitesimally small support, or simply as measurement of the time when the input signal crossed the reference signal. The former representation is more applicable when discussing networks of time encoding machines, similar to networks of spiking neurons, and the latter is typically more appropriate when treating the time encoding machine as a sampler.

Let τ_n be the times the crossing encoder produces, since the reference signal is known, the amplitudes $x_n = x(\tau_n) = r(\tau_n)$ are clearly implicit in the times, τ_n . Thus, the sampled times alone represent a time encoding of the input signal $x(t)$. An example, of such an encoding is given in Fig. 2.2 where the reference signal is chosen to be a sinusoid.

The operation of comparing the two signals is represented by the difference of the two signals followed by a zero crossing detector. Indeed, the reference crossing detector is a generalisation of a zero crossing encoder when the reference signal is $r(t) = 0$. In such a case Logan's theorem can be used to recover bandlimited signals from their zero crossings alone [13]. The combination of difference and zero detection operations implements detection of equality between the input signal and the reference signal. In other formulations, most notably [14], the output time encoding is represented as a signal whose polarity inverts at each sample time. As such, their definition of the crossing encoder uses a comparator instead of a zero crossing detector. This allows a more gen-




Figure 2.2: Example time encodings generated by a crossing encoder for two inputs. Signal **a)** is a bandlimited signal. Samples are recorded at each point the signal $y(t)$ is equal to the reference sinusoid. Signal **b)** is a signal containing an initial period of quiescence. Note how the crossing encoder generates uniform samples regardless.

eral framework that encompasses both crossing TEMs, integrate-and-fire TEMs, and asynchronous sigma-delta modulators. The latter are also a form of time encoding machines but not typically referred to as such explicitly. The advantage of the formulation of Fig. 2.1 is that it makes a more explicit connection between the representation of samples as spikes or time measurements and the output of the encoder.

One advantage of a reference crossing encoder lies in the simplicity of the device. There are a few components and, unlike the integrate-and-fire systems discussed later, there are no accumulators or resetting mechanisms that need special consideration. Another advantage is found in the theoretical analysis of such a sampling strategy. Guarantees that samples are produced in sufficient quantities for reconstruction are more easily obtained when the reference signal is chosen appropriately. For instance, let the reference signal be periodic with a wider dynamic range than the input signal and assume the input signal is continuous, then the system can be guaranteed to produce at least one sample within any two periods of the reference signal. Thus, by selecting an appropriate reference signal, a sufficient density of samples can be enforced. This is the principle used in [15] to produce guarantees that enough samples are generated to allow reconstruction.

However, this approach to creating a time encoding has major disadvantages. A major aspiration of time encoding methods in this thesis is to create systems that generate sparser sets of samples and achieve this by being adaptive to the signal they sample. If we consider the samples generated by a crossing encoder (Fig. 2.2), we can see that the samples occur regularly across the entire timespan of the signal and do not localise around periods of activity. Furthermore, in periods of quiescence, samples are still generated at an almost uniform rate and differ from a uniform sampling by only small degrees. Moreover, this is highly inefficient when sampling signals with compact support. These drawbacks will become more apparent when we consider highly sporadic and sparse signals in Section 2.3.1.

2.1.2 The integrate-and-fire time encoding machine

The focus of this thesis is specifically on integrate-and-fire methods for time encoding. Integrate-and-fire describes a class of samplers where signals are accumulated (integrated) and when the accumulated signal satisfies a condition, typically equality with a known threshold or trigger mark, the time is recorded as a sample and the accumulated value is reset back to zero. There are a few slight variations in the literature on the structure of an integrate-and-fire system which I will enumerate in this section.

The integrate-and-fire approach is inspired by biological neurons which exhibit similar, though significantly more complex, behaviour. Information in a neural network is transported in the form of bursts of electrical activity called ‘action potentials’. This all or nothing approach allows efficient transportation of electrical charge along the neuron cell body and between neurons. Integrate-and-fire time encoding machines (IF-TEM) share similarities with biological neurons but for the most part use more ideal models that can be realised with standard electrical components. The first form of the time encoding machine was introduced by [16] along with a reconstruction algorithm for the bandlimited signals. The time encoding machine was constructed as shown in Fig. 2.3 and resembles an asynchronous version of a delta-sigma modulator. For this reason it is referred to as an asynchronous delta-sigma modulator (ADSM). In fact, the same architecture was initially proposed by [17] and has enjoyed a recent revival and development that is seemingly independent to the efforts of those studying time encoding in the signal processing community [18–20].

The ADSM uses feedback to create an astable circuit whose output, $z(t)$, repeatedly and asynchronously changes polarity. To achieve this, the integrated signal, $f(t)$, is fed through a Schmitt trigger with threshold δ to create an output of $z(t) = \pm b$. This output is then used as negative feedback into the integrator. If the output value is sufficiently high such that $b > \sup |x(t)|$, the resulting integral will always be monotonic and decreasing (increasing) for a positive (negative) output. Consequently, each time the integral reaches the threshold $\pm\delta$, the output polarity flips and the integral starts increasing or decreasing towards the opposite threshold. Let τ_n be the switching times of the




Figure 2.3: A Lazar-style integrate-and-fire time encoding machine in the form of an asynchronous delta-sigma modulator (ADSM).




Figure 2.4: A typical integrate-and-fire time encoding machine after [21]. b is a constant bias, C is the threshold, and τ_n are the spike times.

output from the Schmitt Trigger. The switching times of the signal can be related to the input signal by [16],

$$\int_{\tau_n}^{\tau_{n+1}} x(u) du = (-1)^n (-b(\tau_{n+1} - \tau_n) + 2\kappa\delta) . \quad (2.1)$$

The ADSM is closely related to another more commonly used form of integrate-and-fire time encoding machine, typically referred to as an IF-TEM and first introduced in [22]. The architecture is simpler and the associated reconstruction algorithm is the same, therefore, so I will forego a deeper analysis of the ADSM in favour of the IF-TEM. The main difference in architecture is that the IF-TEM

operates without feedback and hysteresis as illustrated in Fig. 2.4. The operation of the IF-TEM is as follows:

1. The input signal is summed with a constant bias, b , to produce $y(t)$.
2. The resulting signal is integrated, to produce $f(t)$, which is the *unreset* integral.
3. Each time the resulting signal crosses a threshold C , a sample is generated.
4. The integral is reset to zero after each sample is generated. The resulting signal is the reset integral, $z(t)$.

Fig. 2.5 plots the intermediate signals of the IF-TEM for the given example signal. Note that a typical assumption of this architecture is that the bias, b , is strictly greater than the input signal $x(t)$. This ensures the integral is monotonic. The use of a bias has a similar effect to that observed in the crossing encoder in that samples are produced constantly and at relatively uniform times. Indeed, samples are generated even when the input signal $x(t)$ is zero. It is also worth noting that the bias term corresponds to a ramp signal of constant gradient in the integral. As such, this architecture is equivalent to Δ -ramp encoders as investigated by [23–25].

The mechanism used to reset the integral is not stated explicitly but can be considered analogous to the discharging of the membrane potential of a biological neuron or the discharging of a capacitor through a transistor. The combination of bias and the resetting operation serves a similar purpose to the alternating polarity of the Schmitt trigger output of the ADSM in that instead of bouncing between a positive and negative threshold, $\pm\delta$, the IF-TEM repetitively reaches a single threshold, C . It then resets and continues in the same fashion to reach the threshold.

The sample times generated by the IF-TEM are related to the input signal in a similar way to the ADSM through,

$$\int_{\tau_n}^{\tau_{n+1}} f(t) dt = C - b(\tau_{n+1} - \tau_n) , \quad (2.2)$$

which is sometimes referred to as the t -transform [22]. In both the case of the ADSM and the IF-TEM, the resulting time encoding cannot be described simply as a set of samples, (y_n, τ_n) where a




Figure 2.5: The intermediate signals of an IF-TEM: **a)** input signal, **b)** input signal with bias, **c)** integrated signal (this is not manifest in a TEM but an abstract ‘unreset’ signal), **d)** integrated signal with resets where sample times are taken. Note the monotonicity of **c)** and the relatively uniform firing of **d)**.

simple magnitude measurement, y_n , is implicit in its corresponding time, τ_n . Instead, the magnitude measurement must be expressed as a finite difference on the integral of the input signal, that is,

$$\int_{\tau_n}^{\tau_{n+1}} y(t) dt = f(\tau_{n+1}) - f(\tau_n) . \quad (2.3)$$

In conclusion, it is interesting to note the meagre variety in architectures of the single time encoding machine. In particular, the ADSM and IF-TEM architectures are very similar architectures which produce very similar time encodings. Indeed, the choice between the two is likely governed more by practical considerations of implementation and not by the difference in the encodings they produce. This lack of variety motivates the results I present in Chapter 4 to a great extent.

2.2 Reconstruction of bandlimited and shift-invariant signals

2.2.1 Iterative reconstruction

Reconstruction of signals from their time encodings was first investigated in the context of ADSMs and IF-TEMs by Lazar and Tóth [16,22,26]. They showed that if the input signal, $x(t)$, is bandlimited to $[-\Omega, \Omega]$ (and consequently $y(t)$ is bandlimited) then it can be reconstructed perfectly from its time encoding when,

$$\frac{\pi(b - c)}{2C} > \Omega \quad \text{rad/s} , \quad (2.4)$$

where C is the trigger mark, b is the bias of the IF-TEM, and $c = \|y\|_\infty$ is the maximum absolute value of the input signal. Since lower trigger marks and higher biases correspond to higher firing rates, this provides a Nyquist-like sampling criterion where a higher bandwidth signal requires the IF-TEM's parameters to be altered to increase firing rate.

Reconstruction is achieved through an iterative procedure that converges on the solution. Consider an operator $\mathcal{A}$,

$$\mathcal{A}y(t) = \sum_{n \in \mathbb{Z}} \int_{\tau_n}^{\tau_{n+1}} y(t) dt \operatorname{sinc}_{\Omega}(t - s_n) , \quad (2.5)$$

where $s_n = (\tau_{n+1} - \tau_n)/2$ is the midpoint of an interspike interval. Crucially, the value of the inner definite integral can be calculated directly from a time encoding via the t-transform, Eq. (2.2), without knowledge of the signal. Consequently, $\mathcal{A}y(t)$ can be directly calculated as,

$$\mathcal{A}y(t) = \sum_{n \in \mathbb{Z}} (C - b(\tau_{n+1} - \tau_n)) \operatorname{sinc}_{\Omega}(t - s_n) . \quad (2.6)$$

By leveraging the work of [9,27] on sampling averages, the authors showed that $I - \mathcal{A}$ is a contractive operator when Eq. (2.4) is satisfied and where I is the identity operator. As such, one can formulate a Neumann series,

$$\mathcal{A}^{-1} = \lim_{K \rightarrow \infty} \sum_{k=0}^K (I - \mathcal{A})^k , \quad (2.7)$$

so that,

$$I = \lim_{K \rightarrow \infty} \sum_{k=0}^K \mathcal{A}(I - \mathcal{A})^k . \quad (2.8)$$

Expanding the summation gives the following iteration,

$$\begin{aligned} \hat{y}^{(0)}(t) &= \mathcal{A}y(t) \\ \hat{y}^{(k+1)}(t) &= \hat{y}^{(k)}(t) + \mathcal{A}(y(t) - \hat{y}^{(k)}(t)) . \end{aligned} \quad (2.9)$$

The base case of the recursion, $\hat{y}^{(0)}(t) = \mathcal{A}y(t)$, can be calculated by Eq. (2.6) from the time encoding of $y(t)$. An example of reconstruction with the iterative approach of Eq. (2.9) is given in Fig. 2.6.

This iterative approach to reconstructing bandlimited signals from time encodings has been extended to signals in shift-invariant subspaces by [25] and with an alternative reconstruction algorithm called projection onto convex sets (POCS) by [28,29].




Figure 2.6: An example time encoding and reconstruction of bandlimited signal ($\Omega = 5 \text{ rad s}^{-1}$) using the iterative reconstruction algorithm of Eq. (2.9). Intermediate results for $k = 0, 2, 4$ are given.

These reconstruction results have also been extended beyond the single TEM architecture. In particular, multi-TEM sampling architectures have been used to gain advantages in firing rate and tradeoffs in spatial and temporal resolution [14,30,31]. The TEM architecture has also been combined with modulo sampling techniques [32] to provide a more meaningful tradeoff between firing rate and dynamic range [33].

2.2.2 Multichannel time encoding machines

Reconstruction from time encodings has also been built upon with novel multi-TEM architectures for multichannel encodings of correlated signals such as video [31,34,34–36], and multichannel encodings of a single input signal [30,37]. In particular, the work of Adam [14,30,37] looked extensively at how multichannel systems can grant more beneficial tradeoffs between firing rate, spatial, and temporal resolution. They proposed an M -channel system where each channel consists of a single TEM that individually samples the input signal. With this architecture the input signal can still be constructed even when each constituent TEM fires at a rate slower than in the single-channel case. If the integral of each TEM is offset from the integral of the others, the sets of spikes generated by each channel can be used together to reconstruct the input signal. More specifically they show that the signal must satisfy the following more relaxed criterion,

$$\frac{M\pi(b-c)}{2C} > \Omega, \quad (2.10)$$

where, as before, c is the maximum of the input signal. In fact, as can be seen by comparing Eq. (2.10) to Eq. (2.4), one can use a multichannel system to trade off between the number of channels and the spike rate of each channel. That is, in an M -channel system, the firing rate can be M times slower per channel while still allowing perfect reconstruction of the input signal. This is advantageous since practical implementations of IF-TEMs and the biological neurons that they imitate have limited rate at which they can fire. In fact, individual neurons in the human body rarely fire above rates of 100Hz [38]. A tradeoff between the number of channels and the firing rate of each individual channel can be made without affecting the effective global firing rate. In a video time encoding setting, the

channels are analogous to pixels that receive different but correlated signals. Adam goes on to show that in this case a similar tradeoff can be made between the number of pixels (spatial resolution) and the firing rate of individual channels (temporal resolution).

2.2.3 Modulo

An area of active research which bears strong resemblance to time encoding is modulo sampling, sometimes referred to as unlimited sampling. In the modulo sampling paradigm, samples are taken in a conventional manner except that the input signal is wrapped such that it stays within a specific dynamic range. That is, each time the signal crosses an upper or lower threshold, a bias is applied bringing the signals to its respective opposite threshold in much the same way as an IF-TEM. Conventional samples (for example uniform samples) are taken from the wrapped signal.

It has been shown that perfect reconstruction can be achieved for various classes of signals [32]. The main advantage of this approach is that high dynamic range signals can be confined to a lower dynamic range while still allowing sampling and reconstruction. This is also beneficial when quantisation is taken into account since fewer bits are required to sample a narrower dynamic range.

It is apparent that time encoding and modulo sampling have similar characteristics, they both employ some form of wrapping the signal when a threshold is exceeded, however, they differ in how this process is used. In the time encoding case, we have seen that the timing of these crossing points is measured at sampling time and is directly used to reconstruct the input signal. Indeed, the crossing times are the samples themselves. In the modulo case, the timing of the crossings is unknown. Instead, in general terms, the objective is estimate the timing of the crossings and consequently retrieve non-modulo samples that can be used to reconstruct the signal in a conventional manner.

2.2.4 Modulo IF-TEM

The proposed methods for time encoding with an IF-TEM typically include the use of a constant bias, b , to ensure that the input signal is monotonically increasing and that the spike rate is sufficiently

high to achieve reconstruction. This is achieved by ensuring that the bias is greater than the infinity norm of the input signal as expressed through the criteria of Eq. (2.4) & Eq. (2.10). If the input becomes greater than the bias, guarantees that a sufficient density of samples is achieved are lost. In [33], an adaptation of the IF-TEM architecture was proposed that incorporates recent innovations in the area of modulo sampling [32], to address the reliance on a carefully selected bias. Their approach involves applying a non-linear modulo operation to the input signal that wraps the signal within a chosen range, $[-\lambda, \lambda]$. This ensures the signal is always smaller than the bias and thus the standard guarantees for sampling rate apply. They go on to show that the resulting time encoding of a modulo folded signal can be used to reconstruct the input signal up to a constant bias.

2.3 Reconstruction of signals with finite rate of innovation

2.3.1 Finite rate of innovation

The main focus of this thesis is on signals that exhibit sporadic activity. By this I mean signals where regions of interest or activity in are highly localised and interspersed with periods of inactivity or quiescence of varying lengths. For example, one might be interested in sampling and reconstructing signals that consist of pulses occurring at times not aligned to a grid, where the number of pulses may or may not be finite. Since the pulses are localised in time, there can be extended periods of time where the signal is zero. Another example are piecewise constant signals which have discontinuities or groups of discontinuities separated by long periods of quiescence where the signal remains unchanged. The notion of a sporadic signal often coincides with that of a sparse signal but care must be taken not to consider them one and the same. A signal that is sparse is not necessarily sporadic, for example, a train of close but evenly spaced pulses may be sparse but it is not sporadic since regions of activity are evenly spaced throughout time. Similarly, a signal that is sporadic is not necessarily sparse, for example, intermittent pulses of white noise may be sporadic since they constitute highly localised bursts of activity interspersed with periods of inactivity, but each pulse

itself is not sparse. Sporadicity can be seen as the change in sparsity over time or the locality of sparsity.

Bandlimited and shift-invariant signals are not the most appropriate signal model for representing highly sporadic signals. For instance, a strictly bandlimited signal, that is a signal that is well localised in Fourier domain, is necessarily poorly localised in time domain due to the uncertainty principle. In contrast, highly sporadic signals exhibit activity that is highly localised. Furthermore, the sampling strategies for such signals (e.g. uniform sampling) also not the most appropriate when sampling highly sporadic signals. The typical uniform sampling strategy due to Shannon requires measuring the signal at points during the signals' periods of quiescence. It is clear that such a sampling must be sub-optimal.

To model and define sporadic signals I will instead use the framework of signals with finite rate of innovation (FRI) which was first introduced in [11]. FRI signals are a class of signals that can be parameterised by a finite number of parameters over a finite interval. The number of parameters, or degrees of freedom, required to describe a signal on an interval $[-L/2, L/2]$ is called the local rate of innovation,

$$\hat{\rho}_L = \frac{1}{L} C_f \left(-\frac{L}{2}, \frac{L}{2} \right) , \quad (2.11)$$

where $C_f(a, b)$ counts the degrees of freedom in a signal, $f(t)$, in an interval $[a, b]$. When $\hat{\rho}_L < \infty$, we say that the signal has finite local rate of innovation. The limit of $\hat{\rho}_L$ as the interval widens to the entire real line is called the global rate of innovation,

$$\rho = \lim_{L \rightarrow \infty} \hat{\rho}_L . \quad (2.12)$$

When ρ is finite we say the signal has finite (global) rate of innovation. In practice one is usually interested in local rate of innovation as this determines how many samples are required to reconstruct the signal in a given interval, however, finite global rate of innovation is useful since it implies that the signal must have finite local rate of innovation at every point.

FRI signals are of particular interest in the context of integrate-and-fire time encoding since they naturally provide a representation for highly sporadic signals and sparse signals. Consider signals of the form,

$$\begin{aligned} x(t) &= \varphi(t) * \sum_{k \in \mathbb{Z}} x_k \delta(t - t_k) \\ &= \sum_{k \in \mathbb{Z}} x_k \varphi(t - t_k), \end{aligned} \tag{2.13}$$

where $t_k < t_{k+1}$ and there are a finite number of t_k 's in any finite interval or the real line. By definition, such a signal has finite local rate of innovation. Given a choice of $\varphi(t)$ that is well localised in time, such a signal resembles a stream of pulses with highly sporadic and localised behaviour.

By this I mean two things. First, the structure of the signal is made up of events that occur at asynchronous points in time. This is evident in that each component can occur at an arbitrary time, t_k , which is not aligned on a grid. Second, activity in the signal can be highly localised such that certain intervals have a high local rate of innovation whilst others have a low local rate of innovation. In the extreme, the signal consists of intervals with local rate of innovation $\hat{\rho} = 0$, which I shall refer to as quiescence, interspersed with intervals with high local rate of innovation $\hat{\rho} = K$, which I shall refer to as a burst. An example of this is plotted in Fig. 2.7 which shows a highly sporadic signal comprised of periods of quiescence ($x(t) = 0$) interspersed with bursts of overlapping pulses.

Optimal sampling of this signal would produce the same or close to the same number of samples in a given interval as the number of innovations in the interval. Since intervals of quiescence can be arbitrarily long, it is evident that a conventional uniform sampling is suboptimal as it would produce samples within any quiescent interval with $\hat{\rho} = 0$. Indeed, if quiescent intervals can be arbitrarily long, then the number of conventional samples taken in the interval can be arbitrarily large. As discussed earlier, IF-TEMs with a bias and crossing encoders also suffer from this inefficiency.

The signal model also extends to other FRI signals such as filtered piecewise constant signals and




Figure 2.7: Sporadic signal comprised of two bursts of three Diracs with a pulse shape $\varphi(t) = \text{sech}(5t)$. Since the pulse shape is well localised, the bursts are separated by an interval where $y(t) = 0$ and $\hat{\rho} = 0$ (quiescence). $\hat{\rho} = 3$ for the interval of each burst.

more generally filtered piecewise polynomial signals via the integral. In the case of the filtered piecewise constant, the relationship is simple,

$$\begin{aligned}
 y(t) &= \int x(t) dt = \int \varphi(t) * \sum_{k \in \mathbb{Z}} x_k \delta(t - t_k) dt \\
 &= c + \varphi(t) * \sum_{k \in \mathbb{Z}} x_k u(t - t_k) ,
 \end{aligned} \tag{2.14}$$

where c is a constant of integration and $u(t)$ is the Heaviside step function. Intuitively, $u(t)$ is not localised in time, however, the filtered piecewise constant of Eq. (2.14) still has finite local rate of innovation on a given interval, $[t_a, t_b]$, since

$$\begin{aligned}
\int x(t) dt &= c + \varphi(t) * \left(\sum_{t_k < t_a} x_k u(t - t_k) + \sum_{t_a \leq t_k \leq t_b} x_k u(t - t_k) + \sum_{t_b \leq t_k} x_k u(t - t_k) \right) \\
&\stackrel{(a)}{=} c + \underbrace{\|\varphi\|_1 \sum_{t_k < t_a} x_k}_{\xi_a} + \varphi(t) * \sum_{t_a \leq t_k \leq t_b} x_k u(t - t_k), \tag{2.15}
\end{aligned}$$

for $t \in [t_a, t_b]$ and where (a) follows from assuming that $t_a - t_k$ is larger than the support of $\varphi(t)$ such that the first summation is equivalent to the l-1 norm of the kernel. The third summation is zero inside the interval.¹ The third sum is eliminated because for any t_k greater than t_b , $u(t - t_k) = 0$. Thus, any finite interval of a piecewise constant signal is parameterised by $2K + 1$ parameters for $2K$ edges where one of these parameters, ξ_a , represents the value of the signal at the beginning of the interval.

Piecewise constant signals clearly also represent highly sporadic signals except that periods of quiescence, for which $\hat{\rho} = 0$, are not necessarily zero but are regions where the signal remains constant. Bursts are also not made up of pulses but of edges. An example of a filtered piecewise constant is given in Fig. 2.8 which is related through the integral to Fig. 2.7.

