Aalborg Universitet

DOWNSAMPLING OF DFT PRECODED SIGNALS FOR THE AWGN CHANNEL

Tobias Lindstrøm Jensen, Karsten Fyhn, Thomas Arildsen, Torben Larsen

Aalborg University, Faculty of Engineering and Science
Department of Electronic Systems
DK-9220 Aalborg, Denmark

ABSTRACT

In this paper we propose and analyze a method for downsampling discrete Fourier transform (DFT) precoded signals. Since the symbols (in frequency) are in the constellation set, which is a subset of the entire complex plane, it is possible to detect N symbols from the DFT precoded signal when transmitting $M < N$ symbols, where M is not too small. We build our analysis on so-called simple vectors, and show that it is possible to detect in the noise-less case with high probability down to approximately $M \geq N/4$ for BPSK and $M \geq N/2$ for QPSK. We develop extensions from the noise-less to the noisy case, and propose two different detectors for the AWGN channel. Simulations show that using the two proposed detectors in the AWGN channel, we observe empirically a phase transition at $M \approx N/2$ for QPSK. Further, it is shown how downsampled QPSK signals can achieve the same BER and data rate as 8PSK at a lower signal-to-noise-ratio per information bit.

Index Terms— Communication, detection, signal processing, compressed sensing, convex relaxation, semidefinite relaxation

1. INTRODUCTION

Digital communication [1] is a valuable technology and is used extensively. An important technique in digital communication is precoding. We could e.g. precode a digital signal with N symbols using the inverse discrete Fourier transform (DFT). This corresponds to applying a constellation symbol to each sub-carrier, perform the inverse DFT, and then transmit N time symbols. The receiver would then estimate the constellation symbol by performing an N point DFT of the input followed by detection. But is it necessary to transmit N time symbols to detect N constellation symbols?

In general, we need N linear measurements to be able to perfectly reconstruct a general length N signal. The advent of compressed sensing [2, 3] however showed that if a signal obeys certain structure, such as a sparse description, it is not

necessary to take N, but $M < N$ linear measurements and still be able to reconstruct with high probability. Recently, it has been proposed to exploit other structure than sparsity. If, *e.g.*, we only need to categorize a received signal to a small number of different transmitted signals, it was shown how to do such using $M < N$ linear measurements in [4]. Such a technique may be used in IEEE 802.15.4, where it was shown how to detect a transmitted signal using $M = N/2$ samples [5]. Being able to take fewer than N measurements introduces flexibility to the system and possibly enables increasing data rate and/or decreasing transmission energy and/or reducing ADC and DAC sampling rate at the receiver/transmitter.

In this paper we show an approach for downsampling DFT precoded signals. In [4] it is required to investigate all possible signals that could be transmitted, but this is not tractable for a DFT precoded system with even a moderate size of sub-carriers N, since the total number of possible signals grows exponentially in N. Instead, it is in [6, 7] suggested to solve a convex feasibility problem with N variables in the noise-less case. We generalize the noise-less detector to the noisy case and propose two different detectors: one based on convex relaxation (CR) and one on semidefinite relaxation (SDR).

The outline of the paper is as follows: first we present the considered system model and notation in Sec. 2. In Sec. 3 we show how the downsampled DFT precoded system can be described as a special case of [6, 7] and develop two detectors for the case of noise. In Sec. 4 we show simulations of the system and in Sec. 7 we discuss the results.

