Efficient and Robust Detection of GFSK Signals under Dispersive Channel, Modulation Index, and Carrier Frequency Offset Conditions

Charles Tibenderana

Communications Research Group, School of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK
Email: ct02r@ecs.soton.ac.uk

Stephan Weiss

Communications Research Group, School of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK
Email: sw1@ecs.soton.ac.uk

Gaussian frequency shift keying is the modulation scheme specified for Bluetooth. Signal adversities typical in Bluetooth networks include AWGN, multipath propagation, carrier frequency, and modulation index offsets. In our effort to realise a robust but efficient Bluetooth receiver, we adopt a high-performance matched-filter-based detector, which is near optimal in AWGN, but requires a prohibitively costly filter bank for processing of K bits worth of the received signal. However, through filtering over a single bit period and performing phase propagation of intermediate results over successive single-bit stages, we eliminate redundancy involved in providing the matched filter outputs and reduce its complexity by up to 90% (for $K = 9$). The constant modulus signal characteristic and the potential for carrier frequency offsets make the constant modulus algorithm (CMA) suitable for channel equalisation, and we demonstrate its effectiveness in this paper. We also introduce a stochastic gradient-based algorithm for carrier frequency offset correction, and show that the relative rotation between successive intermediate filter outputs enables us to detect and correct offsets in modulation index.

Keywords and phrases: Gaussian frequency shift keying, Bluetooth, constant modulus algorithm, carrier frequency offset correction, modulation index offset correction.

1. INTRODUCTION

Gaussian frequency shift keying (GFSK) is a bandwidth preserving digital modulation technique, which has been used for low-cost transmission standards such as Bluetooth [1]. Relative low cost makes Bluetooth an attractive alternative to expensive high-data-rate wireless local area network (WLAN) services such as IEEE 802.11b. Therefore, in this contribution we aim at deriving GFSK receivers for high integrity data transfer, which can enable their use in inexpensive standards similar to Bluetooth more efficiently.

Maximum likelihood detection of a sequence of GFSK-modulated bits can be achieved by a Viterbi receiver, which correlates the received signal over a bit period with all possible alternatives, before deploying the Viterbi algorithm to penalise illegitimate phase transitions in the filter outputs [2]. However, the use of a Viterbi detector is limited to coherent detection of signals where the modulation index (h) used in the modulation process is a rational number, thereby ensuring a finite number of states. In addition, the Viterbi receiver is very vulnerable to inaccuracies in the modulation index, whereby it has been shown to be robust to variations of about $|\Delta h| \leq 0.01$ only [3]. Even if it were possible to estimate the transmitter modulation index accurately, it would be difficult to compensate for this at the receiver, as the receiver architecture, including the number of states, would have to be changed [4]. The Viterbi receiver therefore seems unsuitable for wireless standards such as Bluetooth, in which significant offsets in modulation index are permitted [1].

On the other hand, near optimal detection of a symbol in the maximum likelihood sense can be effected using a system of filters, which are matched to legitimate waveforms over an observation interval of several bit periods. The filter with the largest output determines the received waveform, and the bit at the centre of the modulating bit-sequence responsible for producing the received waveform is chosen as the received bit. Detailed descriptions of this high-performance receiver are recorded in [5, 6].

It is important to note that the Viterbi receiver estimates the maximum likelihood of a data sequence, while the high-performance detector nearly determines the maximum likelihood estimate of a single data symbol. Consequently, the Viterbi and high-performance receivers are optimal in a different sense [4]. Relatively simple reception techniques for GFSK signals which yield more modest results include FM-AM conversion, phase-shift discrimination, zero-crossing detection, and frequency feedback [7].

The advent of software-defined radios (SDR) has allowed multiple wireless standards to run on common hardware platforms. A platform implementing relatively complex standards, such as IEEE 802.11b WLAN, will have excess computational capacity when running relatively simple standards like Bluetooth which uses GFSK. This motivates us to seek out the best way to utilise the extra capacity that may be available in an SDR to improve reception of GFSK-modulated signals. The high-performance receiver represents a plausible option because of its good reception in additive white Gaussian noise (AWGN), and its relative insensitivity to errors in modulation index when compared to the Viterbi detector. However, an observation interval of K bit periods and N samples per bit will require 2^K filters of length KN. This results in a prohibitively large filter bank for high values of K necessary to ensure best performance. In this paper we assert that the computational cost of realising the high-performance receiver on a digital signal processor (DSP) can be reduced by performing filtering over a single bit interval and propagating the intermediate results over K successive bit periods, thereby eliminating redundancy involved in providing the matched filter outputs.

Many Bluetooth applications are likely to suffer from multipath propagation if pleas to increase the operation range are heeded. Strategies have been suggested to tackle dispersive channels in Bluetooth using decision feedback equalisers [8, 9], but they will be undermined by frequency errors. Other more common equalisation techniques that rely on a training sequence may not be capable of tracking fast changes caused by frequency errors [10] and are not recommended for point-to-multipoint networks such as Bluetooth because of the requirement for the control unit to interrupt transmission to retrain a tributary receiver that may have experienced a change in channel conditions, or that was not online during the initial training procedure [11]. In this paper we demonstrate the effectiveness of the constant modulus algorithm (CMA) [12, 13] as an equalisation technique for Bluetooth in presence of carrier frequency offsets.

The necessity for cheap transceivers motivates the Bluetooth Special Interest Group to allow up to 75 kHz initial frequency errors [1]. Research has shown that performance deteriorates significantly even when operating within this range [14], more so, in the high-performance receiver where frequency errors propagate through an observation interval of K bit periods, thereby trading off robustness to Gaussian noise with immunity to a carrier frequency offset.

Work done to address the problem of frequency errors in continuous phase-modulated signals, of which GFSK forms a subset, can be categorised as being training based or blind. Notable research on blind algorithms is reported in [15], where the phase rotation between two received samples separated by M chip periods is estimated. Propositions in [15] rely heavily on the ability of the receiver to determine M which would represent a maximum phase shift of $\pi/2$ in the transmit signal. This is not easily attained in Bluetooth where the modulation index $h \in (0.28, 0.35)$ [1].

Therefore, in this paper we present a novel algorithm based on gradient descent techniques, which converges under conditions specified in [15], without the necessity for the receiver to know the precise transmitter modulation index. The algorithm is derived analytically and assessed via simulation.

The Bluetooth standard accommodates initial errors in modulation index of $\Delta h \le 0.07$, which can cause significant degradation in performance. However, we demonstrate that the relative rotation between intermediate filter outputs enables the detection and correction of offsets in modulation index.

The structure of this paper is as follows. A signal model is developed in Section 2, while the constant modulus algorithm is presented as a possible solution to multipath propagation in Section 3. A novel stochastic-gradient-based algorithm for carrier frequency offset correction is derived in Section 4. The standard high-performance receiver is introduced in Section 5, before the analytic description of a new reduced complexity realisation of it in Section 6. In Section 7 we explain how the errors in modulation index can be detected and corrected. Our results are contained in Section 8 before we conclude in Section 9.