Sampling of sporadic signals such as these in an optimal way is a problem that is only partially solved by existing FRI theory. Whilst the framework provides an appropriate signal model for sporadic signals, the acquisition methods are largely unchanged, mostly relying on uniform sampling. It is clear that uniform sampling produces suboptimal sample sets since samples are taken even during periods of quiescence when there is no new information to be recorded. In the case of pulses there can be long periods where the signal is zero, and in the case of piecewise constant signals there can be long periods where the signal does not change. Redundant sampling in these intervals is also encountered when using the typical IF-TEM with an input bias.

¹Note, the sums are indexed over k




Figure 2.8: A filtered piecewise constant signal equivalent to the integral of Fig. 2.7.

2.3.2 Reproduction of polynomials and exponentials

The first approach to reconstructing FRI signals from time encodings is due to [15] who demonstrated that FRI signals can be reconstructed when the filter, φ , exhibits a polynomial or exponential reproducing property. This approach was adapted from the work of [39] who leveraged the same polynomial (exponential) reproducing property but in the context of uniform sampling. The reconstruction algorithm consists of estimating the innovation parameters, t_k & x_k , through two key operations. First, the exponential reproducing property of the kernel is used to transform the uniform samples of the signal into uniform samples of an equivalent sum of exponentials or polynomials. Spectral estimation techniques like Prony's method can then be used to recover the parameters. A function $\varphi(t)$ can reproduce polynomials up to a maximum degree R if,

$$\sum_{k \in \mathbb{Z}} c_{m,k} \varphi(t - n) = t^m, \quad (2.16)$$

for $m \leq R$ and some $c \in \mathbb{C}$. That is, there exists a weighted sum of uniform shifts of the function $\varphi(t)$

that reproduces a monomial of maximum degree R exactly. By linearity, there also exists a weighted sum that reproduces any polynomial of maximum degree R . Functions that have this property are known to satisfy the Strang-Fix conditions of order R [40]:

$$\hat{\varphi}^{(q)}(2ni\pi) = 0, \quad (2.17)$$

for $n \in \mathbb{Z}/\{0\}$, $|q| < R$ and $\hat{\varphi}$ is the Laplace transform of φ . Examples of functions that satisfy these conditions are polynomials basis splines of degree R and the scaling function of the Daubechies wavelet. The Strang-Fix conditions can be extended to the generalised Strang-Fix conditions in [41],

$$\hat{\varphi}^{(q)}(a_m + 2ni\pi) = 0. \quad (2.18)$$

A function that satisfies the generalised Strang-Fix conditions can reproduce exponentials so that coefficients $c_{m,k}$ can be found such that,

$$\sum_{k \in \mathbb{Z}} c_{m,k} \varphi(t - n) = e^{a_m t}. \quad (2.19)$$

An example of this kind of reproduction is shown in Fig. 2.9. I will focus on exponential reproduction but the following results apply equally to both polynomial and exponential reproducing functions. Consider an FRI signal consisting of a finite stream of K pulses sampled uniformly,

$$y_n = \sum_{k=1}^K x_k \varphi(n - t_k), \quad (2.20)$$

where $\varphi(t)$ satisfies the Generalised Strang-Fix conditions Eq. (2.18) up to degree $2K$. Taking a linear combination of samples, y_n , to produce signal moments, s_m , we can reproduce exponentials,

$$s_m = \sum_{l \in \mathbb{Z}} c_{m,l} \sum_{k=1}^K x_k \varphi(n - t_k) = \sum_{k=1}^K x_k \sum_{l \in \mathbb{Z}} c_{m,l} \varphi(n - t_k) = \sum_{k=1}^K x_k e^{-a_m t_k}, \quad (2.21)$$

where possible exponents a_m are dependent on the kernel $\varphi(t)$. If the a_m are uniformly spaced such that $a_m = a_0 + m\lambda$ where $\lambda \in \mathbb{R}$ then the parameters t_k can be estimated using spectral estimation




Figure 2.9: Reproduction of three exponentials using an E-Spline which can reproduce the exponents $\alpha = -0.5, 0, 1$.

techniques like Prony's method [11,42] when there are at least $2K$ signal moments. Therefore, the sampling kernel must be able to reproduce at least $2K$ different exponentials and at least $2K$ samples must be taken. Once the innovation times, t_k , are estimated, the parameters x_k can then be estimated through a least squares approach or, in the case of hyperbolic functions, by direct calculation as shown in Chapter A.

This approach to reconstruction demonstrates another advantage of the FRI approach. The sampling is optimal in the sense that only $2K$ samples are required to reconstruct the $2K$ innovation parameters of the FRI signal. In this case I have assumed that a signal has a finite number of innovations globally across the entire real line. This assumption can be relaxed when we assume that $\varphi(t)$ is well localised or of compact support so that the signal has finite local rate of innovation,

$$y_n = \sum_{k \in \mathbb{Z}} x_k \varphi(n - t_k) . \quad (2.22)$$

Apart from Prony's method, other techniques have also been proposed to improve the robustness of the parameter estimation such as sparse optimisation methods [10], subspace methods [41,43–46], and model fitting approaches [47–50]. Different kernels have also been proposed including sum-of-sinc kernels (SoS) [51] and exponential maximum order minimum support kernels (E-MOMS) [41,52] as well as approximate reproduction of exponentials with arbitrary kernels [41,52].

The generalised Strang-Fix conditions ensure that polynomials and exponentials can be reproduced globally with uniform shifts of the kernel. Unfortunately, time encodings naturally map to non-uniform samples. Consequently, the exponential reproduction of Eq. (2.19) does not hold. From the t -transform, Eq. (2.2), we can relate the spike times generated from an IF-TEM to the trigger mark and the FRI signal parameters,

$$C - b(\tau_n - \tau_{n-1}) = \int_{\tau_{n-1}}^{\tau_n} y(t) dt = \sum_{k \in \mathbb{Z}} x_k \int_{\tau_{n-1}}^{\tau_n} \varphi(u - t_k) du . \quad (2.23)$$

From this representation, [15] made two observations: Firstly, when $\varphi(t)$ can reproduce R exponentials, the exponential reproducing property is maintained and the function $\psi_n(t) = \int_{\tau_{n-1}}^{\tau_n} \varphi(u -$

t_k) du can also reproduce the same exponentials. Secondly, they noted that whilst it is not possible to globally reproduce exponentials using these non-uniform samples, it is still possible to reproduce exponentials locally within certain intervals. In particular exponential reproduction can still be achieved when non-uniform samples fall within the same knotless region of the input signal. By knots I mean points x where a signal is only finitely differentiable. Local reproduction is possible because a compact sampling kernel that satisfies the Generalised Strang-Fix conditions is a piecewise function where each piece is a sum of exponentials. Let $\tau_n \in \mathbb{R}$ then each piece of the shifted kernel can be expressed as,

$$\varphi_L(t - \tau_n) = \sum_{m=1}^R d_{m,L} e^{-a_m \tau_n} e^{a_m t}, \quad (2.24)$$

where $\tau_n \in \mathbb{R}$ is the non-uniform shift of the kernel, $t - \tau_n \in L$ and L is a knotless interval of the kernel. Thus, a linear combination of shifts can be found to reproduce exponentials $e^{a_m t}$ within the given interval and parameters can be estimated with the usual spectral estimation technique.

The need for samples to occur in the same knotless regions is a downside of this approach to reconstruction from time encodings since it makes it hard to present simple criteria that guarantee reconstruction is possible. This problem is compounded further when one considers that kernels which can reproduce larger numbers of exponentials have larger numbers of knots. To mitigate this, [15] proposed a multichannel system where different filters are used for each channel. In this way each channel can reproduce different sets of exponentials thereby reducing the order of individual filters.

2.3.3 Reconstruction from Fourier series

In [53], another approach to reconstructing FRI signals from time encodings is proposed. This was similarly based on a previous work, [11,54], which is also predated by the work of [55], regarding the reconstruction of streams of pulses from uniform samples. Consider a periodic stream of pulses,




Figure 2.10: Examples of non-uniform reproduction on the interval $[2.4, 3]$ from non-uniformly spaced E-splines that can reproduce exponents $\alpha = -0.5, 0, 1$

$$y(t) = \varphi(t) * \sum_{n \in \mathbb{Z}} \sum_{k=1}^K x_k \delta(t - t_k - nT) , \quad (2.25)$$

where x_k and t_k are the usual innovation parameters and T is the period of the signal. The signal is clearly an FRI signal with a finite number of $2K$ innovations. One can rewrite Eq. (2.25) with the Fourier series representation of the sum of Diracs such that,

$$y(t) = \varphi(t) * \sum_{n \in \mathbb{Z}} c_n e^{2i\pi n t / T} , \quad (2.26)$$

where,

$$c_n = \frac{1}{T} \sum_{k=1}^K x_k e^{2i\pi n t_k / T} . \quad (2.27)$$

Eq. (2.27) is a sum of exponentials similar to Eq. (2.21). The innovations can therefore be retrieved through the same spectral estimation techniques such as Prony's method in Fourier domain when we know at least $2K$ consecutive Fourier coefficients, c_n . To determine $2K + 1$ coefficients [53] propose using a bandlimited filter, $g(t)$, to remove the extra Fourier coefficients outside the range $n \in [-K, \dots, K]$. Alternatively, one may use a more general filter that satisfies the following alias cancellation property [56],

$$\hat{g}(n\omega_0) \begin{cases} \neq 0 & \text{for } n \in \{-K, \dots, K\} \\ = 0 & \text{otherwise} \end{cases} . \quad (2.28)$$

In fact, filters that satisfy this alias cancellation property are precisely those that satisfy the Strang-fix properties Eq. (2.17). Note that in both cases, a slightly suboptimal $2K + 1$ Fourier coefficients is taken so that the filter $g(t)$ remains real. The Fourier coefficients can then be easily estimated from the filtered signal through standard interpolation techniques on uniform or non-uniform samples. This algorithm was extended by [53] to time encoding by noting that the non-uniform samples of

a time encoding generated from a Lazar style IF-TEM and the Fourier coefficients, c_n , are related through a known matrix transformation, $\mathbf{Q}$, derived from the t-transform, Eq. (2.2),

$$\underbrace{\begin{pmatrix} \int_{\tau_0}^{\tau_1} y(t) dt \\ \int_{\tau_1}^{\tau_2} y(t) dt \\ \vdots \end{pmatrix}}_{\mathbf{y}} = \underbrace{\begin{pmatrix} e^{-i\omega_0 K\tau_1} - e^{-i\omega_0 K\tau_0} & e^{-i\omega_0 K\tau_2} - e^{-i\omega_0 K\tau_1} & \dots \\ \vdots & \vdots & \dots \\ e^{-i\omega_0 \tau_1} - e^{-i\omega_0 \tau_0} & e^{-i\omega_0 \tau_2} - e^{-i\omega_0 \tau_1} & \dots \\ \tau_1 - \tau_0 & \tau_2 - \tau_1 & \dots \\ e^{i\omega_0 \tau_1} - e^{i\omega_0 \tau_0} & e^{i\omega_0 \tau_2} - e^{i\omega_0 \tau_1} & \dots \\ \vdots & \vdots & \dots \\ e^{i\omega_0 K\tau_1} - e^{i\omega_0 K\tau_0} & e^{i\omega_0 K\tau_2} - e^{i\omega_0 K\tau_1} & \dots \end{pmatrix}}_{\mathbf{Q}}^T \underbrace{\begin{pmatrix} c_{-K} \\ \vdots \\ c_0 \\ \vdots \\ c_K \end{pmatrix}}_{\mathbf{c}}, \quad (2.29)$$

where $\omega_0 = 2\pi/T$. [53] conjectured and [57] went on to show that $\mathbf{Q}$ is left invertible when $L \geq 2K + 2$ spike times are known. Thus, an FRI signal as defined in Eq. (2.25) with K innovations can be reconstructed from a time encoding if there are at least $2K + 2$ spikes. This can be guaranteed when the parameters of the IF-TEM satisfy [57,58],

$$\frac{b - c}{C} \geq \frac{2K + 2}{T}, \quad (2.30)$$

where $c = \|y * g\|_\infty$ and $c < b$. Whereas the Nyquist-like criteria of Eq. (2.4) & Eq. (2.10) related the trigger mark, C , to the bandwidth of the signal, in this case the trigger mark is constrained by the number of innovations.

2.4 Summary

In this chapter I have surveyed the state of the art time encoding architectures, mainly the integrate-and-fire time encoding machine and the asynchronous delta sigma modulator, as well as the algorithms that are used to perform reconstruction for signals whose samples are time encoded in this way. The iterative approach of Lazar et al. and its derivatives like POCS show that time encodings

can be used to sample and reconstruct bandlimited signals. The work of Rudresh et al. further showed that FRI signals in the form of a train of pulses can be reconstructed when an appropriate filter is used. In particular, they make use of bandlimited filters and filters that satisfy the Strang-Fix (alias cancellation) properties.

I also discussed another quite distinct approach to reconstruction of signals from their time encodings where exponential reproducing filters were exploited to transform the problem into a standard FRI reconstruction problem. These filters could reproduce exponentials because they satisfied the Strang-Fix properties. In Chapter 3, I will build upon this work to show that the α -synaptic filter, a filter that does not satisfy the Strang-Fix conditions, can also be used to reproduce exponentials locally and used for reconstruction.

Chapter 3

Time encoding and recovery of signals with finite rate of innovation

In this chapter I will look at two classes of sampling kernels and reconstruction strategies associated with them. First, I will consider the α -synaptic function and its related kernels. The α -synaptic function provides some benefits due to its greater smoothness in comparison to exponential splines and polynomial splines. The minimum number of samples required for reconstruction will be derived as well as constraints on FRI signals that ensure an integrate-and-fire system will produce the required number of samples and a criterion to discard ‘invalid’ samples. I also show that FRI signals with intrinsic structure, such as power signals, can be reconstructed with the α -synaptic kernel.

The second class of kernels I will consider are hyperbolic kernels, in particular, the derivatives of the hyperbolic secant and hyperbolic secant squared. These kernels extend the smoothness of the α -synaptic to the entire real line, thus eliminating concerns of knotless regions entirely. Reconstruction methods for time encodings derived from hyperbolic sampling kernels will be developed.




Figure 3.1: Two alternative depictions of the prototype integrate-and-fire time encoding machine. **Top:** The representation used in [15]. The spike-triggered reset is shown explicitly without specifying the physical mechanism by which the reset is achieved. **Bottom:** A mathematical representation of the operation of an IF-TEM that defines an intermediate ‘non-reset’ version of the integral, $f(t)$.

These are extended to sampling when the initial conditions of the IF-TEM are unknown and to the case of direct reconstruction from differential samples.

3.1 The prototype IF-TEM

Before proceeding I will first give an explicit and precise definition of the prototypical IF-TEM that we will be considering in this chapter and building upon in later chapters. This is to avoid any confusion with the definitions of time encoding machines discussed in the review of the literature.

The IF-TEM is illustrated in Fig. 3.1 in two equivalent representations. The first is the most intuitive and shows the spike-triggered resetting mechanism of the IF-TEM without specifying its implementation. The operation of the IF-TEM is as follows:

1. The input signal $x(t)$ is filtered with a sampling kernel $\varphi(t)$ to produce the intermediate signal $y(t)$.

2. The filtered signal $y(t)$ is integrated to produce the signal $z(t)$.
3. Each time $z(t)$ exceed the trigger mark, C , it is reset back to zero through the spike-triggered reset.

The latter representation is a more mathematically tractable representation that shows the feedback as an accumulated offset that is subtracted from a ‘non-reset’ integrated signal, $f(t)$. I include this representation mainly to more easily define $f(t)$ for the purposes of analysis. In general, the reader should refer to this representation bearing in mind that a hardware implementation of the system more closely resembles the former representation.

A hardware implementation of a resetting integral would most likely be implemented as an accumulator of charge (a capacitor) which is then discharged each time a reset occurs, thus resembling the first representation. In this sense, the non-reset integral, $f(t)$, present in the latter representation is fictitious and used for analytical convenience.

3.2 The α -synaptic function

As discussed in Section 2.3.2, [15] extended the application of exponential reproduction guaranteed by the Strang-Fix conditions to non-uniformly sampled signals by noting that while global reproduction was no longer guaranteed, reproduction within a specific interval of the real line could still be achieved. In fact, the appeal to the Strang-Fix conditions is not really applicable in this setting since it is a result that is intended to apply to an infinite sum of evenly spaced shifts of a kernel. Indeed, there are other kernels that do not satisfy the Strang-Fix conditions but can still locally reproduce exponentials. In truth, the E-spline kernels allow local reproduction of exponentials because each piece of the spline lies in the same space spanned by a basis of exponentials. Thus, a suitable number of shifts of the same or different pieces that are independent form a basis for the same space of exponentials.

The property of having an interval that lies in a space spanned by exponentials is not unique to kernels which conform to the Strang-Fix conditions. In fact, by considering a broader set of kernels




Figure 3.2: The α -synaptic function with $a = 1$.

that do not conform to the Strang-Fix conditions it is possible to find kernels with more desirable characteristics. In particular, we are interested in kernels which result in simpler constraints on the signal to guarantee reconstruction is possible. An example of a kernel that provides more relaxed restrictions is the α -synaptic kernel which can *locally* reproduce exponentials. In what follows I will demonstrate that the proposed kernel can reproduce exponentials despite not conforming to the Strang-Fix conditions. I will then show how this can be used to sample and reconstruct FRI signals with a time encoding approach. Finally, I provide analysis of the conditions required to guarantee that the samples generated by an IF-TEM are sufficient to reconstruct the signal and a method by which appropriate samples can be selected based on these conditions.

3.2.1 Exponential reproduction with the α -synaptic function

The α -synaptic function originates in the field of neuroscience where it has been used to characterise the shape of a post-synaptic impulse, that is, an action potential [59] and has seen some limited use in the area of signal processing and machine learning [14,60]. It is defined as,

$$\varphi(t) = \begin{cases} te^{-at} & t > 0 \\ 0 & \text{otherwise} \end{cases}, \quad (3.1)$$

where $a \in \mathbb{R}^+$. It is plotted in Fig. 3.2. The main advantage of the α -synaptic function is that it is smooth for $t > 0$ and causal. This is in contrast to the E-spline kernel which has knots at N points of the kernel where N is the degree of the E-spline. In effect, the α -synaptic function has a single knot at $t = 0$. These points will be leveraged in Section 3.3.1 to produce constraints on the input signal that guarantee reconstruction can be achieved and that the necessary samples can be selected.

The kernel does not conform to the Strang-Fix properties. This is evident from the Fourier transform which has no zeros,¹

$$\hat{\varphi}(\omega) = \frac{1}{(a - i\omega)^2}. \quad (3.2)$$

I now show that the exponential reproduction method can be used to reconstruct a signal sampled with an α -synaptic kernel. Consider an FRI signal consisting of a stream of Diracs,

$$x(t) = \sum_{k=1}^K x_k \delta(t - t_k). \quad (3.3)$$

After sampling with an IF-TEM using the α -synaptic kernel the resulting sample times, τ_n , and sample amplitudes, $z_n = C$, are related by,

$$\begin{aligned} z_n &= \sum_{k=1}^K x_k \int_{\tau_{n-1}}^{\tau_n} (t - t_k) e^{-a(t-t_k)} dt \\ &= \sum_{k=1}^K x_k \left(e^{at_k} \underbrace{\frac{(e^{a\tau_{n-1}}(a\tau_{n-1} + 1) - e^{a\tau_n}(a\tau_n + 1))}{a^2}}_{\beta_n} + t_k e^{at_k} \underbrace{\frac{(e^{a\tau_n} - e^{a\tau_{n-1}})}{a}}_{\gamma_n} \right), \end{aligned} \quad (3.4)$$

where τ_n are the spike times. I also assume that $t_k < \tau_n$, that is, every spike time is after the last input

¹In fact, the α -synaptic function is the Fourier dual of a Cauchy wavelet [61].

Dirac. This ensures that the sum contains contributions from each shift of the sampling kernel. As can be seen from Eq. (3.4), for each sample, the summand of the summation is a linear combination of e^{at_k} and $t_k e^{at_k}$. Moreover, given two consecutive inter-spike intervals, we can show that the summands are linearly independent. For linear independence we require that,

$$\begin{pmatrix} \beta_1 \\ \gamma_1 \end{pmatrix} \neq b \begin{pmatrix} \beta_2 \\ \gamma_2 \end{pmatrix}, \quad (3.5)$$

for $0 \neq b \in \mathbb{R}$. Without loss of generality, assume $a = 1$ such that,

$$\begin{aligned} b = \frac{\beta_1}{\beta_2} &= \frac{e^{\tau_1}(\tau_1 + 1) - e^{\tau_0}(\tau_0 + 1)}{e^{\tau_1}(\tau_1 + 1) - e^{\tau_2}(\tau_2 + 1)} \\ &= \frac{\gamma_1}{\gamma_2} = \frac{e^{\tau_0} - e^{\tau_1}}{e^{\tau_1} - e^{\tau_2}}. \end{aligned} \quad (3.6)$$

The equality of Eq. (3.6) between the two lines does not hold for any times $\tau_0 < \tau_1 < \tau_2$. This can be shown by considering the monotonicity of the exponential function. Note, that $e^{\tau_0} < e^{\tau_1} < e^{\tau_2}$. Thus, β_1/β_2 must always have a positive numerator and a negative denominator such that,

$$b = \frac{\beta_1}{\beta_2} < 0. \quad (3.7)$$

Similarly, γ_1/γ_2 must have negative numerator and denominator such that,

$$0 < \frac{\gamma_1}{\gamma_2} = b. \quad (3.8)$$

Thus, there is no value of b that satisfies Eq. (3.5) and therefore the summands of any two consecutive samples are independent. More generally, any two samples consisting of non-overlapping inter-spike intervals are independent. Consequently, a linear combination of two consecutive samples can reproduce the exponential in the sum,

$$s_a = Az_n + Bz_{n+1} = \sum_{k=1}^K x_k e^{at_k}, \quad (3.9)$$

where the reproduction coefficients A and B are given by,

$$\begin{pmatrix} \beta_n & \beta_{n+1} \\ \gamma_n & \gamma_{n+1} \end{pmatrix} \begin{pmatrix} A \\ B \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (3.10)$$

$$\begin{pmatrix} A \\ B \end{pmatrix} = \frac{1}{\beta_n \gamma_{n+1} - \beta_{n+1} \gamma_n} \begin{pmatrix} \gamma_{n+1} \\ -\gamma_n \end{pmatrix}. \quad (3.11)$$

Note, since the signal moment, s_a , consists of two samples, we require three consecutive spike times. The first of these spikes can be viewed intuitively as a spike that resets the IF-TEM to a known initial state where the integral is zero. An example of reproduction of an exponential using three shifts of an appropriate α -synaptic function is given in Fig. 3.3.

3.2.2 Multichannel approach

A set of samples derived from an integrate-and-fire sampling with a single α -synaptic function is capable of reproducing a single exponential of the same exponent, a . In order to reconstruct a signal with $2K$ innovations we must reproduce moments of $2K$ uniformly spaced exponentials. This can be achieved with a multichannel approach where each of the $2K$ channels has an α -synaptic kernel with a different exponent. This architecture is in a similar spirit to that used by [15] where lower order exponential splines were used to reproduce exponentials in separate channels.