2. SYSTEM MODEL

Fig. 1 shows the considered downsampled DFT precoded system. First an *i.i.d.* bit stream $d[k] \in \{0,1\}$ with $P(d[k] = 1) = P(d[k] = 0) = \frac{1}{2}$ is passed trough a mapping from information bits to constellations symbols $\tilde{x}[n] \in \mathcal{S}^N \subset \mathbb{C}^N$ in the complex plane where $|\mathcal{S}| = L$ is the number of possible constellation points for each symbol. The constellation symbols have $E[|\tilde{x}[n]|^2] = 1$ where $E[\cdot]$ is the expectation operator. Each block of N symbols is inverse DFT'ed using the normalized inverse DFT matrix $D \in \mathbb{C}^{N \times N}$ with $\|Dz\|_2 = \|z\|_2, \forall z \in \mathbb{C}^N$. The information of each sym-
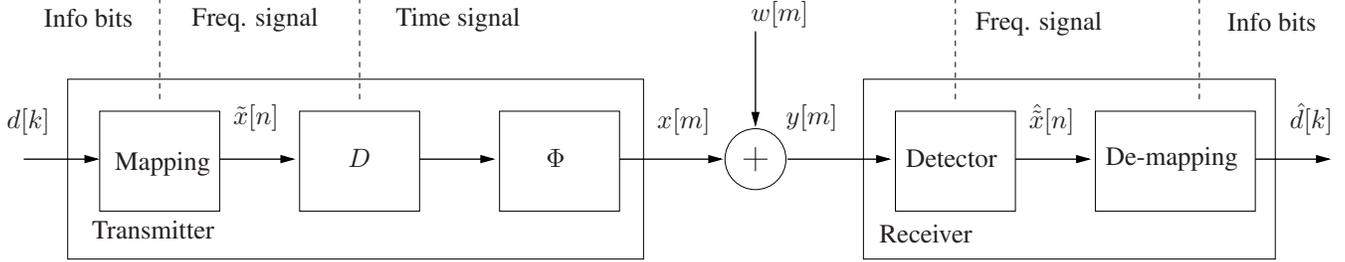
Fig. 1. Transmitter, channel and receiver structure for the considered downsampled DFT precoded system.

bol is then spread over the whole time domain since D is a dense matrix. The time symbols are then *"randomly punctured"* by block processing of N to M time symbols with Φ. This gives the *downsampling factor*

$$\kappa = \frac{M}{N}. \tag{1}$$

The Φ block can also be seen as a matrix-vector multiplication with a random $\Phi \in \{0,1\}^{M \times N}$, $M \leq N$, a matrix where each row of Φ contains exactly one element with a one. One possible realization is

$$\Phi = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \ddots & & \vdots & \vdots \\ 0 & \cdots & & & 1 & 0 \end{bmatrix}. \tag{2}$$

The selection matrix Φ can be generated by a random uniform selection of M unique rows of the $N \times N$ identity matrix I. Note that $\Phi = I$ corresponds exactly to standard DFT precoded transmission and since $\Phi D = D$ is an orthogonal transform, precoding has no effect compared to non-precoding. Due to Φ and the normalized D we have $E[\,|x[m]|^2\,] = 1$. We use an AWGN channel with $w[n] \in \mathbb{C}$ a circular complex zero-mean white Gaussian noise variable $w[n] \sim \mathcal{CN}(0, \sigma^2)$. The received signal $y[m]$ is processed by the detector to obtain an estimate $\hat{\tilde{x}}[n] \approx \tilde{x}[n]$. These estimates are further processed using demapping to obtain an estimation of the transmitted bit stream. Note the different time indexes k, m, n which might run at different frequencies, *e.g.*, with QPSK modulation, k runs "two times faster than" n.

The signal-to-noise-ratio (SNR) is given by

$$\frac{E_{\mathrm{s}}}{N_0} = \frac{E[\,|x[m]|^2\,]}{E[\,|w[m]|^2\,]} = \frac{1}{\sigma^2}. \tag{3}$$

The SNR per information bit is given by

$$\frac{E_{\mathrm{b}}}{N_0} = \frac{E_{\mathrm{s}}}{N_0} \frac{1}{R} = \frac{\kappa}{\sigma^2 \log_2(L)}, \quad R = \frac{\log_2(L)}{\kappa}. \tag{4}$$

Note that the rate R both depends on the constellation size L as well as the downsampling factor, such that the energy

per bit is reduced by κ. With $\Phi = I$, we have the *maximum likelihood* (ML) detector [1]

$$\hat{\tilde{x}}_{\mathrm{ML}} = \underset{\tilde{x} \in \mathcal{S}^N}{\operatorname{argmin}} \|D\tilde{x} - y\|_2 = \underset{\tilde{x} \in \mathcal{S}^N}{\operatorname{argmin}} \|\tilde{x} - \tilde{y}\|_2 = h(\tilde{y}) \tag{5}$$

where $\tilde{y} = D^H y$, $\tilde{x} = [\tilde{x}[n'], \cdots \tilde{x}[n' + N - 1]]^T$, and the other symbols are formed in a similar way due to the block processing. The solution in (5) is simply the closest point in the constellation to $\tilde{y}$ measured in Euclidean distance.

