EFFICIENT SEQUENCE DETECTION OF MULTI-CARRIER TRANSMISSIONS OVER DOUBLY DISPERSIVE CHANNELS

Sung-Jun Hwang and Philip Schniter

Dept. of ECE, The Ohio State University, Columbus, OH, USA.
{hwangsu,schniter}@ece.osu.edu

ABSTRACT

Multicarrier modulation (MCM) over a doubly dispersive (DD) channel yields complicated inter-carrier interference (ICI) and inter-symbol interference (ISI) responses. With appropriately designed MCM pulse shapes, however, ISI can be mostly suppressed, as can ICI outside a small subcarrier radius. In this case, the channel can be well described by a quasi-banded subcarrier coupling matrix. Several sequence detectors (SDs) have been proposed to leverage this quasi-banded structure, including linear, decision feedback (DF), and maximum likelihood (ML) schemes. Relative to linear and DF schemes, the ML schemes offer superior performance, but are significantly more complex, even when efficient Viterbi or sphere-detection algorithms are used. In this paper, we propose a new SD algorithm for the quasi-banded application with a frame error rate (FER) that is nearly indistinguishable from ML and an average complexity that is on par with DF SD.[1]

1. INTRODUCTION

Multi-carrier modulation (MCM) has been extensively studied as a practical method for communication over channels which are both time dispersive and frequency dispersive, i.e., doubly dispersive (DD). (See, e.g., the many references in [1].) The principle challenge faced when using MCM over these channels is effectively combating a rich and quickly varying inter-symbol interference (ISI) plus inter-carrier interference (ICI) response. Traditionally, MCM systems have been designed to make ISI negligible, enabling block demodulation of ICI-corrupted multicarrier symbols. More recently, MCM systems have been designed to also ensure a sparse ICI response, so that a given subcarrier sees significant interference only from a few neighboring subcarriers [2–5]. In these systems, the DD channel is well described by a "quasi-banded" subcarrier coupling matrix of the form shown in Fig. 1(a). Furthermore, by turning edge subcarriers off, the subcarrier coupling matrix can be made banded (rather than quasi-banded).

A number of sequence detection (SD) algorithms have been designed to exploit this banded or quasi-banded channel structure. These schemes include linear [6–8], decision feedback (DF) [5, 9, 10], iterative [3, 4], and maximum likelihood (ML) [11, 12] schemes. The linear and DF schemes have the advantage of low complexity, the ML schemes have the advantage of excellent performance,and the iterative schemes fall somewhere in-between in both performance and complexity.

In this paper, we propose a new approach to SD for quasi-banded channels that yields a frame error rate (FER) close to that of MLSD with a complexity close to that of DFSD. Our SD algorithm combines a novel sequential decoding (SqD) algorithm with a novel pre-processing algorithm, both specifically designed for the quasi-banded channel. Specifically, we propose a new channel-adaptive T-algorithm [13] and a fast implementation of the MMSE-GDFE pre-processor [14]. As we discuss in Section 2.2, other well-known SqD algorithms, such as the Schnorr-Euchner sphere decoder (SE-SpD) and the Fano algorithm, do not behave well on these quasi-banded channels. Numerical experiments are conducted to evaluate the efficacy of the proposed SqD relative to other algorithms.

The paper is organized as follows. Section 2 reviews MCM and SqD, Section 3 presents our proposed algorithms, and Section 4 presents numerical results. We use $(\cdot)^T$ to denote the transpose, $(\cdot)^*$ the conjugate, and $(\cdot)^H$ the conjugate transpose. $\boldsymbol{I}_L$ denotes the $L \times L$ identity matrix, and $[\boldsymbol{B}]_{m,n}$ denotes the element in the m^{th} row and n^{th} column of matrix $\boldsymbol{B}$, where row/column indices begin with zero. Similarly, $[\boldsymbol{b}]_m$ denotes the m^{th} entry of vector $\boldsymbol{b}$. Expectation is denoted by $\mathrm{E}\{\cdot\}$, the ℓ_2 norm by $\|\cdot\|$, the Kronecker delta by δ_l, and the modulo-N operation by $\langle\cdot\rangle_N$. Finally, $\mathbb{R}$ denotes the real field, $\mathbb{C}$ the complex field, and $\mathbb{Z}$ the integers.

2. BACKGROUND

2.1. System Model

Equations (1)-(4) describe the baseband-equivalent operation of an N-subcarrier QAM-based MCM system in a DD channel.

$$s(t) = \sum_{n=-\infty}^{\infty} \sum_{k=0}^{N-1} s_{k,n} a(t - nT_s) e^{j2\pi k F_s (t - nT_s)} \quad (1)$$

$$x(t) = \int_0^{T_h} h(t,\tau) s(t-\tau) d\tau + z(t) \quad (2)$$

$$x_{l,m} = \int_{-\infty}^{\infty} x(t)\, b^*(t - mT_s) e^{-j2\pi l F_s t} dt,\ 0 \le l < N \quad (3)$$

$$= \sum_{n=-\infty}^{\infty} \sum_{k=0}^{N-1} h_{l,m,k,n} s_{k,n} + z_{l,m}. \quad (4)$$

The MCM transmitter uses time-frequency shifts of the pulse $a(t)$ to modulate the QAM data $\{s_{k,n}\}$ onto the transmitted waveform $s(t)$. In (1), T_s denotes the symbol spacing and F_s the subcarrier spacing. The channel, characterized by the time-varying impulse response $h(t,\tau)$ and the noise waveform $z(t)$, produces the received signal $x(t)$. The receiver then uses time-frequency shifts of the pulse $b(t)$ to generate the subchannel outputs $\{x_{l,m}\}$. Equation

[1]This work was supported by the National Science Foundation CAREER grant CCR-0237037.

