

Low-complexity DFT-pair Carrier Acquisition

S. Jayasimha, P. Jyothendar and B. Satish Bhargav

Signion Systems
Hyderabad, India

Abstract— When the maximum frequency offset to be acquired is a small fraction of the symbol-rate, a DFT-pair based carrier acquisition method (a frequency-domain analog of the early-late gate synchronizer) provides low-complexity frequency-offset acquisition using a modest number of symbols. Several new modulation and coding standards (e.g., DVB-S2, allowing receivers to operate at lower E_b/N_0 's) require that this DFT-based acquisition scheme be modified in order to allow acquisition at a range of signal-to-noise ratios. As the signal-to-noise ratio is unknown, an intuitively appealing method, explored in this paper, takes the smaller of the two DFT bins as the bin with more noise.

Keywords— frequency acquisition; DFT; DVB-S2; carrier-to-noise;

I. INTRODUCTION

Frequency acquisition may be either pilot-based [1, 2] or unaided [3-5]. In this paper we consider unaided (or “blind”) frequency acquisition.

For continuous mode MPSK, carrier acquisition described in [4, 5] removes modulation (exponentiating the band-pass filtered signal by M , implemented as $\log_2 M$ stages of squaring), while for burst-mode, an un-modulated preamble field is exploited for faster acquisition. M is 16 for 16-QAM, since its phases are approximately a subset of 12 of 16 equally-spaced phases. This approximation, in addition to amplitude modulation, results in a less localized spectral peak after exponentiation.

Thereafter, a traditional method (as in [5]) then obtains a sampled spectrum (computed by an FFT) and identifies the sample (or FFT bin) that consistently peaks to obtain coarse frequency acquisition. When the maximum frequency offset is a small fraction of the symbol rate (i.e., for a high symbol rate), a DFT-pair discriminator allows the frequency offset to be estimated as described in [4, 5]. We describe a refinement to that method that allows a single discriminator to perform well at all signal-to-noise ratios. The improvement described applies to both continuous and burst modes.

A Chebyshev window [6] with parameter a , $1.5 \leq a \leq 3$, suppresses out-of-band noise and interference by 30-60dB, while incurring processing losses of 3-3.23dB. DFT-based acquisition is robust in many noise/ interference scenarios. In order to acquire a digital carrier's frequency offset, a receiver

applies a Chebyshev window and calculates DFT envelopes, y_1 and y_2 , at two equidistant bins (this, as stated in pp. 239 of [4] is the frequency domain analog of the early-late gate synchronizer described pp.226-228 of [4]) from the nominal carrier. Averaging periodograms over K 50% overlapped segments improves SNR by almost K . The frequency offset is obtained as a function of DFT ratio (as described by Figure 6 in section III for various modulation schemes). This estimator enables an acquisition procedure which decreases the offset progressively and rapidly; in contrast, a PLL's phase estimates become random once the offset exceeds its limited tracking range.

When the signal-to-noise ratio is poor, the DFT ratio approaches unity (particularly for small frequency offsets that are greater than the coherent demodulator's PLL tracking range). In such a case, there is no frequency correction that is made and acquisition failure results. We describe a method (that automatically adapts to the prevailing signal-to-noise ratio) to modify the DFT ratio in Section II.

II. IMPROVED ACQUISITION

A better acquisition procedure (compared to that reported in [5]) is to find the minimum of the two DFT bin energies and subtract a fraction α of it from both DFT bins before computing their ratio.

The qualitative rationale is that the signal-to-noise ratio in the smaller DFT bin depends on a) the carrier-to noise ratio (C/N) and b) the frequency offset. With moderate to large frequency offsets, the smaller DFT bin contains more noise (at lower C/N 's), thus providing a good estimate of N , while making an incorrect N estimate at small frequency offsets does change significantly the computed DFT ratio.