3. DOWNSAMPLING AND DETECTION

The maximum likelihood detector for the downsampled signal is

$$\hat{\tilde{x}}_{\mathrm{ML}} = \underset{\tilde{x} \in \mathcal{S}^N}{\operatorname{argmin}} \|\Phi D\tilde{x} - y\|_2 \tag{6}$$

which is not easily calculated for the case $M < N$. A way is to investigate all L^N combinations and select the one with the smallest objective. This however requires L^N evaluations. In general, (6) is NP-hard [8]. We now relate the downsampling problem (6) to known results in the noise-less case and present two different ways of approximately solving (6) in the noisy case.

3.1. Downsampling in the noise-less case

In [6, 7], so-called s-*s*imple[1] vectors are described which are vectors $z \in [0, 1]^J$ with $J - s$ elements equal to exactly 0 or 1. Notice that this vector is structured but does not necessarily have a sparse representation. For such structured descriptions, it is not necessary to take J linear measurements of z to exactly reconstruct z. In-fact, it is possible to downsample using

$$b = Az \tag{7}$$

---

[1][6, 7] uses the term k-simple, but we use s-simple due to notations

where $A \in \mathbb{R}^{H \times J}$, $H \leq J$ are new variables which will be related to M, N later. It is then possible to reconstruct using the convex feasibility problem[2]

$$\hat{z} = \underset{Az=b, 0 \leq z_i \leq 1, i=1, \cdots, J}{\text{argmin}} \ 0 \,. \tag{8}$$

Then $\hat{z} = z$ with high probability for certain settings of the triplet (s, H, J) and A [6, 7]. The so-called *phase-transition* for this problem occurs at $\frac{s}{H} = \max(0, 2 - \frac{J}{H})$ [6]. The phase-transition divides the plot of $\frac{s}{J}$ and $\frac{H}{J}$ in two distinct regions. A region where it is expected that $\hat{z} = z$ and a region where it is expected that $\hat{z} \neq z$. Fig. 2 shows this for a 0-simple vector, where we plot the probability of correct recovery [6, 7]

$$P(z = \hat{z}) = 2^{-J+1} \sum_{l=0}^{H-1} \binom{J-1}{l} \tag{9}$$

for two different values of J. We see that as we approach $\frac{H}{J} = \frac{1}{2}$, the reconstruction will start to fail. We also observe that for larger J, the transition region becomes more narrow.
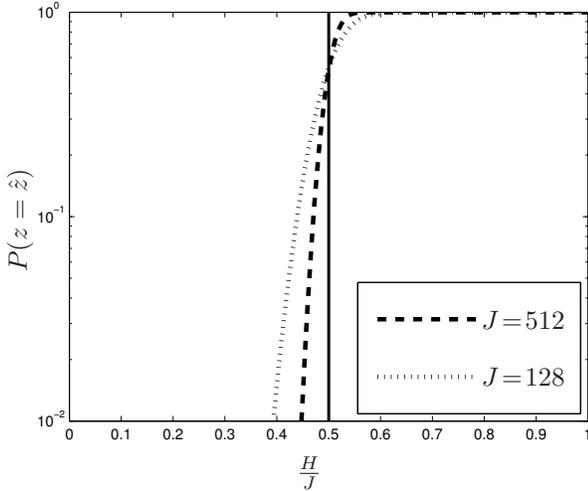

Fig. 2. Probability of correct recovery for $J = 128$ and $J = 512$. Full line marks $\frac{H}{J} = \frac{1}{2}$.

This use of so-called s-simple vectors is related to compressed sensing in the sense that we exploit the structure of the problem at hand to reduce the number of linear measurements. In the setting here we do however not exploit sparse representations.

3.2. Relation to downsampled DFT precoded signals

We show how the DFT precoded system in Fig. 1 can be described in terms of downsampling of 0-simple vectors for

---

[2] A convex feasibility problem is just a convex optimization problem with objective $f(x) = 0$.

the noise-less case $w[m] = 0$ ($\sigma^2 = 0$).