Each channel is required to spike at least 3 times to produce 2 samples per channel. Since there are at least $2K$ channels the total number of spikes generated by this architecture is $6K$.

The α -synaptic multichannel approach has some obvious practical problems, most prominently the need to have multiple heterogenous filters in the same system, however, it has several theoretical advantages:




Figure 3.3: Reproduction of the exponential e^{-t} using three shifts of the integral of the α -synaptic function, $\psi(t)$.

Each channel can fire at a lower rate. Since each channel need only fire 3 times, the firing rate of each channel can be lower and spread over a greater interval of the observed signal. This may be useful when implementing samplers that have a limited firing rate. For example, in biological neurons this can be due to refractory periods and the slow propagation velocity of electrical charge. In electronic circuitry, constraints such as slew rate may play similar roles.

Samples in different channels are independent. By this I mean that samples in different channels can independently reproduce exponentials. Because of this, there is no need to estimate the relationship between spikes in different channels or the integrals of the constituent IF-TEMs. Indeed, each channel can have IF-TEMs that are heterogenous in their trigger marks, resetting characteristics, and other non-idealities.

Constraints to guarantee recovery are more tractable. Since each channel reproduces its exponential independently and because each channel's filter contains a single exponential, guarantees that a channel will produce enough spikes become more tractable. This is elaborated on in greater detail in Section 3.3.1. The same analysis is not attempted for the higher order α -synaptic kernels of Section 3.5

Finally, it is worth noting that the reconstruction method described, assumes that the number of pulses, K , is known a priori. If K is not known, multiple values of K up to a maximum can be tried and compared using a model selection criterion to determine the appropriate model order. The simplest criterion can be calculated simply by resynthesising and resampling the estimated signal. By comparing the error between the estimated samples and the true samples one can select the most appropriate value of K .

In a multi-channel system, the maximum K for which a reconstruction can be achieved is restricted by the number of channels or equivalently the number of exponentials that can be reproduced. Thus, when designing such a multi-channel sampler, care must be taken to decide on a maximum value of K beforehand and implement an appropriate number of channels.




Figure 3.4: An FRI signal which consists of pulses with an α -synaptic response ($\alpha = 1$). The Dirac magnitudes have been scaled by 0.5 in the plot for readability.

3.3 Recovery guarantees with the α -synaptic kernel

3.3.1 Constraints to guarantee recovery

We have seen that a certain number of spikes must be generated per channel in order to allow reconstruction of the input signal. In an integrate-and-fire system without bias an arbitrary input is not guaranteed to produce enough spikes. I now analyse constraints on the input signal that guarantee that enough valid spikes are generated to allow us to reconstruct the input signal. The minimum number of channels needed is constrained by the number of innovations or equivalently the number of exponentials that need to be reproduced. This means our analysis can focus on the constraints on a single channel to produce a given number of spikes, N .

In Fig. 3.4, we can see that a signal filtered by an α -synaptic has a knotless region that extends from the position of the last Dirac to the end of the real line. Since this knotless region is more easily defined than in the E-spline case, the problem of determining when spikes are sufficient becomes




Figure 3.5: A two pulse input signal, $x(t)$, filtered with an α -synaptic kernel ($\alpha = 1$) and integrated to get $f(t)$. The corresponding quantity D is denoted with a red brace.

more tractable. Moreover, because the kernel is causal, this is the only region that has contribution from all input Diracs and thus, if simultaneous reconstruction of all Diracs is desired, spikes necessarily must be taken from after the last Dirac. Another way of expressing this is that spikes must occur within the support of the kernel otherwise Eq. (3.4) does not hold and exponential reproduction cannot be achieved. We call spikes that occur after the last Dirac *valid* spikes, because they can be used for reconstruction, and those that occur before as *invalid* spikes.

We would like to characterise the number of valid spikes generated for a given input signal and trigger mark C . Furthermore, we would like these valid spikes to be generated within a t^* -long interval following the last input Dirac. This is to ensure we don't have to wait too long for samples and to increase stability of the reconstruction. To do this I first derive an expression for the difference between the integral of the signal at the position of the last Dirac, $f(t_0)$, and a later time $f(t_0 + t^*)$. Note that for convenience I have reindexed and reversed the typical indexing of the input Diracs such that t_{K-1} is the time of the first Dirac and t_0 is the time of the last Dirac. I call the difference in the integral,

$$D = f(t_0 + t^*) - f(t_0) . \quad (3.12)$$

Since $f(t)$ is a continuous function it must take any given value somewhere in the interval $[f(t_0), f(t_0 + t^*)]$ for some t . Therefore, if D is non-zero we must be able to find a trigger mark C that is small enough to generate M spikes. The number of spikes generated is equal to the number of times the integral crosses a multiple of the trigger mark. Note that due to the architecture of a bipolar IF-TEM, the integral may take a value close to the negative trigger mark, that is $f(t_0) = \epsilon - C$ where ϵ is arbitrarily small. In such a circumstance, to cross the positive trigger mark it must increase by $2C$. An example of this degenerate case of detecting level sets is plotted in Fig. 3.6. The level crossing at $t = 2.5$ is not detected because the integral is reset to zero and not the negative trigger mark. As a consequence the number of samples generated is,

$$M \geq \left\lfloor \frac{D - C}{C} \right\rfloor . \quad (3.13)$$

Let $\hat{M}$ be the desired number of spikes then the criterion by which we can set C such that at least $\hat{M}$ spikes are generated is,

$$C \leq \frac{D}{\hat{M} + 1} . \quad (3.14)$$

We can verify this by substitution into Eq. (3.13),

$$M \geq \left\lfloor \frac{D - C}{C} \right\rfloor \geq \hat{M} . \quad (3.15)$$

The number of spikes is dependent on the value of D , which can be arbitrarily large or small for an unconstrained input signal. We constrain our input signal by enforcing upper and lower bounds on the amplitude innovation parameters, x_k , and a minimum bound on the separation between any two consecutive Diracs. The latter constraint is important since without this constraint, the signal can be made arbitrarily small by closely separated but opposite sign Diracs. Under these constraints,




Figure 3.6: A degenerate case of detecting the level set. The filtered input to the IF-TEM is $y(t)$ is plotted scaled by 0.2 for readability. The IF-TEM resets and generates a spike each time the integral of this signal, $f(t)$, crosses a multiple of the trigger mark. The correct detection of these level sets is marked in green. The level crossing marked in red is not detected since the reset integral $z(t)$ has not crossed a trigger mark.

the quantity D can be constrained to be non-zero as summarised in Proposition 1. Given $D \neq 0$, one can set an appropriate trigger mark, C , by Eq. (3.15).

Proposition 1. *Let $x_{\min}$ and $x_{\max}$ be the minimum and maximum absolute values of the amplitude innovation parameters such that the amplitudes of the input Diracs satisfy $x_{\min} \leq |x_k| \leq x_{\max}$ and let $\Delta_{\min}$ be the minimum separation between any two input Diracs such that $\Delta_{\min} \leq t_k - t_{k+1}$. The following are sufficient conditions such that $D \neq 0$:*

$$1 > \frac{x_{\min}}{x_{\max}} > \sum_{k=1}^{K-1} e^{-ak\Delta_{\min}} (ak\Delta_{\min} + 1 - e^{-at^*} (ak\Delta_{\min} + at^* + 1)) , \quad (3.16)$$

and

$$\frac{t^* e^{-at^*}}{1 - e^{-at^*}} < \Delta_{\min} . \quad (3.17)$$

Proof. In what follows I will first prove that Eq. (3.16) ensures $D \neq 0$ under the assumption that the input Diracs are uniformly spaced with a separation of Δ between any two consecutive Diracs. The extension to non-uniform spacing with $\Delta_{\min}$ minimum separation is then simple. First we derive an expression for the integral at $t \geq t_0$:

$$\begin{aligned} f(t) &= \int_{-\infty}^t y(t) dt = \sum_{k=0}^{K-1} x_k \int_{-\infty}^t \varphi(t - t_k) dt \\ &\stackrel{(a)}{=} \sum_{k=0}^{K-1} x_k \int_0^t (t - t_k) e^{-a(t-t_k)} dt \\ &\stackrel{(b)}{=} \sum_{k=0}^{K-1} \frac{x_k}{a^2} (e^{at_k} (1 - at_k) - e^{-a(t-t_k)} (a(t - t_k) + 1)) , \end{aligned} \quad (3.18)$$

where (a) follows from the $\varphi(t)$ being causal and (b) follows by integration by parts. Substituting into Eq. (3.12) we obtain an expression in terms of Δ and t^* ,

$$\begin{aligned}
D &= \frac{1}{a^2} \sum_{k=0}^{K-1} x_k \left(e^{-a(t_0-t_k)} (a(t_0-t_k) + 1) - e^{-a(t_0+t^*-t_k)} (a(t_0+t^*-t_k) + 1) \right) \\
&\stackrel{(a)}{=} \frac{1}{a^2} \sum_{k=0}^{K-1} x_k \underbrace{e^{-ak\Delta} (ak\Delta + 1 - e^{-at^*} (ak\Delta + at^* + 1))}_{\xi_k} \\
&= \frac{1}{a^2} \left(x_0 + \sum_{k=1}^{K-1} x_k \xi_k \right), \tag{3.19}
\end{aligned}$$

where (a) follows from assuming, without loss of generality, that the first Dirac is at $t_{K-1} = 0$, and therefore $t_0 = (K-1)\Delta$ and $t_k = (K-1-k)\Delta$. We can say that $D \neq 0$ if the following is satisfied:

$$|x_0| > \left| \sum_{k=1}^{K-1} x_k \xi_k \right|. \tag{3.20}$$

Appealing to the triangle inequality, we can derive a constraint that ensures Eq. (3.20) is satisfied:

$$|x_0| > \sum_{k=1}^{K-1} |x_k| \xi_k > \left| \sum_{k=1}^{K-1} x_k \xi_k \right|,$$

which holds provided:

$$\begin{aligned}
x_{\min} &> \sum_{k=1}^{K-1} x_{\max} \xi_k \\
\frac{x_{\min}}{x_{\max}} &> \sum_{k=1}^{K-1} \xi_k. \tag{3.21}
\end{aligned}$$

The constraint in equation Eq. (3.21) can be extended to the case of non-uniform spacing with a minimum separation, $\Delta_{\min}$, when ξ_k is monotonically decreasing. Consider the derivative of ξ_k with respect to Δ ,

$$\frac{d\xi_k}{d\Delta} = ake^{-ak\Delta} (ak\Delta(e^{-at^*} - 1) + ate^{-at^*}). \tag{3.22}$$

One can observe that ξ_k is monotonically decreasing with Δ when,

$$ake^{-ak\Delta} (ak\Delta(e^{-at^*} - 1) + ate^{-at^*}) < 0$$

$$\frac{t^*e^{-at^*}}{1 - e^{-at^*}} < k\Delta, \quad (3.23)$$

where I have used the fact that the quotient of Eq. (3.22) is always positive. Taking the most strictest constraint of Eq. (3.23), that is for $k = 1$ and $\Delta = \Delta_{\min}$, we have a condition for ξ_k to be monotonically decreasing in Eq. (3.17). Under this condition, the separation of any two Diracs can be increased and Eq. (3.21) still holds.

□

These conditions simplify when we let $t^* \rightarrow \infty$. Indeed, Eq. (3.17) is satisfied since we only need $0 < \Delta_{\min}$ and Eq. (3.21) simplifies to,

$$\frac{x_{\min}}{x_{\max}} > \sum_{k=1}^{K-1} e^{-ak\Delta_{\min}} (ak\Delta_{\min} + 1). \quad (3.24)$$

This asymptotic constraint ensures that any IF-TEM will fire some time after t_0 . The constraint is plotted for different values of K in Fig. 3.7. Each trace plots the lower bound on the ratio between minimum and maximum Dirac amplitude. Each new input Dirac makes the constraint more stringent but with diminishing returns. This is because as the number of Diracs increases, the contribution from the first Diracs becomes negligible.

The quantity $|D|$ in Eq. (3.19) can also be minimised by substituting $x_0 = -x_{\min}$ and $x_{i \neq 0} = x_{\max}$. This can then be used to determine the appropriate trigger mark according to Eq. (3.15) to generate a given number of samples within the interval $[t_0, t_0 + t^*]$.

An example of the samples generated using these constraints is plotted in Fig. 3.8. In this example it can be observed that these constraints tend to underestimate the number of samples generated. This is because we assumed that the integral, $f(t)$, only passes through each multiple of C once. This is often not the case as can be seen from the generation of both positive and negative spikes after t_0 .




Figure 3.7: Comparison of the constraint on ratio $x_{\min}/x_{\max}$ with respect to $\Delta_{\min}$ imposed for different numbers of Diracs, K , where $a = 1$ and $t^* \rightarrow \infty$. The ratio must be higher than the lower boundary curves, which is plotted for different K .

3.3.2 Selection of valid samples for reconstruction

The output of an integrate-and-fire neuron consists of a sequence of samples. The first of these, up to the position of the last Dirac, t_0 , are invalid and cannot be used in reconstruction. The rest are valid samples which can be used in reconstruction. By imposing a final maximum constraint on the separation of any two consecutive Diracs, one can determine a rule by which M valid samples can always be selected for reconstruction.

Proposition 2. *Let $\Delta_{\max}$ be the maximum separation between consecutive Diracs such that the first and last Diracs are separated by at most $(K-1)\Delta_{\max}$, that is, $t_0 - t_{K-1} \leq (K-1)\Delta_{\max}$. Let $x_{\min}$, $x_{\max}$, $\Delta_{\min}$, and C be calculated by Proposition 1 and Eq. (3.15) such that at least $\hat{M}$ samples are generated in the interval $[t_0, t_0 + t^*]$. The sequence of samples with times $T = (\tau_0, \dots, \tau_N)$ is generated by an integrate-and-fire neuron with trigger mark C . The last $\hat{M}$ samples in the interval $[\tau_0, \tau_0 + L]$ are valid samples that can be used for reconstruction where $L = (K-1)\Delta_{\max} + t^*$.*

Proof. First, note that if $\hat{M}$ valid samples are generated in the interval $[a, b]$ then at least $\hat{M}$ valid samples are generated in the interval $[a, b + \epsilon]$, for $\epsilon > 0$. The conditions in Proposition 1 and Eq. (3.15) guarantee that at least $\hat{M}$ valid samples are generated before $t_0 + t^*$. Moreover, since $t_0 \leq t_{K-1} + (K - 1)\Delta_{max}$, at least $\hat{M}$ valid samples are generated before $t_{K-1} + (K - 1)\Delta_{max} + t^*$. Finally, by noting that the first sample, τ_0 , occurs after the first Dirac, t_{K-1} , we know that at least $\hat{M}$ valid samples are generated before $\tau_0 + (K - 1)\Delta_{max} + t^*$, or in other words in the interval $[\tau_0, \tau_0 + L]$. □

Leveraging Proposition 2, a set of M consecutive valid samples can be determined from the output of an integrate-and-fire neuron by selecting the last M samples in the interval $[\tau_0, \tau_0 + L]$. In general, we require $M = 3$. In this manner, the conditions derived in this section allow us not only to guarantee valid samples are generated but also that the valid samples can be determined in a simple and efficient manner. Consequently perfect reconstruction of the input Diracs is always possible.

In Fig. 3.8, a sufficient number of valid samples can be seen to be generated. Furthermore, at least three samples occur before $\tau_0 + L$ (marked by a red cross). This being the case, it is simple to select three valid samples with which to reproduce an exponential.

Bursts of pulses

Finally I note that the method for sampling and reconstructing a finite number of Diracs described above can be extended to the case of sampling and reconstructing a stream of bursts of Diracs. In this context, I define a burst of Diracs to be a set of Diracs who satisfy the constraints on separation Δ_{min} and Δ_{max} as well as the constraints on the magnitude of the Diracs, x_{min} and x_{max} . If the first Dirac of each burst is separated from the next burst by at least $L = (K - 1)\Delta_{max} + t^*$ then we know that enough samples are generated by a burst before the subsequent burst to allow reconstruction of the first burst. The influence of the first burst can then be removed from the samples of the subsequent burst. In this way bursts can be sequentially reconstructed.

Furthermore, since the alpha synaptic function decays to zero, if t^* is sufficiently large, the Diracs of two bursts will be sufficiently separated such that the influence of the first burst on the samples




Figure 3.8: Two examples of the unreset integral, $f(t)$, and the samples generated by integrate-and-fire neuron where $a = 1$ and $t^* = 5$. The input Diracs are constrained such that $\Delta_{\min} = 2.5$, $\Delta_{\max} = 4$ and $x_{\min}/x_{\max} = 0.589$. The trigger mark C is set so that at least 3 valid samples are generated. The position $\tau_0 + L$ is marked by a red cross. Samples have been scaled up for readability. Left: case for $\Delta = \Delta_{\max}$, Right: case for $\Delta = \Delta_{\min}$.

generated by the subsequent burst is negligible. In this case it can be shown empirically that each burst can be reconstructed independently with good accuracy.

3.4 Piecewise constant and polynomial signals by sampling the derivative

As in previous works, sampling and reconstruction from the α -synaptic function can be extended to piecewise constant and polynomial signals by modifying the kernel. We can reconstruct piecewise constant signals by using the simple approach of sampling with derivative of the sampling kernel. The derivative of the α -synaptic function as plotted in Fig. 3.9 is,

$$\frac{d\varphi}{dt}(t) = e^{-at}(1 - at) . \quad (3.25)$$




Figure 3.9: The derivative of the α -synaptic function $\alpha = 1$.

The filtered signal, $y(t)$, before integration can be expressed as,

$$y(t) = \left(x * \frac{d\varphi}{dt} \right) (t) \stackrel{(a)}{=} \left(\frac{dx}{dt} * \varphi \right) (t) , \quad (3.26)$$

where (a) follows from integration by parts. In this way, our samples are equivalent to samples of the derivative of the input signal. If $x(t)$ is a piecewise constant signal expressed as a superposition of Heaviside functions, $u(t)$, then,

$$x(t) = \sum_{k=1}^K x_k u(t - t_k) \Rightarrow \frac{dx}{dt}(t) = \sum_{k=1}^K x_k \delta(t - t_k) . \quad (3.27)$$

The innovation parameters can then be estimated in the same manner as before using spectral estimation techniques like Prony's method followed by least squares approximation.

3.5 Higher order α -synaptic kernel

Throughout this section I have assumed the use of a multichannel system where each channel is configured to allow reproduction of a single exponential. This affords the advantages of multichannel system, however, the realisation of such a system may not be practical. Each channel requires its own integrator, reset circuitry, and heterogenous filter. An alternative or complimentary approach is to employ higher order α -synaptic kernels. That is we consider a kernel constructed as the sum of multiple α -synaptic functions with different exponents,

$$\varphi(t) = t \left(\sum_{r=1}^R c_r e^{\alpha_r t} \right), \quad (3.28)$$

where R is the order of the α -synaptic kernel and $c_r \in \mathbb{R}$. Appealing to Eq. (3.4) one can see that the integral of shifts of the kernel lie in a space spanned by,

$$\int_{\tau_{n-1}}^{\tau_n} \varphi(t - t_k) dt \in \text{span}\{t_k e^{a_1 t_k}, e^{a_1 t_k}, t_k e^{a_2 t_k}, e^{a_2 t_k}, \dots, t_k e^{a_R t_k}, e^{a_R t_k}\}. \quad (3.29)$$

Similar to the first order case, given $2R$ samples ($2R + 1$ spikes), one can reproduce R different exponentials that lie in the given basis, however, this is dependent on the assumption that the samples are independent like in the first order kernel case. The proof of linear independence of Eq. (3.6) to Eq. (3.8) does not apply. I supply no proof for the higher order case except to say that exponential reproduction is successful in an empirical setting. An example of exponential reproduction with a second order alpha-synaptic function is given in Fig. 3.10.

If the signal being sampled has $2K$ innovations, then we must reproduce $2K$ exponentials to estimate the innovation parameters, therefore, a multichannel sampler using a R^{th} order kernel must have at least $L = 2K/R$ channels where each must generate at least $2R + 1$ spikes. The minimum number of spikes required to allow reconstruction is therefore,

$$\frac{2K}{R}(2R + 1) = 4K + \frac{2K}{R}. \quad (3.30)$$




Figure 3.10: Reproduction of two different exponentials with the integral of the same second order α -synaptic function, $t(e^{-\frac{3t}{2}} + e^{-3t})$. In the plot, $\frac{d}{dt}\psi(t) = \varphi(t)$.

Increasing the order of the kernel reduces the minimum number of spikes required to reconstruct the signal. In the extreme case of a mono-channel sampler with a kernel of order $R = 2K$, the number of spikes required to reconstruct the signal is $4K + 1$. The number of samples required by a mono-channel system is much closer to the optimum number of samples unlike that of the fully multichannel system where $6K$ spikes are required. The reason for this is because in the multichannel setting each channel uses a unique set of three spikes. In the mono-channel system, each set of three spikes overlap by a single spike.

3.6 Signals with structure - recovery of power signals

In some cases, it may be of interest to reconstruct FRI signals which have a structure beyond the typical linear sum. An example of this is the power signal of an FRI signal. By power signal I mean a signal that is the point wise square of an underlying signal,

$$p(t) = \left| \sum_{k=1}^K x_k \varphi(t - t_k) \right|^2 = \sum_{k=1}^K \sum_{l=1}^K x_k x_l \varphi(t - t_k) \varphi(t - t_l) . \quad (3.31)$$

The power signal is also an FRI signal. Of course, its rate of innovation, $2K$, remains unchanged, however, one can model such a signal as having K^2 innovation parameters, $\chi_k \in \{x_l x_m | l, m \in 1, \dots, K\}$ and $\xi_k \in \{t_l + t_m | l, m \in 1, \dots, K\}$. This allows an alternative approach to reconstruction where the χ_k and ξ_k parameters are first estimated before the x_k and t_k are subsequently inferred.

Sampling such a power signal in the context of integrate-and-fire may yield some advantages. Foremost, unipolar IF-TEMs can be used without having to resort to an input bias that results in constant firing even when there is no signal to observe. Since the power signal must always be positive, no negative trigger mark is required. Secondly, since the power signal is positive, the integral must be monotonic, therefore, the IF-TEM will always come closer to firing. It cannot remain quiescent indefinitely even when the input signal is small. In this way, the pointwise squaring operation may be an alternative to the bias typically applied to the input of an IF-TEM. The squaring operation

itself can be implemented in hardware through the use of mixers or log amplifiers. In some cases measuring the power of a process directly may also be necessary.

