AN EFFICIENT, PSP-BASED DIGITAL MODULATION CLASSIFIER

Prokopios Panagiotou
Helix Communication Technologies, Inc.
Seal Beach, CA

Andreas Polydoros
TrellisWare Technologies, Inc.
San Diego, CA

Chul Oh
US Army RDECOM, CERDEC
Fort Monmouth, NJ

ABSTRACT

This paper presents several methods that can be used for the design of an efficient, PSP-based digital modulation classifier. Such methods include: (a) an efficient Decision-Making Procedure that reaches a final decision in stages, (b) a reduced-complexity Feature-Generation Algorithm that generates the necessary discriminating metrics, (c) a method to calculate approximate thresholds that uses polynomial fitting techniques to establish a functional relationship between the threshold and the Signal-to-Noise Ratio. Selected performance simulation results presented herein demonstrate the validity of these methods for digital, memoryless modulation menus.

1. INTRODUCTION

Radios that can autonomously and efficiently characterize the operational environment are of increasing importance due to a general current trend for "cognitive" commercial tasks (e.g., operating in a short-term leased band) [1]. When done in a non-cooperative manner, many parameters and features must be identified quickly, efficiently and accurately. Of those, we focus here on the task of Modulation Classification (MC), which renders efficient decisions on the exact transmit modulation type out of a given menu of possibilities.

In previous work [2], the authors have addressed the problem of MC for digital, memoryless modulations in arbitrary propagation environments. It was shown therein that the so-called waveform-likelihood framework [3] leads to an assortment of Per-Survivor-Processing [4] (PSP)-based classifiers. The approach provides reliable MC performance in both Frequency-nonselective (FNS) and Frequency-selective (FS) channels, even for very small data records (e.g., 100 bauds). However, the brute-force exhaustive version is too complex for large modulation menus involving higher-order modulation types and large channel spreads. A proposed architecture for significantly reducing this complexity is the subject of the current paper.

In order to design an efficient, PSP-based digital MC engine, it is important to understand that all modulation classifiers reach a final decision as to which modulation type was transmitted by following the following steps:

a. generating features (such as likelihood-based metrics) associated with particular modulation types (or, more generally, classes of modulations);
b. forming decision statistics based on these features;
c. comparing each decision statistic to an appropriately selected threshold.

The number and order of comparisons, as well as the modulation types (or classes) involved in each comparison are determined by the Decision-Making Procedure (DMP). The DMP that was considered in [2] was using hypothesis-exhaustive, sequential, pair-wise comparisons to reach a final decision. Although this procedure is straightforward to implement, it is very inefficient due to its exhaustive-search aspect.

In Section 2, we propose a more efficient DMP that reaches a final decision in stages, in other words in a tree-like search procedure. Such a procedure can offer significant reduction in the overall computational complexity, because the number of required pair-wise comparisons is much lower than that of the exhaustive sequential procedure. The comparisons in the multi-stage DMP are not necessarily between individual modulation types but include (in the first stages) comparisons between classes of modulations. The idea, then, is to gradually narrow the class to which a hypothesized modulation belongs and then, finally, to home in to the specific elementary modulation type.

Further reduction in the computational complexity of the modulation classifier can be achieved by designing more efficient algorithms that generate the statistical features involved in the decisions, namely the Feature-Generation Algorithms (FGA). The Full-Record PSP algorithms that were presented in [2] utilized the whole data-record length in order to produce two quantities for each modulation type involved in a pair-wise comparison: (a) a metric, which is used for the formation of the decision statistic, and (b) a Signal-to-Noise Ratio (SNR) estimate, which is used to select the appropriate threshold. In the following, these two quantities for each hypothesized modulation type will be collectively referred as the feature vector.

Section 3 presents an efficient FGA, which is motivated by the rapid, blind Overall-Channel-Impulse-Response

(OCHIR) acquisition capabilities of PSP. This algorithm replaces each Full-Record PSP algorithm with a cascade consisting of a PSP algorithm, using few of the available observation samples for blind OCHIR acquisition, followed by an appropriate Data Detector which incorporates OCHIR tracking, using either the remaining (i.e., after acquisition) or all of the observation samples to generate the required feature vector. It reduces the overall computational complexity by reducing the processing volume, processing time and memory requirements of the modulation classifier.