2. SIGNAL AND SYSTEM MODEL

GFSK generally modulates a multilevel symbol stream $p[k]$, which here is assumed to be binary, $p[k] \in \{\pm 1\}$ with bit index k. This bit sequence is expanded by a factor of N and passed through a Gaussian filter with a bandwidth-time product of K_{BT} and impulse response $g[n]$ of length LN, thus having a support of L bit periods, yielding its discrete instantaneous angular frequency signal

$$\omega[n] = 2\pi h \sum_{k=-\infty}^{\infty} p[k]g[n - kN], \tag{1}$$

where n represent the chip index. The phase of the baseband version of the transmitted signal,

$$s[n] = \exp\left\{j \sum_{\nu=-\infty}^{n} \omega[\nu]\right\} = \prod_{\nu=-\infty}^{n} e^{j\omega[\nu]}, \tag{2}$$

is determined as the cumulative sum over all previous frequency values $\omega[n]$.

According to the baseband system model in Figure 1, the signal is dispersed by a stationary channel impulse response
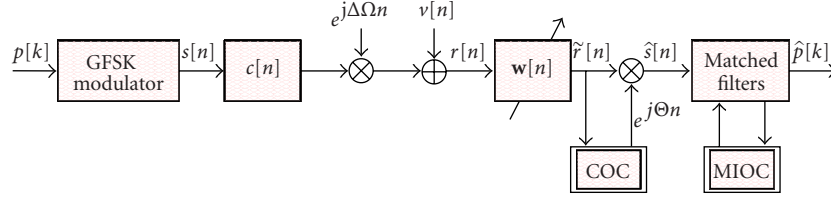
FIGURE 1: System model.

(CIR) $c[n]$, and subjected to a carrier frequency offset $\Delta\Omega$ and AWGN $v[n]$. The received signal can be expressed as

$$r[n] = \sum_{\lambda=0}^{L_c-1} c[\lambda]s[n-\lambda]e^{j\Delta\Omega n} + v[n], \tag{3}$$

with L_c being the length of the CIR. Functionalities of the system blocks for the receiver portrayed in Figure 1 will be detailed later. At this point it suffices to say that $c[n]$ and $\mathbf{w}[n]$ are the channel and equaliser blocks, respectively, while COC and MIOC stand for carrier and modulation index offset correction, respectively.

3. CHANNEL EQUALISATION

Due to the constant modulus nature of the GFSK signal $s[n]$ in (2), the CMA lends itself as a candidate for blindly equalising the dispersive channel $c[n]$. Due to its phase invariance, the CMA is known to be tolerant towards frequency offsets, and we exploit this trait here.

Let the equaliser coefficients be written in vector notation

$$\mathbf{w}[n]^{\mathrm{T}} = [w_0[n], w_1[n], \ldots, w_{L_w-1}[n]], \tag{4}$$

then equalisation of the CIR takes place via

$$\tilde{r}[n] = \sum_{\lambda=0}^{L_w-1} w_\lambda[n]r[n-\lambda], \tag{5}$$

ideally resulting in a version of the received signal $\tilde{r}[n]$, which is corrupted only by an arbitrary phase offset, a frequency offset, and AWGN. The nonconvex CMA cost function used to achieve this is

$$J = \mathcal{E}\left\{ \left| \tilde{r}[n]^2 - 1 \right|^2 \right\}, \tag{6}$$

whereby $\mathcal{E}\{\cdot\}$ is the expectation operator, $\tilde{r}[n]$ is the equaliser output, and the expected magnitude of the received signal samples is 1. The equaliser coefficients can be adjusted via a stochastic gradient search [16] according to

$$\mathbf{w}[n+1] = \mathbf{w}[n] - \mu_{\mathrm{w}}\hat{\nabla}[n], \tag{7}$$

where μ_{w} is a step size and $\hat{\nabla}[n]$ is the instantaneous estimate

of the gradient of $J[n]$, given by

$$\hat{\nabla}[n] = \mathbf{r}^*[n]\tilde{r}[n]\left(\left| \tilde{r}[n] \right|^2 - 1 \right), \tag{8}$$

in which $\mathbf{r}[n]$ is a vector of received signal samples

$$\mathbf{r}[n]^{\mathrm{T}} = [r_0[n], r_1[n], \ldots, r_{L_w-1}[n]]. \tag{9}$$

4. CARRIER FREQUENCY OFFSET CORRECTION

From the discussions in Section 3 and assuming perfect equalisation, the remaining signal components in (5) will be

$$\tilde{r}[n] = s[n] + \tilde{v}[n], \tag{10}$$

where

$$\tilde{v}[n] = w[n] * v[n] \tag{11}$$

is the AWGN filtered by the equaliser's impulse response $w[n]$. Therefore an estimation of the carrier frequency offset can be based on $\tilde{r}[n]$ by denoting

$$\begin{aligned}
\mathcal{E}&\{\tilde{r}[n]\tilde{r}^*[n-M]\} \\
&= \mathcal{E}\{s[n]s^*[n-M]\}\, e^{j\Delta\Omega M} \\
&\quad + \mathcal{E}\{s[n]\tilde{v}^*[n-M]\}\, e^{j\Delta\Omega n} \\
&\quad + \mathcal{E}\{\tilde{v}[n]s^*[n-M]\}\, e^{-j\Delta\Omega(n-M)} \\
&\quad + \mathcal{E}\{\tilde{v}[n]\tilde{v}^*[n-M]\} \\
&= e^{j\Delta\Omega M}.
\end{aligned} \tag{12}$$
$$\tag{13}$$

Due to the independence and zero mean of $s[n]$ and $\tilde{v}[n]$, the second and third term in (12) will be zero. By selecting M sufficiently large, the autocorrelation term of the noise in (12) vanishes. Since the instantaneous frequency accumulated over M samples of the transmitted signal $s[n]$ will either rotate in a positive or negative direction but on average be zero, we have $\mathcal{E}\{s[n]s^*[n-M]\} = 1$. To minimise the fourth term in (12) we select $M > L_w$, thereby reducing the autocorrelation term to zero and enabling the simplification in (13). Note that the detection of the carrier frequency offset is independent of any other receiver functions.
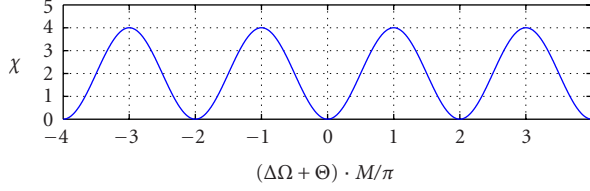
FIGURE 2: Cost function for carrier frequency offset correction.

4.1. Cost function

From the equaliser output $\tilde{r}[n]$, we create a received signal

$$\hat{s}[n] = \tilde{r}[n]e^{j\Theta n}, \tag{14}$$

that is, modulating by Θ, to match the carrier offset $\Delta\Omega$. In order to determine Θ, we can use the following constant modulus cost function:

$$\chi = \left| \mathcal{E}\{\hat{s}[n]\hat{s}^*[n-M]\} - 1 \right|^2. \tag{15}$$

Inserting (14) and (13) into (15) yields

$$\begin{aligned} \chi &= \left(e^{j(\Delta\Omega+\Theta)M} - 1\right)\left(e^{j(\Delta\Omega+\Theta)M} - 1\right)^* \\ &= 2 - 2\cos\left((\Theta + \Delta\Omega)M\right) \end{aligned} \tag{16}$$

$$\text{with } \chi = 0 \longleftrightarrow \Theta = \frac{2\pi k}{M} - \Delta\Omega.$$

Figure 2 confirms the assertion in (16), however we are interested in the solution for $k = 0$ only, for which the cost function provides a unique minimum under the condition

$$-\pi < (\Theta + \Delta\Omega)M < \pi, \tag{17}$$

similar to [15]. Hence, a trade-off exists for the selection of M between decorrelating the noise in the receiver and not exceeding the bounds in (17).