We shall see that the α -synaptic kernel allows us to reproduce exponentials as before. The resulting signal moments can be used to recover the original power signal, or equivalently the signal up to a sign change. Consider the integrate-and-fire samples generated from the power signal in Eq. (3.31) for channel m of a multi-channel sampler (as before, each channel reproduces a different exponential):

$$z_n^{(m)} = C = \int_{\tau_{n-1}}^{\tau_n} p(t) dt = \int_{\tau_{n-1}}^{\tau_n} \sum_{k=1}^K \sum_{l=1}^K x_k x_l \varphi(t - t_k) \varphi(t - t_l) dt. \quad (3.32)$$

Substituting the expression for the α -synaptic function into Eq. (3.32) yields,

$$\begin{aligned} z_n^{(m)} &= \int_{\tau_{n-1}}^{\tau_n} \sum_k \sum_l x_k x_l (t - t_k) e^{-\alpha_m(t-t_k)} (t - t_l) e^{-\alpha_m(t-t_l)} dt \\ &= \sum_k \sum_l x_k x_l e^{\alpha_m(t_k+t_l)} \int_{\tau_{n-1}}^{\tau_n} e^{-2\alpha_m t} (t^2 - (t_k + t_l)t + t_k t_l) dt. \end{aligned} \quad (3.33)$$

Similar to the case with the non-power signal, we can calculate the square of a signal moment by taking a weighted sum of samples, $z_n^{(m)}$:

$$s_m^2 = \sum_n c_{m,n} z_n^{(m)} = \sum_k \sum_l x_k x_l e^{\alpha_m(t_k+t_l)} \underbrace{\sum_n c_{m,n} \int_{\tau_{n-1}}^{\tau_n} e^{-2\alpha_m t} (t^2 - (t_k + t_l)t + t_k t_l) dt}_{I_n}, \quad (3.34)$$

where,

$$I_n = \underbrace{\int_{\tau_{n-1}}^{\tau_n} t^2 e^{-2\alpha_m t} dt}_{\gamma_{1,n}} - \underbrace{(t_k + t_l)}_{a_{k,l}} \underbrace{\int_{\tau_{n-1}}^{\tau_n} t e^{-2\alpha_m t} dt}_{\gamma_{2,n}} + \underbrace{t_k t_l}_{b_{k,l}} \underbrace{\int_{\tau_{n-1}}^{\tau_n} e^{-2\alpha_m t} dt}_{\gamma_{3,n}}, \quad (3.35)$$

where $\gamma_{1,n}, \gamma_{2,n}, \gamma_{3,n}$ are known functions of n and where $a_{k,l}$ and $b_{k,l}$ are unknown constants. To

reproduce the exponentials we must find $c_{m,n}$ such that the innermost summation of Eq. (3.34) to equals one,

$$\sum_n c_{m,n} (\gamma_{1,n} + \gamma_{2,n} a_{k,l} + \gamma_{3,n} b_{k,l}) = 1 . \quad (3.36)$$

This can be represented simply in matrix form:

$$\Gamma \mathbf{c}_m = \begin{pmatrix} 1 \\ a_{k,l} \\ b_{k,l} \end{pmatrix} , \quad (3.37)$$

where $\{\Gamma\}_{i,j} = \gamma_{i,j}$ and $\{\mathbf{c}_m\}_i = c_{m,i}$. Hence, by solving Eq. (3.37), we can determine coefficients, $c_{m,n}$, such that:

$$s_m^2 = \sum_k \sum_l x_k x_l e^{\alpha_m(t_k+t_l)} . \quad (3.38)$$

This can be achieved if we know the values of $\gamma_{i,n}$ for at least 3 values of n . Since each $\gamma_{i,n}$ is determined from two consecutive sample times, τ_n and τ_{n-1} , we need at least 4 consecutive samples to determine Γ (though this may include the first “sample” at $\tau_0 = -\infty$).

3.6.1 Determining Dirac times

Normally, when determining Dirac times from signal moments, we use Prony’s method to directly retrieve the Dirac times, $\{t_k\}$, in one step. This is not possible in this case because we lose the sign of the signal moments, however, we can retrieve the times of a related problem. Collecting like terms of the summation in Eq. (3.38) we can define a new problem which can be solved in the normal way:

$$s'_m = s_m^2 = \sum_{i=1}^N \chi_i e^{\alpha_m \xi_i} \quad \chi_i \in \{x_k x_l \mid k \leq l\} \quad \xi_i \in \{t_k + t_l \mid k \leq l\} , \quad (3.39)$$

where $N = K(K + 1)/2$ and thus requires $K(K + 1)$ signal moments to solve via Prony's method. Applying Prony's method we can retrieve the times, $\{\xi_i\}$. At this point we need to match each ξ_i with its corresponding sum of t_k 's. In the following I show that there is terminating algorithm that can perform this task. First I consider the ordered of the set $\{\xi_i\}$ such that if $\xi_i < \xi_j$ then $i < j$. Also note the following inequality on the Dirac times:

$$2t_1 < t_1 + t_2 < t_i + t_j \quad \text{for} \quad i, j > 2. \quad (3.40)$$

From Eq. (3.40) and the ordering of $\{\xi_i\}$ we can determine values for t_1 and t_2 ,

$$\xi_1 = 2t_1 \quad \xi_2 = t_1 + t_2. \quad (3.41)$$

For $K \leq 4$ it is possible to determine all values of t_k by considering a similar inequality to Eq. (3.40) that enforces the ordering of the top two elements of $\{\xi_i\}$. However, for completeness, we can show that the values of t_k can be determined for any value of K . Consider all the possible pair-wise combinations of t_k that comprise all the elements of $\{\xi_i\}$ arranged as the upper triangular elements of a matrix:

$$\Xi = \begin{pmatrix} \mathbf{2t_1} & \mathbf{t_1 + t_2} & t_1 + t_3 & \dots & t_1 + t_K \\ & 2t_2 & t_2 + t_3 & \dots & t_2 + t_K \\ & & 2t_3 & \dots & t_3 + t_K \\ & & & \ddots & \vdots \\ & & & & 2t_K \end{pmatrix}. \quad (3.42)$$

The first two elements, marked in bold, of Ξ can be determined to be ξ_1 and ξ_2 respectively as previously shown using Eq. (3.40). The rest of the matrix can then be determined algorithmically by noting the following properties:

1. Each element is smaller than all elements to the right and below. That is $\{\Xi\}_{i,j} \leq \{\Xi\}_{m,n}$ if $i \leq m$ and $j \leq n$

2. If we know the top element in a column and the diagonal elements to its left then we can determine all the elements in the column. That is if we know the elements $\{\Xi\}_{1,j}$ for some fixed j and $\{\Xi\}_{m,m}$ for $m < j$ then we can determine the values of $\{\Xi\}_{i,j}$

Using property 2 we can immediately determine the missing value in the second column, $2t_2$. Now, excluding the three known values from the set $\{\xi_i\}$ we can now determine the top element of column 3:

$$\{\Xi\}_{1,3} = t_1 + t_3 = \min \{\xi_i\} \setminus \{2t_1, t_1 + t_2, 2t_2\} , \quad (3.43)$$

where $A \setminus B$ indicates the set of elements of A not in B . Eq. (3.43) follows directly from property 1 since we exclude all values to the left of $\{\Xi\}_{1,3}$ from the set $\{\xi_i\}$. By alternately applying the two properties and Eq. (3.43) it can be seen that the entire matrix can be determined. The Dirac times can then be extracted from the diagonal entries.

3.6.2 Determining Dirac Amplitudes, $\{x_k\}$

The final step in reconstruction is to determine the values of x_k . This is done by retrieving the Gram matrix corresponding to the values of x_k . Let the column vector $\mathbf{x} = \{x_1, x_2, \dots, x_K\}$ and let the column vector $\mathbf{a}_m = \{e^{\alpha_m t_1}, e^{\alpha_m t_2}, \dots, e^{\alpha_m t_K}\}$, the squared moments in Eq. (3.38) can now written as quadratic matrix form.

$$\begin{aligned} s_m^2 &= \mathbf{a}_m^T X \mathbf{a}_m \\ &= \text{vec}(\mathbf{a}_m \mathbf{a}_m^T)^T \text{vec}(X) , \end{aligned} \quad (3.44)$$

where X is the Gram matrix $\mathbf{x} \mathbf{x}^T$ and where $\text{vec}(\cdot)$ converts a matrix into a column vector. Since X is symmetric, given $K(K+1)/2$ such equations we can solve for the matrix X . The magnitudes $|x_k|$

can then be calculated from the squared roots of the diagonal entries of X and their signs (up to a global sign) can be determined from the signs of any row or column of X .

3.7 Reconstruction with hyperbolic and rational kernels

The α -synaptic kernel does not fully resolve the issue of samples having to occur in certain known regions to be used for reconstruction. Although it is smooth for $t > 0$, the knot at $t = 0$ means that some consideration for where samples occur still has to be made as demonstrated in Section 3.3.1. In fact, these constraints can be very stringent due to the front loaded shape of the kernel as could be seen in Section 3.3.1.

In this section I will consider an alternative class of kernels, namely, hyperbolic functions, and an alternative, but related, reconstruction strategy based on rational approximation. This class of kernels has the advantage of being smooth across the entire real line whilst still yielding signals that have finite rate of innovation and can be reconstructed using rational approximation techniques. An example of a hyperbolic kernel, $\varphi(t) = \text{sech}(t)$, is illustrated in Fig. 3.11 along with its derivative.

The approach to reconstruction is largely based on [62] in which a rational approximation approach was used to achieve reconstruction of FRI signals from non-uniform samples. I will first review this method for a non-uniformly sampled signal. I will then consider how the integrate-and-fire samples can be related to these samples, how many samples are required, and how the choice of hyperbolic kernel affects these. Consider an FRI signal of the form,

$$x(t) = \sum_{k=1}^K x_k \delta(t - t_k) . \quad (3.45)$$

The output of a non-uniform sampling process with kernel $\varphi(t)$ is:

$$y_n = y(\tau_n) = \sum_{k=1}^K x_k \varphi(\tau_n - t_k) , \quad (3.46)$$

where τ_n are the sampling times. The ability to reconstruct a signal sampled in this manner depends




Figure 3.11: The hyperbolic secant and its derivative.

on the choice of sampling kernel, $\varphi(t)$. For demonstration purposes, let us consider the hyperbolic secant kernel. The kernel can be expressed in terms of exponentials, as follows:

$$\varphi(t) = \text{sech}(at) = \frac{2e^{at}}{1 + e^{2at}}, \quad (3.47)$$

where $a \in \mathbb{R}$. Substituting Eq. (3.47) into Eq. (3.46) shows that when $\varphi(t) = \text{sech}(at)$ each sample is the quotient of two polynomials evaluated at τ_n :

$$\begin{aligned} y_n = y(\tau_n) &= \sum_{k=1}^K x_k \text{sech}(a\tau_n - at_k) \\ &= \sum_{k=1}^K \frac{2x_k e^{-at_k} e^{a\tau_n}}{1 + e^{-2at_k} e^{2a\tau_n}} \\ &= \frac{e^{a\tau_n} P(e^{2a\tau_n})}{Q(e^{2a\tau_n})}, \end{aligned} \quad (3.48)$$

where $P(x)$ and $Q(x)$ are polynomials of order $K - 1$ and K respectively. Furthermore, information of the locations, t_k , of the input Diracs is contained in the $Q(x)$ polynomial through the factorisation:

$$Q(x) = \prod_{k=1}^K (1 + e^{-2at_k}x) . \quad (3.49)$$

By multiplying both sides of Eq. (3.48) by $Q(x)$ we can form the following system of linear equations:

$$Q(e^{2a\tau_n})y_n = e^{a\tau_n}P(e^{2a\tau_n}) , \quad (3.50)$$

or,

$$y_n \sum_{k=0}^K q_k e^{2ka\tau_n} = \sum_{k=0}^{K-1} p_k e^{(2k+1)a\tau_n} , \quad (3.51)$$

where q_k and p_k are the coefficients of the polynomials $Q(x)$ and $P(x)$ respectively. This can be expressed in the following matrix form:

$$\underbrace{\begin{bmatrix} y_1 & 0 & \dots & 0 \\ 0 & y_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & y_N \end{bmatrix}}_{\mathbf{Y}} \underbrace{\begin{bmatrix} 1 & e^{2a\tau_1} & \dots & e^{2Ka\tau_1} \\ 1 & e^{2a\tau_2} & \dots & e^{2Ka\tau_2} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & e^{2a\tau_N} & \dots & e^{2Ka\tau_N} \end{bmatrix}}_{\mathbf{V}} \begin{bmatrix} q_0 \\ q_1 \\ \vdots \\ q_K \end{bmatrix} = \underbrace{\begin{bmatrix} e^{a\tau_1} & e^{3a\tau_1} & \dots & e^{(2K-1)a\tau_1} \\ e^{a\tau_2} & e^{3a\tau_2} & \dots & e^{(2K-1)a\tau_2} \\ \vdots & \vdots & \ddots & \vdots \\ e^{a\tau_N} & e^{3a\tau_N} & \dots & e^{(2K-1)a\tau_N} \end{bmatrix}}_{\mathbf{W}} \begin{bmatrix} p_0 \\ p_1 \\ \vdots \\ p_{K-1} \end{bmatrix} ; \quad (3.52)$$

This resembles the form of a Padé approximant except that in this case we are performing an interpolation on a set of non-uniformly spaced points, e^{τ_n} , and not equating coefficients. The system can be solved up to a multiplicative constant when $N \geq 2K$. Using the coefficients of $Q(x)$ we can then compute its roots, ζ_k . The locations of the input Diracs are then retrieved using,

$$t_k = -\frac{\ln(-\zeta_k)}{2a} , \quad (3.53)$$

where we note that the roots, ζ_k , are necessarily negative due to Eq. (3.49). By substituting values of t_k and τ_n into Eq. (3.46) we can now form a system of linear equations, which can be solved using a least squares approach to retrieve the magnitudes, x_k , of the input Diracs. The amplitudes, x_k , can also be retrieved directly without a least squares approach as detailed in Chapter A.

Finally, we note that the system in Eq. (3.52) is always solvable. This property was not explicitly stated or proved by [62] and so I give the following brief proof. Assuming the case using a minimum number of samples $N = 2K$, the product $\mathbf{YV}$ has $2K$ rows and $K + 1$ columns. Therefore, in order to have full rank, a maximum of $K - 1$ rows of $\mathbf{YV}$ may be zero. The number of zero rows corresponds to the number of cumulative samples, y_n , that are zero. These samples are equivalent to the samples of a sum of scaled and shifted hyperbolic secants in Eq. (3.48). From Eq. (3.48) we can see that $y(t)$ has the same zeros as the polynomial $P(x)$. This polynomial is of order K and therefore has $K - 1$ roots. Consequently, assuming that each sample is unique, only $K - 1$ samples can be zero. $\mathbf{YV}$ is thus guaranteed to be full rank. The argument extends to $N = 2K + L$ for $L > 1$ where at most $K - 1 + L$ samples can be zero.

3.7.1 Time encoding and reconstruction using an integrate-and-fire system

The non-uniform sampling algorithm due to [62] described in the previous section can be used in the context of integrate-and-fire sampling with the prototypical time encoder. First, consider the samples generated from an FRI signal by the prototypical time encoder,

$$z_n = z(\tau_n) = \sum_k x_k \int_{\tau_{n-1}}^{\tau_n} \varphi(t - t_k) dt . \quad (3.54)$$

Note that τ_0 represents the time that the IF-TEM starts integrating from. For now I will assume that the IF-TEM started integrating from the $t = -\infty$ with an initial integral of $f(t) = 0$ but this will be examined in a later section. We would like to map samples of the form in Eq. (3.54) back to the non-uniform sampling setting. Currently, the samples each represent a finite difference between

two consecutive spike times, τ_{n-1} and τ_n . To achieve this we can take a cumulative sum of samples to generate a new set of *augmented* samples,

$$\begin{aligned}
y_n &= \sum_{m=1}^n z_m = \sum_{m=1}^n \sum_{k=1}^K x_k \int_{\tau_{n-1}}^{\tau_n} \varphi(t - t_k) dt \\
&= \sum_{k=1}^K x_k \int_{\tau_0}^{\tau_n} \varphi(t - t_k) dt \\
&= \sum_{k=1}^K x_k (\psi(\tau_n - t_k) - \psi(\tau_0 - t_k)) ,
\end{aligned} \tag{3.55}$$

where $\psi(t)$ is the integral of $\varphi(t)$. The augmented samples, y_n , now have a common term related to τ_0 . Given our assumption that the IF-TEM has been integrating from the start of time and its initial integral is zero, this term is eliminated leaving a non-uniform sampling of an FRI signal with kernel $\psi(t)$,

$$y_n = \sum_{k=1}^K x_k \psi(\tau_n - t_k) . \tag{3.56}$$

Clearly, the innovation parameters x_k and t_k can be determined with the reconstruction strategy reviewed previously in the case where we let $\psi(t) = \text{sech}(at)$. The choice of sampling kernel, $\varphi(t)$, to facilitate this form of reconstruction is then the derivative of the hyperbolic secant,

$$\varphi(t) = \frac{d}{dt} \psi(t) = -a \text{sech}(at) \tanh(at) . \tag{3.57}$$

Given this sampling kernel, we can reconstruct the input signal given at least $2K$ augmented samples generated from the cumulative sum of $2K$ samples of the form Eq. (3.54). An example sampling and reconstruction using a time encoder with the hyperbolic kernel is plotted in Fig. 3.12. The filtered signal is smooth with no knots and so all spikes, regardless of position, can be used for reconstruction. Note also that since the kernel is a derivative of the hyperbolic secant, each Dirac results in the signal having both intervals where it is positive and intervals where it is negative. This




Figure 3.12: Example time encoding and reconstruction of a two Dirac burst using the hyperbolic secant kernel. Intermediate signals are shown: input signal consisting of two Diracs, $x(t)$. Filtered signal, $y(t)$. Integral with resets $z(t)$. As can be seen, the estimation is exact.

is problematic as it makes deriving firing guarantees largely intractable, even for the case of only two Diracs.

3.7.2 Reconstruction with unknown starting conditions

The aforementioned method for reconstruction relies on knowledge of the starting conditions of the time encoding machine. That is, I assumed that $\tau_0 = -\infty$ and that $f(-\infty) = 0$. In this section I relax this assumption and present a reconstruction scheme tailored to the following practical considerations. Our knowledge of the starting conditions of the encoder can be lacking in three main ways:

1. *The encoder may have spiked one or more times before recording started* - The method in Section 3.7.1 requires that we know and use the first output spike, τ_1 . Ideally we would like to be able to reconstruct using a number of different subsets of the generated spikes and not just the

first $2K$ spikes. For example, the first M spikes may not have been recorded and we must use a later subset for reconstruction.

2. *The kernel may be truncated in some way resulting in an initial bias to $f(t)$* - I have considered kernels of infinite support. In principle there is no reason why truncated or similarly windowed kernels could not be used subject to certain conditions on the input signal. However, truncating the kernel introduces a (negative) bias into the integral $f(t)$ which must be taken into account when reconstructing.
3. *The integral $f(t)$ may have started with a non-zero value* - This is important in the case where we want to perform reconstruction of localised portions of the input signal. For example, reconstruction of streams of bursts of Diracs was demonstrated in [15] using kernels of compact support. The first spike to follow a burst of Diracs in the input signal was influenced not only by the current burst but also by the residual integral remaining from the previous burst. This was accounted for by realising that the influence of the previous burst only affects the first spike not the subsequent spikes. Since this method requires the first sample be used, a different approach is needed to account for the residual integral.

All three of these cases are equivalent to adding a constant term to the first sample we use. For the first and third case I show this by modifying the expression for the first sample, z_1 , with a residual integral, b , and the starting time of the IF-TEM, τ_0 ,

$$z_1 = \pm C = b + \int_{\tau_0}^{\tau_1} dt . \quad (3.58)$$

Now consider the samples starting from $n = M$ for $M > 0$. For example, if $M = 2$ we are discarding the first sample. The first sample, z_M , is then,

$$\begin{aligned}
z_M &= \pm C_T = \int_{\tau_{M-1}}^{\tau_M} y(t) dt \\
&= \int_{\tau_0}^{\tau_M} y(t) dt - \int_{\tau_0}^{\tau_{M-1}} y(t) dt \\
&\stackrel{(a)}{=} \int_{\tau_0}^{\tau_M} y(t) dt + b - \sum_{n=1}^{M-1} z_n ,
\end{aligned} \tag{3.59}$$

where (a) follows from Eq. (3.58) and b is constant. Substituting Eq. (3.59) into Eq. (3.55), we can express an augmented sample y_n for $n > M$ as,

$$\begin{aligned}
y_n &= \sum_{m=1}^n z_m \\
&= \sum_{m=1}^{M-1} z_m + z_M + \sum_{m=M+1}^n z_m \\
y_n + c &= \sum_{m=M+1}^n z_m ,
\end{aligned} \tag{3.60}$$

where c is a constant bias common to all augmented samples. Thus, when the first M samples are omitted, a non-zero initial integral is assumed and an unknown starting time is considered, the cumulative samples are equivalent to the original cumulative samples except with a constant unknown bias,

$$y_n + c = \sum_{k=1}^K x_k \operatorname{sech}(a\tau_n - at_k) . \tag{3.61}$$

The linear system in Eq. (3.52) can also be modified to take this constant into account,

$$(\mathbf{Y} + c\mathbf{I})\mathbf{V}\mathbf{q} = \mathbf{W}\mathbf{p} . \tag{3.62}$$

Given knowledge of the constant, c , Eq. (3.62) can be solved up to a multiplicative constant using

$2K$ samples. In the case when c is unknown, the system can be solved using $3K + 1$ samples by linearising the problem as follows,

$$\begin{bmatrix} \mathbf{y}\mathbf{V} & \mathbf{v} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ c\mathbf{q} \end{bmatrix} = \mathbf{w}\mathbf{p} . \quad (3.63)$$

3.7.3 Piecewise constant signals

Reconstruction of sums of scaled and shifted Diracs can be further extended to the case of piecewise constant signals. Let $x(t)$ be a piecewise constant signal:

$$x(t) = \sum_{k=1}^K x_k u(t - t_k) , \quad (3.64)$$

where $u(t)$ is the Heaviside function. By defining a new sampling kernel,

$$\varphi'(t) = \frac{d\varphi(t)}{dt} = a^2 \left(\tanh^2(at) \operatorname{sech}(at) - \operatorname{sech}^3(at) \right) , \quad (3.65)$$

we can reformulate the problem of sampling a piecewise constant signal into the problem of sampling a sum of scaled and shifted Diracs:

$$y(t) = \left(x * \frac{d\varphi}{dt} \right) (t) \stackrel{(a)}{=} \left(\frac{dx}{dt} * \varphi \right) (t) ,$$

where (a) follows from integration by parts and,

$$\frac{dx}{dt} = \sum_{k=1}^K x_k \delta(t - t_k) . \quad (3.66)$$

Sampling a piecewise constant using the filter in Eq. (3.65) is therefore equivalent to sampling a sum of scaled and shifted Diracs with our original filter in Eq. (3.47). An example of sampling and reconstructing a piecewise constant signal is shown in Fig. 3.13.




Figure 3.13: Reconstruction of a piecewise constant signal from output samples generated by integrate-and-fire method. $\hat{x}(t)$ is the estimated signal.

3.7.4 Squared hyperbolic secant kernel

In the previous section I looked at how a particular hyperbolic function, the derivative of the hyperbolic secant can be used to recover signals from integrate-and-fire time encodings. The development of this particular sampling kernel is largely because it is a very natural extension of the secant kernel used in [62]. Of course, the method is amenable to other hyperbolic kernels, some of which may be more advantageous, appropriate or plausible for a given context. A good example of this is the square of the hyperbolic secant ($\text{sech}^2(x)$) as plotted in Fig. 3.14. This kernel is similar to the derivative of the hyperbolic secant because it conforms to the same general shape. It is also positive for the entire real line.