Regardless of which particular FGA is employed, all DMPs require the determination of appropriate thresholds, which usually involves considerable off-line processing. In Section 4, we provide a discussion of threshold-related issues, such as setting, calculation, and storage requirements, and we present a method that can reduce the off-line processing and storage needs for PSP-based modulation classifiers.

The issue of reliable SNR estimation methods is discussed in Section 5. The paper concludes with Section 6, which presents the conclusions that we have drawn from this work.

2. DECISION-MAKING PROCEDURES

As its name suggests, a DMP determines the steps by which the modulation classifier reaches a final decision as to which modulation type was transmitted. It is a term denoting the particular search algorithm involved and does not address the specific signal processing of the received observation record. When the multi-hypothesis MC problem is treated as a sequence of pair-wise hypothesis-testing problems, a final decision is reached by a DMP that employs exhaustive, sequential, pair-wise comparisons. Such a procedure is described in the following.

A PSP-based FGA generates a feature vector consisting of a metric plus an SNR estimate for each modulation type in the menu. The difference of the metrics of the first two modulation types in the menu is taken in order to form the decision statistic, whereas the two SNR estimates are utilized to select the appropriate threshold from a file containing SNR values and corresponding thresholds. Comparing the decision statistic to the selected threshold yields the winning modulation type (either the first or the second one in the total menu). The metric and SNR estimate associated with the winning modulation type are then utilized to compare it to the third one. This procedure continues until the total menu is exhausted and a final decision is reached.

We have implemented this DMP for a modulation menu consisting of 2ASK, 4ASK, 4PSK, and 8PSK. The assumed operating environment is an FNS, Random-Phase channel, whereas the FGA is a Full-Record PSP algorithm utilizing 100 baud-spaced samples to generate the necessary discriminating metrics [2]. The operating SNR value is assumed to be known (an ideal case) and the corresponding Empirical Ensemble Threshold (EET) [2] has been used. The confusion matrix for this modulation menu and an SNR of 3dB is presented in Table 1. It is clear from this matrix that the MC performance when a one-dimensional constellation (2ASK or 4ASK) is transmitted is not significantly affected by the addition of a two-dimensional constellation (4PSK or 8PSK) in the modulation menu and vice-versa. A DMP employing exhaustive, sequential, pair-wise comparisons would not take advantage of this fact and, therefore, would be very inefficient because it performs a large number of essentially unnecessary pair-wise comparisons.

TABLE 1
CONFUSION MATRIX, SNR=3dB

	SELECTED MODULATION (%)			
TX MODULATION	*2ASK*	*4ASK*	*4PSK*	*8PSK*
2ASK	***95.6127***	4.3850	0.0003	0.0020
4ASK	4.3550	***95.6427***	0.0016	0.0007
4PSK	0.0007	0.0005	***73.0147***	26.9840
8PSK	0.0015	0.0006	20.7857	***79.2122***

For such a modulation menu consisting of a mixture of one- and two-dimensional constellations, an efficient DMP would utilize the fact that it is relatively easy to distinguish between one- and two-dimensional constellations and would then reach a final decision in two stages. The first stage determines the dimensionality of the constellation, whereas the second stage determines the actual transmitted modulation type. The main issue in the implementation of such a DMP is the choice of the modulation types that can be used to represent one- and two-dimensional constellations (in other words, we introduce the notion of a representative constellation per modulation class).

The chosen pair of representative modulations must provide both good performance in terms of a low Probability of False Dimensionality Classification, as well as minimum computational complexity. We have considered 4PSK as the representative of all two-dimensional constellations (since its alphabet size is the smallest possible among all quadrature-symmetric two-dimensional constellations) and we have examined the two cases where either 2ASK or 4ASK can be used as the representative of one-dimensional constellations.

Extensive simulation results [5] have shown that the pair 4ASK-4PSK provides better performance when an ASK modulation (one-dimensional constellation) with alphabet size greater than two is transmitted, whereas the pair 2ASK-4PSK provides better performance when a QAM modulation (two-dimensional constellation) is transmitted.

Both pairs provide equivalent performance when a PSK modulation is transmitted.