4.2. Stochastic gradient method

Within the bounds of (17), Θ can be iteratively adapted over time based on gradient descend techniques [16] according to

$$\Theta[n+1] = \Theta[n] - \mu_\Theta \frac{\partial\hat{\chi}[n]}{\partial\Theta}, \tag{18}$$

with a suitable step size parameter μ_Θ and the substitution of the exact cost function by an instantaneous estimate $\hat{\chi}[n]$. The latter is obtained by dropping expectations in (15) and assuming small changes in Θ such that $\Theta[n] \approx \Theta[n-M]$,

$$\begin{aligned} \frac{\partial\hat{\chi}[n]}{\partial\Theta} &= \frac{\partial}{\partial\Theta}\left(\hat{s}[n]\hat{s}^*[n-M] - 1\right) \cdot \left(\hat{s}[n]\hat{s}^*[n-M] - 1\right)^* \\ &= -2M\mathfrak{I}\{\hat{s}[n]\hat{s}^*[n-M]\left(\hat{s}[n]\hat{s}^*[n-M] - 1\right)^*\}. \end{aligned} \tag{19}$$

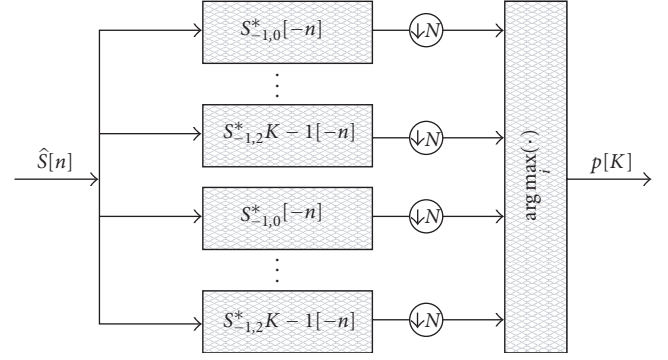Therefore the COC block in Figure 1 is defined by (14), (18), and (19).

FIGURE 3: Standard high-performance GFSK receiver via a bank of matched filters and a detector according to (20).

5. MATCHED FILTER BANK RECEIVER

A standard high-performance receiver, which achieves near-optimum maximum likelihood noncoherent estimation of a bit in AWGN is discussed in [4, 5, 6, 17, 18]. The method is based on a filter bank containing all possible transmitted sequences $s[n]$ over a duration of K bit periods. Over this observation interval, due to the support length of the Gaussian filter, 2^{K+L-1} possible sequences exist apart from an initial phase shift. However, the effects of the marginal bits of the modulating data sequence, on the K bit long sequence $s[n]$, may be neglected without significant loss in performance, resulting in 2^K matched filters [5]. The filter with the largest output determines the received sequence, and the bit at the centre of the modulating bit sequence responsible for producing the received sequence is chosen as the detected bit, assuming K to be odd. The resulting scheme is depicted in Figure 3, where $s_{i,j}[n]$ are possible transmitted sequences with $i \in \{\pm 1\}$ indicating the value of the middle bit and $j = 0(1)2^{K-1}$ indexing the possible combinations of the remaining $K-1$ bits. The detector selects the largest magnitude value, determining the output $\hat{p}[k]$ as

$$\hat{p}[k] = \arg\max_i \left| \sum_{n=0}^{KN-1} \hat{s}[kN-n] \cdot s_{i,j}^*[-n] \right|, \tag{20}$$

where $s_{i,j}^*[-n]$ are the 2^K matched filter responses. Note that the detector imposes a delay such that ideally $\hat{p}[k] = p[k - (K+1)/2]$. The performance of this receiver improves with an increase in the observation interval K. However, despite its performance merits and neglecting $L-1$ marginal bits, the computational complexity of

$$C_{\text{standard}} = 4NK2^K = NK2^{K+2} \tag{21}$$

real-valued multiply accumulates (MACs), where a complex MAC accounts for 4 real-valued ones, is prohibitive. Therefore, in the following we seek a low-complexity implementation of this receiver.
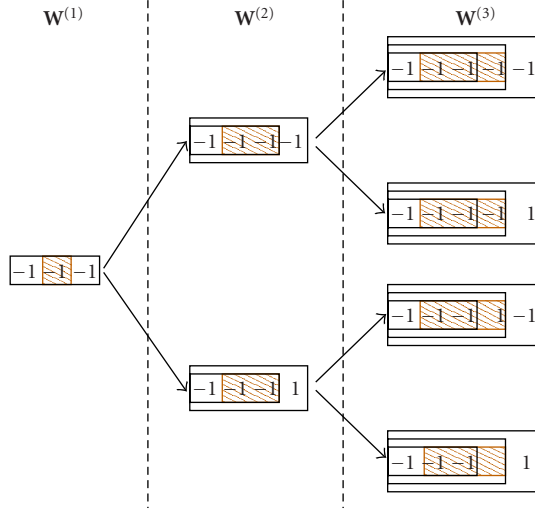
FIGURE 4: Example for $L = 3$ of the evolution from the first row of $\mathbf{W}^{(1)}$, representing a sequence of -1 data pulses, to the first two and four rows of $\mathbf{W}^{(2)}$ and $\mathbf{W}^{(3)}$ respectively.

6. LOW-COMPLEXITY RECEIVER

We will first inspect the matched filter responses in Section 6.1, and thereafter develop a recursive scheme for their representation in Section 6.2, leading to an analysis of its complexity in Section 6.3.

6.1. Received signals

For simplicity, we assume in the following that the channel distortion and carrier offset have been eliminated by preceding signal processing blocks. Therefore we assume that K bit periods of the received signal $\hat{s}[n] = s[n]$ are held in a tap delay line (TDL) vector $\hat{\mathbf{s}}_k$, synchronised with the kth bit to be the most recent datum:

$$
\hat{\mathbf{s}}_k = \begin{bmatrix} \grave{\mathbf{s}}_k \\ \grave{\mathbf{s}}_{k-1} \\ \vdots \\ \grave{\mathbf{s}}_{k-K+1} \end{bmatrix} = \mathbf{s}_k + \mathbf{v}_k = \underbrace{\begin{bmatrix} \tilde{\mathbf{s}}_k \\ \tilde{\mathbf{s}}_{k-1} \\ \vdots \\ \tilde{\mathbf{s}}_{k-K+1} \end{bmatrix}}_{\mathbf{s}_k} + \mathbf{v}_k, \quad (22)
$$

where $\mathbf{v}_k \in \mathbb{C}^{NK}$ holds the noise samples. The vector $\tilde{\mathbf{s}}_k$ is defined as