Previously I used the fact that the integral of the kernel was the hyperbolic secant. The integral of the new kernel is the hyperbolic tangent with a constant of integration which can be expressed as,

$$\psi(t) = \int_{-\infty}^t \text{sech}^2(x) dx = \tanh(t) + 1 = \frac{2e^{2t}}{e^{2t} + 1} . \quad (3.67)$$




Figure 3.14: The hyperbolic secant squared, $\text{sech}^2(x)$, and its integral $\tanh(x) + 1$.

Note that the degree of the numerator and denominator of $\tanh$ are equal. Unlike with the derivative of the hyperbolic secant, the constant of integration is significant since it is non-zero for an integrate-and-fire mechanism starting with an initial state of zero. A simple way to circumvent this problem is to treat this constant of integration as an unknown starting condition of the integrator as done in Section 3.7.2. Substituting Eq. (3.67) for $\text{sech}(\cdot)$ in Eq. (3.55) we can see that the result is samples, y_n , each with a constant but unknown bias,

$$y_n = c + \sum_{k=1}^K x_k (\tanh(a\tau_n - at_k) + 1)$$

$$y_n - c - \underbrace{\sum_{k=1}^K x_k}_b = \sum_{k=1}^K x_k \tanh(a\tau_n - at_k), \quad (3.68)$$

where c is an unknown constant representing the initial condition of the sampler. The situation is the same as that in Eq. (3.61) and as such we can use the same approach to achieve reconstruction.

Once again we linearise the problem in the same way as Eq. (3.63). Furthermore, note that the polynomials P and Q contain the same monomials ($\mathbf{W} = \mathbf{V}$) since the integrated kernel has equal degree in numerator and denominator. Thus, the linearised problem is simplified,

$$\begin{aligned} \begin{bmatrix} \mathbf{YV} & \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ -(c+b)\mathbf{q} \end{bmatrix} &= \mathbf{Vp} \\ \begin{bmatrix} \mathbf{YV} & \mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ -(c+b)\mathbf{q} - \mathbf{p} \end{bmatrix} &= 0. \end{aligned} \quad (3.69)$$

The new linearised system can be solved with $2K + 1$ samples. Thus the coefficients of Q and consequently the times t_k can be determined with $2K + 1$ samples with or without unknown starting conditions. As before, a least squares approach can then be used to determine both the initial starting condition and the values of x_k . Note that the number of samples required is the same as the number of degrees of freedom in the signal when the starting condition of the sampler is unknown, $2K + 1$. As such, the reconstruction method uses the sparsest set of sample points possible.

In a sense, the squared hyperbolic secant kernel results in an integrated kernels whose numerator and denominator polynomials are of a lower degree than the original integrated kernel, the hyperbolic secant. As shown in Eq. (3.69), the constant arising from the unknown initial conditions of the sampler, c , and the constant of integration of the kernel, b , are subsumed within the same quantity during the reconstruction. Furthermore, the hyperbolic tangent with integration constant of Eq. (3.67) generalises all kernels of minimum degree when the constants b and c are accounted for. Consider any rational function of degree $[1/1]$, that is with a degree of 1 in the numerator and the denominator,

$$r(x) = \frac{ax + b}{cx + d}. \quad (3.70)$$

Such a rational function can always be decomposed into an equivalent rational function,

$$r(t) = \frac{1}{p(x)} + c; , \quad (3.71)$$

where c is a constant. The kernel with minimum allowable degree has a degree of $[0/1]$. When constants of integration and unknown initial conditions are considered, these are equivalent to kernels of degree $[1/1]$ through addition of the constant term.

3.8 Recovery from differential samples

Typically, time encoded samples can be considered to represent finite differences on the integral of the input signal. The effect of the cumulative samples approach (Section 3.7) is to produce samples that are finite differences between a common point, τ_0 , and a set of non-uniform points $\tau_{n>0}$. The advantage of this approach is that it maps a time encoding into a non-uniform samples which can be reconstructed with an existing method for reconstructing FRI signals from non-uniform samples. The required number of samples is also optimal in the sense that it is the same or close to degrees of freedom of the input signal.

A major pitfall of the cumulative samples approach is that it assumes that samples are generated from IF-TEMs with ideal characteristics. In particular, I assumed that the resetting time was instantaneous and that the IF-TEM reset to a known value. When this is not the case, cumulative samples are no longer equivalent to differential samples with a common starting point, τ_0 (Eq. (3.55) no longer holds).

If we sacrifice optimality in sampling complexity, reconstruction can instead be achieved directly from the samples by approximating the interpolating function with a bivariate rational function. This approach is also clearly significantly more unstable since the degree of the estimated rational function is significantly higher.

Finally, another major motivation for using this form of reconstruction is that it allows samples from two well-matched IF-TEMs to be combined to cancel out non-idealities in their resetting char-

acteristics. ‘Differential mode’ sampling of this variety is discussed in greater detail in Chapter 5. The comparative merits of the two approaches are tabulated in Table 3.1

Table 3.1: Comparison of reconstruction from cumulative samples or differential samples

Cumulative samples	Differential Sampling
- Optimal number of samples, $\mathcal{O}(K)$	- Sub-optimal number of samples, $\mathcal{O}(K^2)$
- More stable	- Less stable
- Assumes Ideal resetting characteristics	- Resetting characteristics can be non-ideal or unknown
- Samples generated by one IF-TEM	- Samples from one or many IF-TEMs

In this section, I will focus specifically on how reconstruction is achieved from differential samples. Recall the expression for a sample derived from an IF-TEM,

$$\begin{aligned}
z_n = z(\tau_n) &= \sum_k x_k \int_{\tau_{n-1}}^{\tau_n} \varphi(t - t_k) dt \\
&= \sum_k x_k (\psi(\tau_n - t_k) - \psi(\tau_{n-1} - t_k)) \\
&= \sum_k x_k \tanh(\tau_n - t_k) - \sum_k x_k \tanh(\tau_{n-1} - t_k),
\end{aligned} \tag{3.72}$$

where I have assumed that $\varphi(t) = \text{sech}^2(t)$. I will call such a sample a differential sample. Substituting the usual rational function (Eq. (3.67)) expression of two summands into Eq. (3.72) we have the difference of two rational functions,

$$z_n = \frac{P(e^{2\tau_n})}{Q(e^{2\tau_n})} - \frac{P(e^{2\tau_{n-1}})}{Q(e^{2\tau_{n-1}})}, \tag{3.73}$$

where P and Q are polynomials of order K . As before, the input Dirac times, t_k , are contained within the roots of Q . If the Q polynomial can be determined, the input Dirac times can be determined from

its roots. To this end we can consider how to determine the polynomials P and Q directly from a set of differential samples. First we combine the quotients to express the finite difference of Eq. (3.73) as a bivariate rational function,

$$\begin{aligned} z_n &= \frac{P(e^{2\tau_n})Q(e^{2\tau_{n-1}}) - P(e^{2\tau_{n-1}})Q(e^{2\tau_n})}{Q(e^{2\tau_n})Q(e^{2\tau_{n-1}})} \\ &= \frac{R(e^{2\tau_n}, e^{2\tau_{n-1}})}{S(e^{2\tau_n}, e^{2\tau_{n-1}})} , \end{aligned} \quad (3.74)$$

where R and S are bivariate polynomials which are antisymmetric and symmetric polynomials respectively by construction. That is, $R(x, y) = -R(y, x)$ and $S(x, y) = S(y, x)$. Let $x = e^{2\tau_n}$, $y = e^{2\tau'_{n-1}}$ and let $r_{m,n}$ and $s_{m,n}$ be the terms of the R and S bivariate polynomials respectively. We then note that, due to their symmetry, R and S have the following expressions,

$$R(x, y) = \sum_{m=0}^K \sum_{n=m}^K r_{m,n} (x^m y^n - x^n y^m) , \quad (3.75)$$

and,

$$S(x, y) = \sum_{m=0}^K \sum_{n=m}^K s_{m,n} (x^m y^n + x^n y^m) . \quad (3.76)$$

Given Eq. (3.75) and Eq. (3.76), there are $N = K^2 + 2K + 1$ unknown coefficients. We construct a linear system similar to Eq. (3.51),

$$z_n \sum_{m=0}^K \sum_{n=m}^K s_{m,n} (x^m y^n + x^n y^m) = \sum_{m=0}^K \sum_{n=m}^K r_{m,n} (x^m y^n - x^n y^m) . \quad (3.77)$$

The system can be solved for the coefficients of the R and S given N finite differences, z_n . This system can be expressed in matrix form as,

$$\mathbf{ZV} = \mathbf{W} \mathbf{r} , \quad (3.78)$$

where $\mathbf{s}$ is the vector of $s_{m,n}$ coefficients, $\mathbf{r}$ is the vector of $r_{m,n}$ coefficients, $\mathbf{V}$ and $\mathbf{W}$ are matrices of terms corresponding to $\mathbf{s}$ and $\mathbf{r}$ respectively, and $\mathbf{Z}$ is the diagonal matrix of samples, z_n .

Having determined the coefficients of S we can now determine the roots of Q as follows. From Eq. (3.49) we can see that the roots of S are repeated in x and y such that,

$$S(x, y) = Q(x)Q(y) = \prod_{k=1}^K (1 + e^{-2t_k}x)(1 + e^{-2t_k}y) . \quad (3.79)$$

Therefore, using the substitution $x = y$, we can transform $S(x, y)$ into a univariate polynomial, $S(x)$, whose doubled roots coincide with the roots of the original bivariate polynomial,

$$S(x) = \prod_{k=1}^K (1 + e^{-2t_k}x)^2 . \quad (3.80)$$

If α_k are the K double roots of $S(x)$ then the input Dirac times can be determined as,

$$t_k = \frac{\ln(-\alpha_k)}{2} . \quad (3.81)$$

The magnitudes of the Diracs, x_k , can then be determined in the standard least squares sense from Eq. (3.72).

An example of the samples generated by a nine channel system where each channel is comprised of a single differential pair is plotted in Fig. 3.15 along with the resulting reconstruction.

3.9 Summary

In this chapter I have looked at time encoding strategies for signals with finite rate of innovation with novel classes of sampling kernels, namely the α -synaptic and hyperbolic kernels. The various pros and cons of these two classes of kernels are tabulated in Table 3.2.




Figure 3.15: The spikes generated by multiple channels of differential samplers (aggregated) and the resulting reconstruction. The input Diracs are $x(t)$ and the estimated Diracs are $\hat{x}(t)$. The aggregated output spikes generated by 9 differential pair channels are z_n . Each channel produces 1 finite difference (2 spikes).

Table 3.2: Comparison of FRI sampling kernels for time encoding

Spline kernels	α -Synaptic kernels	Hyperbolic kernels
- Multiple knots	- Single knot	- No knots
- Exponential reproducing	- Exponential reproducing	- Rational interpolation
- Compact support	- Supported $t > 0$, vanishing $t \rightarrow \infty$	- Infinite support, vanishing $ t \rightarrow \infty$
- Samples must be in same knotless region	- Samples must be after last Dirac	- Samples can be in any region
- Complicated signal conditions	- Restrictive signal conditions	

The case of the α -synaptic kernels is significant for several reasons. Firstly, I have shown that signals can be reconstructed with exponential reproducing techniques even when the sampling kernel does not conform to the generalised Strang-Fix conditions. Furthermore, in the original case of polynomial or exponential splines, samples had to be generated in the same knotless region of the kernel to allow reconstruction. In the α -synaptic case, the kernel has a single knot, and consequently it becomes tractable to derive conditions (Proposition 1) on the input signal that guarantee enough useful samples are generated in a knotless region to allow reconstruction.

The exponential reproducing approach to reconstruction was also shown to be applicable to FRI signals with more complicated structure. I developed an algorithm to reconstruct the power signal of an FRI signal with the α -synaptic sampling kernel. Power signals may be significant in the use of IF-TEMs since they are intrinsically positive and their integral is always monotonic. IF-TEMs used to sample such signals thus only require a positive threshold and are guaranteed to sample all level sets of the integrated signal.

I then showed that signals sampled with hyperbolic sampling kernels can be reconstructed using a univariate rational interpolation technique. This was first demonstrated with the derivative of the hyperbolic tangent sampling kernel. The method was further extended to time encodings where the

initial conditions of the encoder were unknown. In both cases the number of samples required to enable reconstruction was shown to be optimal or nearly optimal, depending on the choice of kernel. Such kernels are of interest because, unlike splines or α -synaptic functions, they do not contain any knots, therefore, there is no restriction on when samples may or may not be generated. On the other hand, they are infinitely supported, therefore, reconstruction results for multiple bursts of pulses must appeal to the fact that they vanish as $|t| \rightarrow \infty$ and thus we can approximate the kernel as being zeroed values for large values of $|t|$.

Finally, I showed that a bivariate rational interpolation approach can be used to reconstruct the input signal directly from differential samples. Integrate-and-fire time encoded samples can be expressed as finite differences and so rational interpolation can be used to directly reconstruct the signal without first mapping the samples to equivalent non-uniform samples as is typical in time encoding. This has some advantages over the univariate approach including relaxed assumptions on the non-idealities of an IF-TEM implementation. This is achieved at the expense of robustness and optimality. The bivariate rational interpolation approach from differential samples will be leveraged in Chapter 5 allowing pairs of IF-TEMs to be combined in a single network with the benefit of nullifying the non-idealities common within a pair of IF-TEMs.

In some cases, the kernels described in this chapter may match closely with the pulse shapes encountered in some real world signals. In particular, the α -synaptic function has been used to model post-synaptic potentials in biological neurons [59]. However, in practice, other signals that are comprised of pulses are unlikely to have pulse shapes that precisely conform to the pulse shapes used in this chapter. A potential way around this is to use higher order α -synaptic functions and higher order hyperbolic functions that can be fitted to the desired pulse shape.

Chapter 4

Amplitude-domain and the time dilating integrate-and-fire encoder (TDIF)

The majority of work on time encoding, including the contributions of Chapter 3, take an approach that maps the samples of the time encoding into their equivalent non-uniform samples. The signal can then be reconstructed using standard reconstruction algorithms for non-uniform samples. For example, I have used rational approximation in the context of hyperbolic sampling kernels. As such, the reconstruction algorithms for time encoded samples are typically handled entirely in the time-domain.

As alluded to in Chapter 2, conventional samples and time encoded samples are two sides of the same coin. Briefly, I say this because the time measurements of the former are signal independent or implicit while the amplitude measurements of the latter are signal independent or implicit. In this chapter I will consider a similar question in the context of reconstruction algorithms. If time-domain

reconstruction is the conventional side of the coin, is there an ‘amplitude-domain’ reconstruction strategy on the other side of the coin?

The organisation of this chapter is as follows. In Section 4.1 I will show how time encoding can be seen as sampling in an equivalent domain. More specifically, how integrate-and-fire time encodings perform uniform sampling in this domain. The definitions and properties of time encoding presented in this section are then directly used in Section 4.2 to demonstrate reconstruction of a restricted class of signals via uniform sampling. In Section 4.3 I will introduce a new sampling architecture that uses dilation of time measurements to allow us to use arbitrary filters with compact support. I do this with a view to achieving reconstruction of shift-invariant and FRI signals in Section 4.4 with a multichannel system of TEMs that implement the new architecture.

4.1 Time encoding is a form of uniform sampling

Integrate-and-fire time encoding is a form of uniform sampling in amplitude is an oft repeated description of the mechanism by which time encoding is achieved. This observation is relatively intuitive. Samples are produced at non-uniform instants in time every time the integrated signal crosses some known, constant threshold. This implies that the time samples encode the level sets of the integrated signal.

This said, the prevailing sampling and reconstruction strategies are dependent on the equivalence of time encodings and non-uniform samples. Typically, reconstruction strategies for time encodings have leveraged the equivalence of time encoded samples and conventional non-uniform samples through the t-transform which provides a piecewise constant initial approximation [16,25,29,30]. In other works such as [63], the t-transform is used to relate the cumulative sum of samples to non-uniform samples of the integral. Once time encodings are mapped to their equivalent non-uniform samples, the reconstruction is achieved with standard non-uniform reconstruction strategies in time domain, such as those contained in [9].

Notable exceptions to this are the works of [64] and [23,24] who consider the reconstruction prob-




Figure 4.1: Time domain signal, $h(t)$, and its integral, $f(t)$, plotted alongside the amplitude-domain signal, $g(c)$, along the vertical axis. The time domain signal, $h(t)$, is bandlimited and periodic.

lem from the perspective of uniform sampling in amplitude. The work of [64] showed that given a bandlimited signal in time-domain, the equivalent uniform sampled signal is not bandlimited, however, with a sufficient bias, an effective approximation can be made in the form of the essential bandwidth. In a similar vein, the work of [24] compared the operation of an integrate-and-fire TEM with delta-ramp encoders and analysed the rate of decay of the spectrum of the equivalent signal (amplitude-time signal). Some of the fundamental observations presented in these works are also made in this section but from a different perspective. I will provide a concrete setting to the phrase *integrate-and-fire time encoding is a form of uniform sampling in amplitude* and some observations from which later sections will build upon.

First, consider a monotonically increasing function $f(t)$, its derivative $h(t) = f'(t)$, and the derivative of its inverse $g(c) = [f^{-1}]'(c)$. These three functions are plotted together in Fig. 4.1 where $g(c)$ is plotted rotated 90 degrees to emphasise its relation with the inverse of $f(t)$.

The function $h(t)$ is equivalent to the time-domain input signal of an integrate-and-fire encoder and the function $f(t)$ its integral. As such, the $f(t)$ represents the accumulated integral at a time t and $f^{-1}(c)$ represents the time t at which the system reaches an accumulated integral of c . I note that $f(t)$ must be one-to-one for its inverse to be well-defined and consequently $g(c)$ and $h(t)$ must also be strictly positive. I will refer to the function $g(c)$ as the amplitude-domain equivalent of the time domain signal $h(t)$. The foundation for the work in this chapter rests in the following property:

Property 1 (Inter-trigger timespan). *Let the inter-trigger timespan $\theta_n = \tau_{n+1} - \tau_n$ where $\{\tau_n\}$ are the trigger times of the integrate-and-fire system, C is the threshold of the TEM, $b < C$ is the initial integral of the TEM, and $\tau_0 = 0$ is the initial time then,*

$$\theta_n = \int_{nC-b}^{(n+1)C-b} g(x) dx = \langle \Pi_C(x + b - nC), g(x) \rangle, \quad (4.1)$$

where $\Pi_C(x)$ is the rectangular function supported on $[0, C]$. That is,

$$\Pi_C(x) = \begin{cases} 1 & 0 \leq x \leq C, \\ 0 & \text{otherwise} . \end{cases} \quad (4.2)$$

Proof. The trigger times, τ_n and τ_{n+1} , are related to the initial integral, b , and threshold, C , through f :

$$f(\tau_n) = nC - b \quad f(\tau_{n+1}) = (n+1)C - b, \quad (4.3)$$

therefore,

$$\begin{aligned} \int_{nC-b}^{(n+1)C-b} g(x) dx &= f^{-1}((n+1)C - b) - f^{-1}(nC - b) \\ &= \tau_{n+1} - \tau_n = \theta_n. \end{aligned} \quad (4.4)$$

□

Property 1 can be viewed as an equivalence between the timespans of the time domain signal, $h(x)$, and the area under the curve for the corresponding signal in the amplitude-domain, $g(x)$. An illustrative example of this is plotted in Fig. 4.2. Alternatively, from the inner-product form of Eq. (4.1), one can see that the inter-trigger timespans correspond to uniform samples of a filtered version of $g(x)$. The support of the filter is equal to the threshold, C , as is the sampling period, therefore, by modifying the threshold one can change the equivalent sampling rate in amplitude-domain. I formalise the transform from the time-domain signal to the amplitude-domain signal with the transform denoted $\mathcal{T}$ (Definition 1).

Definition 1 (Domain swap transform). *Let $\gamma(x) : \mathbb{R} \rightarrow \mathbb{R}$ be a function with $\gamma(0) = 0$. Let $\mathcal{I}$ be the integral operator;*

$$\mathcal{I}\gamma = \int_0^x \gamma(\tau) d\tau; \quad (4.5)$$

.

Let $\mathcal{D}$ be the derivative operator such that by the fundamental theorem of calculus,




Figure 4.2: Illustration of Property 1 and how it relates a timespan θ_n in the time domain signal to the area under the curve of the amplitude-domain signal. In this case, the amplitude-domain sampling period is determined by a threshold $C = 0.1$.

$$\mathcal{I} \circ \mathcal{D} = \mathcal{I} \circ \mathcal{D} = I , \quad (4.6)$$

where I is the identity. Then the transform $\mathcal{T}$ is,

$$\mathcal{T} = \mathcal{D} \circ (\cdot)^{-1} \circ \mathcal{I} , \quad (4.7)$$

where $(\cdot)^{-1}$ is an operator that takes a function to its inverse.

From Definition 1 one can determine that the transform is in fact self inverse (Property 2). Thus, properties that hold for the forward (reverse) transform hold in the reverse (forward) transform. For instance, by Property 1 and Property 2, a span in the amplitude-domain is equivalent to the area under the corresponding curve in the time-domain. That is, $\int_{\tau_n}^{\tau_{n+1}} h(x) dx = C$ where C is a span in the amplitude-domain. In fact this is a restatement of the t-transform coined by [16] and used throughout the literature on time encoding.

Property 2 (Involution). *The operator $\mathcal{T}$ defined in Definition 1 is self-inverse.*

Together, these two properties suggest an alternative approach to reconstructing from time encoded samples. First, Property 1 indicates that one can treat inter-trigger timespans, θ_n , as uniform samples of the equivalent amplitude-domain signal, $g(c)$. Thus, by applying standard reconstruction techniques for uniform samples one can recover the amplitude-domain signal. The invertibility of the transform $\mathcal{T}$ (Property 2), then allows us to regain the time-domain signal, $h(t)$. This is the fundamental procedure that I now build upon.

4.2 Bandlimited and piecewise constant signals in amplitude-domain

From Eq. (4.1) it can be seen that an integrate-and-fire time encoding is equivalent to a uniform sampling of an equivalent amplitude-domain signal using a rectangular sampling kernel with support

equal to the threshold, C . The sampling period is also the threshold and the signal itself is given by the transform $\mathcal{T}h$. Immediately, we can note that when the signal, in amplitude-domain, belongs to certain subspaces, we can reconstruct by interpolation. I will refer to the subspace that the amplitude-domain signal belongs to as the representation of the time-domain signal. In general, the representation of the signal in amplitude-domain is not from the same class of signals as the time-domain signal. For example, when the amplitude-domain signal is bandlimited the time-domain signal is necessarily non-bandlimited [24]. Theorem 1 formalises the reconstruction for the case of bandlimited representations.

Theorem 1 (Bandlimited reconstruction). *Let h be a signal in time-domain and $g = \mathcal{T}h$. Let PW_B be the set of signals bandlimited to B then if $g \in PW_B$, h can be reconstructed from an integrate-and-fire time encoding if*

$$C < \frac{\pi}{B}. \quad (4.8)$$

Proof. Given an integrate-and-fire time encoding, $\{\theta_n\}$, of h , then by Property 1 the time encoding is equivalent to a uniform sampling of $G(x) = g(x) * \varphi(x)$ where $\varphi(x) = \Pi_C(x)$. If $g(x) \in PW_{\pi/C}$ then $G(x)$ is also bandlimited to π/C and can be reconstructed by the Shannon-Whittaker sampling theorem. Furthermore, the Fourier transform of $\varphi(x)$,

$$\Pi_C(x) \overset{FT}{\leftrightarrow} C \text{sinc}\left(\frac{\omega C}{2}\right) e^{-\frac{i\omega C}{2}}, \quad (4.9)$$

does not vanish on the interval $(-\pi/C, \pi/C)$. Consequently, $g(x)$ can be recovered from $G(x)$ by the convolution theorem. Finally, $h(x)$ can be recovered by applying the transform, $\mathcal{T}$ (Property 2).