(4) relates $x_{l,m}$ to $s_{k,n}$ using the pulse-shaped channel coefficients $\{h_{l,m,k,n}\}$, which are functions of $h(t,\tau)$, $a(t)$ and $b(t)$.

We assume a square QAM constellation of size Q^2, with real and imaginary components chosen from the Q-ary PAM constellation $\mathcal{S} := \{-\frac{Q-1}{2}, -\frac{Q-1}{2}+1, \ldots, \frac{Q-1}{2}\}$. By splitting the complex-valued elements $\{x_{l,m}\}_{l=0}^{N-1}$, $\{s_{k,m}\}_{k=0}^{N-1}$, $\{z_{l,m}\}_{l=0}^{N-1}$, and $\{h_{l,m,k,n}\}_{l,k=0}^{N-1}$ from (4) into their real and imaginary components, we obtain the real-valued vector model (5), which will be more convenient for SqD implementation.

$$\boldsymbol{x}_m \;=\; \sum_{n=-\infty}^{\infty} \boldsymbol{H}_{m,n}\boldsymbol{s}_{m-n} + \boldsymbol{z}_m. \tag{5}$$

In particular, the vector $\boldsymbol{x}_m \in \mathbb{R}^{2N}$ is constructed so that $[\boldsymbol{x}_m]_{2l} = \mathrm{Re}(x_{l,m})$ and $[\boldsymbol{x}_m]_{2l+1} = \mathrm{Im}(x_{l,m})$ for $0 \le l < N$, while $\boldsymbol{s}_m \in \mathbb{R}^{2N}$, $\boldsymbol{z}_m \in \mathbb{R}^{2N}$, and $\boldsymbol{H}_{m,n} \in \mathbb{R}^{2N\times 2N}$ are constructed in a similar manner. Note that the matrix sequence $\{\boldsymbol{H}_{m,n}\}_{n=-\infty}^{\infty}$ specifies the impulse response relating the transmitted multicarrier-symbol sequence $\{\boldsymbol{s}_n\}_{n=-\infty}^{\infty}$ to the time-m modulator output $\boldsymbol{x}_m$; the coefficients $\{\boldsymbol{H}_{m,n}\}_{n\neq 0}$ characterize the inter-symbol interference (ISI) while the off-diagonal elements of $\boldsymbol{H}_{m,0}$ characterize the inter-carrier interference (ICI). For brevity, we define $L := 2N$.

While much of the theoretical MCM literature assumes continuous pulse shapes as in (1)-(3), practical MCM implementations use pulse sequences $\{a_k\}$ and $\{b_k\}$ to modulate a chip-waveform $p(t)$ with approximate time support $T_c = (NF_s)^{-1}$ and approximate frequency support NF_s, i.e., $a(t) = \sum_k a_k p(t-kT_c)$ and $b(t) = \sum_k b_k p(t-kT_c)$. In this case, the significant entries in $\boldsymbol{H}_{m,0}$ lie within the "quasi-banded" region shown in Fig. 1(a), where the "ICI radius" D depends on the pulse designs and channel spreading characteristics. This phenomenon motivates the partition $\boldsymbol{H}_{m,0} = \boldsymbol{H}_m^{(D)} + \bar{\boldsymbol{H}}_m^{(D)}$, where $\boldsymbol{H}_m^{(D)}$ extracts the coefficients of $\boldsymbol{H}_{m,0}$ inside the shaded region in Fig. 1(a), and $\bar{\boldsymbol{H}}_m^{(D)}$ extracts those outside. Using this partition, we rewrite (5) as

$$\boldsymbol{x}_m \;=\; \boldsymbol{H}_m^{(D)}\boldsymbol{s}_m + \underbrace{\bar{\boldsymbol{H}}_m^{(D)}\boldsymbol{s}_m + \sum_{n\neq 0}\boldsymbol{H}_{m,n}\boldsymbol{s}_{m-n} + \boldsymbol{z}_m}_{:=\,\boldsymbol{w}_m}, \tag{6}$$

where $\boldsymbol{H}_m^{(D)}\boldsymbol{s}_m$ contains the signal and "significant-ICI," and where $\boldsymbol{w}_m$ contains the noise, ISI, and "insignificant-ICI." With the proper choice of pulse shapes (e.g., the max-SINR transmitter pulse shaping scheme from [4]), we find $\mathrm{E}\{\boldsymbol{w}_m\boldsymbol{w}_m^T\} \approx \sigma_z^2\boldsymbol{I}_L$, even with a highly dispersive channel at relatively high SNR. In the sequel, we suppress the m and D notation in (6) to yield

$$\boldsymbol{x} \;=\; \boldsymbol{H}\boldsymbol{s} + \boldsymbol{w}, \tag{7}$$

where $\boldsymbol{H}$ exhibits the quasi-banded structure in Fig. 1(a) and $\boldsymbol{w}$ is white Gaussian noise.