We observe

$$y = \Phi D \tilde{x}, \quad y \in \mathbb{C}^M \tag{10}$$

which can be written as

$$\breve{y} = \begin{bmatrix} \Re(y) \\ \Im(y) \end{bmatrix} = \begin{bmatrix} \Re(\Phi D) & -\Im(\Phi D) \\ \Im(\Phi D) & \Re(\Phi D) \end{bmatrix} \begin{bmatrix} \Re(\tilde{x}) \\ \Im(\tilde{x}) \end{bmatrix} = \breve{T}\breve{x} \tag{11}$$

where $\Re$ and $\Im$ are the real and imaginary part and with an implicit definition of $\breve{T}$ and $\breve{x}$. For BPSK, we have $\Re(\tilde{x}) \in \{+1, -1\}$, $\Im(\tilde{x}) = 0$, and the above system reduces to

$$\breve{y} = \begin{bmatrix} \Re(y) \\ \Im(y) \end{bmatrix} = \begin{bmatrix} \Re(\Phi D) \\ \Im(\Phi D) \end{bmatrix} \begin{bmatrix} \Re(\tilde{x}) \end{bmatrix} . \tag{12}$$

Using the transformation $\Re(\tilde{x}) = 2z - \mathbb{1}$, $z \in \{0, 1\}^N$, $\mathbb{1} = [1, \cdots, 1]^T$ we have

$$\breve{y} + \begin{bmatrix} \Re(\Phi D) \\ \Im(\Phi D) \end{bmatrix} \mathbb{1} = 2 \begin{bmatrix} \Re(\Phi D) \\ \Im(\Phi D) \end{bmatrix} z \tag{13}$$

which corresponds to the linear system $b = Az$ with $b \in \mathbb{R}^{2M}$ and $A \in \mathbb{R}^{2M \times N}$. Since $s = 0$, we have phase transition for this problem at $0 = \max(0, 2 - \frac{J}{H}) \Rightarrow \frac{H}{J} = \frac{1}{2}$,

$$\frac{H}{J} = \frac{2M}{N} = \frac{1}{2} \Leftrightarrow \kappa = \frac{M}{N} = \frac{1}{4} \,. \tag{14}$$

That is, in the noise-less case, the downsampling procedure works down to approximately $\kappa = \frac{1}{4}$, approximately one quarter of the time symbols are actually needed to reconstruct. However, the transition region has a certain width, see Fig. 2, so at $\kappa = \frac{1}{4}$ we would expect that errors still occur, even in the noise-less case. Similarly, it can be shown that for QPSK the phase transition occurs at $\kappa = \frac{1}{2}$.

3.3. A convex relaxation for the noisy case

The material in [6, 7] does not consider the noisy case. Since the noisy case is the most important case in communication systems, we quantify the impact of noise by simulation in Sec. 4. We also need to develop another detector, since (8) requires consistency (equality) between the transmitted and received signal. Consider the following ML reconstruction of linear measurements $b = Az + e$ with e *i.i.d.* Gaussian noise

$$\hat{z}_{\text{ML}} = \underset{z_i \in \{0,1\}, i=1, \cdots, N}{\text{argmin}} \|Az - b\|_2 \,. \tag{15}$$

Then the closest CR of the above ML detection problem is

$$\hat{z}_{\text{ML-CR}} = \underset{0 \leq z_i \leq 1, i=1, \cdots, N}{\text{argmin}} \|Az - b\|_2 \,. \tag{16}$$

This formulation could also be naturally extracted from formulation (8). In terms of the original variable $\tilde{x}$, the above problem is equivalent to solving

$$\hat{\tilde{x}}_{\text{ML-CR}} = \underset{|\Re(\tilde{x}_i)| \leq 1,\, \Im(\tilde{x}_i)=0,\, i=1,\cdots,N}{\text{argmin}} \|\Phi D \tilde{x} - y\|_2 \quad (17)$$

for BPSK and

$$\hat{\tilde{x}}_{\text{ML-CR}} = \underset{|\Re(\tilde{x}_i)|+|\Im(\tilde{x}_i)| \leq 1,\, i=1,\cdots,N}{\text{argmin}} \|\Phi D \tilde{x} - y\|_2 \quad (18)$$

for QPSK, where the constellation and the CR are shown in Fig. 3.
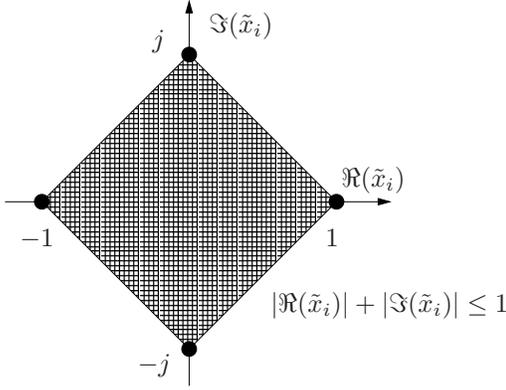