□

An example of an integrate-and-fire sampling where the input signal is represented by a bandlimited and periodic function in the amplitude-domain is given in Fig. 4.3.

Similarly, signals in amplitude-domain that are shift invariant where the basis functions are dis-




Figure 4.3: Sampling and reconstruction of a signal with a bandlimited and periodic representation in amplitude-domain. The integrate-and-fire samples in time-domain (bottom) are mapped to uniform samples in the amplitude-domain (left).

joint can also be sampled and reconstructed. For example, any signal of the form $\sum_{n=-\infty}^{\infty} c_n \varphi(x - nC)$ where $\text{supp}(\varphi) = [0, C]$.

This is a very restricted class of signals and so I will only consider the case of piecewise constant signals. Piecewise constant signals are unique in that the transform $\mathcal{T}$ maps positive (negative) piecewise constant signals in time-domain to positive (negative) piecewise constant signals in amplitude domain and vice versa as summarised in Property 3. This property will also be of significance in Section 4.5 when I revisit sampling and reconstructing piecewise constant signals with a new architecture.

Property 3 (Automorphism on piecewise constant signals). *Let $\mathcal{H}^+$ be a strictly positive piecewise constant signal. The operator is an automorphism on $\mathcal{H}^+$. That is, if $h(t) \in \mathcal{H}^+$ then $\mathcal{T}h \in \mathcal{H}^+$.*

Proof. Since $h(t)$ is piecewise constant, its integral must be piecewise linear. The inverse of a piecewise linear is also piecewise linear. Thus, $g(x) = [f^{-1}]'(x)$ is piecewise constant.

□

Using this property, Theorem 2 shows that a time-domain piecewise constant signal can be reconstructed from its time encoding using a uniform sampling approach when the integral of each piece is equal to a multiple of the trigger mark, C . This is equivalent to the edges of amplitude-domain signal lying on a uniform grid of period C . This form of reconstruction is equivalent to the initial piecewise constant approximations used to achieve bandlimited reconstructions in other works [16,29] and is therefore not in itself novel, however, I have formalised it for the sake of completeness and intuition.

Theorem 2. *A piecewise constant signal in time-domain can be reconstructed from its amplitude-domain samples (inter-trigger timespans) if the time-domain signal has pieces whose integral is a multiple of the threshold.*

Proof. Let the first piece be supported on the interval $[0, t_1]$. If the area of said piece is a multiple of the threshold, that is,

$$\int_0^{t_1} h(t) dt = nC, \quad (4.10)$$

where $n \in \mathbb{Z}^+$, then by Property 1, $f^{-1}(t_1) = nC$. The edge of the amplitude-domain piecewise constant signal must therefore align on a grid of period C in the amplitude-domain. The same argument applies for all subsequent and preceding pieces. Since the edges in amplitude-domain lie on a grid with sampling period C , the amplitude-domain signal can be expressed as,

$$g(x) = \sum_{n \in \mathbb{Z}} c_n \Pi_C(x - nC), \quad (4.11)$$

for $c_n \in \mathbb{R}$. Substituting Eq. (4.11) into Eq. (4.1) shows that inter-trigger time spans correspond to the coefficients of the signal, c_n , when the initial integral $b = 0$,

$$\theta_m = \left\langle \Pi_C(x - mC), \sum_{n \in \mathbb{Z}} c_n \Pi_C(x - nC) \right\rangle = c_m$$

□

4.3 Sampling kernel shaping via time measurement dilation

Previously we have seen that time domain signals represented by a limited subset of shift-invariant signals in the amplitude-domain can be reconstructed via a uniform sampling approach. Unfortunately, this family of signals is highly restrictive, in particular because the basis functions of these signals, $\varphi(\cdot - nC)$, must be disjoint and the sampling kernel used is restricted to a rectangular function. To broaden the set of signals one must tackle these two problems. In this section I will modify the sampling architecture to allow sampling with kernels other than the rectangular function. I do this with a view to sampling and reconstructing a much broader set of shift-invariant signals and signals with finite rate of innovation with multichannel systems.

The fundamental approach to modifying the sampling kernel is dilation of the time measurements, τ_n . The architecture by which one can achieve dilation of time measurements and how it

fits in the broader context of conventional samplers and integrate-and-fire samplers is presented in the following Section 4.3. The effect of this particular form of time dilation is then analysed in Section 4.3.

Generalised samplers and the time dilating integrate-and-fire sampler

First, I would like to introduce the new architecture within the broader context of conventional and time encoding sampling architectures. Consider the sampling architectures illustrated in Fig. 4.4. The sampler depicted in Fig. 4.4(a) is an equivalent model for a classical uniform sampler where the role of the timing circuitry is emphasised. The lower branch is a timing circuit used to determine when measurements of the signal should be taken. The input to the timing circuitry is a scale factor, in this case equal to 1. This is then integrated to provide a signal representing time. Whenever this time signal crosses a multiple of the period, T , a sample is taken. The integration operation is emphasised to show the commonality between this conventional sampling architecture and the later architectures in (b) and (c). Although this seems like a gross abstraction of the real life hardware, it is not dissimilar to the operation of a digital circuit comprising an oscillator, counter, and comparator in series.

The typical integrate-and-fire time encoding machine illustrated in Fig. 4.4(b) is very similar in construction except that the roles of the branches are reversed. Measurements of the time signal are now made and the timing of these samples are dictated by when the integral of the input signal crosses multiples of a threshold, C . The purpose of this exposition is to emphasise that in both cases, whilst the input to the top branch is signal driven and time-varying, the input to the timing branch is constant. It is also important to emphasise that the figure represents the integrate-and-fire mechanism as a level set detection for analytical convenience. In practice the level set detection would be implemented with an integral resetting mechanism as in Fig. 3.1.

In the final incarnation of the sampler, Fig. 4.4(c), the output of the upper branch is fed through






Figure 4.4: Generalised sampling architectures: **a)** A generalised model for conventional sampling of an integrated signal. **b)** A generalised model for integrate-and-fire sampling. **c)** A time-dilating integrate-and-fire sampler (TDIF). In all examples, the typical comparator with reset mechanism has been replaced with a level detector ($\cdot = nC$). Note, φ is a filter and β is a non-linear amplifier.

a non-linear amplifier¹ β into the lower branch providing a signal-driven time signal, $T(t)$. Assume that the response of the amplifier, β , is periodic such that,

$$\beta(x) = \varphi(x + nC) \quad \text{for} \quad n \in \mathbb{Z}, \quad (4.12)$$

where $\varphi(x) : [0, C] \rightarrow \mathbb{R}$ is a real valued function of compact support. This is also an abstraction for analytical convenience. The function $\varphi(x)$ represents the response of the amplifier in a practical implementation that uses an integral resetting mechanism. In this case, we have unfolded the resetting mechanism into a level set detector, and so the equivalent amplifier response is $\beta(x)$, the C -periodic extension of $\varphi(x)$.

In the simplest case, if the amplifier is constant, $\beta(x) = 1$, then the original integrate-and-fire setting is retrieved. On the other hand, one can design β to act as a sort of filter in amplitude-domain. The amplifier β acts to achieve a form of time dilation of the time measurements. At different values of the integral $f(t)$, the value of the time scale, $\beta(f(t))$, determines how fast or slow ‘time’ is accumulated. For this reason I call this sampler a time dilating integrate-and-fire sampler (TDIF) and $\varphi(x)$ the *dilation function*. The effects of time dilated measurements are analysed in the next section.

The effect of dilation of time measurements

The function $T(t) : \mathbb{R} \rightarrow \mathbb{R}$ is a many-to-one function that represents the mapping from (undilated) time to the dilated time. It is many-to-one because the time scale, $T'(t)$, can be positive and negative valued. The implication is that the measured time can run either forwards or in reverse. From Fig. 4.4, one can see that,

$$\begin{aligned} T'(t) &= \beta(b + f(t)) \\ &\stackrel{(a)}{=} \varphi(b + nC + f(t)), \end{aligned} \quad (4.13)$$

¹Non-linear amplifiers are realisable circuits, for example the reader can refer to implementations with operational amplifiers [65].

where b is the initial condition of the integrator, and (a) follows from Eq. (4.12). We wish to see the effect of time dilation on the observed signal in the amplitude domain. Let $G(x) = [F^{-1}]'(x)$ represent the new amplitude-domain signal after time dilation, that is, $G(x)$ is the equivalent amplitude-domain signal if $f(t)$ were dilated according to $T(t)$. $T(t)$ maps the non-dilated time measurements to the dilated measurements and so,

$$F^{-1}(x) = T(f^{-1}(x)) . \quad (4.14)$$

Note that I use $F^{-1}(x)$ for notational consistency, however, the $F^{-1}(x)$ does not necessarily have a well defined inverse. If $T(t)$ is many-to-one (time dilation can be both positive and negative) then $F^{-1}(x)$ is many-to-one. Now taking the derivative of Eq. (4.14) and applying the chain rule to obtain the amplitude-domain signal after dilation,

$$\begin{aligned} G(x) &= [F^{-1}]'(x) = [f^{-1}]'(x) T'(f^{-1}(x)) \\ &\stackrel{(a)}{=} g(x) \cdot T'(f^{-1}(x)) \\ &\stackrel{(b)}{=} g(x) \cdot \varphi(x + b + nC) , \end{aligned} \quad (4.15)$$

where (a) follows from definition of the non-dilated amplitude-domain signal and (b) follows from substitution of Eq. (4.13). From Eq. (4.15) we can see that the signal in amplitude-domain is a product of the original amplitude-domain signal, $g(x)$, and the dilation function, $\varphi(x)$. In fact, the second term of the product represents the time measurement dilation of the signal at the time corresponding to the value of the integral. By designing an appropriate dilation function we can shape the amplitude-domain signal as we see fit. Moreover, due to the timespan property (Property 1), the inter-trigger timespans now equal an inner product between the dilation function and the amplitude-domain signal as formalised in Property 4. An example of this is plotted in Fig. 4.5 which plots both the original signal and the equivalent signal under time dilation by a linear spline. Note that while the




Figure 4.5: An example of a time-dilated signal to affect a multiplication between the amplitude-domain signal and a known sampling kernel. The red area indicates the integral that is equivalent to an inter-trigger timespan, θ , in the time-domain.

dilation of the integral is plotted, in practice this effect is only observed in the measured times and is not a physical effect, that is, I'm not proposing time travel.

Property 4 (TDIF inter-trigger timespan). *Let the inter-trigger timespan of a TDIF be $\theta_n = \tau_n - \tau_{n-1}$ where $\{\tau_i\}$ are the trigger times of the TDIF then,*

$$\theta_n = \langle \varphi(x + b - nC), g(x) \rangle, \quad (4.16)$$

where C is the threshold of the TDIF, b is the initial integral, and $\varphi(x)$ is the dilating function with compact support $B \subseteq [0, C]$.

Proof. Using the timespan property (Property 1) substitute Eq. (4.15) into Eq. (4.1),

$$\begin{aligned}
\theta_n &= \int_{b+nC}^{b+(n+1)C} G(x) dx \\
&= \int_{b+nC}^{b+(n+1)C} \varphi(x+b-nC)g(x) dx \\
&\stackrel{(a)}{=} \int_{-\infty}^{\infty} \varphi(x+b-nC)g(x) dx \\
&= \langle \varphi(x+b-nC), g(x) \rangle,
\end{aligned} \tag{4.17}$$

where (a) follows from the assumption that $\varphi(x)$ is supported on the interval $[0, C]$.

□

Finally, I would like to note some of the symmetry between the conventional sampler and the TDIF sampler. Firstly, referring to Fig. 4.4, the filter φ in the conventional sampler and the amplifier β in the TDIF (a periodic extension of φ for analytical convenience) perform similar roles but for different branches. Namely, they define the sampling kernel in the time-domain and amplitude domain respectively. This is apparent from Eq. (4.16). Also note that in the TDIF, the time-domain signal is restricted to be positive, such that $f(t)$ is increasing. On the other hand the timescale, $T'(t)$, can be positive and negative valued depending on the choice of β . This is the converse of the conventional sampling architecture where the timescale is positive while the time domain signal can be positive and negative valued.

4.4 Interleaved time-dilating integrate-and-fire samplers

Multichannel systems for FRI and shift-invariant representations

In Section 4.3 I demonstrated how to shape the sampling kernel, however, there is a final problem that prevents sampling of a wider range of signals, namely, the shifts or the sampling kernel are disjoint. This is inherent to the sampling mechanism since the system does not filter the amplitude-domain signal but instead mimics filtering through point-wise multiplication of the amplitude-




Figure 4.6: Multichannel sampler with 3 channels. The time-evolving integral, $z(t)$, is plotted on each TDIF along with the output spike pattern of each channel, $z_n^{(l)}$, and the resulting combined output, z_n .

domain signal with the dilation function. Thus, in order to sample with overlapping basis functions, we must employ a multichannel system where each channel consists of a single TDIF, each with a different initial integral. Multichannel systems have been employed in various works in the area of time encoding [30,31,35,66].

Referring to Property 1 and Property 4, the initial integral, b , equates to a shift by b of the sampling kernel in amplitude-domain. The effect of using a multichannel system with different initial integrals is plotted in Fig. 4.7.

In particular, if one considers a multichannel sampler with L channels and threshold C and lets the initial integral of the l^{th} channel be $b^{(l)} = \frac{lC}{L}$ then $\theta_n^{(l)}$, the n^{th} inter-trigger timespan of the l^{th} channel, is given by:

$$\theta_n^{(l)} = \left\langle \varphi \left(x + \frac{C}{L}(l - nL) \right), g(x) \right\rangle \quad \Longleftrightarrow \quad \tilde{\theta}_m = \left\langle \varphi \left(x + m \frac{C}{L} \right), g(x) \right\rangle. \quad (4.18)$$

Since $\{l - nL : n \in \mathbb{Z}, l \in \{1, 2, \dots, L\}\} = \mathbb{Z}$, the multichannel system admits a uniform sampling, $\tilde{\theta}_m$, of the amplitude-domain signal with period C/L under a reordering of the multichannel inter-trigger timespans, $\theta_n^{(l)}$. Specifically, the reordering is taken first in descending channel number, l ,




Figure 4.7: Integral and amplitude-domain signals after dilation. The dilated signals are plotted aligned in the amplitude-domain to emphasise the decomposition of the amplitude-domain signal. Therefore, in this plot, the thresholds are at C and $2C$ for the blue TDIF and $C/2$ and $3C/2$ for the red TDIF and the dilation function is a linear spline as plotted in Fig. 4.8. In practice, offsets in the TDIF integral result in sampling a shifted version of the amplitude-domain signal and not a shift in the sampling kernel.

then ascending in n as shown in Fig. 4.6. Note that the task of setting and maintaining the correct offsets between channels is similar to the problem of maintaining synchronised clocks in a conventional multichannel time-domain system.

Shift invariant signals in amplitude-domain

Theorem 3. Let a time domain signal, $h(t)$, be represented by a shift-invariant signal in amplitude-domain,

$$g(x) = \sum_{n \in \mathbb{Z}} c_n \psi(x - nT) , \quad (4.19)$$




Figure 4.8: A linear spline dilating function, $\varphi(x)$, supported on $[0, C]$. and its periodic extension $\beta(x)$.

The time domain signal, $h(t)$, can be reconstructed from a time encoding derived from a multichannel integrate-and-fire sampler when $\varphi(x) = \tilde{\psi}(x)$ is the dual basis element of $\psi(x)$, has compact support $M \subset [0, LT]$, $L \in \mathbb{Z}^+$ and there are L channels with initial integrals $b_l = lC/L$ for $l \in [0, \dots, L-1]$ and $C = LT$.

Proof. From Eq. (4.18) we know that there is an ordering of inter-trigger timespans that produces the time encoding,

$$\begin{aligned}
 \tilde{\theta}_m &= \left\langle \varphi\left(x - m\frac{C}{L}\right), g(x) \right\rangle \\
 &= \left\langle \varphi\left(x - m\frac{C}{L}\right), \sum_{n \in \mathbb{Z}} c_n \psi(x - nT) \right\rangle \\
 &= \sum_{n \in \mathbb{Z}} c_n \left\langle \varphi\left(x - m\frac{C}{L}\right), \psi(x - nT) \right\rangle .
 \end{aligned} \tag{4.20}$$

Now substituting in $C = LT$ and $\varphi(x) = \tilde{\psi}(x)$,

$$\begin{aligned}\tilde{\theta}_m &= \sum_{n \in \mathbb{Z}} c_n \langle \tilde{\psi}(x - mT), \psi(x - nT) \rangle \\ &= c_m\end{aligned}\tag{4.21}$$

where I have used $\langle \tilde{\psi}(x - mT), \psi(x - nT) \rangle = \delta_{m-n}$. From the coefficients c_m the amplitude-domain signal can be reconstructed and the time-domain signal can be recovered through the transform $\mathcal{T}$.

□

4.5 Piecewise-constant signals

In Section 4.1, the reconstruction of time-domain piecewise constant signals was restricted to those that lay on a grid in their amplitude-domain representation. Furthermore, the signals were required to be strictly positive. With the new architecture, both these issues can be alleviated when viewing these signals as ones with finite rate of innovation (FRI). Archetypal examples of an FRI signals include sums of shifted and scaled Dirac deltas as well as piecewise constant signals and their filtered versions. The latter is related to the former through the integral and is the focus of this section. Sampling and reconstruction of FRI signals from time encodings has already been considered by various authors. For example, [35,53,58] considered periodic FRI signals and [15,63,67] considered reconstruction of pulses and piecewise constant signals using various kernels.

Of particular interest, [15] leveraged the work of [39] which showed how a uniform sampling problem on FRI signals can be formulated as one of sparse high resolution spectral estimation when the sampling kernel satisfies the Strang-Fix conditions and so can globally reproduce exponentials. [15] extended this idea to integrate-and-fire time encoding by noting that local exponential reproduction was still possible even in a non-uniform sampling context. The main caveat is that relatively strict constraints have to be placed on the signal to ensure samples are triggered within knotless

regions. If an approximate reconstruction is acceptable, a relaxation of these constraints can be achieved through the use of approximate reproduction of exponentials as used in [68].

As stated in Property 3, signals that are positive piecewise constant in time-domain are also piecewise constant in amplitude-domain. Thus, the output of a multichannel TDIF sampler, when the input signal is a piecewise constant signal, is a uniform sampling of an equivalent piecewise constant signal in amplitude-domain. Formulating the reconstruction from a time encoding in this manner means that one can directly apply the results of [39] to achieve exact reconstruction without appealing to local only reproduction of exponentials or approximate reproduction of exponentials.

Theorem 4. *Let $h(t) \in \mathcal{H}^+$ be a positive (negative) piecewise constant signal and θ_n be the time encoding of $h(t)$ generated by a multichannel TDIF where C is the threshold of the constituent TDIFs and L is the number of channels, and $\varphi(x) = \psi'(x)$, where $\psi(x)$ can reproduce $M \geq 2K - 1$ exponentials and has compact support C . The signal $h(t)$ can be exactly reconstructed when the integral of $h(t)$ on the interval between any K edges is bounded below as,*

$$\int_{t_n}^{t_{n+K}} h(t) dt > 2KC, \quad (4.22)$$

where t_n are the locations of the edges of the piecewise constant signal.

Proof. From Property 3, we know that if $h \in \mathcal{H}^+$ then, $\mathcal{T}h = g \in \mathcal{H}^+$ as well so that,

$$g(x) = \sum_{k \in \mathbb{Z}} a_k u(x - x_k), \quad (4.23)$$

where $u(x)$ is the Heaviside step function and $a_k, x_k \in \mathbb{R}$. When the time-domain signal is sampled with a L -channel TDIF with a threshold of C and $\varphi(x) = \psi'(t)$ the equivalent amplitude-domain uniform sampling is,

$$\begin{aligned}
g_n &= \left\langle \sum_{k \in \mathbb{Z}} a_k u(x - x_k), \psi'(x - nT) \right\rangle \\
&\stackrel{(a)}{=} \left\langle \sum_{k \in \mathbb{Z}} a_k \delta(x - x_k), -\psi(x - nT) \right\rangle,
\end{aligned} \tag{4.24}$$

where, $T = C/L$ is the sampling period and (a) follows from integration by parts. The amplitude-domain signal, $g(x)$, can be recovered from the samples g_n . This is done by using the exponential or polynomial reproducing property of $\psi(x)$ followed by high resolution spectral estimation techniques or model fitting approaches, such as Prony's method, to retrieve the innovation parameters x_k and a_k . By exponential reproduction I mean that a linear sum of the shifts of the kernel reproduce an exponential:

$$-\sum_{n \in \mathbb{Z}} c_{m,n} \psi(t - nT) = e^{\alpha_m t} \tag{4.25}$$

where $\alpha_m = \alpha_0 + m\lambda \in \mathbb{C}$, $m \in [0, \dots, M]$, and $c_{m,n} \in \mathbb{C}$. Polynomial reproduction can also be used but for the sake of brevity I will focus on exponential reproduction. Exponential and polynomial reproduction can be achieved when $\psi(t)$ satisfies the generalised Strang-Fx conditions [41]. One can leverage the exponential reproduction property of the kernel to generate signal moments,

$$\begin{aligned}
s_m &= \sum_{n \in \mathbb{Z}} c_{m,n} g_n = \left\langle \sum_{k=1}^K a_k \delta(x - x_k), -\sum_{n \in \mathbb{Z}} c_{m,n} \psi(x - nT) \right\rangle \\
&= \left\langle \sum_{k=1}^K a_k \delta(x - x_k), e^{\alpha_m x} \right\rangle \\
&= \sum_{k=1}^K a_k e^{\alpha_m x_k},
\end{aligned} \tag{4.26}$$

where we consider a signal with a finite number of edges, K . The parameters x_k and a_k can be recovered through various methods including Prony's method [11,42], sparse optimisation methods

[10], subspace methods [41,43–46], and model fitting approaches [47–50]. Using Prony’s method, $g(x)$ can be uniquely determined when $M \geq 2K - 1$ by noting that Eq. (4.26) is annihilated by a filter,

$$H(z) = \prod_{k=1}^K (1 - e^{\lambda x_k} z^{-1}) , \quad (4.27)$$

such that $h_n * s_n = 0$ where h_n are the coefficients of the filter. The coefficients h_n can be determined from a Yule-Walker system and the parameters x_k determined from the roots of the filter. The parameters a_k can then be determined by least squares approach.

If there are an infinite number of edges [39] states that $g(x)$ is uniquely determined when the extra condition that there are at most K Diracs in any $2KC$ interval is satisfied. From Property 1 we know that a bound on an interval in amplitude-domain corresponds to a bound on the integral over the corresponding interval in time-domain.

$$\int_{t_n}^{t_{n+K}} h(t) dt > 2KC . \quad (4.28)$$

Given $g(x)$, the time-domain signal, $h(t)$, can then be retrieved by application of the transform $\mathcal{T}$. Note that the transformation between domains can be achieved parametrically for piecewise constant signals.