2.2. Sequential Decoding

By definition, the MLSD solution to (7) for known $\boldsymbol{H}$ is given by

$$\hat{\boldsymbol{s}}_{\mathsf{ML}} \;=\; \arg\min_{\boldsymbol{s}\in\mathcal{S}^L} \left\|\boldsymbol{x} - \boldsymbol{H}\boldsymbol{s}\right\|^2. \tag{8}$$

The brute-force approach to finding $\hat{\boldsymbol{s}}$ requires $\mathcal{O}(Q^L)$ operations, which is impractical for large L. If $\boldsymbol{H}$ was banded with a band radius of D, then the Viterbi algorithm could be used to solve (8) with a complexity of $L(2D+1)Q^{(2D+1)}$ real MACs per frame [12]. But, since $\boldsymbol{H}$ is only quasi-banded, a different approach is needed. One could instead use, for example, a "tail-biting" MLSD which hypothesizes an initial state at an arbitrary location within the frame, runs the standard Viterbi algorithm from that state, and forces a termination back to that state. Exhaustively searching among the Q^{2D} possible hypotheses yields a MLSD algorithm with a complexity of $L(2D+1)Q^{(4D+1)}$ real MACs per frame. These Viterbi-based algorithms, while much cheaper than brute force search, will still be impractical in many applications.

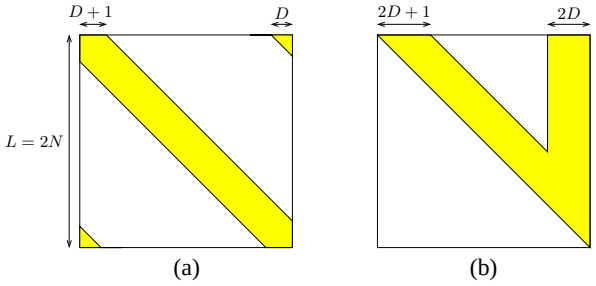

Fig. 1. Channel matrices associated with MCM: (a) "quasi-banded" channel matrix, (b) "V-shaped" channel matrix.

As an alternative to brute-force and Viterbi MLSD, one might consider sphere decoding (SpD) [15–17] or more general forms of sequential decoding (SqD) [18]. For instance, modern SpD algorithms have been claimed to yield ML (or nearly ML) estimates with approximately $\mathcal{O}(L^3)$ MACs per frame *on average*, and at *high* SNR, regardless of constellation size Q [17]. In fact, we will show that, by leveraging the quasi-banded structure of $\boldsymbol{H}$ in (6), it is possible to coax near-ML performance out of SqD with approximately $\mathcal{O}(L^2)$ MACs per frame at *any* SNR.

SqD consists of a pre-processing step which triangularizes $\boldsymbol{H}$ and a tree-search step which searches for $\hat{\boldsymbol{s}}_{\mathsf{ML}}$. We provide background on these two steps below.

2.2.1. SqD Pre-Processing

The traditional SqD pre-processing stage uses the QR decomposition $\boldsymbol{H} = \boldsymbol{Q}\boldsymbol{R}$ to transform (7) into the equivalent system $\boldsymbol{x}' = \boldsymbol{Q}^T\boldsymbol{x} = \boldsymbol{R}\boldsymbol{s} + \boldsymbol{w}'$, where $\boldsymbol{R}$ is upper triangular and $\boldsymbol{w}'$ is statistically equivalent to $\boldsymbol{w}$. In this case, the detection problem (8) can be equivalently restated as $\hat{\boldsymbol{s}}_{\mathsf{ML}} = \arg\min_{\boldsymbol{s}\in\mathcal{S}^L}\|\boldsymbol{x}' - \boldsymbol{R}\boldsymbol{s}\|^2$. When the pre-processed channel matrix $\boldsymbol{R}$ is ill-conditioned, however, the complexity of SqD is known to grow significantly [19].

MMSE-GDFE pre-processing [14] was recently proposed as an alternative. Intuitively, MMSE-GDFE pre-processing yields better "conditioning," and, practically, it has been observed to reduce SqD search time significantly. We now describe the standard MMSE-GDFE pre-processing algorithm. Under the assumption that $\boldsymbol{s}$ and $\boldsymbol{w}$ are zero-mean uncorrelated with covariance matrices $\sigma_s^2\boldsymbol{I}_L$ and $\sigma_z^2\boldsymbol{I}_L$, respectively, we define $\gamma := \sigma_s^2/\sigma_z^2$ and the augmented channel matrix $\tilde{\boldsymbol{H}}$:

$$\tilde{\boldsymbol{H}} \;:=\; \begin{pmatrix}\boldsymbol{H}\\ \frac{1}{\sqrt{\gamma}}\boldsymbol{I}_L\end{pmatrix} \;=\; \tilde{\boldsymbol{Q}}\tilde{\boldsymbol{R}} \;=\; \begin{pmatrix}\boldsymbol{Q}_1\\ \boldsymbol{Q}_2\end{pmatrix}\tilde{\boldsymbol{R}}. \tag{9}$$

As shown in (9), we take the QR decomposition $\tilde{\boldsymbol{H}} = \tilde{\boldsymbol{Q}}\tilde{\boldsymbol{R}}$ and

partition $\tilde{\boldsymbol{Q}}$ into two square matrices. The transformed observation $\boldsymbol{\rho} := \boldsymbol{Q}_1^T \boldsymbol{x}$ is then used in the pre-processed detection problem

$$\hat{\boldsymbol{s}}_{\mathsf{PP}} \;=\; \arg\min_{\boldsymbol{s}\in\mathcal{S}^L} \left\|\boldsymbol{\rho} - \tilde{\boldsymbol{R}}\boldsymbol{s}\right\|^2. \tag{10}$$