Quantitatively, the smaller bin DFT energy is $D(f_c + f_{offset}) \cdot C + N$ and the larger bin is $D(f_c - f_{offset}) \cdot C + N$, where D is the Dolph-Chebyshev window energy response, f_c is the nominal carrier frequency (or center frequency, midway between the two DFT bins). Subtracting α (to be selected in $N/C \leq \alpha < 1$, the upper limit being selected to avoid an indeterminate ratio) times the smaller of the DFT energies from both and computing the ratio yields:

$$\frac{(1-\alpha) \cdot D(f_c + f_{offset}) \cdot C + N}{\alpha \cdot [D(f_c + f_{offset}) \cdot C + N]} \quad (1)$$

At low SNRs, the raised noise floor (due to exponentiation by M) drives the DFT ratio close to unity, causing the DFT ratio to lose discrimination of frequency offset, i.e., frequency offsets are estimated as near zero offsets causing incorrect acquisition. The subtraction shown in (1) reduces the ratio to less than unity (i.e., subtracting noise), sufficient so that large frequency offsets are acquired by iteratively accumulating smaller offsets, until the DFT ratio is truly near unity indicating that the offset has been acquired. Hence, this method improves acquisition at the expense of increase in acquisition time. Though the choice of α is based on C/N , pulse shaping and the modulation-type of the signal, it is more-or-less unvarying over the operating E_s/N_0 range.

Frequently, with fixed-point arithmetic on typical DSP processors, one likes to work on DFT envelopes (rather than energies). The noise subtraction method remains the same (with a different scale factor on α as compared to energy-based computations) as in a finite-dimensional complex vector space, all norms are topologically equivalent [7]. A popular approximation [8] for the envelope is also equivalent to the L^2 norm: suppose df_{tl} and df_{th} are the smaller and larger DFT envelopes before alteration respectively and corresponding primed envelopes after modification, then $df'_{tl}=df_{tl}\cdot(1-\alpha^2)^{1/2}$ and $df'_{th}=(df_{th}^2-\alpha^2\cdot df_{tl}^2)^{1/2}$. We use the approximation that the square root of the sum of two squared quantities is approximately (with up to a maximum error of 6%) the larger quantity plus $(\sqrt{2}-1)$ times the smaller quantity. Since $(1-\alpha^2)^{1/2}\approx[1-(\sqrt{2}-1)\cdot\alpha]$, $df'_{tl}=df_{tl}\cdot(\sqrt{2}-1)\cdot\alpha\cdot df_{tl}$ and $df'_{th}=df_{th}-(\sqrt{2}-1)\cdot\alpha\cdot df_{tl}$. Thus, the scale factor for envelopes (as compared to the energies) is $(\sqrt{2}-1)\cdot\alpha/\alpha^2=(\sqrt{2}-1)/\alpha$.

Figure 1 shows a family of curves, parameterized by α , of the fraction of DFT ratio for each original DFT ratio. As can be seen, the DFT ratio does not change significantly for small offsets, but rapidly increases before saturating (particularly for higher α 's) for medium to large frequency offsets.




Figure 1. Fraction of DFT ratio computed after noise subtraction as a function of original DFT ratio

We can see that with large offsets α should be large (but less than unity, because the smaller bin energy is almost entirely noise), while with small offsets, only a fraction of the smaller bin energy should be subtracted (since it has both signal and noise components). This suggests a fixed non-negative, decreasing schedule, with unbounded sum, for α regardless of the unknown frequency offset [9].