□

Theorem 4 demonstrates one of the desirable properties of an integrate-and-fire approach when taken from an FRI perspective. The maximum number of edges that can be resolved by a particular multichannel TDIF is dependent on the magnitude of the signal. The larger the signal the more edges can be resolved and the smaller the signal the fewer the edges that can be resolved. This adaptive nature is present in previous work on piecewise signals [15,63,67], however, it is often obscured by stricter conditions that must be met simultaneously.

Non-negative piecewise constant signals in time-domain

Another desirable characteristic of an integrate-and-fire time encoding framework is that when a signal is quiescent, the sampler does not produce any samples [15,63,67]. Up until now I have assumed that our signals are strictly positive. This assumption can be relaxed to allow intervals where $h(t) = 0$. Such intervals correspond to Dirac delta functions in the signal in the opposite domain. This result is summarised in Property 5.

Property 5 (Automorphism on non-negative piecewise constants). *Let $h(x) \in \mathcal{H}$ where $\mathcal{H}$ is the set of sums of a piecewise constant and Dirac deltas of the form,*

$$h(x) = \sum_k c_k u(x - \tau_k) + d_k \delta(x - \tau_k) , \quad (4.29)$$

and where $h(t) \geq 0$ is supported on a disconnected subset of the real line where $u(x)$ is the Heaviside step function and $\delta(x)$ is the Dirac delta function and $d_k \geq 0$. The transform $\mathcal{T}$ is then an automorphism on $\mathcal{H}$ such that $\mathcal{T}h(x)$ is also a sum of a piecewise constant and Dirac deltas of the form of Eq. (4.29). When $h(x)$ is a simple piecewise constant ($d_k = 0$), the support of $\mathcal{T}h(x)$ is the entire real line.

Proof. The correspondence of intervals of zero value in one domain and Dirac deltas in the other domain can be shown either by restricting the domain of $f(t)$ to the support of $h(t)$ or by appealing to a limiting argument based on Property 4. I will use the latter argument.

Consider, a single piece of a positive piecewise constant signal on an interval with length d and with height ϵ . According to Property 4, the area under the corresponding piece in amplitude-domain is d and has a support of length $d\epsilon$. In the limit $\epsilon \rightarrow 0$ this is a standard representation of the Dirac delta function scaled by d . Thus, intervals where $h(t) = 0$ correspond to Dirac deltas scaled by the length of the interval. By the symmetry of the transform, the reverse is also true. Dirac deltas in $h(t)$ correspond to intervals where $g(x) = 0$.

□

An example of a piecewise constant signal which equals zero for a period is plotted in Fig. 4.9 along with its amplitude-domain equivalent. Intuitively, if the time-domain signal had Dirac deltas




Figure 4.9: An example of a non-negative piecewise constant time-domain signal and its resulting amplitude-domain signal that is the sum of a strictly positive piecewise constant signal and Dirac deltas.

as well, these would correspond to periods of zero in the amplitude-domain equivalent. Property 5 yields a sampling theorem for non-negative piecewise constant signals in Theorem 5 where I once again leverage [39].

Theorem 5. *Let $h(t) \in \mathcal{H}$ be a non-negative piecewise constant signal. Given a time encoding θ_n of $h(t)$ generated by a multichannel TDIF, $h(t)$ can be exactly reconstructed when the dilation function $\varphi(x) = \psi'(x)$, where $\psi(x)$ can reproduce $N \geq 4K - 1$ exponentials, and the integral between any K edges, $n_k, \dots, t_{n+K}$ is at least,*

$$\int_{t_n}^{t_{n+K}} h(t) dt > 4KC, \quad (4.30)$$

where C is the threshold of the TDIFs and N is the number of channels.

Proof. From Property 5, we know that if $h(t) \in \mathcal{H}$ then $\mathcal{T}h(x) = g(x) \in \mathcal{H}$. Let,

$$g(x) = \sum_k a_{0,k} u(x - x_k) + a_{1,k} \delta(x - x_k) , \quad (4.31)$$

where $a_{r,k} \in \mathbb{R}$. Similar to Theorem 4, the amplitude-domain uniform samples generated by a L -channel TDIF with threshold C and $\varphi(x) = \psi'(t)$ are,

$$\begin{aligned} g_n &= \left\langle \sum_{k=1}^K a_{0,k} u(x - x_k) + a_{1,k} \delta(x - x_k), \psi'(x - nT) \right\rangle \\ &\stackrel{(a)}{=} \left\langle \sum_{k=1}^K a_{r,0} \delta(x - x_k) + a_{1,k} \delta^{(1)}(x - x_k), -\psi(x - nT) \right\rangle , \end{aligned} \quad (4.32)$$

where $T = C/L$ is the sampling period and (a) follows from integration by parts, and $\delta^{(1)}$ is a differentiated Dirac. Taking a linear combination of signal moments, we can produce a sparse estimation problem as before:

$$g_n = \sum_{k=1}^K a_{r,0} e^{\alpha_m t_k} - a_{r,1} \alpha_m e^{\alpha_m t_k} . \quad (4.33)$$

$g(n)$ can be determined using Prony's method when the annihilating filter has doubled roots,

$$H(z) = \prod_{k=1}^K (1 - e^{\lambda x_k} z^{-1})^2 , \quad (4.34)$$

and when the $\psi(x)$ can reproduce exponentials up to a degree $M \geq 4K - 1$. When there are an infinite number of edges, [39] states that $g(x)$ can be uniquely determined from the samples g_n if there are at most K differentiated Diracs of the same order in any $4KC$ interval is met. From Property 1 we know that a bound on an interval in amplitude-domain corresponds to a bound on the integral over the corresponding interval in time-domain.

$$\int_{t_n}^{t_{n+K}} h(t) dt > 4KC . \quad (4.35)$$

Given $g(x)$, the time-domain signal, $h(t)$, can then be determined by application of the transform $\mathcal{T}$.

□

A TDIF sampling architecture cannot produce samples during intervals where the input signal is zero since the accumulated integral will never surpass the trigger mark. Due to the dilating nature of the sampling process, the samples generated are also not an accurate measure of when samples are generated. Despite these two behaviours, Theorem 5 gives a method for reconstructing piecewise constant signals that have intervals of arbitrary length where the signal is zero. This is possible because time measurement ($T(t)$ in Fig. 4.4) is still accumulated during the zero-valued intervals. In fact, each channel of the sampler accumulates at different rates determined by their own value of $f(t)$ which are different due to their offsets. In a typical multi-channel integrate-and-fire setting [15,30], time is accumulated at an equal rate by each channel.

I note that the reconstruction of a non-negative piecewise constant is less stable due to the requirement to resolve doubled roots in the annihilating filter. This can be ameliorated to some extent by methods more stable than Prony's but this is beyond the scope of this work.

A natural question to ask is whether the reconstruction method can be extended to piecewise constant signals that run both positive and negative. A simple extension of the sampling architecture can be made to accommodate this when one considers that any piecewise constant signal can be decomposed into a non-negative and a non-positive piecewise constants signal. Thus, two multi-channel TDIF samplers can be used in parallel, one which only samples the positive component and another which only samples the negative component. The positive and negative components of the signal must then satisfy the conditions from Theorem 5 separately.

4.6 Summary

In summary, in this chapter I have considered time encoding from the perspective of uniform sampling in an equivalent 'amplitude-domain' representation.

State-of-the-art methods for time encoding and reconstruction from time encodings revolve around mapping the time encoding problem to one of non-uniform sampling in *time-domain*. Indeed, this is the approach followed in Chapter 3. To some degree, one could argue that the apple has not fallen far from the tree and, consequently, the potential benefits of a time encoding approach are diminished. With the TDIF architecture we can now move away from reconstructing signals in time-domain with non-uniform sampling approaches and towards reconstructing signals directly in amplitude-domain.

This chapter has shown that with the appropriate hardware implementation where time samples are manipulated during acquisition, one can not only uniformly sample in amplitude-domain with a fixed kernel, as in [23,24,64], but also with any arbitrary kernel of compact support (Theorem 3). This significantly broadens the uses of the amplitude-domain model.

The full benefits of such an approach are as yet largely unexplored but Theorem 4 and Theorem 5 demonstrate that there are simple examples where the state of the art in time encoding can be improved upon. They show that in the case of time encodings of piecewise constant signals, exact reconstructions can be achieved with simple constraints on the piecewise constant in contrast to the complicated restrictive constraints and approximate reconstructions of Chapter 3 and [15]. Moreover, reconstruction is achieved whilst maintaining the desirable characteristics of IF-TEMs, in particular that of quiescence where samples are not generated when the input signal is zero.

As to the unexplored, they present several interest academic questions and avenues for further research. There are several potential benefits of the TDIF that are immediately apparent. Firstly, it can be clearly seen that the class of a signal is not generally the same in time-domain and amplitude-domain since the transformation between the two domains is highly non-linear. For example, a bandlimited signal in amplitude-domain is necessarily non-bandlimited in time-domain. This begs the question, what practical time-domain signals can be represented by standard signal models in amplitude-domain. Is the amplitude-domain approach a way to model highly non-linear classes of signals?

Secondly, since the TDIF produces samples that are uniform in amplitude-domain, one could read-

ily apply multi-resolution decompositions to the amplitude-domain samples. The meaning of such decompositions in time-domain is intuitive, but the meaning in amplitude-domain is an open question.

Finally, the TDIF may present a more natural way to leverage the modulo sampling paradigm [32]. A deeper explanation of this is out of the scope of this work, but suffice it to say, the hardware implementation of timing circuitry could be naturally adapted to implement the modulo operation itself. For example, the simplest of timing circuits, a clock driving a counter, can be turned into a modulo sampler simply by restricting the bit depth of the counter.

Chapter 5

Networks of time encoding machines

In Chapter 3 I looked at how signals with finite rate of innovation can be reconstructed from their time encodings derived from a prototypical IF-TEM. I then showed in Chapter 4 how by using a modified architecture, namely the TDIF, we can now use alternative reconstruction strategies leveraging existing uniform sampling theory. In both cases I used multichannel encoding machines where each channel consisted of a single TEM. In this chapter I will move beyond simple multichannel systems and look at ways networks of multiple TEMs can be used to gain desirable sampling behaviours.

In Section 3.8, I showed that FRI signals can be recovered directly from finite differences of the input signal. Although less stable than reconstruction from cumulative samples, this method may be required when the resetting characteristics of the TEMs are unknown. In Section 5.2 I will consider how pairs of matched TEMs can be used to generate finite differences while mitigating the effect of non-ideal and unknown resetting characteristics.

I will also propose a network in Section 5.3 for the purpose of synchronisation of TEMs in amplitude-domain. This can have a couple of potential applications. First, the multichannel approach of the TDIF architecture as well as multichannel approaches to sampling with hyperbolic

kernels require synchronisation of sampling in the amplitude domain. This is a similar problem to synchronising digital clocks in synchronous circuits. Second, even when synchronisation is not required, coarser synchronisation can be used to ensure different TEMs do not fire simultaneously or in quick succession.

5.1 Inhibitory connections

One phenomenon that can be observed in biological neurons is that some synaptic connections are inhibitory. This means that they further polarise the membrane potential of the neuron increasing the amount of excitatory input required to depolarise a neuron and cause it to generate an action potential. It seems that inhibitory connections form a crucial component of a neural network's computational capabilities. I will use analogous mechanisms of inhibition in the construction of the networks presented in this chapter. I model inhibition through two different mechanisms:

1. **Simple inhibition** - Inhibiting spikes are subtracted from the accumulated integral of the TEM.
2. **Resetting inhibition** - For each inhibiting spike, the accumulated integral of the TEM is reset to zero as if it had fired.

Simple inhibition treats each inhibitory spike as a discrete packet of energy that is removed from the accumulated integral. On the other hand, resetting inhibition triggers the TEMs integral reset mechanism. Both models are illustrated in Fig. 5.1.

Note, both forms of inhibition are non-linear. This is obvious in the case of resetting inhibition. For simple inhibition, this is because I defined the accumulated integral to be non-negative. As such, if the inhibiting spike exceeds the accumulated integral then the integral accumulated will bottom out at $f(\tau_{\text{inhib}}) = 0$ as can be seen in Fig. 5.1(c).




Figure 5.1: Types of inhibition: **a)** an example uninhibited cumulative integral $f(t)$, **b)** a pair of inhibiting spikes of different amplitude, **c)** the original cumulative integral after *simple* inhibition with the spikes of b), **d)** the original cumulative integral after *resetting* inhibition with the spike of b). Changes in the integral in the simple case are equal to the spike amplitudes.

5.2 Differential pairs

Up to now, I have largely considered only ideal IF-TEMs. Specifically, I have assumed that the integral gets reset instantaneously, begin integrating immediately following a reset, and reset to a known value, typically zero. In a practical implementation of an IF-TEM it is likely that the resetting process is not instantaneous nor is the reset value known. This could be due to physical constraints such as the rate at which a capacitor can be discharged, or due to manufacturing uncertainty. Indeed, biological neurons exhibit a behavior called a refractory period where the neuron becomes desensitised to further stimulus following an action potential. Taking into account these non-idealities, each pair of consecutive spikes from an IF-TEM can represent a finite difference of the form,

$$z_n = C = c + \int_{\tau_{n-1}+r}^{\tau_n} y(t) dt = c + f(\tau_n) - f(\tau_{n-1} + r) , \quad (5.1)$$

where r is the length of time the IF-TEM takes to reset and c is the value of the integral immediately following a reset. The parameters r and c are non-idealities that may not be known. If this is the case, these finite differences cannot be used for reconstruction.

The non-ideal parameters of Eq. (5.1) can be eliminated using a network of two IF-TEMs that I call a differential pair. Such a network is illustrated in Fig. 5.2 and consists of two IF-TEMs with trigger marks C_a and C_b where $C_b > C_a$. I will refer to the IF-TEMs by their respective trigger marks. I also assume that the IF-TEMs are unipolar in that they only have positive trigger marks.

The integrals of the IF-TEMs for a constant input signal are plotted in Fig. 5.3 against time for the case where both IF-TEMs start from the same initial integral. When this is the case, it is obvious that C_a will always fire first as its trigger mark is lower. After firing, C_a resets and will continue integrating starting from the reset value c and may fire before C_b . In fact, C_a will fire M times before C_b according to,

$$M = \left\lfloor \frac{C_b}{C_a} \right\rfloor , \quad (5.2)$$




Figure 5.2: A simple differential pair consisting of two IF-TEMs with trigger marks C_a and C_b . The filtering component of the IF-TEMs has been separated and combined into φ for clarity in later developments of the circuit. Inhibitory connections are represented by connection with circles.




Figure 5.3: The integrals of a differential pair with global reset against time.

where $\lfloor \cdot \rfloor$ rounds down to the nearest integer. In general, multiple spikes from C_a can be handled by selecting only the first of the M spikes either after the fact or through a post processing network of IF-TEMs. Otherwise, if $C_b < 2C_a$ then C_a will only spike once before C_b and such considerations need not be made. I will assume that $C_b < 2C_a$ and C_a spikes only once before C_b .

When C_b fires, its output is fed back C_a as a *resetting* inhibitory input that forces C_a to reset simultaneous to C_b resetting. This returns the network to its original starting condition where the integrals of the two IF-TEMs are equal. In this case I have shown the inhibiting feedback only to C_a . One could also feed the inhibiting signal to both IF-TEMs simultaneously. This could be warranted in cases where there is a delay between the neuron resetting and outputting a spike.

When C_a and C_b have the same non-ideal parameters, c and r , I call them well matched. When they are well matched the non-ideal parameters can be eliminated by taking the difference of their samples,

$$\begin{aligned} C_b - C_a &= C_{\text{diff}} = f(\tau_n) - f(\tau_{\text{reset}}) - c - (f(\tau'_n) - f(\tau_{\text{reset}}) - c) \\ &= f(\tau_n) - f(\tau'_n) , \end{aligned} \tag{5.3}$$

where τ_n and τ'_n are the firing times of C_a and C_b respectively, and $\tau_{\text{reset}} = \tau_{n-1} + r = \tau'_{n-1} + r$ (both IF-TEMs have the same resetting time). Eq. (5.3) can be graphically verified in Fig. 5.3 as the integral of the signal clearly changes by $C_b - C_a$ over the interval $[\tau_n, \tau'_n]$. The example considers only a constant input signal, however, the result is general to any input signal.

The differential pair allows uncertainties in the non-ideal resetting characteristics of IF-TEMs to be eliminated when the constituent IF-TEMs are well-matched. The use of a differential output to reject the common mode (the resetting value c) is a common technique in analogue circuit design. Matching of components, for example in differential amplifiers and current mirrors, is also common practice [69]. This is typically a consideration for the manufacturing process. For example, transistors manufactured on the same dye in close proximity are likely to be well matched.




Figure 5.4: Two ring synchronisers with single and triple feedback. Arrows indicate inhibitory feedback.

5.3 Amplitude domain synchronisation

The multichannel of TDIFs presented in Chapter 4 relies on each constituent TDIF maintaining a particular offset in the integral, $f(t)$, from the other channels. This translates to uniform spacing of the amplitude-domain signals measured by each channel and is analogous to maintaining synchronisation between clocks in digital circuits. This kind of synchronisation in amplitude-domain can be achieved by means of a network of IF-TEMs called a ring synchroniser.

The operation of the synchroniser is simple and intuitive but relies on a few assumptions about the operation of the IF-TEMs. The M constituent IF-TEMs of a ring synchroniser are labelled in order $[0, \dots, M - 1]$ and have the same trigger mark, C . I also assume that when an IF-TEM spikes, it generates a spike with energy equivalent to half its trigger mark, $C/2$. Finally, I assume that the integrals of the IF-TEMs are always positive as is the case with TDIFs.

When the m^{th} IF-TEM fires, its output is subtracted from its $M - 2$ predecessors to prevent them firing in quick succession. By the n^{th} predecessor of the m^{th} IF-TEM I mean the IF-TEM with label $m - n \bmod M$. Consequently, we can say that if the m^{th} IF-TEM has fired, then its neighbours $M - 2$




Figure 5.5: Integrals of constituent IF-TEMs in ring synchronisers for a constant input signal. The aggregated spike pattern is shown on the bottom row in red.

predecessors must be at least $C/2$ below their trigger marks. On the other hand, its $M-1$ predecessor, which is also its successor, will not have been inhibited and thus can be closer than $C/2$ to its trigger mark. In such a condition, the successor IF-TEM is then the next to fire. Thus, the firing of IF-TEMs proceeds in sequence clockwise around Fig. 5.4.

Examples of how the integrals of the IF-TEMs develop over time are shown in Fig. 5.5. The initial conditions of the constituent IF-TEMs are set randomly and yet the network quickly converges on a uniform and ordered firing steady state. It can be shown that this is the only stable state of the network through two properties. It is important to note that each IF-TEM is inhibited $M-1$ times keeping the integrals of each IF-TEM below $C/2$.

Property 6 (Ring Synchroniser Order). *Given any signal, the ring synchroniser will converge on a steady state where the constituent IF-TEMs will fire in a sequence increasing order.*

Proof. Let the $f_m(t)$ be the integrals of M constituent IF-TEMs. First, note that if the IF-TEMs fire in sequence, then they will remain firing in this sequence. When the i^{th} IF-TEM fires, all IF-TEMs except

the $i+1 \bmod M$ IF-TEM will have integrals $f_m(t) < C/2$ due to inhibition. IF-TEM will have integrals $f_m(\tau_1 + \epsilon) < C/2$ due to inhibition. The $i+1 \bmod M$ IF-TEM will have an integral $f_m(\tau_1 + \epsilon) > C/2$ and thus will be the next to fire.

Now we can show that if two IF-TEMs fire out of sequence, the order of their integrals is swapped. Let the i^{th} IF-TEM fire at τ_1 such that $f_i(\tau_1 + \epsilon) = 0$ then all IF-TEMs except the $j = i+1 \bmod M$ IF-TEM will have integral $f_m(\tau_1 + \epsilon) < C/2$. If the $k = i+n \bmod M$ IF-TEM is the next to fire at τ_2 , where $1 < n < M-1$, then $f_k(\tau_2 + \epsilon) = 0$. Moreover, $f_j(\tau_2 - \epsilon) > C/2$ and $f_j(\tau_2 + \epsilon) > 0 = f_k(\tau_2 + \epsilon)$. Therefore, the j^{th} IF-TEM will fire before the k^{th} IF-TEM and the order of firing has swapped. $\square$

Property 7 (Ring Synchroniser Uniform Firing). *In the steady state, the ring synchroniser must fire uniformly. That is, consecutive spikes in the aggregate spike pattern are equally spaced for a constant input signal.*

Proof. Note that from Property 6 that in the steady state, the IF-TEMs must fire in sequence. Also, if the input signal is constant $y(t) = a$ then a uniform firing of an M IF-TEM ring synchroniser has spikes with a separation of C/Ma . Without loss of generality, assume $a = 1$. Let the i^{th} IF-TEM fire prematurely at $\tau_n + \delta$ where $\delta < C/M$ and τ_n is the last spike time. The $k = i+n \bmod M$ IF-TEMs, where $1 < n < M-1$ will have integrals $f_k(\tau_n + \delta - \epsilon) = C/2 - \delta$ and $f_k(\tau_n + \delta + \epsilon) = 0$. Thus the integrals of these IF-TEMs is reduced by $C/2 - \delta$ and will therefore fire earlier by δ compensating for the premature firing. $\square$

5.4 Summary

In this chapter, I have briefly presented two networks of time encoding machines that are directly applicable to time encodings. The first, a pair of TEMs operating in a differential mode, provide a way to mitigate the non-idealities that could potentially affect a single TEM in practice. By operating in a differential mode, issues such as slow resets, resetting to an unknown value, or refractory

periods can be cancelled out if the constituent TEMs are well-matched. Reconstruction from the samples generated by such a network leverages the algorithm described in Section 3.8.

The second TEM network, which I call a ring synchroniser, presents a way to achieve synchronisation in amplitude-domain similar to how it is common to need clock synchronisation in time-domain. This may be important in the implementation of the TDIF described in Chapter 4 where it is assumed that strict synchronisation in the amplitude-domain is maintained across the channels of a multi-channel system.

Chapter 6

Conclusion

In this thesis I have made contributions to several aspects of time encoding. In particular, in Chapter 3 we looked at reconstruction of FRI signals from time encodings using novel sampling kernels including the α -synaptic function and hyperbolic kernels. I showed that the α -synaptic kernel allows local exponential reproduction despite not satisfying the Generalised Strang-Fix properties. Exponential spline kernels used in previous works contain multiple non-smooth knots. Consequently, time encoded samples need to be generated in highly restrictive knotless intervals in order to allow reconstruction. On the other hand, the α -synaptic kernel has a single knot meaning these constraints become more tractable. In particular, I derived constraints on the input signal that guarantee reconstruction even when the input consists of a mix of positive and negative Diracs.

I showed that sampling and reconstruction of FRI signals with a hyperbolic kernel can be achieved through a rational approximation approach. I also demonstrated that this can be achieved even when the initial conditions of the IF-TEM are unknown. The hyperbolic kernels represent entirely smooth kernels and so samples generated by an IF-TEM are not restricted to any particular intervals like in the exponential spline and α -synaptic kernel cases.