Because $\boldsymbol{Q}_1 \in \mathbb{R}^{L\times L}$ is generally non-orthogonal, we cannot claim that $\hat{\boldsymbol{s}}_{\mathsf{PP}} = \hat{\boldsymbol{s}}_{\mathsf{ML}}$. However, it is interesting to note that the error $\boldsymbol{n} := \boldsymbol{\rho} - \tilde{\boldsymbol{R}}\boldsymbol{s}$, while signal dependent and non-Gaussian, is white with covariance $\sigma_z^2 \boldsymbol{I}_L$. Finally, it is important to note that, when $\boldsymbol{H}$ has the quasi-banded structure in Fig. 1(a), $\tilde{\boldsymbol{R}}$ will have the "V-shaped" structure in Fig. 1(b). In Section 3.1, we propose a fast MMSE-GDFE implementation suitable for quasi-banded $\boldsymbol{H}$.

2.2.2. Tree Search

After pre-processing, the MLSD problem (10) corresponds to a tree search over a tree with depth L, where every tree node has Q children. A brute-force approach to tree search would entail the examination of the Euclidean metric in (10) at each of the Q^L leaf nodes. We are interested in search algorithms which prune branches that are unlikely to contain the ML path, thus drastically reducing the search complexity. Unlike their ML counterparts, *quasi*-ML tree search algorithms can, in some cases, accidentally discard the ML path, and hence return a suboptimal sequence estimate. Thus, each quasi-ML algorithm achieves a particular tradeoff between performance and complexity.

Tree search algorithms can be categorized [18, 19] as breadth-first search (BrFS), depth-first search (DFS), or best-first search (BeFS). BrFS includes, e.g., the M-algorithm [18], T-algorithm [13], and Pohst sphere decoder [20]. DFS includes, e.g., the Schnorr-Euchner sphere decoder (SE-SpD) and its variants [15–17]; and BeFS includes, e.g., the stack and Fano algorithms [19]. A recent comprehensive comparison [19] found that a properly-designed Fano algorithm achieved a better complexity/performance tradeoff than other SqD algorithms when $\tilde{\boldsymbol{R}}$ has a fully populated upper triangle. When $\tilde{\boldsymbol{R}}$ is V-shaped, however, we have found that the Fano algorithm—and BeFS and DFS algorithms in general—do not give a good complexity/performance tradeoff [1]. A thorough explanation for this behavior is given in [1].

Some BrFS algorithms have a complexity that is relatively insensitive to SNR and the structure of $\tilde{\boldsymbol{R}}$, suggesting that BrFS might be advantageous in our application. The M-algorithm, for example, has complexity that is *invariant* to both SNR and $\tilde{\boldsymbol{R}}$. At high SNR, however, the M-algorithm is more expensive than DFS and BeFS because it is not aggressive enough in branch pruning. Hence, a better complexity/performance tradeoff could be achieved by a BrFS algorithm that varies the number of branches taken at each level. For example, at level i, the T-algorithm only extends paths from nodes whose Euclidean metrics are within T_i of best Euclidean metric found at that level, where T_i is chosen to achieve a particular complexity/performance tradeoff. While several approaches to the design of T_i have been proposed, we are not aware of any that make use of the channel realization as well the SNR. In Section 3.2, we propose such an algorithm.

3. MCM SEQUENCE DETECTION

3.1. Fast MMSE-GDFE Pre-Processing

The MMSE-GDFE pre-processing originally proposed in [14] involves QR decomposition with complexity $\mathcal{O}(L^3)$. Here we propose an $\mathcal{O}(D^2L)$ implementation of MMSE-GDFE pre-processing that leverages the quasi-banded structure of $\boldsymbol{H}$. We note connections to the fast MMSE-DFE in [5], which was formulated for the banded (as opposed to quasi-banded) $\boldsymbol{H}$ that occurs with inactive edge subcarriers.

Recall the augmented channel matrix $\tilde{\boldsymbol{H}}$ and its QR decomposition in (9). Note that, while $\boldsymbol{H}$ is quasi-banded, $\tilde{\boldsymbol{H}}$ is not. However, the matrix $\tilde{\boldsymbol{H}}^T\tilde{\boldsymbol{H}}$, which can be computed in $(4D^2+4D+2)L$ multiply-accumulate (MAC) operations, *is* quasi-banded with $4D+1$ active diagonals. Now, since $\tilde{\boldsymbol{Q}}$ is an orthogonal matrix, we know $\tilde{\boldsymbol{H}}^T\tilde{\boldsymbol{H}} = \tilde{\boldsymbol{R}}^T\tilde{\boldsymbol{R}}$. Hence, $\tilde{\boldsymbol{R}}$ can be obtained via Cholesky factorization [21] of $\tilde{\boldsymbol{H}}^T\tilde{\boldsymbol{H}}$ in $O(D^2L)$ operations. Table 1 details the fast Cholesky factorization $\boldsymbol{A} = \boldsymbol{G}\boldsymbol{G}^T$, where $\boldsymbol{A} := \tilde{\boldsymbol{H}}^T\tilde{\boldsymbol{H}}$ and where $\boldsymbol{G} := \tilde{\boldsymbol{R}}^T$ is the lower triangular Cholesky factor. This fast computation of $\tilde{\boldsymbol{R}}$ can be shown to consume $(10D^2+11D+2)L - \frac{1}{3}(74D^3+133D^2+44D+3)$ MACs [1].