Fig. 3. Example of the feasible set in dimension i in the case of QPSK.

3.4. A semidefinite relaxation for the noisy case

Using (15), we formed the closest CR (16), but it is also possible to form a SDR of this problem. We can represent the QPSK constellation $\tilde{x} \in \mathcal{S}^N = \{+1, -1, +j, -j\}^N$ as [9]

$$\Re(\tilde{x}_i)^2 + \Im(\tilde{x}_i)^2 = 1, \Re(\tilde{x}_i)\Im(\tilde{x}_i) = 0, \forall i = 1, \cdots, N. \tag{19}$$

We can then formalize the problem (6) in real variables using (11), as

$$
\begin{aligned}
\text{minimize} \quad & \|\breve{T}\breve{x} - \breve{y}\|_2^2 \\
\text{subject to} \quad & \breve{x}_i^2 + \breve{x}_{i+N}^2 = 1, \quad \forall i = 1, \cdots N \\
& \breve{x}_i \breve{x}_{i+N} = 0, \quad \forall i = 1, \cdots N
\end{aligned} \tag{20}
$$

or equivalently in a homogeneous form

$$
\begin{aligned}
\text{minimize} \quad & [\breve{x}^T\ t] \begin{bmatrix} \breve{T}^T\breve{T} & -\breve{T}^T\breve{y} \\ -\breve{y}^T\breve{T} & \breve{y}^T\breve{y} \end{bmatrix} \begin{bmatrix} \breve{x} \\ t \end{bmatrix} \\
\text{subject to} \quad & \breve{x}_i^2 + \breve{x}_{i+N}^2 = 1, \quad \forall i = 1, \cdots N \\
& \breve{x}_i \breve{x}_{i+N} = 0, \quad \forall i = 1, \cdots N \\
& t^2 = 1.
\end{aligned} \tag{21}
$$

Since the feasible set only contains $t^2 = 1$, the term $t^2 \breve{y}^T \breve{y}$ is a constant, and we have

$$
\begin{aligned}
\text{minimize} \quad & [\breve{x}^T\ t]\, C \begin{bmatrix} \breve{x} \\ t \end{bmatrix} \\
\text{subject to} \quad & \breve{x}_i^2 + \breve{x}_{i+N}^2 = 1, \quad \forall i = 1, \cdots N \\
& \breve{x}_i \breve{x}_{i+N} = 0, \quad \forall i = 1, \cdots N \\
& t^2 = 1
\end{aligned} \tag{22}
$$

with

$$C = \begin{bmatrix} \breve{T}^T\breve{T} & -\breve{T}^T\breve{y} \\ -\breve{y}^T\breve{T} & 0 \end{bmatrix}. \tag{23}$$

We can then form a SDR of the above quadratically constrained quadratic program (QCQP) as [10]

$$
\begin{aligned}
\text{minimize} \quad & \mathbf{tr}\,(CX) \\
\text{subject to} \quad & X_{i,i} + X_{i+N,i+N} = 1,\ \forall i = 1, \cdots, N \\
& X_{i,i+N} = 0,\ \forall i = 1, \cdots, N \\
& X_{2N+1,2N+1} = 1 \\
& X \succeq 0
\end{aligned} \tag{24}
$$

with solution $X^\star$ and $\mathbf{tr}$ the trace function. A popular approach to improve SDR is to use randomization see, *e.g.*, [10] for an overview. With randomization we approximate the solution of the QCQP (22) by generating $T = 500$ realizations with $\zeta^{(r)} \sim \mathcal{N}(0, X^\star)$ and obtain feasible points (constellation points) as

$$\bar{\tilde{x}}^{(r)} = h\Big(\Re\big([\zeta_1^{(r)}, \cdots, \zeta_N^{(r)}]^T\big) + j\Im\big([\zeta_{N+1}^{(r)}, \cdots, \zeta_{2N}^{(r)}]^T\big)\Big) \tag{25}$$

where h is given in (5), and then select

$$r^\star = \underset{r=1,\cdots,T}{\text{argmin}} \|\Phi D \bar{\tilde{x}}^{(r)} - y\|_2 \tag{26}$$

and $\bar{\tilde{x}}^{(r^\star)} = \hat{\tilde{x}}_{\text{ML-SDR}}$ as our estimate.