In Chapter 4, a perspective on time encoding was taken which considered integrate-and-fire time encoding of a time-domain signal as a uniform sampling of an equivalent amplitude domain sig-

nal. Building upon this I showed that a novel time encoding architecture, that we call a time dilating integrate-and-fire encoder, can be used to sample and reconstruct a number of different signal classes from their amplitude-domain samples. This was achieved modifying the typical IF-TEM architecture so that time measurements are manipulated (dilated) in a signal driven manner. By appropriate choice of this signal driven dilation, I showed that this is equivalent to shaping the sampling kernel in amplitude-domain. By incorporating multiple such TDIFs into a single sampler, one can then sample and reconstruct signals in shift-invariant spaces in amplitude-domain as well as piecewise constant signals in time-domain. The latter example of piecewise constant signals demonstrated the suitability of the TDIF architecture to sampling signals which have periods of quiescence.

Finally, I introduced some networks of time encoding machines that address some of the non-idealities and difficulties of real implementations of TEMs. In particular, I developed a differential pair network that uses two well-matched TEMs to eliminate non-ideal resetting characteristics from the measured samples. I also presented a synchronising network that allows synchronising of the channels of a multichannel TEM.

6.1 Future directions

The time dilating integrate-and-fire approach, as presented in this thesis, is still very much in its infancy and many questions still are unanswered both in terms of the theoretical points of reconstruction from amplitude-domain and potential improvements of the sampling architecture. I will enumerate several areas that may be worth exploring below.

Noise and jitter analysis

Throughout this thesis I have shied away from explicitly analysing the robustness of these approaches to noise. This, to some degree, is deliberate since many of the techniques of Chapter 3 leverage existing non-uniform reconstruction algorithms which have been the subject of significant study regarding their robustness to noise. Moreover, it is my belief that advances in architectures

of time encoding are a more prescient concern. Finding areas where time encoding can give us new capabilities, such as quiescence in sampling and new signal representations will further the cause of time encoding more considerably than marginal gains in robustness.

This thinking motivated the introduction of the TDIF. In the case of the TDIF, analysis of noise and its effect on sampling may be of greater interest than simply quantifying the robustness of the reconstruction. To see this, consider the effect of noise on the input signal to a TDIF. Deviations in the amplitude of the signal accumulate in the integrated measure of time (the dilated time measurement) as well as in the timing of the spike and reset mechanism. Both effects represent not noise but jitter in the amplitude-domain. Similarly, distortions in the timing of the input signal (perhaps we might call this jitter in a continuous domain) represents noise in the amplitude-domain. Analysis of these characteristics and what that might mean for the robustness of such a system is an open question to be pursued.

Modulo reconstructions in time domain

A recent trend in signal processing is the modulo or unlimited sampling approach [32]. The advantage of this approach is that it provides an advantageous tradeoff between quantisation resolution and dynamic range. This is achieved by folding the signal before sampling such that the amplitude of samples is always within a fixed interval. In fact, the folding operation is similar to the resetting of the integral of an IF-TEM. It has been shown that some signals can be reconstructed from these folded measurements.

The folding of an amplitude-domain signal, that is folding time measurements, is potentially a very natural operation. In a practical TDIF implementation, the time measurements would be made by a digital counter that counts up and down depending on the clock frequency (voltage controlled frequency in the case of a TDIF). Folded time measurements in amplitude domain can be obtained simply by using a smaller number of bits. For example, for a 10.24kHz clock, a 10-bit counter will wrap around every 0.1s. Due to this, the TDIF architecture may be well suited to the modulo sampling approach.

Amplitude domain approximation and non-linear preprocessing

Bandlimited signals in time-domain are necessarily represented by non-bandlimited signals in amplitude domain. Similarly, shift-invariant signals in time-domain are not shift-invariant in amplitude-domain. As such perfect reconstruction cannot be achieved with standard reconstruction algorithms for uniform samples, however, approximate results are still achievable. For example, in the case of bandlimited signals, [64] shows that good reconstruction can be achieved when the essential bandwidth is sufficient. It would be interesting to analyse the approximation properties of amplitude-domain reconstruction in finer detail and to extend it to shift-invariant signals.

This could also inform investigation into non-linear preprocessing techniques. As indicated by the use of the essential bandwidth criterion in [64], the amplitude-domain approach is related to frequency modulation. Companding is a non-linear amplification technique used to maximise the dynamic range of a message signal in frequency modulation. Perhaps such non-linear preprocessing steps could be used to similarly minimise the essential bandwidth of an amplitude-domain signal.

True amplitude domain filtering

The TDIF of Chapter 4 imitates filtering through the combination of non-amplification and multiple channels. This has several drawbacks. Principal amongst these are implementation complexity, due to the requirement for multiple identical channels, and the restriction to kernels with compact support. A natural question that arises is whether amplitude-domain filtering can be directly implemented in hardware in the similar way to how filters can be implemented in time-domain. If this is achievable, it may render multichannel systems unnecessary.

Interpretation of multi-resolution analysis in amplitude-domain

Multi-resolution techniques such as wavelet decompositions are readily applicable to amplitude-domain uniform samples generated by a TDIF. However, the interpretation of such decompositions is less obvious. For instance, signals are decomposed across different amplitude scales and each the

size of each coefficient represents the length of the corresponding timespan. The precise meaning of this is unclear.

Appendix A

Innovation amplitudes & phase without least-squares

In an FRI problem, once the time parameters have been estimated, the most straightforward way to estimate the amplitude parameters is via a least-squares estimation. It is worth noting that this is not the only way to achieve the estimation. In this appendix I briefly demonstrate that since the problem is one of interpolation, the amplitude parameters can be calculated directly from the time parameters.

In the simplest case, consider a stream of Diracs filtered with a hyperbolic tangent:¹

$$y(\tau) = \sum_{k=1}^K x_k \tanh(\tau - t_k) = \sum_{k=1}^K 2x_k \frac{e^{-2t_k} e^{2\tau} - 1}{e^{-2t_k} e^{2\tau} + 1}. \quad (\text{A.1})$$

For simplicity, let $\gamma_k = 2x_k$, $\alpha_k = e^{2t_k}$, and $z = e^{2\tau}$, then:

$$y(z) = \sum_{k=1}^K \gamma_k \frac{z - \alpha_k}{z + \alpha_k} = \frac{P(z)}{Q(z)}, \quad (\text{A.2})$$

where,

¹Equivalent to a signal filtered with $\varphi(t) = \text{sech}^2$ and integrated

$$Q(z) = \prod_{k=1}^K (z + \alpha_k) , \quad (\text{A.3})$$

and,

$$P(z) = \sum_{k=1}^K \gamma_k (z - \alpha_k) \prod_{l \neq k} (z + \alpha_l) . \quad (\text{A.4})$$

Similar to Lagrange interpolation, Eq. (A.4) shows that the values of $P(z)$ at any of the pole positions, $-\alpha_k$, are dependent only on a single term of the summation. For any $P(-\alpha_k)$ all but one term is equal to zero:

$$P(-\alpha_k) = -2\alpha_k \gamma_k \prod_{l \neq k} (\alpha_l - \alpha_k) , \quad (\text{A.5})$$

and,

$$\gamma_k = \frac{-P(-\alpha_k)}{2\alpha_k \prod_{l \neq k} (\alpha_l - \alpha_k)} . \quad (\text{A.6})$$

Thus, given an estimate of the time parameters, α_k , derived by some method such as rational approximation or a Prony-like method, the amplitude parameters can be directly calculated by Eq. (A.6).

References

- [1] V.A. Kotel'nikov, On the transmission capacity of the “ether” and wire in electrocommunications, in: J.J. Benedetto, P.J.S.G. Ferreira (Eds.), *Modern Sampling Theory: Mathematics and Applications*, Birkhäuser Boston, Boston, MA, 2001: pp. 27–45. https://doi.org/10.1007/978-1-4612-0143-4_2.
- [2] C.E. Shannon, Communication in the presence of noise, *Proceedings of the IRE* 37 (1949) 10–21. <https://doi.org/10.1109/JRPROC.1949.232969>.
- [3] M. Unser, Sampling-50 years after Shannon, *Proceedings of the IEEE* 88 (2000) 569–587. <https://doi.org/10.1109/5.843002>.
- [4] B. Alberts, *Molecular biology of the cell*, 6th ed., Garland Publishing, New York, NY, 2014.
- [5] Delbruck, T, Frame-free dynamic digital vision, (2008). <https://doi.org/10.5167/UZH-17620>.
- [6] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A.J. Davison, J. Conradt, K. Daniilidis, D. Scaramuzza, Event-based vision: A survey, *IEEE Transactions on Pattern Analysis & Machine Intelligence* 44 (2022) 154–180. <https://doi.org/10.1109/TPAMI.2020.3008413>.
- [7] H.J. Landau, Necessary density conditions for sampling and interpolation of certain entire functions, *Acta Mathematica* 117 (1967) 37–52. <https://doi.org/10.1007/BF02395039>.
- [8] A. Beurling, P. Malliavin, On the closure of characters and the zeros of entire functions, *Acta Mathematica* 118 (1967) 79–93. <https://doi.org/10.1007/bf02392477>.

- [9] A. Aldroubi, K. Gröchenig, Nonuniform sampling and reconstruction in shift-invariant spaces, *SIAM Review* 43 (2001) 585–620. <https://doi.org/10.1137/s0036144501386986>.
- [10] E.J. Candes, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, *IEEE Transactions on Information Theory* 52 (2006) 489–509. <https://doi.org/10.1109/TIT.2005.862083>.
- [11] M. Vetterli, P. Marziliano, T. Blu, Sampling signals with finite rate of innovation, *IEEE Transactions on Signal Processing* 50 (2002) 1417–1428. <https://doi.org/10.1109/TSP.2002.1003065>.
- [12] Y. Chi, M. Ferreira Da Costa, Harnessing sparsity over the continuum: Atomic norm minimization for superresolution, *IEEE Signal Processing Magazine* 37 (2020) 39–57. <https://doi.org/10.1109/MSP.2019.2962209>.
- [13] B.F. Logan, Information in the zero crossings of bandpass signals, *The Bell System Technical Journal* 56 (1977) 487–510. <https://doi.org/10.1002/j.1538-7305.1977.tb00522.x>.
- [14] K. Adam, Timing is everything, PhD thesis, École polytechnique fédérale de Lausanne, 2022.
- [15] R. Alexandru, P.L. Dragotti, Reconstructing classes of non-bandlimited signals from time encoded information, *IEEE Transactions on Signal Processing* 68 (2020) 747–763. <https://doi.org/10.1109/TSP.2019.2961301>.
- [16] A.A. Lazar, L.T. Toth, Time encoding and perfect recovery of bandlimited signals, in: 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)., 2003: pp. VI–709. <https://doi.org/10.1109/ICASSP.2003.1201780>.
- [17] C. Kikkert, D.J. Miller, Asynchronous delta sigma modulation, *Proc. IREE* 36 (1975) 83–88.
- [18] W. Chen, Asynchronous sigma delta modulators for data conversion, PhD thesis, Imperial College London, 2013.
- [19] W. Chen, C. Papavassiliou, A new decoding solution for the asynchronous sigma delta modulator, in: Proceedings of the 2013 9th Conference on Ph.d. Research in Microelectronics and Electronics (PRIME), 2013: pp. 49–52. <https://doi.org/10.1109/PRIME.2013.6603115>.

- [20] J. Daniels, W. Dehaene, M. Steyaert, A. Wiesbauer, A/d conversion using an asynchronous delta-sigma modulator and a time-to-digital converter, in: 2008 IEEE International Symposium on Circuits and Systems, 2008: pp. 1648–1651. <https://doi.org/10.1109/ISCAS.2008.4541751>.
- [21] A.A. Lazar, L.T. Toth, Perfect recovery and sensitivity analysis of time encoded bandlimited signals, IEEE Transactions on Circuits and Systems I: Regular Papers 51 (2004) 2060–2073. <https://doi.org/10.1109/TCSI.2004.835026>.
- [22] A.A. Lazar, Time encoding with an integrate-and-fire neuron with a refractory period, Neurocomputing 58-60 (2004) 53–58. <https://doi.org/10.1016/j.neucom.2004.01.022>.
- [23] P. Martínez-Nuevo, H.-Y. Lai, A.V. Oppenheim, Amplitude sampling, in: 2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2016: pp. 17–22. <https://doi.org/10.1109/ALLERTON.2016.7852205>.
- [24] P. Martínez-Nuevo, H.-Y. Lai, A.V. Oppenheim, Delta-ramp encoder for amplitude sampling and its interpretation as time encoding, IEEE Transactions on Signal Processing 67 (2019) 2516–2527. <https://doi.org/10.1109/TSP.2019.2904027>.
- [25] D. Gontier, M. Vetterli, Sampling based on timing: Time encoding machines on shift-invariant subspaces, Applied and Computational Harmonic Analysis 36 (2014) 63–78. <https://doi.org/10.1016/j.acha.2013.02.002>.
- [26] A.A. Lazar, E.K. Simonyi, L.T. Toth, Fast recovery algorithms for time encoded bandlimited signals, in: Proceedings. (ICASSP '05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005., 2005: pp. iv/237–iv/240 Vol. 4. <https://doi.org/10.1109/ICASSP.2005.1415989>.
- [27] K. Gröchenig, A new approach to irregular sampling of band-limited functions, in: J.S. Byrnes, J.L. Byrnes (Eds.), Recent Advances in Fourier Analysis and Its Applications, Springer Netherlands, Dordrecht, 1990: pp. 251–260. https://doi.org/10.1007/978-94-009-0665-5_16.

- [28] N.T. Thao, D. Rzepka, Time encoding of bandlimited signals: Reconstruction by pseudo-inversion and time-varying multiplierless FIR filtering, *IEEE Transactions on Signal Processing* 69 (2021) 341–356. <https://doi.org/10.1109/TSP.2020.3043809>.
- [29] N.T. Thao, D. Rzepka, M. Miśkowicz, Bandlimited signal reconstruction from leaky integrate-and-fire encoding using POCS, *IEEE Transactions on Signal Processing* 71 (2023) 1464–1479. <https://doi.org/10.1109/TSP.2023.3256269>.
- [30] K. Adam, A. Scholefield, M. Vetterli, Sampling and reconstruction of bandlimited signals with multi-channel time encoding, *IEEE Transactions on Signal Processing* 68 (2020) 1105–1119. <https://doi.org/10.1109/TSP.2020.2967182>.
- [31] A.A. Lazar, E.A. Pnevmatikakis, Video time encoding machines, *IEEE Transactions on Neural Networks* 22 (2011) 461–473. <https://doi.org/10.1109/TNN.2010.2103323>.
- [32] A. Bhandari, F. Krahmer, R. Raskar, On unlimited sampling and reconstruction, *IEEE Transactions on Signal Processing* 69 (2021) 3827–3839. <https://doi.org/10.1109/TSP.2020.3041955>.
- [33] D. Florescu, F. Krahmer, A. Bhandari, Event-driven modulo sampling, in: *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021: pp. 5435–5439. <https://doi.org/10.1109/ICASSP39728.2021.9414152>.
- [34] K. Adam, A. Scholefield, M. Vetterli, Asynchrony increases efficiency: Time encoding of videos and low-rank signals, *IEEE Transactions on Signal Processing* 70 (2022) 105–116. <https://doi.org/10.1109/TSP.2021.3133709>.
- [35] A.J. Kamath, C.S. Seelamantula, Multichannel time-encoding of finite-rate-of-innovation signals, in: *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023: pp. 1–5. <https://doi.org/10.1109/ICASSP49357.2023.10096150>.
- [36] R. Alexandru, P.L. Dragotti, Time encoding and decoding of multidimensional signals with finite rate of innovation, in: *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021: pp. 842–846. <https://doi.org/10.1109/IEEECONF53345.2021.9723165>.

- [37] K. Adam, A. Scholefield, M. Vetterli, Multi-channel time encoding for improved reconstruction of bandlimited signals, in: ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2019.
- [38] S. Thorpe, A. Delorme, R. Van Rullen, Spike-based strategies for rapid processing, *Neural Networks* 14 (2001) 715–725. [https://doi.org/10.1016/S0893-6080\(01\)00083-1](https://doi.org/10.1016/S0893-6080(01)00083-1).
- [39] P.L. Dragotti, M. Vetterli, T. Blu, Sampling moments and reconstructing signals of finite rate of innovation: Shannon meets strang-fix, *IEEE Transactions on Signal Processing* 55 (2007) 1741–1757. <https://doi.org/10.1109/TSP.2006.890907>.
- [40] A. Cohen, *Numerical analysis of wavelet methods*, Elsevier, 2003.
- [41] J.A. Urigüen, T. Blu, P.L. Dragotti, FRI sampling with arbitrary kernels, *IEEE Transactions on Signal Processing* 61 (2013) 5310–5323. <https://doi.org/10.1109/TSP.2013.2278152>.
- [42] L. Weiss, R.N. McDonough, Prony's method, *Z-transforms*, and pade approximation, *SIAM Review* 5 (1963) 145–149. <http://www.jstor.org/stable/2027478> (accessed August 4, 2023).
- [43] R. Schmidt, Multiple emitter location and signal parameter estimation, *IEEE Transactions on Antennas and Propagation* 34 (1986) 276–280. <https://doi.org/10.1109/TAP.1986.1143830>.
- [44] R. Roy, A. Paulraj, T. Kailath, Estimation of signal parameters via rotational invariance techniques - ESPRIT, in: *MILCOM 1986 - IEEE Military Communications Conference: Communications-Computers: Teamed for the 90's*, 1986: pp. 41.6.1–41.6.5. <https://doi.org/10.1109/MILCOM.1986.4805850>.
- [45] Y. Hua, T.K. Sarkar, Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise, *IEEE Transactions on Acoustics, Speech, and Signal Processing* 38 (1990) 814–824. <https://doi.org/10.1109/29.56027>.
- [46] B. Halder, T. Kailath, Efficient estimation of closely spaced sinusoidal frequencies using subspace-based methods, *IEEE Signal Processing Letters* 4 (1997) 49–51. <https://doi.org/10.1109/97.554470>.

- [47] Y. Bresler, A. Macovski, Exact maximum likelihood parameter estimation of superimposed exponential signals in noise, *IEEE Transactions on Acoustics, Speech, and Signal Processing* 34 (1986) 1081–1089. <https://doi.org/10.1109/TASSP.1986.1164949>.
- [48] R. Guo, Y. Li, T. Blu, H. Zhao, Vector-FRI recovery of multi-sensor measurements, *IEEE Transactions on Signal Processing* 70 (2022) 4369–4380. <https://doi.org/10.1109/TSP.2022.3204402>.
- [49] Y. Li, R. Guo, T. Blu, H. Zhao, Generic FRI-based DOA estimation: A model-fitting method, *IEEE Transactions on Signal Processing* 69 (2021) 4102–4115. <https://doi.org/10.1109/TSP.2021.3092344>.
- [50] H. Pan, T. Blu, M. Vetterli, Towards generalized FRI sampling with an application to source resolution in radioastronomy, *IEEE Transactions on Signal Processing* 65 (2017) 821–835. <https://doi.org/10.1109/TSP.2016.2625274>.
- [51] R. Tur, Y.C. Eldar, Z. Friedman, Innovation rate sampling of pulse streams with application to ultrasound imaging, *IEEE Transactions on Signal Processing* 59 (2011) 1827–1842. <https://doi.org/10.1109/TSP.2011.2105480>.
- [52] J.A.U. Garaizabal, Exact and approximate strang-fix conditions to reconstruct signals with finite rate of innovation from samples taken with arbitrary kernels, PhD thesis, Imperial College London, 2013.
- [53] S. Rudresh, A.J. Kamath, C.S. Seelamantula, A time-based sampling framework for finite-rate-of-innovation signals, in: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020: pp. 5585–5589. <https://doi.org/10.1109/ICASSP40776.2020.9053120>.
- [54] T. Blu, P.-L. Dragotti, M. Vetterli, P. Marziliano, L. Coulot, Sparse sampling of signal innovations, *IEEE Signal Processing Magazine* 25 (2008) 31–40. <https://doi.org/10.1109/MSP.2007.914998>.
- [55] R.J.P. De Figueiredo, C.-L. Hu, Waveform feature extraction based on tauberian approximation, *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-4* (1982) 105–116. <https://doi.org/10.1109/TPAMI.1982.4767214>.

- [56] S. Mulleti, C.S. Seelamantula, Paley–wiener characterization of kernels for finite-rate-of-innovation sampling, *IEEE Transactions on Signal Processing* 65 (2017) 5860–5872. <https://doi.org/10.1109/TSP.2017.2733484>.
- [57] A.J. Kamath, S. Rudresh, C.S. Seelamantula, Time encoding of finite-rate-of-innovation signals, (2021). <https://doi.org/10.48550/ARXIV.2107.03344>.
- [58] H. Naaman, S. Mulleti, Y.C. Eldar, FRI-TEM: Time encoding sampling of finite-rate-of-innovation signals, *IEEE Transactions on Signal Processing* 70 (2022) 2267–2279. <https://doi.org/10.1109/TSP.2022.3167146>.
- [59] W. Rall, Distinguishing theoretical synaptic potentials computed for different somadendritic distributions of synaptic input., *Journal of Neurophysiology* 30 (1967) 1138–1168. <https://doi.org/10.1152/jn.1967.30.5.1138>.
- [60] I.M. Comsa, K. Potempa, L. Versari, T. Fischbacher, A. Gesmundo, J. Alakuijala, Temporal coding in spiking neural networks with alpha synaptic function, in: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020: pp. 8529–8533.
- [61] M. Holschneider, *Wavelets*, Clarendon Press, Oxford, England, 1995.
- [62] X. Wei, B. Thierry, P.L. Dragotti, Finite rate of innovation with non-uniform samples, in: *2012 IEEE International Conference on Signal Processing, Communication and Computing (ICSPCC 2012)*, 2012: pp. 369–372. <https://doi.org/10.1109/ICSPCC.2012.6335674>.
- [63] M. Hilton, R. Alexandru, P.L. Dragotti, Time encoding using the hyperbolic secant kernel, in: *2020 28th European Signal Processing Conference (EUSIPCO)*, 2021: pp. 2304–2308. <https://doi.org/10.23919/Eusipco47968.2020.9287806>.
- [64] D. Florescu, D. Coca, A novel reconstruction framework for time-encoded signals with integrate-and-fire neurons, *Neural Computation* 27 (2015) 1872–1898. https://doi.org/10.1162/NECO_a_00764.
- [65] G.B. Clayton, *Operational amplifiers*, 2nd ed., Butterworth-Heinemann, Oxford, England, 1979.

- [66] M. Hilton, P.L. Dragotti, Sparse asynchronous samples from networks of tems for reconstruction of classes of non-bandlimited signals, in: ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2023: pp. 1–5. <https://doi.org/10.1109/ICASSP49357.2023.10095620>.
- [67] M. Hilton, R. Alexandru, P.L. Dragotti, Guaranteed reconstruction from integrate-and-fire neurons with alpha synaptic activation, in: ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2021: pp. 5474–5478. <https://doi.org/10.1109/ICASSP39728.2021.9414759>.
- [68] J. Oñativia Bravo, Sampling and reconstruction of finite rate of innovations signals with applications in neuroscience and sparse representation, PhD thesis, Imperial College London, 2016.
- [69] A.S. Sedra, K.C.A. Smith, Microelectronic circuits, 4th ed., Oxford University Press, New York, NY, 1997.