$$
\tilde{\mathbf{s}}_k = [s[kN]s[kN-1]\cdots s[(k-1)N+1]]^{\mathrm{T}}, \quad (23)
$$

and $\grave{\mathbf{s}}_k$ is defined analogously. According to (2), $\tilde{\mathbf{s}}_k$, holding N samples within a bit period, can be expanded as

$$
\tilde{\mathbf{s}}_k = \underbrace{\begin{bmatrix} \prod_{\nu=(k-1)N+1}^{kN} e^{j\omega[\nu]} \\ \prod_{\nu=(k-1)N+1}^{kN-1} e^{j\omega[\nu]} \\ \vdots \\ e^{j\omega[(k-1)N+1]} \end{bmatrix}}_{\mathbf{u}_k} \cdot \prod_{\nu=-\infty}^{(k-1)N} e^{j\omega[\nu]}, \quad (24)
$$

whereby for the samples in $\mathbf{u}_k$ the instantaneous frequency is only accumulated from the start of the kth bit period. Inserting (24) into $\mathbf{s}_k$ yields

$$
\mathbf{s}_k = \begin{bmatrix} \mathbf{u}_k & \cdot e^{j\left(\theta_{k-K+1}+\cdots+\theta_{k-2}+\theta_{k-1}\right)} \\ \mathbf{u}_{k-1} & \cdot e^{j\left(\theta_{k-K+1}+\cdots+\theta_{k-2}\right)} \\ \vdots & \\ \mathbf{u}_{k-K+2} & \cdot e^{j\theta_{k-K+1}} \\ \mathbf{u}_{k-K+1} & \cdot 1 \end{bmatrix} \cdot e^{j\alpha} \quad (25)
$$

with

$$
\theta_k = \sum_{\nu=(k-1)N+1}^{kN} \omega[\nu], \qquad \alpha = \sum_{\nu=-\infty}^{(k-K)N} \omega[\nu]. \quad (26)
$$

Firstly, note that each vector $\mathbf{u}_m$ can take on the shape of 2^L different waveforms, whereby L was the support length of the Gaussian window in bit periods. Secondly, observe that a phase correction term $e^{j\theta_k}$ contains the instantaneous frequency values accumulated over the kth bit period, which is held in the top element of $\mathbf{u}_k$ in (24) and is applied to all subsequent bit periods. The initial phase of $s[n]$ entering the TDL is α.

6.2. Recursive matched filter formulation

The matched filter responses $s^*_{i,j}[-n]$ are designed from the transmitted signal $s[n]$ in (2). Utilising the previous observation that $\mathbf{u}_k$ only takes on 2^L basic waveforms independent of k, we will construct a matched receiver in steps.

Case 1 ($K = 1$). Consider a matched filter for $K = 1$ covering the kth bit period. The 2^L matched filter outputs are given by

$$
\mathbf{y}_k^{(1)} = \mathbf{W}^{(1)}\grave{\mathbf{s}}_k \quad (27)
$$

with $\mathbf{W}^{(1)} \in \mathbb{C}^{2^L \times N}$ containing the possible complex conjugated waveforms in its rows. The superscript (1) indicates that only a single bit period $K = 1$ is observed. The first column of $\mathbf{W}^{(1)}$, denoted by $\mathbf{w}$, holds the 2^L possible values for $e^{j\theta_k}$. We assume that the first row of $\mathbf{W}^{(K)}$ is the matched filter for K bits at the centre of $K+L-1$ bits, all of value -1, binary coded decimally down to the last row with $K + L - 1$ bits of value $+1$. Hence, rows of $\mathbf{W}^{(1)}$ are formed from the central N samples obtained via GFSK modulation of L bits of value -1, binary coded decimally down to the last row with L bits of value $+1$.

Case 2 ($K = 2$). Expanding to $K = 2$, we can denote

$$
\mathbf{y}_k^{(2)} = \mathbf{W}^{(2)} \begin{bmatrix} \grave{\mathbf{s}}_k \\ \grave{\mathbf{s}}_{k-1} \end{bmatrix}. \quad (28)
$$

In constructing the 2^{L+1} matched filter responses in $\mathbf{W}^{(2)}$, only one extra bit needs to be considered compared to the responses in $\mathbf{W}^{(1)}$.
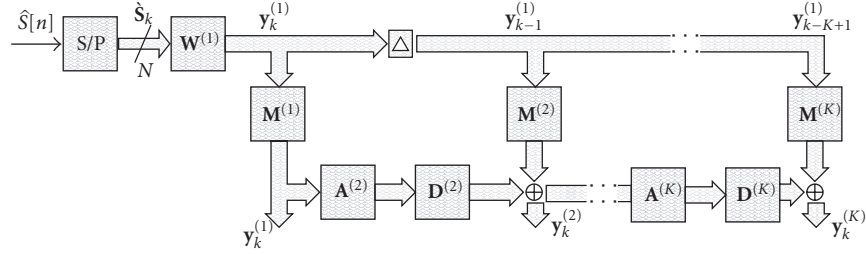
FIGURE 5: Low-complexity implementation of a matched filter bank high-performance GFSK receiver. The received GFSK signal $\hat{s}[n]$ is passed through a serial/parallel converter and a filter bank $\mathbf{W}^{(1)}$ with a single bit duration. Processed over K stages, the matched filter bank outputs are contained in $\mathbf{y}_k^{(K)}$.

Example 1. For $L = 3$, the rows in $\mathbf{W}^{(1)}$ should contain the central N samples of the responses to the sequences $\{-1, -1, -1\}$ to $\{1, 1, 1\}$, while $\mathbf{W}^{(2)}$ would cater for an additional bit, hence covering the middle $2N$ samples of combinations $\{-1, -1, -1, -1\}$ to $\{1, 1, 1, 1\}$. So for each possible sequence contained in $\mathbf{W}^{(1)}$, two new possibilities arise in $\mathbf{W}^{(2)}$, and so on for higher values of K. This formulation is portrayed in Figure 4. Thus, 2^L outputs of filter bank $\mathbf{W}^{(1)}$ from the previous bit period can be used with its current results to compute the 2^{L+1} outputs of $\mathbf{W}^{(2)}$, enabling us to write

$$\mathbf{y}_k^{(2)} = \mathbf{D}^{(2)}\mathbf{A}^{(2)}\mathbf{W}^{(1)}\hat{\mathbf{s}}_k + \mathbf{M}^{(2)}\mathbf{W}^{(1)}\hat{\mathbf{s}}_{k-1}$$
$$= \mathbf{D}^{(2)}\mathbf{A}^{(2)}\mathbf{y}_k^{(1)} + \mathbf{M}^{(2)}\mathbf{y}_{k-1}^{(1)}, \tag{29}$$

whereby $\mathbf{y}_{k-1}^{(1)}$ are the single bit matched filter outputs for the $(k-1)$th bit. The matrix $\mathbf{A}^{(2)}$,

$$\mathbf{A}^{(2)} = \text{block diag} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \in \mathbb{Z}^{2^{L+1} \times 2^L}, \tag{30}$$

produces an extra copy of each response in $\mathbf{W}^{(1)}$, while

$$\mathbf{D}^{(2)} = \begin{bmatrix} \text{diag}\,\mathbf{w} & \mathbf{0} \\ \mathbf{0} & \text{diag}\,\mathbf{w} \end{bmatrix} \in \mathbb{C}^{2^{L+1} \times 2^{L+1}} \tag{31}$$

applies the phase correction term $e^{j\theta_k}$, and the matrix

$$\mathbf{M}^{(2)} = \begin{bmatrix} \mathbf{I}_{2^L} \\ \mathbf{I}_{2^L} \end{bmatrix} \in \mathbb{Z}^{2^{L+1} \times 2^L} \tag{32}$$

is assigning the expansion by the extra bit considered for $K = 2$, whereby $\mathbf{I}_{2^L}$ is a $2^L \times 2^L$ identity matrix.