Once a decision about the dimensionality of the constellation is made, the second stage of this DMP determines the actual transmitted modulation type within the subset of one- or two-dimensional constellations by using sequential, pair-wise comparisons. In the case where the first stage determines that a two-dimensional constellation was transmitted, first distinguishing between PSK and QAM modulations and then determining the actual transmitted modulation can achieve further reduction in computational complexity.

Once again, the main implementation issue is the choice of the constellation types that will be used to represent PSK and QAM modulations. The natural candidates are 4PSK and 16QAM, respectively. Several simulations that we have conducted have revealed that when 4PSK is used as the representative of the PSK modulations, all PSK modulations with alphabet size greater than 4 are classified as QAM modulations; hence, it is not the right representative choice for PSKs. These simulations have also shown that 16PSK is the best representative of the PSK modulations. However, from an implementation standpoint, it is preferable to use 8PSK as the representative of the PSK modulations, because the PSP trellis size and processing requirements of 8PSK are significantly lower than that of 16PSK, whereas the discriminating performance suffers only a minor degradation. A block diagram of this multi-stage DMP for modulation menus consisting of ASK, PSK, and QAM modulation types is shown in Figure 1.

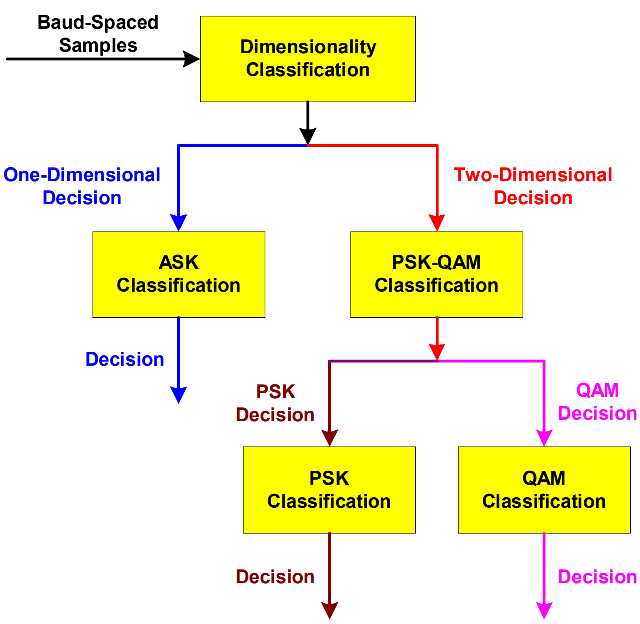

Figure 1. Multi-Stage Decision-Making Procedure

3. FEATURE-GENERATION ALGORITHMS

In [2], the authors presented several PSP-based FGAs that utilized the whole data record length in order to generate the required feature vector for each modulation type of interest. A block diagram of these Full-Record PSP FGAs is shown in Figure 2. These algorithms were found to provide very good MC performance in both FNS and FS environments. Their computational complexity however can be prohibitive for their implementation, especially for large modulation menus involving higher-order modulation types.