4. SIMULATIONS

We divide the simulations into two cases. The first case is downsampling at transmitter. The second case is undersampling at receiver. For the simulations we use solver [11] with the interface [12] to formulate the convex optimization problems. The AWGN channel is formed with a given SNR (3) and we obtain the SNR per information bit via (4).

5. DOWNSAMPLING AT TRANSMITTER

For each block transmission, we generate a new random Φ by random selection of M rows from the identity matrix. Since it is required that the receiver also knows Φ, practical implementation of this system requires that transmitter and receiver

seed a random generator with the same number, similar to code division multiple access (CDMA).

Fig. 4 shows the (uncoded) BER versus $\frac{E_b}{N_0}$ of the QPSK based CR detector (18). We observe that it is indeed possible to detect the symbols for downsampled DFT precoded signals in the AWGN channel. We also observe that as we decrease κ (more downsampling), the BER performance decreases, and at $\kappa = 0.55$, it is almost impossible to detect correctly in the considered SNR range.
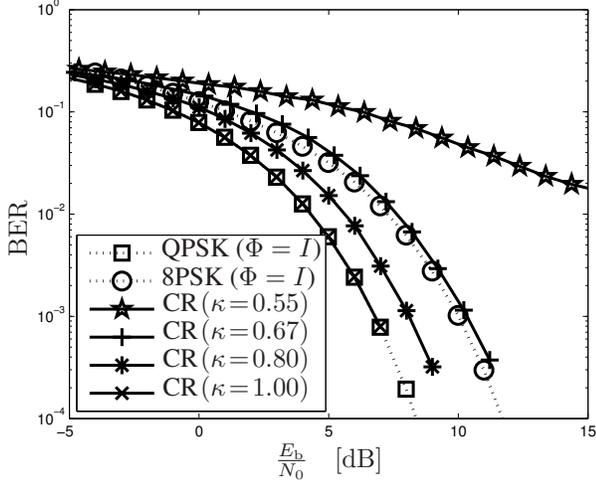

Fig. 4. The (uncoded) BER for: *i)* QPSK, 8PSK using $\Phi = I$ and optimal detection (5), *ii)* QPSK based CR detectors with (18) and downsampling factor κ. $N = 128$ sub-carriers.
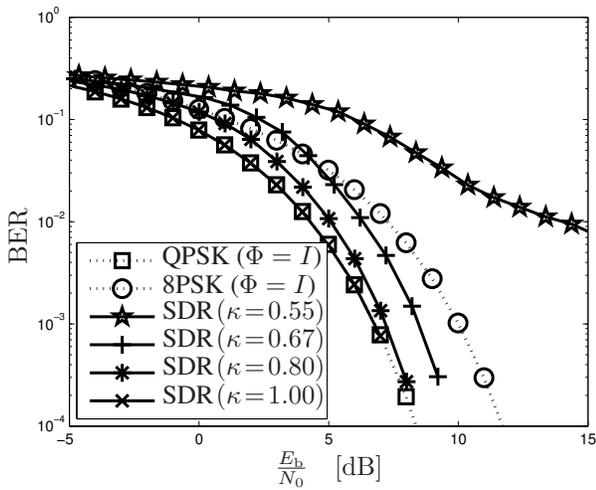

Fig. 5. The (uncoded) BER for: *i)* QPSK, 8PSK using $\Phi = I$ and optimal detection (5), *ii)* QPSK based SDR detectors with (24)-(26) and downsampling factor κ. $N = 128$ sub-carriers.