Case 3 (K arbitrary). Generalising from the previous cases, we formulate recursively for $\mathbf{y}_k^{(K)} \in \mathbb{C}^{2^{K+L-1}}$,

$$\mathbf{y}_k^{(K)} = \mathbf{D}^{(K)}\mathbf{A}^{(K)}\mathbf{y}_k^{(K-1)} + \mathbf{M}^{(K)}\mathbf{y}_{k-K+1}^{(1)}, \tag{33}$$

where

$$\mathbf{M}^{(K)} = \begin{bmatrix} \mathbf{M}^{(K-1)} \\ \mathbf{M}^{(K-1)} \end{bmatrix} \quad \text{with } \mathbf{M}^{(1)} = \mathbf{I}_{2^L},$$

$$\mathbf{A}^{(K)} = \begin{bmatrix} \mathbf{A}^{(K-1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}^{(K-1)} \end{bmatrix}$$
$$\text{with } \mathbf{A}^{(1)} = \text{block diag} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \in \mathbb{Z}^{2^L \times 2^{L-1}},$$

$$\mathbf{D}^{(K)} = \begin{bmatrix} \mathbf{D}^{(K-1)} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^{(K-1)} \end{bmatrix} \quad \text{with } \mathbf{D}^{(1)} = \text{diag}\,\mathbf{w}. \tag{34}$$

This form of the matched filter bank receiver is depicted by the flow graph in Figure 5. A detector similar to (20) selecting the index of the largest element would operate on $\mathbf{y}_k^{(K)}$ to determine the correct output bit.

6.3. Computational complexity

Inspecting the operations in Figure 5, per bit period 2^L matched filter operations of length N have to be performed. With the matrices $\mathbf{M}^{(k)}$ and $\mathbf{A}^{(k)}$ only performing indexing, the only arithmetic operations required are multiplications with the diagonal elements of the phase correction matrices $\mathbf{D}^{(K)}$, yielding a total of

$$C_{\text{efficient}} = 2^{L+2}N + 4 \sum_{k=1}^{K-1} 2^{L+k} \lessapprox 2^{L+2}N + 2^{L+K+2}. \tag{35}$$

If marginal bits are disregarded analogously to the matched filter receivers in [5, 6] as discussed in Section 5, then desired outputs $\widetilde{\mathbf{y}}_k^{(K)} = \mathbf{S}^{(K)}\mathbf{y}_k^{(K)}$ can be extracted. As an example for $L = 3$, the extraction matrix $\mathbf{S}^{(K)}$ takes the form

$$\mathbf{S}^{(K)} = \begin{bmatrix} \mathbf{G}^{(K)} & \mathbf{0}_{2^K \times 2^K + 1} \end{bmatrix} \tag{36}$$

with

$$\mathbf{G}^{(K)} = \text{block diag} \begin{bmatrix} 1 & 0 \end{bmatrix} \in \mathbb{Z}^{2^K \times 2^{K+1}}. \tag{37}$$
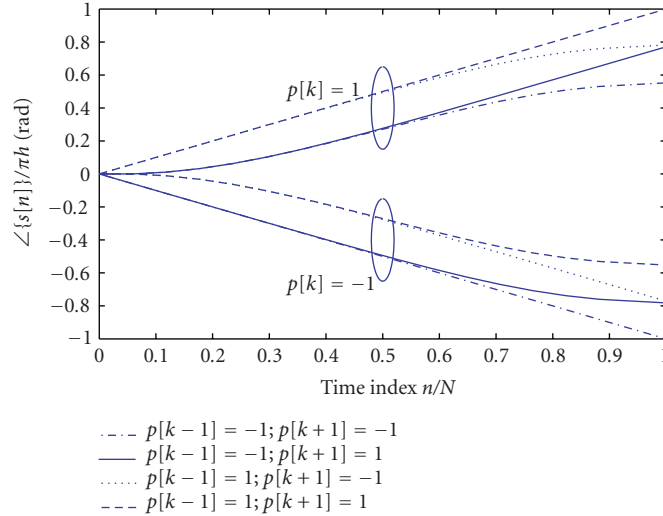
FIGURE 6: Example of the possible phase trajectories of $s[n]$ for one bit period for $L = 3$.

The extraction matrices can be appropriately absorbed into (34), yielding a reduced complexity of

$$\tilde{C}_{\text{efficient}} \lesssim 2^{L+2}N + 2^{K+3}. \tag{38}$$

The matched filter bank detector discussed in this section works accurately if the exact modulation index h, used in the transmitter, is also known to the receiver. In this case, the Viterbi algorithm could be employed instead of the proposed structure [2]. However, in case of a mismatch in modulation index, Viterbi decoding fails [3], while the above iterative matched filter bank offers a basis for the detection and correction of this offset, which will be addressed in the following section.

7. MODULATION INDEX OFFSET CORRECTION

Some GFSK systems permit a tolerance in the modulation index, such as Bluetooth with $h \in (0.28, 0.35)$. Thus, for example, for Bluetooth the assumed modulation index in the receiver, $\hat{h}$, may vary by a maximum offset of $|\Delta h| = |h - \hat{h}| \leq 0.07$ from the modulation index h in the transmitter. In the following we propose a method to adapt the modulation index $\hat{h}$ in the receiver, and thus $\mathbf{W}^{(1)}$ from Section 6. Note from (1) and (2) that $|\angle\{s[n]\}| \propto h$. As an example, for $L = 3$ the resulting $2^L = 8$ possible phase trajectories of $s[n]$ are depicted in Figure 6, assuming that the phase at the beginning of the bit period was zero. With an increase in h, the phase trajectories will fan further out. Since both the transmitted signal and $\mathbf{W}^{(1)}$ are derived from this set of curves, the phase term

$$\eta = \angle\mathcal{E}\left\{\left(y^{(1)}_{k-(K-1)/2} \cdot \left(y^{(1)}_{k-(K-3)/2}\right)^*\right) \cdot p\left[k - \frac{k+1}{2}\right]\right\} \propto \Delta h \tag{39}$$

can be verified to be proportional to the mismatch in modulation index Δh. In (39), the quantity $y^{(1)}_k$ refers to the element of $\mathbf{y}^{(1)}_k$ in Figure 5 associated with the correct bit sequence leading to the detection of the middle bit $p[k - (k + 1)/2]$. The complex conjugate term in (39) ensures that the phase is measured relative to zero, while $p[k-(k+1)/2]$ compensates for the sign change imposed by the middle bit onto the phase as evidenced in Figure 6. An example for $\eta(\Delta h)$ is given in Figure 7 for $L = 3$.