Next, we consider the implementation of the pre-processing operation $\boldsymbol{\rho} = \boldsymbol{Q}_1^T\boldsymbol{x}$. Multiplication of this equality by $\tilde{\boldsymbol{R}}^T$ yields

$$\tilde{\boldsymbol{R}}^T\boldsymbol{\rho} \;=\; \tilde{\boldsymbol{R}}^T\boldsymbol{Q}_1^T\boldsymbol{x} \;=\; \boldsymbol{H}^T\boldsymbol{x} \;:=\; \boldsymbol{b}. \tag{11}$$

Due to quasi-banded $\boldsymbol{H}$, the vector $\boldsymbol{b}$ can be computed in $(2D+1)L$ MAC operations. From $\boldsymbol{b}$ we can solve (11) for $\boldsymbol{\rho}$ using forward substitution in $\mathcal{O}(DL)$ additional operations, because $\tilde{\boldsymbol{R}}^T$ has the sparse "V-shaped" structure in Fig. 1(b). In total, this consumes $(6D+2)L - 6D^2 - 3D$ MAC operations. Combining forward substitution with fast Cholesky decomposition, our fast MMSE-GDFE preprocessing requires $\big(14D^2+21D+6\big)L - \frac{76}{3}D^3 - 53D^2 - \frac{53}{3}D - 1$ real MAC operations.

Though SqD pre-preprocessing normally includes lattice reduction and column ordering [17], these operations would destroy the quasi-banded structure of $\boldsymbol{H}$ and preclude the opportunity for fast MMSE-GDFE pre-processing. Still, we might consider an n-place circular shift in the column ordering, since this preserves the quasi-banded nature of $\tilde{\boldsymbol{H}}^T\tilde{\boldsymbol{H}} = \tilde{\boldsymbol{R}}^T\tilde{\boldsymbol{R}}$. Specifically, we order the columns so that the last one has maximum norm. This ensures that the PAM symbol contributing the most energy to $\boldsymbol{x}$ is placed at the root of the tree. Evaluating the column norms requires $\mathcal{O}(DN)$ operations. We have observed, numerically, that this "circular ordering" scheme yields a modest improvement in terms of the performance/complexity tradeoff.

3.2. Channel-Adaptive T-algorithm

The T-algorithm [13] is a BrFS algorithm which, at the i^{th} level, discards any partial path $\boldsymbol{s}^{(i)} := [s_i, s_{i+1}, \ldots, s_{L-1}]^T \in \mathcal{S}^{L-i}$ whose partial-path metric

$$\mathcal{M}(\boldsymbol{s}^{(i)}) \;:=\; \sum_{k=i}^{L-1}\Big|\rho_k - \sum_{l=k}^{L-1}\tilde{r}_{k,l}s_l\Big|^2. \tag{12}$$

exceeds that of the "best" partial path $\boldsymbol{s}_\star^{(i)} := \arg\min_{\boldsymbol{s}^{(i)}} \mathcal{M}(\boldsymbol{s}^{(i)})$ by an amount $\geq T_i$. (Here, the root node corresponds to the L^{th} level and the leaf nodes to the 0^{th} level.) Clearly, the T-algorithm will make a frame error if the *true* partial path $\boldsymbol{s}_{\mathsf{T}}^{(i)}$ is discarded at any level $i \in \{L-1, L-2, \ldots, 0\}$.

In our adaptive T-algorithm, we set T_i so that, *when the true path is not the best partial path*, the true path will be discarded with probability less than ϵ_o

$$\Pr\big\{\mathcal{M}(\boldsymbol{s}_{\mathsf{T}}^{(i)}) > \mathcal{M}(\boldsymbol{s}_\star^{(i)}) + T_i \;\big|\; \mathcal{M}(\boldsymbol{s}_{\mathsf{T}}^{(i)}) > \mathcal{M}(\boldsymbol{s}_\star^{(i)})\big\} < \epsilon_o \tag{13}$$

Table 1. Fast Cholesky factorization of quasi-banded $\boldsymbol{A}$.

Say that $\boldsymbol{A} = \boldsymbol{G}\boldsymbol{G}^T$, where $\boldsymbol{G}$ is lower triangular and
$\boldsymbol{A} \in \mathbb{R}^{L\times L}$ is quasi-banded with $4D+1$ active diagonals.

```
for j = 0 : L − 4D − 1
    v_{j:L−1} = [A]_{j:L−1,j}
    m_1 = max{0, j − 2D − 1}
    m_2 = j + 2D − 1
    for i = m_1 : j − 1
        v_{j:m_2} = v_{j:m_2} − [G]_{j,i}[G]_{j:m_2,j}
        v_{L−2D−1:L−1} = v_{L−2D−1:L−1} − [G]_{j,i}[G]_{L−2D−1:L−1,j}
    end
    [G]_{j:m_2,j} = v_{j:m_2}/√v_j
    [G]_{L−2D−1:L−1,j} = v_{L−2D−1:L−1}/√v_j
end
for j = L − 4D : L − 2D − 1
    v_{j:L−1} = [A]_{j:L−1,j}
    m_1 = max{0, j − 2D − 1}
    for i = m_1 : j − 1
        v_{j:L−1} = v_{j:L−1} − [G]_{j,i}[G]_{j:L−1,j}
    end
    [G]_{j:L−1,j} = v_{j:L−1}/√v_j
end
for j = L − 2D : L − 1
    v_{j:L−1} = [A]_{j:L−1,j}
    for i = 0 : j − 1
        v_{j:L−1} = v_{j:L−1} − [G]_{j,i}[G]_{j:L−1,j}
    end
    [G]_{j:L−1,j} = v_{j:L−1}/√v_j
end
```

Note that this is different from simply setting T_i so that true paths are discarded with probability ϵ_o. In the latter case, T_i will increase—thereby increasing search complexity—at low SNR. Intuition, however, tells us that it is not worthwhile to search extensively at low SNR because, even if found, the ML path is likely to be in error.