In Fig. 5 we show the (uncoded) BER versus $\frac{E_b}{N_0}$ of the QPSK based SDR detector (24). We generally observe the same behaviour as for Fig. 4, that noise as well as downsampling introduce errors in the detector. To compare with a standard modulation scheme, we also show Gray encoded 8PSK with 3 bits per symbol calculated via [13]. The modulation scheme 8PSK has the same number of bits per symbol as QPSK modulation with downsampling $\kappa = \frac{2}{3}$, *i.e.*, $\frac{\log_2(L)}{\kappa} = 3$ bits per. symbol. In Fig. 4 we observe that QPSK at $\kappa = \frac{2}{3}$ and 8PSK has approximately the same BER at the same SNR, with 8PSK slightly better. In Fig. 5 we observe that QPSK at $\kappa = \frac{2}{3}$ has a lower BER than 8PSK at the same SNR. At $\text{BER} = 10^{-3}$ there is an improvement of ≈ 1.8 dB.

6. UNDERSAMPLING AT RECEIVER

We note an important simple extension of the proposed signal model Fig. 1. Consider the case of undersampling where the transmitter still sends a full signal and the receiver just takes fewer samples than normally required. The undersampling system can be obtained by moving the $D\Phi$ to the other side of the AWGN channel in Fig. 1. Specifically, let

$$y = \Phi D(x + w) = \Phi Dx + \tilde{w}, \quad \tilde{w} = \Phi Dw. \quad (27)$$

Then $E[\tilde{w}\tilde{w}^H] = E[\Phi Dww^H D^H \Phi^H] = \sigma^2 I$, *i.e.*, the noise is still *i.i.d.* Gaussian. For this system the SNR is the same

$$\frac{E_s}{N_0} = \frac{E[\,|x[m]|^2\,]}{E[\,|w[m]|^2\,]} = \frac{1}{\sigma^2} \quad (28)$$

but the SNR per information bit is instead

$$\frac{E_b}{N_0} = \frac{E_s}{N_0} \frac{1}{\tilde{R}} = \frac{1}{\sigma^2 \log_2(L)} \quad (29)$$

since we still transmit all time symbols (and only undersample at the receiver).

For each block transmission, we generate a new random Φ by random selection of M rows from the identity matrix. With undersampling at the receiver it is not required that the transmitter also knows Φ. Due to the system design, undersampling at the receiver then comes with a SNR penalty but it becomes possible to reduce the ADC sampling rate at the receiver by a factor of κ.

We see the result of undersampling at the receiver in Figs. 6 and 7. We observe that as we increase undersampling, the SNR penalty for detection with same BER increases when compared to standard QPSK with $\Phi = I$. Again we observe that the SDR detector is performing better than the CR detector, *i.e.*, the curves in Fig. 7 lie closer to the QPSK curve than in Fig. 6. Using the SDR receiver at $\text{BER} \approx 10^{-3}$ we can undersample with a factor of $\kappa = 0.8$ and $\kappa = 0.67$ with an SNR penalty of approximately 1.5 dB and 3.5 dB, respectively.
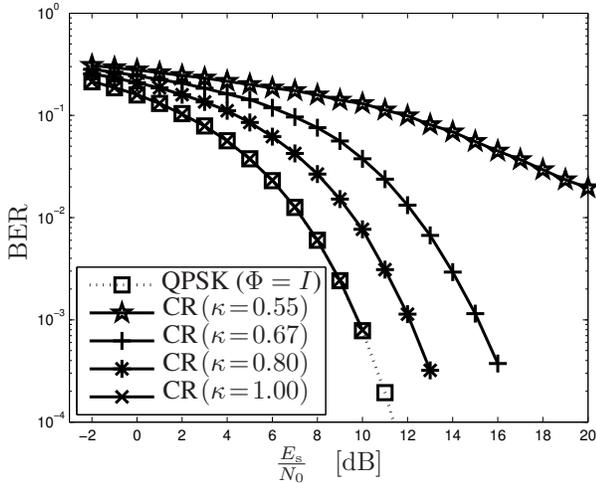
Fig. 6. The (uncoded) BER for: *i)* QPSK using $\Phi = I$ and optimal detection (5), *ii)* QPSK based CR detectors with (18) and undersampling factor κ. $N = 128$ sub-carriers.
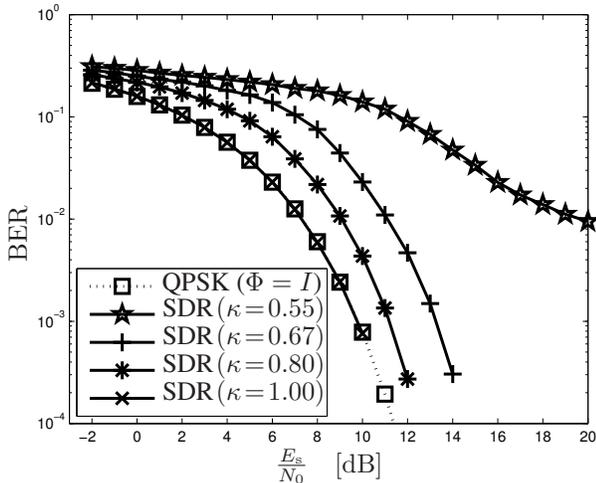