To adapt the modulation index estimate $\hat{h}$ and therefore $\mathbf{W}^{(1)}$ in the receiver, we employ an iterative technique

$$\hat{h}[k + 1] = \hat{h}[k] + \mu_h \cdot \hat{\eta}[k], \tag{40}$$

where $\hat{\eta}[k]$ is an instantaneous estimate of the term in (39) based on a single bit period,

$$\hat{\eta}[k] = \angle\left\{y^{(1)}_{k-(K-1)/2} \cdot \left(y^{(1)}_{k-(K-3)/2}\right)^*\right\} \cdot \hat{p}[k]. \tag{41}$$

Note that in (41) the element $y^{(1)}_k$ is based on the estimated bit sequence and the estimated middle bit $\hat{p}[k]$ rather than the true quantities assumed in (39), and therefore the substitution of η by the iterative estimate $\hat{\eta}[k]$ appears justified.

8. SIMULATION RESULTS

This section contains a description of experiments that were performed and the results obtained.

8.1. Simulation model

Figure 1 shows a flow graph of our simulation model. The transmitter produces a GFSK modulated signal $s[n]$, as specified in Section 2, with parameters $K_{\text{BT}} = 0.5$, $h = 0.35$, and $N = 2$, thereby simulating a Bluetooth signal [1].
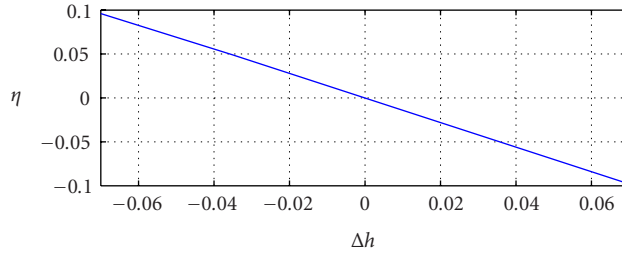
FIGURE 7: Term η for $L = 3$, independent of the modulation index offset.
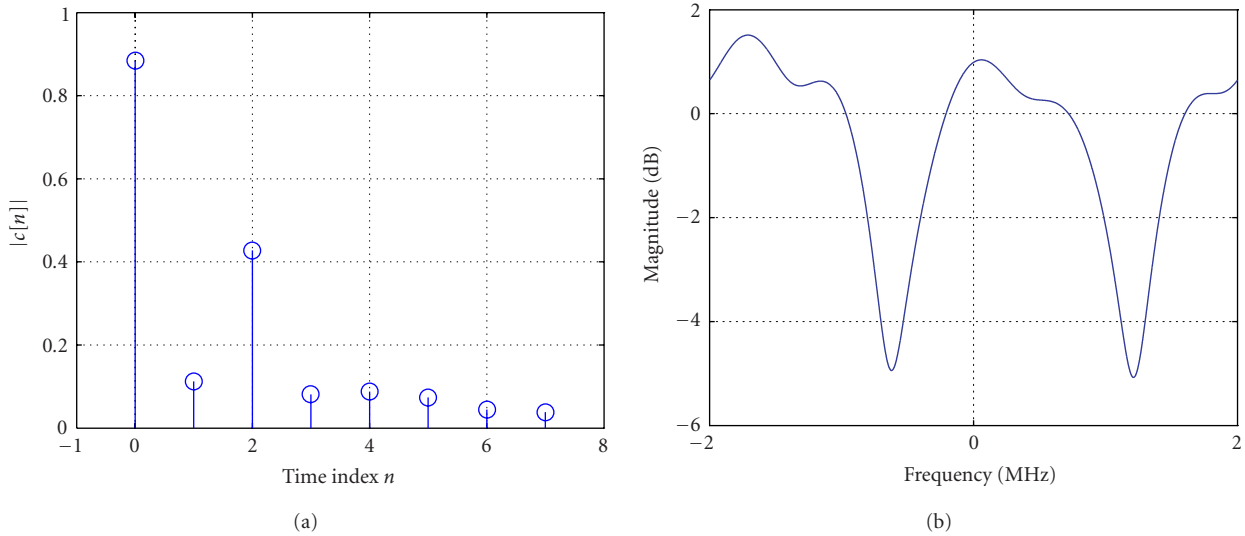

(a)

(b)

FIGURE 8: Example for a dispersive baseband channel of the generated ensemble: (a) CIR modulus $|c[n]|$ and (b) its magnitude response.
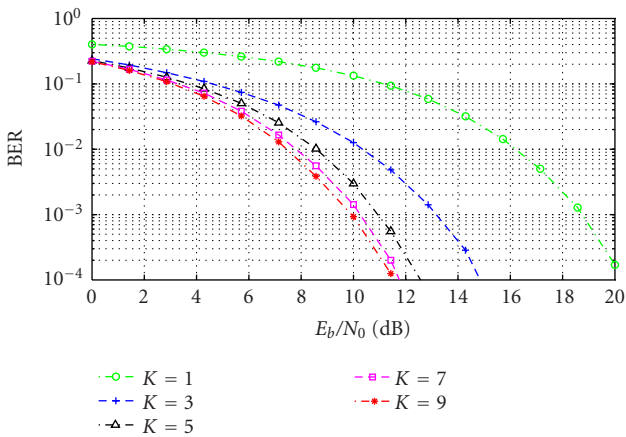

FIGURE 9: BER performance in AWGN with different observation intervals K, and parameters $K_{BT} = 0.5$, $h = 0.35$.

A channel $c[n]$, with impulse response and spectrum shown in Figure 8, was derived via discretisation of a Saleh-Valenzuela indoor propagation model [19]. The cluster and ray arrival rates for this model were set to 50 MHz and 200 MHz, while cluster and ray power decay constants were 400×10^{-9} and 3×10^{-3}, respectively. These parameters were used to derive 100 channel models, whose ensemble was used in our simulations. The resulting channel root mean square (RMS) value was approximately 500 nanoseconds—or half a bit period for Bluetooth—and typifies a large-sized office [20, 21], in which Bluetooth transceivers would normally operate. Output of the channel is subjected to a carrier offset term $e^{j\Delta\Omega n}$, and AWGN $v[n]$.

The receiver was the reduced complexity high-performance detector derived in Section 6, with $L_w = 32$, $\Delta\Omega = 0.3\pi$ rad, and $\Delta h = 0.07$. Note that the values of $\Delta\Omega$ and Δh define the worst-case scenario for initial carrier frequency and modulation index offsets between a transceiver pair adhering to the Bluetooth standard.

In our derivation, perfect timing synchronisation was assumed. In practice, we have utilised the maximum likelihood synchronisation method in [4] with good results.

8.2. High-performance receiver under AWGN channel conditions

If the channel does not impose any impairments other than AWGN, the matched filter receiver as suggested by [5] and implemented according to Section 6 approaches the performance of a maximum likelihood estimation of a data bit. This is exemplified in Figure 9 by the fact that it attains a

TABLE 1: GFSK high-performance receiver complexity with $L = 3$ and $N = 2$.