With $\mu^{(i)} := \mathcal{M}(\boldsymbol{s}_{\mathsf{T}}^{(i)}) - \mathcal{M}(\boldsymbol{s}_{\star}^{(i)})$ we can rewrite (13) as

$$\Pr\{\mu^{(i)} > T_i \mid \mu^{(i)} > 0\} < \epsilon_o. \tag{14}$$

We now analyze the random variable $\mu^{(i)}$. To do this, we define $\boldsymbol{\rho}^{(i)} := [\rho_i, \rho_{i+1}, \ldots, \rho_{L-1}]^T$ and construct $\tilde{\boldsymbol{R}}^{(i)} \in \mathbb{R}^{(L-i)\times(L-i)}$ from the last $L-i$ rows and columns of $\tilde{\boldsymbol{R}}$, i.e., $[\tilde{\boldsymbol{R}}^{(i)}]_{j,k} = [\tilde{\boldsymbol{R}}]_{j+i,k+i}$. This way, (12) can be written as $\mathcal{M}(\boldsymbol{s}^{(i)}) = \|\boldsymbol{\rho}^{(i)} - \tilde{\boldsymbol{R}}^{(i)}\boldsymbol{s}^{(i)}\|^2$. Defining the error vector $\boldsymbol{e}^{(i)} := \boldsymbol{s}_{\star}^{(i)} - \boldsymbol{s}_{\mathsf{T}}^{(i)}$ and the interference vector $\boldsymbol{n}^{(i)} := \boldsymbol{\rho}^{(i)} - \tilde{\boldsymbol{R}}^{(i)}\boldsymbol{s}_{\mathsf{T}}^{(i)}$, we find

$$\mu^{(i)} \;=\; 2\boldsymbol{n}^{(i)T}\tilde{\boldsymbol{R}}^{(i)}\boldsymbol{e}^{(i)} - \|\tilde{\boldsymbol{R}}^{(i)}\boldsymbol{e}^{(i)}\|^2. \tag{15}$$

Since the statistics of $\boldsymbol{e}^{(i)}$ are difficult to characterize, we approximate $\boldsymbol{e}^{(i)}$ by the simple error event most likely to occur at the i^{th} level, i.e., an error of the form $\boldsymbol{e}^{(i)} = [0, \ldots, 0, \pm 1, 0, \ldots, 0]^T$. The partial metric $\mathcal{M}(\boldsymbol{s}^{(i)}) = \|\boldsymbol{\rho}^{(i)} - \tilde{\boldsymbol{R}}^{(i)}\boldsymbol{s}^{(i)}\|^2$ suggests that this error will occur at the index of the "weakest" column of $\tilde{\boldsymbol{R}}^{(i)}$. Thus we assume $[\boldsymbol{e}^{(i)}]_l = \pm\delta_{l-l_i}$ for $l_i := \arg\min_l \|\tilde{\boldsymbol{r}}_l^{(i)}\|$, where $\tilde{\boldsymbol{r}}_l^{(i)} \in \mathbb{R}^{L-i}$ denotes the l^{th} column of $\tilde{\boldsymbol{R}}^{(i)}$. In this case, $\mu^{(i)} = \pm 2\boldsymbol{n}^{(i)T}\tilde{\boldsymbol{r}}_{l_i}^{(i)} - \|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|^2$. Recall from Section 2.2 that $\boldsymbol{n}$ is zero-mean, white, and non-Gaussian, where the non-Gaussianity is due to a contribution from not-yet-detected PAM symbols. To proceed further, we approximate $\boldsymbol{n}$ as Gaussian with covariance $\sigma_z^2\boldsymbol{I}_L$, so that $\mu^{(i)} \sim \mathcal{N}\big(-\|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|^2, 4\|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|^2\sigma_z^2\big)$. Using the Gaussian cdf, we can solve (14) for T_i given a particular ϵ_o. Using Bayes rule and $\mathcal{Q}(x) := \frac{1}{\sqrt{2\pi}}\int_x^\infty e^{-\frac{x^2}{2}}dx$, we find [1]

$$T_i \;=\; 2\sigma_z\|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|\mathcal{Q}^{-1}\left(\epsilon_o\mathcal{Q}\left(\frac{\|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|}{2\sigma_z}\right)\right) - \|\tilde{\boldsymbol{r}}_{l_i}^{(i)}\|^2. \tag{16}$$

4. NUMERICAL RESULTS

Uncoded QPSK symbols (i.e., $Q = 2$) were communicated over $N = 64$ MCM subcarriers (i.e., $L = 128$), and the demodulator output $\boldsymbol{x}_m$ from (4) was used to detect the transmitted sequence $\boldsymbol{s}_m$. The max-SINR transmitter pulse (MSTP) MCM scheme from [4] was used since it does an excellent job to ensure $\mathrm{E}\{\boldsymbol{w}_m\boldsymbol{w}_m^T\} \approx \sigma_z^2\boldsymbol{I}_L$ at the high spectral efficiency of 1 QPSK-symbol/sec/Hz. Realizations of a wide sense stationary uncorrelated scattering (WSSUS) Rayleigh fading channel were generated using Jakes method. The channel had a uniform delay-profile with normalized delay spread $N_h = T_h/T_c = 16$ and normalized single-sided Doppler spreads $f_dT_c \in \{0.001, 0.003\}$. The receiver was assumed to have perfect knowledge of ± 3 subcarriers of local ICI (i.e., $D = 6$); ISI and residual ICI were treated as unknown interference.