Fig. 7. The (uncoded) BER for: *i)* QPSK using $\Phi = I$ and optimal detection (5), *ii)* QPSK based SDR detectors with (24)-(26) and undersampling factor κ. $N = 128$ sub-carriers.

7. CONCLUSION AND DISCUSSION

We have presented an approach for downsampling of DFT precoded signals and shown that it is possible to obtain the same BER and data rate for a downsampled QPSK mapped signal at $\approx 1.8\,\text{dB}$ lower SNR compared to a fully sampled 8PSK mapped signal. The penalty of this downsampling scheme is that the detector is more computationally demanding, *i.e.*, it is required to solve a semidefinite problem and generate many realizations of a certain distribution. How computational demanding depend on the problem structure and algorithm but it can be shown that certain semidefinite problems can be solved fast, see overview in [10] or [14, 15].

We also show a setup where it is possible to undersample at the receiver. Undersampling at the receiver gives a SNR penalty but makes it possible to run the receiver ADC at a lower rate.

8. REFERENCES

[1] J. Proakis, *Digital Communications*, McGraw-Hill Higher Education, 4 edition, 2000.

[2] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[3] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[4] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk, "Signal processing with compressive measurements," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 445–460, 2010.

[5] K. Fyhn, T. Arildsen, T. Larsen, and S. H. Jensen, "Demodulating subsampled direct sequence spread spectrum signals using compressive signal processing," in *Proc. IEEE European Signal Process. Conf. (EUSIPCO)*, Bucharest, Romania, Aug. 2012, pp. 2556–2560.

[6] D. L. Donoho and J. Tanner, "Precise undersampling theorems," *Proc. of the IEEE*, vol. 98, no. 6, pp. 913–924, 2010.

[7] D. L. Donoho and J. Tanner, "Counting the faces of randomly-projected hypercubes and orthants, with application," *Discrete Comput. Geom.*, vol. 43, no. 3, pp. 522–541, 2010.

[8] S. Verdú, "Computational complexity of optimum multiuser detection," *Algorithmica*, vol. 4, no. 1-4, pp. 303–312, 1989.

[9] J. Dahl, B. H. Fleury, and L. Vandenberghe, "Approximate maximum-likelihood estimation using semidefinite programming," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Hong Kong, Apr. 2003, pp. VI 721–724.

[10] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Process. Mag., Special issue: Advances in Convex Optimization*, vol. 27, no. 3, pp. 20 – 34, May 2010.

[11] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optim. Methods Softw.*, vol. 11-12, pp. 625–653, 1999.

[12] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," http://cvxr.com/cvx/, Apr. 2011.

[13] P. J. Lee, "Computation of the bit error rate of coherent M-ary PSK with Gray code bit mapping," *IEEE Trans. Comm.*, vol. 34, no. 5, pp. 488 – 491, May 1986.

[14] C. Helmberg, F. Rendl, R. J. Vanderbei, and H. Wolkowicz, "An interior-point method for semidefinite programming," *SIAM J. Opt.*, vol. 6, pp. 342–361, 1996.

[15] W.-K. Ma, C.-C. Su, J. Jaldén, and C.-Y. Chi, "Some results on 16-QAM MIMO detection using semidefinite relaxation," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. (ICASSP)*, Las Vegas, Nevada, USA, Apr. 2008, pp. 2673–2676.