K	3	5	7	9
$C_{\text{standard}}/[\text{MAC}]$	192	1280	7168	36864
$\widetilde{C}_{\text{efficient}}/[\text{MAC}]$	96	288	1056	4128

BER of 10^{-3}—the maximum permissible in Bluetooth [1]—at 9.8 dB when $K = 9$, while the best of the relatively simple algorithms require 14.8 dB for the same feat [22], and some practitioners even assume that 21 dB will be required [23]. Performance also improves with increase in the observation interval, up to around $K = 9$. This was observed by the original proponent of this reception technique [5]. Hence, by reducing the computational penalty involved in increasing K, within limits, we can afford to increase the observation interval and improve BER. The complexity reduction attained by the algorithm derived in Section 6, with respect to its standard realisation, is epitomised in Table 1. When $K = 9$, the value for which increase in K does not yield significant improvement, the complexity is reduced by approximately 90%.

Example 2. As an indication of the implication of the above results, consider that for example, a Texas Instruments C6000 series DSP supports approximately 1 GMAC/s. Given the Bluetooth data rate of 1 Mbps with $M = 2$ and $N = 2$, Table 1 and Figure 9 show that a standard receiver implementation can only afford $K = 3$ and would require 13.2 dB for a BER of 10^{-3}, while the efficient realisation would enable $K = 7$ needing only 10.3 dB.

8.3. High-performance receiver under hostile channel conditions

Additionally to AWGN, channel impairments now include a dispersive channel, and a potential carrier and modulation index offset according to Section 8.1.

The magnitude response of the CIR in Figure 8b exhibits a 6 dB drop approximately every 2 MHz. Results portrayed in Figure 10 indicate that this degrades the system to the tune 11.6 dB at a BER of 10^{-3}.

Offsets in carrier frequency are manifested as offsets in the phase of the transmitted signal, which accumulates over the K bit observation interval. Figure 10 illustrates that a frequency offset of $\Delta\Omega = 0.3\pi$ is sufficient to cause the system to collapse when $K = 9$.

Loss in performance due to errors in modulation index of $\Delta h = 0.07$ is not critical. Instead, a 3 dB loss is experienced at a BER of 10^{-3}.

8.4. High-performance receiver under hostile channel conditions with compensation

We now assume the channel impairments as in Section 8.3, but operate the proposed receiver scheme with a blind equaliser, carrier offset correction, matched filter bank, and modulation index offset correction.
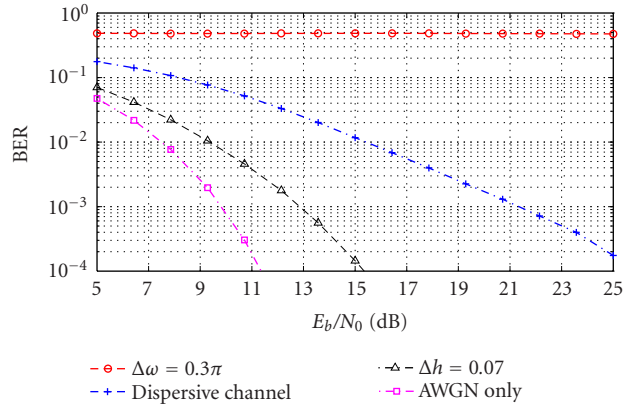

FIGURE 10: BER performance under isolated disturbances of a dispersive channel, a carrier frequency offset of $\Delta\Omega = 0.3\pi$, and a modulation index offset $\Delta h = 0.07$, with parameters $K = 9$, $N = 2$, $K_{\text{BT}} = 0.5$, $h = 0.35$.

Learning curves of the COC and MIOC in Figure 11 demonstrate their efficacy at combating errors in carrier frequency and modulation index, with convergence to the ideal in less than 2000 iterations for step sizes $\mu_\Theta^{\text{initial}} = 0.002$ and $\mu_h = 0.001$. Noting the sensitivity of the steady state error of the COC to μ_Θ, we adjust its value according to

$$\mu_\Theta = \mu_\Theta^{\text{initial}} \cdot \bar{\nabla}_{\Theta\chi}, \tag{42}$$

where $\bar{\nabla}_{\Theta\chi}$ is the mean result of (19), obtained over the last 1000 iterations. Figure 12 shows that when these algorithms are deployed with the CMA to combat the combined effects of a dispersive channel, carrier frequency and modulation index offsets of $\Delta\Omega = 0.3\pi$ and $\Delta h = 0.07$, respectively, they bring the performance of the system to 4 dB from that in AWGN and only 1 dB from the theoretical MMSE solution at BER of 10^{-3}.

9. CONCLUSION

The advent of software-defined radios (SDR) has allowed multiple wireless standards to run on common hardware platforms. A platform implementing relatively complex standards, such as IEEE 802.11b WLAN, will have excess computational capacity when running relatively simple standards like Bluetooth which uses GFSK. It is therefore important to seek the best way to utilise the extra capacity that may be available in an SDR to improve reception of GFSK modulated signals.

Towards this end we have suggested the use of the high-performance receiver because of its performance and relative insensitivity to errors in modulation index when compared to the Viterbi detector. However, in order to make this a practical option, and reduce battery consumption time, we have derived an algorithm that takes advantage of the nature of the GFSK-modulated signals to reduce the computational cost of the receiver by 90% when $K = 9$.
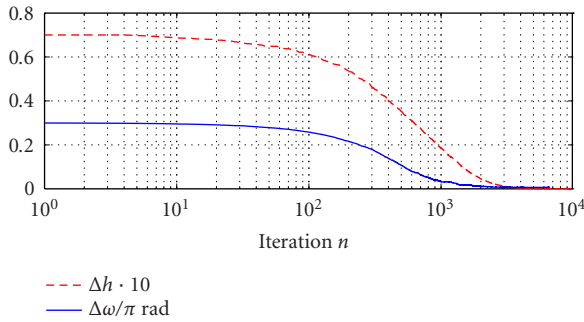
FIGURE 11: Learning curves for carrier frequency (COC), ω, and modulation index (MIOC), h, offset correction $\mu_\Theta^{\text{initial}} = 0.002$ and $\mu_h = 0.001$.
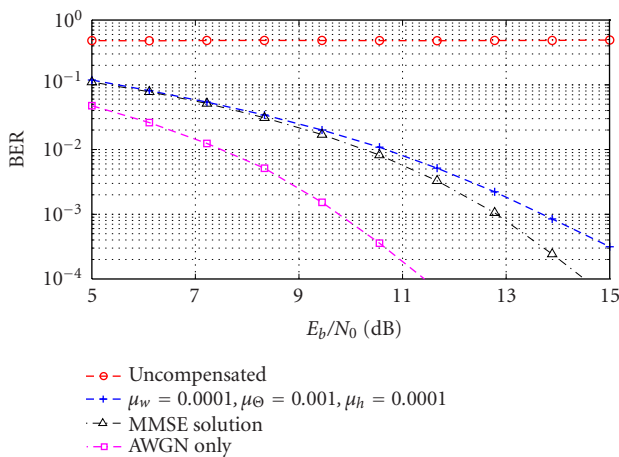

FIGURE 12: BER performance under combined disturbances of a dispersive channel, a carrier frequency offset of $\Delta\Omega = 0.3\pi$, and a modulation index offset $\Delta h = 0.07$, with compensation by the CMA, COC, and MIOC, and with parameters $K = 9$, $N = 2$, $K_{\text{BT}} = 0.5$, $h = 0.35$.