In order to compare the varying schedule with two fixed schedules, Figure 2 shows typical convergence behavior (as a function of DFT-block count) at an E_b/N_0 where all three schedules converge (i.e., for 512ksymb/sec., 8-PSK at $E_b/N_0=9$ dB) to the actual frequency offset. The three schedules are: a) $\alpha=0$ (i.e., no subtraction), b) $\alpha=0.87$, (without any schedule) and c) α -scheduling of $\alpha=0.87$ for 30 DFT blocks and $\alpha=\epsilon$ (a small positive constant near 0) thereafter. While avoiding terminal oscillations of $\alpha=0.87$ (constant), α -scheduling has the fastest rise (to within fine frequency acquisition range indicated by the dashed line) and overall convergence times. As E_b/N_0 is reduced, first the $\alpha=0$ (fixed) schedule fails to converge and then the $\alpha=0.87$ (fixed) schedule fails to converge; thus, the varying schedule evaluated here gains both in terms of converging at lower E_b/N_0 's (this effect is quite modest in the scenarios tested) convergence and in rapidity of convergence (this effect being of substantial significance). Regardless of the actual frequency offset, a decreasing schedule is always better than a fixed schedule.




Figure 2. 8-PSK α -scheduling's faster rise and convergence time

Figure 3 shows typical convergence behavior (as a function of DFT-block count) at an E_b/N_0 where all three schedules converge (i.e., for 512ksymb/sec., 16-QAM at $E_b/N_0=15$ dB). The oscillatory convergence seen is due to spectral peak being data-dependent and less localized (as the exponentiation does not completely remove the modulation). However, we still see that a decreasing α -schedule obtains the most rapid convergence (as in 8-PSK).




Figure 3. 16-QAM α -scheduling's faster rise and convergence time

A further benefit of a decreasing schedule is in a no-noise scenario (such as in loop-back testing of a modem) where, ideally, α should equal 0. Using a decreasing schedule allows oscillating convergence (as shown in Figure 4, for an 8-PSK example, α -scheduling of $\alpha=0.87$ for 30 DFT blocks and $\alpha=\epsilon$ thereafter), provided the initial frequency offset is not greatly overestimated.




Figure 4. 8-PSK α -scheduling's convergence time with no noise

III. FREQUENCY ACQUISITION IN CONTINUOUS-MODE




Figure 5. Carrier acquisition section of a high symbol-rate demodulator

Figure 5 shows the method for a 1024k, 512k, 256k and 128k symbol rate demodulator carrier acquisition including the DFT noise subtraction technique described in Section II. The complex base-band modulated signal is band-pass filtered and decimated, such that the signal is at 6 samples/ symbol prior to exponentiation. In Figure 5, for 1024k, 512k, 256k and 128k symbol rates, n (the number of times the first decimation filter section is iterated)=1, 1, 2 and 3 respectively and P (the decimation factor)=1, 2, 2 and 2 respectively. For BPSK, QPSK, 8-PSK and 16-QAM, the exponentiation parameter $m=2, 4, 8$ and 8 respectively. The DFT ratio vs. frequency offset functional relationship depends, due to pulse-shaping- and ISI- induced spectral envelope fluctuations, upon modulation type and window length as shown by Figure 6 (where f_{bin} is the optimum DFT bins' position that minimizes the probability that frequency offset error, over a desired acquisition range and SNR, exceeds the PLL's frequency tracking range, normalized to the DFT window duration's reciprocal). Table I gives the Chebyshev window design parameters and $[DFTsize/(6 \times symbol\ rate)]$ is the window duration.

TABLE I. Chebyshev acquisition windows

Modulation type	Chebyshev window side-lobe attenuation (dB)	
	Coarse window (DFT size = 32)	Fine window (DFT size = 64)
BPSK	33	40
QPSK	38	45
8-PSK	27	40
16-QAM	40	52




Figure 6. Relationship between frequency offset and DFT ratio, solid and dashed curves correspond to coarse and fine window lengths respectively

Table II summarizes performance obtained with or without α -scheduling with α selected, and performance measured, on a DSP-based demodulator (Figure 7). DFT block-averaging of 512, 512, 1024 and 8192 are selected for BPSK, QPSK, 8-PSK and 16-QAM respectively. As already suggested in section II, α -scheduling reduces acquisition time (rather than improve tolerance to a higher level of noise).