MLSD, quasi-ML SqD, and MMSE-DFE were examined. In each case, we first applied circular ordering and fast MMSE-GDFE pre-processing to arrive at (10). This is justified for MLSD since it has been shown [22] that $\hat{\boldsymbol{s}}_{\mathsf{ML}} = \boldsymbol{s} \Rightarrow \hat{\boldsymbol{s}}_{\mathsf{PP}} = \boldsymbol{s}$ with uncoded QPSK. For MLSD, we solve (10) via SE-SpD, while for quasi-ML SqD, we obtain an approximate solution to (10) via suboptimal tree search. For the M-algorithm, we found $M = 8$ to be the lowest value that yielded near-ML performance over the SNR range of interest, while, for the adaptive T-algorithm, we set $\epsilon_o = 10^{-5}$ and limited the maximum list size to 8. For the Fano algorithm from [19], we used a bias of $\sigma_z^2/2$ and a step size of σ_z^2.

Figure 2 shows the FER performance of various SqD algorithms. With the exception of MMSE-DFE, all SqD algorithms give near-ML performance. Note that the ML and SE-SpD traces are identical. The MMSE-DFE error floor is consistent with that observed in [5] for a max-SINR receiver-pulse MCM scheme.

Figure 3 compares the average complexity of the SqD algorithms and the Viterbi algorithm (as used in [12]). There, "complexity" is plotted on a log base-L scale, as in other near-ML SqD studies (e.g., [17, 19]). For the SqD algorithms, we plot the average number of real MACs per frame needed to achieve the FERs in Fig. 2, including those required for MSTP-MCM demodulation, circular ordering, and fast MMSE-GDFE pre-processing.

Due to the V-shaped structure of $\tilde{\boldsymbol{R}}$, the SE-SpD and Fano algorithms exhibit DFE-like complexity at high SNR but explosive complexity at low SNR, while the M-algorithm has the same complexity at all SNRs. Remarkably, the adaptive T-algorithm yields DFE-like complexity at high SNR and better-than-M-algorithm complexity at low SNR. This is a consequence of the fact that the T-algorithm uses channel knowledge to intelligently guide its search. Note that the Viterbi complexity is much larger than that of BrFS and MMSE-DFE. Furthermore, the Viterbi complexity plotted in Fig. 3, i.e., $L(2D+1)Q^{(4D+1)} = L^{3.39}$, corresponds to the case where D edge subcarriers are inactive; the "tail-biting" MLSD proposed in Section 2.2, suitable for the general case, would require $L(2D+1)Q^{(4D+1)} = L^{5.10}$ MACs per frame.

In conclusion, Fig. 3 shows that, by sacrificing a fraction-of-a-dB in performance relative to MLSD, SqD can be implemented with near-MMSE-DFE average complexity, even when all subcarriers are active.

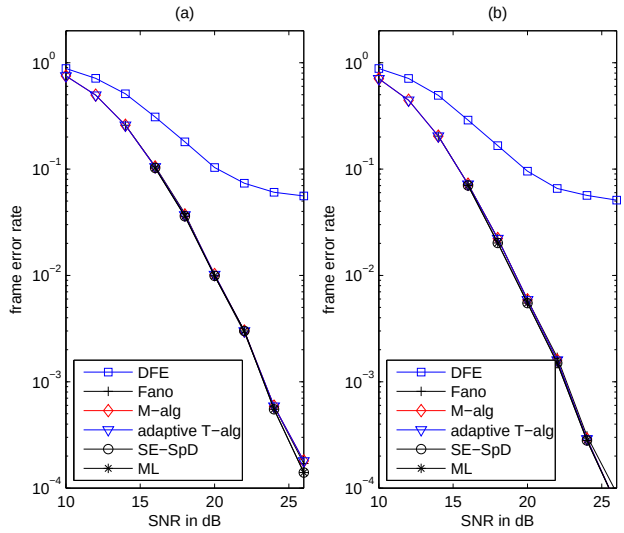

Fig. 2. Performance of several SqDs on doubly dispersed MSTP-MCM with perfect knowledge of local ICI (i.e., $D = 6$) at (a) $f_d T_c = 0.001$; (b) $f_d T_c = 0.003$.