Apart from AWGN, other channel adversities exist. These include multipath propagation, and significant carrier frequency and modulation index offsets, which are permitted in the Bluetooth standard.

To tackle dispersive channels, we suggest the CMA for two reasons. Firstly, the CMA does not require a transmitted training signal, this feature is highly desirable in point-to-multipoint networks like Bluetooth because otherwise transmission will have to be interrupted by the control unit whenever it needs to retrain a tributary node that has experienced a change in channel conditions. Secondly, the CMA is insensitive to frequency error that may exist.

For the carrier frequency offset correction, we derived a stochastic-gradient-based algorithm, that allows us to correct such errors by multiplication of received samples by a derotating phasor.

To cater for modulation index offsets, we have shown that the relationship between the phase gained over a bit period

and the transmitted bit provides a means to detect errors in modulation index, and these can be corrected by recomputing the coefficients of the relatively small filter bank $\mathbf{W}^{(1)}$.

REFERENCES

[1] Specification of the Bluetooth System, *Bluetooth Special Interest Group*, February 2002, core.

[2] T. Aulin, N. Rydbeck, and C.-E. W. Sundberg, "Continuous phase modulation—Part II: partial response signaling," *IEEE Trans. Commun.*, vol. 29, no. 3, pp. 210–225, 1981.

[3] A. Soltanian and R. E. Van Dyck, "Performance of the Bluetooth system in fading dispersive channels and interference," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '01)*, vol. 6, pp. 3499–3503, San Antonio, Tex, USA, November 2001.

[4] J. B. Anderson, T. Aulin, and C.-E. W. Sundberg, *Digital Phase Modulation*, Plenum Press, New York, NY, USA, 1986.

[5] W. P. Osborne and M. B. Luntz, "Coherent and noncoherent detection CPFSK," *IEEE Trans. Commun.*, vol. 22, no. 8, pp. 1023–1036, 1974.

[6] T. A. Schonhoff, "Symbol error probabilities for M-ary CPFSK: coherent and noncoherent detection," *IEEE Trans. Commun.*, vol. 24, no. 6, pp. 644–652, 1976.

[7] B. A. Carlson, *Communication Systems*, McGraw-Hill, New York, NY, USA, 3rd edition, 1986.

[8] M. Nafie, A. Gatherer, and A. Dabak, "Decision feedback equalization for Bluetooth systems," in *Proc. Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing (ICASSP '01)*, vol. 2, pp. 909–912, Salt Lake City, Utah, USA, May 2001.

[9] M. Nafie, A. G. Dabak, T. M. Schmidl, and A. Gatherer, "Enhancements to the Bluetooth specification," in *Proc. 35th Asilomar Conference on Signals, Systems, and Computers*, vol. 2, pp. 1591–1595, Pacific Grove, Calif, USA, November 2001.

[10] M. Rupp, "LMS tracking behavior under periodically changing systems," in *Proc. 9th European Signal Processing Conference (EUSIPCO '02)*, pp. 737–738, Rhodes Island, Greece, September 1998.

[11] G. D. Forney, S. U. H. Qureshi, and C. K. Miller, "Multipoint networks: advances in modem design and control," in *National Telecommunications Conference*, Dallas, Tex, USA, November 1976.

[12] D. N. Godard, "Self-recovering equalization and carrier tracking in two-dimensional data communication systems," *IEEE Trans. Commun.*, vol. 28, no. 11, pp. 1867–1875, 1980.

[13] J. R. Treichler and B. Agee, "A new approach to multipath correction of constant modulus signals," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 31, no. 2, pp. 459–472, 1983.

[14] C. Robinson, A. Purvis, A. Lechner, and M. Hoy, "Characterisation of Bluetooth carrier frequency errors," in *Proc. 9th IEEE International Mixed Signal Testing Workshop (IMSTW '03)*, pp. 119–124, Seville, Spain, June 2003.

[15] P. Spasojevic and C. N. Georghiades, "Blind frequency compensation for binary CPM with $h = 1/2$ and a positive frequency pulse," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '98)*, vol. 1, pp. 573–578, Sydney, Australia, November 1998.

[16] B. Widrow and S. D. Stearns, *Adaptive Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1985.

[17] W. Hirt and S. Pasupathy, "Suboptimal reception of binary CPSK signals," in *IEE Proceedings Part F: Radar and Signal Processing*, vol. 128, no. 3, pp. 125–134, London, UK, June 1981.

[18] J. G. Proakis, *Digital Communications*, McGraw-Hill, New York, NY, USA, 3rd edition, 1995.

[19] A. A. M. Saleh and R. A. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Select. Areas Commun.*, vol. 5, no. 2, pp. 128–137, 1987.

[20] T. S. Rappaport, S. Y. Seidel, and K. Takamizawa, "Statistical channel impulse response models for factory and open plan building radio communicate system design," *IEEE Trans. Commun.*, vol. 39, no. 5, pp. 794–807, 1991.

[21] H. Hashemi, "The indoor radio propagation channel," *Proc. IEEE*, vol. 81, no. 7, pp. 943–968, 1993.

[22] R. Schiphorst, F. W. Hoeksema, and C. H. Slump, "Bluetooth demodulation algorithms and their performance," in *Proc. 2nd Karlsruhe Workshop on Software Radios*, pp. 99–106, Karlsruhe, Germany, March 2002.

[23] Ericsson Bluetooth Development Kit documentation, Ericsson, October 1999.

Charles Tibenderana obtained a B.S. degree in electrical engineering in 1999 from Makerere University in Uganda. He then joined the Uganda Communications Commission, where his main responsibility was spectrum management. In 2002 he completed an M.S. in radio frequency communication Systems at the University of Southampton and went on to pursue a Ph.D. in communications at the same institution.
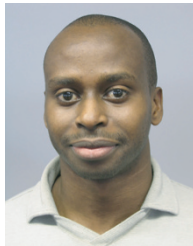
Stephan Weiss received the Dipl. Ing. degree from the University of Erlangen-Nürnberg, Erlangen, Germany, in 1995, and the Ph.D. degree from the University of Strathclyde, Glasgow, Scotland, in 1998, both in electronic and electrical engineering. In 1999 he joined the Communications Research Group within the School of Electronics and Computer Science at the University of Southampton, where he is a Senior Lecturer. Prior to this appointment he held a Visiting Lectureship at the University of Strathclyde. In 1996/1997 he was a Visiting Scholar at the University of Southern California. His research interests are mainly in adaptive and array signal processing, multirate systems, and signal expansions, with applications in communications, audio, and biomedical signal processing. For his work in biomedical signal processing, he was a corecipient of the 2001 Research Award of the German society on hearing aids. He is a Member of the VDE and EURASIP, a Senior Member of the IEEE, and a Member of the IEE Signal Processing Professional Network Executive Team.