Figure 7. DSP-based low power demodulator

TABLE II. Performance summary of carrier acquisition

Modulation type	Symbol rate (symbols/sec)	Acquisition range (Hz)	Threshold E_b/N_0 (dB)	Acquisition blocks
IESS-308 BPSK rate $\frac{1}{2}$ Viterbi RS(126,112)	1024k	$\pm 80k$	4.0	60
	512k	$\pm 40k$	4.0	
	256k	$\pm 20k$	4.0	
	128k	$\pm 10k$	4.0	
IESS-308 QPSK rate $\frac{1}{2}$ Viterbi RS(126,112)	1024k	$\pm 40k$	4.0	700
	512k	$\pm 20k$	4.0	
	256k	$\pm 10k$	4.0	
	128k	$\pm 5k$	4.0	
IESS-310 8-PSK TCM-2/3 RS(219,201)	1024k	$\pm 15k$	8.0	350 ($\alpha=0.87$), 250 ($\alpha=0.87$ for 30 blocks and ϵ thereafter)
	512k	$\pm 7.5k$	8.0	
	256k	$\pm 3.75k$	8.0	
	128k	$\pm 1.87k$	8.0	
DVBS 16-QAM TCM-3/4 RS(204,188)	1024k	$\pm 3k$	13.5	100 ($\alpha=0.8$), 75 ($\alpha=0.8$ for 12 blocks and ϵ thereafter)
	512k	$\pm 1.5k$	14	
	256k	± 750	14.5	
	128k	± 375	14.5	

IV. CONCLUSION

We have shown that a low-complexity two-bin DFT method (as opposed to computing many DFT's over a range) can be modified, in some cases, to provide acquisition at the thresholds that match (subsequent) coherent demodulator performance. This modification consists of subtracting a positive fraction of the minimum of the DFT-pair energy from both energies. While this fraction can be constant or varying, using a decreasing schedule for this fraction yields rapid convergence to the true frequency offset. The additional benefit of α -scheduling is that acquisition is unimpaired when there is very little noise (such as in loopback testing).

REFERENCES

- [1] Alan Barbieri and Giulio Colavolpe, "On Pilot-Symbol-Assisted Carrier Synchronization for DVB-S2 Systems," *IEEE Transactions on Broadcasting*, vol.53, no.3, pp.685-692, Sept. 2007.
- [2] Balasubramanian Ramakrishnan, "Maximum-likelihood based rapid carrier acquisition architectures," *IEEE Military Communications Conference*, pp.1-6, 16-19 Nov. 2008.
- [3] Xiao Yan, Qian Wang and Kaiyu Qin, "Frequency pre-estimation aided carrier recovery algorithm for high-speed M-PSK communication," *11th IEEE Singapore International Conference on Communication Systems*, pp.1339-1343, 19-21 Nov. 2008.
- [4] J.J. Stiffler, *Theory of Synchronous Communications*, Prentice Hall, Englewood Cliffs, New Jersey, 1971, pp. 261, problem 8.2.
- [5] S. Jayasimha, P. Jyothendar and S. Pavanalatha, "SDR Framework for burst/continuous MPSK/ 16-QAM modems," *Proceedings of SPCOM-2004*, ISBN 0-7803-8675-2, IEEE catalog no. 04EX926C.
- [6] F.J. Harris, "On the use of windows for harmonic analysis with the Discrete Fourier Transform," *Proceedings of the IEEE*, vol.66, no.1, pp.51-83, January 1978.
- [7] Nicolas Bourbaki, *Topological vector spaces*, Springer, ISBN 3-540-13627-4, 1987, chapters 1-5.
- [8] W.T. Adams and J. Brady, "Magnitude Approximations for Microprocessor Implementation," *IEEE Micro*, vol.3, no.5, Oct. 1983.
- [9] M. Metivier and P. Priouret, "Applications of a Kushner and Clark Lemma to General Classes of Stochastic Algorithms," *IEEE Transactions on Information Theory*, vol.IT-30, no.2, pp.140-151, March 1984.