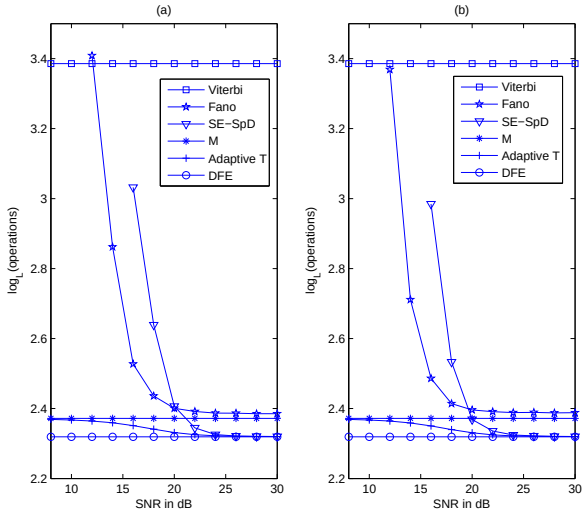

Fig. 3. Number of real MAC operations per frame for doubly dispersed MSTP-MCM reception at (a) $f_d T_c = 0.001$; (b) $f_d T_c = 0.003$. Viterbi complexity above assumes inactive edge subcarriers (as in [12]); for active subcarriers, the tail-biting version would require $L^{5.10}$ MACs/frame.

5. REFERENCES

[1] S.-J. Hwang and P. Schniter, "Efficient sequence detection of multicarrier transmissions over doubly dispersive channels," *EURASIP J. Appl. Signal Processing*. submitted June 2005.

[2] P. Schniter, "A new approach to multicarrier pulse design for doubly-dispersive channels," in *Proc. Allerton Conf. Commun., Control, and Computing*, Oct. 2003.

[3] P. Schniter, "Low-complexity equalization of OFDM in doubly-selective channels," *IEEE Trans. Signal Processing*, vol. 52, pp. 1002–1011, Apr. 2004.

[4] S. Das and P. Schniter, "A new pulse shaped frequency division multiplexing technique for doubly dispersive channels," in *Proc. Asilomar Conf. Signals, Systems and Computers*, Nov. 2004.

[5] L. Rugini, P. Banelli, and G. Leus, "Block DFE and windowing for Doppler-affected OFDM systems," in *Proc. IEEE Workshop Signal Processing Adv. in Wireless Commun.*, June 2005.

[6] W. G. Jeon, K. H. Chang, and Y. S. Cho, "An equalization technique for orthogonal frequency-division multiplexing systems in time-variant multipath channels," *IEEE Trans. Commun.*, vol. 47, pp. 27–32, Jan. 1999.

[7] A. Stamoulis, S. N. Diggavi, and N. Al-Dhahir, "Intercarrier interference in MIMO OFDM," *IEEE Trans. Signal Processing*, vol. 50, pp. 2451–2464, Oct. 2002.

[8] I. Barhumi, G. Leus, and M. Moonen, "Time-domain and frequency-domain per-tone equalization for OFDM in doubly-selective channels," *Signal Processing*, vol. 84, pp. 2055–2066, Nov. 2004.

[9] Y.-S. Choi, P. J. Voltz, and F. A. Cassara, "On channel estimation and detection for multicarrier signals in fast and selective Rayleigh fading channels," *IEEE Trans. Commun.*, vol. 49, pp. 1375–1387, Aug. 2001.

[10] X. Cai and G. B. Giannakis, "Bounding performance and suppressing inter-carrier interference in wireless mobile OFDM," *IEEE Trans. Commun.*, vol. 51, pp. 2047–2056, Dec. 2003.

[11] K. Matheus and K.-D. Kammeyer, "Optimal design of a multicarrier system with soft impulse shaping including equalization in time or frequency direction," in *Proc. IEEE Global Telecommunications Conf.*, vol. 1, pp. 310–314, 1997.

[12] S. Ohno, "Maximum likelihood inter-carrier interference suppression for wireless OFDM with null subcarriers," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, vol. 3, pp. 849–852, 2005.

[13] S. J. Simmons, "Breadth-first trellis decoding with adaptive effort," *IEEE Trans. Commun.*, vol. 38, pp. 3–12, Jan. 1990.

[14] M. O. Damen, H. El Gamal, and G. Caire, "MMSE-GDFE lattice decoding for under-determined linear channels," in *Proc. Conf. Inform. Science and Systems*, 2004.

[15] E. Viterbo and J. Boutros, "A universal lattice code decoder for fading channels," *IEEE Trans. Inform. Theory*, vol. 45, pp. 1639–1642, July 1999.

[16] E. Agrell, T. Eriksson, A. Vardy, and K. Zeger, "Closest point search in lattices," *IEEE Trans. Inform. Theory*, vol. 48, pp. 2201–2214, Aug. 2002.

[17] M. O. Damen, H. El Gamal, and G. Caire, "On maximum-likelihood detection and the search for the closest lattice point," *IEEE Trans. Inform. Theory*, vol. 49, pp. 2389–2402, Oct. 2003.

[18] J. B. Anderson and S. Mohan, "Sequential decoding algorithms: A survey and cost analysis," *IEEE Trans. Commun.*, vol. 32, pp. 169–172, 1984.

[19] A. Murugan, H. El Gamal, M. O. Damen, and G. Caire, "A unified framework for tree search decoding: Rediscovering the sequential decoder," *IEEE Trans. Inform. Theory*, to appear.

[20] U. Fincke and M. Pohst, "Improved methods for calculating vectors of short length in a lattice," *Math. of Comput.*, vol. 44, pp. 463–471, Apr. 1985.

[21] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD: John Hopkins University Press, 3rd ed., 1996.

[22] S.-J. Hwang and P. Schniter, "On the optimality of MMSE-GDFE pre-processed sphere decoding," in *Proc. Allerton Conf. Commun., Control, and Computing*, Oct. 2005.