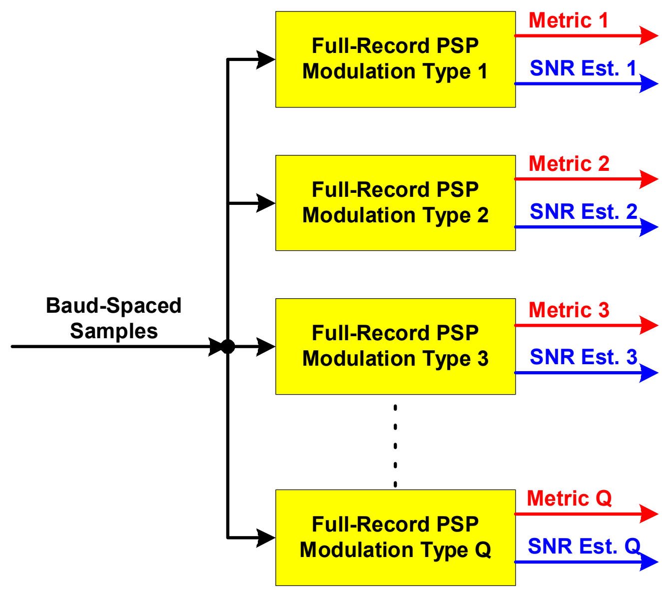

Figure 2. Full-Record PSP FGA

In order to reduce the computational complexity of these FGAs, we can take advantage of the fact that PSP can acquire the OCHIR very rapidly (within few baud-spaced samples) when a Recursive Least-Squares (RLS) or Channel-Weighted RLS (CWRLS) algorithm [2] is used for OCHIR acquisition. Thus, a reduced-complexity FGA can be developed by replacing each Full-Record PSP algorithm with a cascade consisting of a PSP algorithm, using few of the available observation samples for blind OCHIR acquisition, followed by an environment-appropriate Data Detector with or without OCHIR tracking, which uses the remaining or all observation samples to generate the required feature vector. For FNS environments, the appropriate Data Detector is a Symbol-By-Symbol (SBS) detector, whereas for FS environments it is the Viterbi algorithm. A block diagram of this FGA, which has been termed *PSP Cascade*, is illustrated in Figure 3.

We have assessed the performance of PSP Cascade FGAs for a variety of modulation pairs. In Figures 4 and 5, we present Probability of False Classification (PFC) results for the pair 4PSK-8PSK when 4PSK is transmitted; the results for the case where 8PSK is transmitted are similar. The assumed propagation environment is an FNS, Random-Phase channel, whereas the record length is $N_{tot} = 100$ baud-spaced samples.

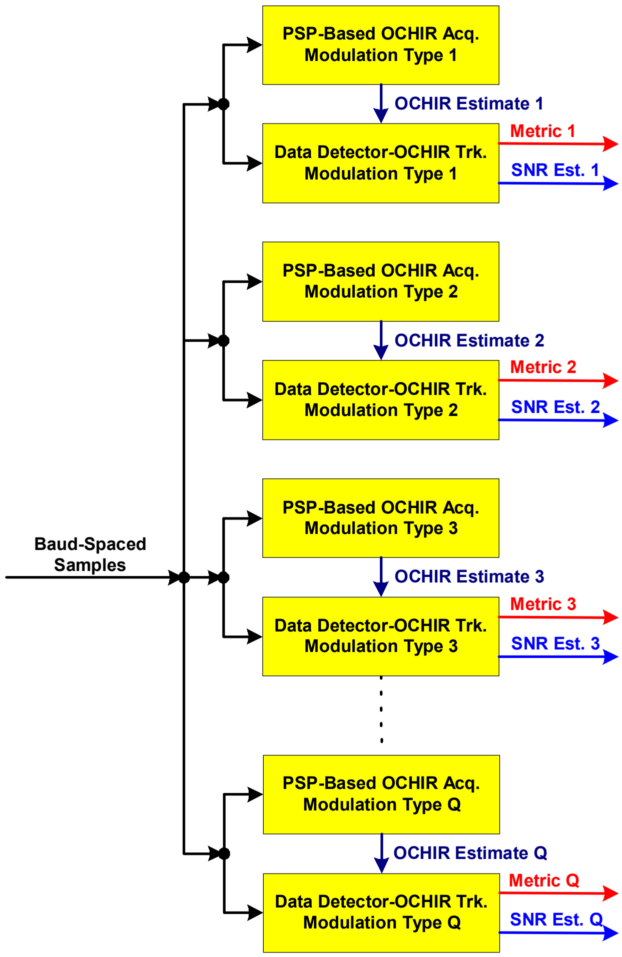

Figure 3. PSP Cascade FGA

We have examined the cases where $N_{ac} = 10$ and 25 baud-spaced samples are used for OCHIR acquisition and all available baud-samples are used for MC (i.e., $N_{mc} = 100$). The PFC was calculated using the EET [2]. For OCHIR acquisition, we have used a PSP-(E-RLS, CWRLS) algorithm with a forgetting factor equal to 0.90 and 1 symbol in the PSP trellis memory (i.e., $N_m = 1$), whereas for MC, two algorithms were considered, namely an SBS detector and an SBS detector with LMS OCHIR tracking. The step-size choices for the latter algorithm were either 0.03 or 0.1. The performance of these PSP Cascade FGAs was compared to that of the Full-Record FGA that provides the best PFC performance [2], namely the PSP-(E-RLS, CWRLS) algorithm with a forgetting factor equal to 0.90 and $N_m = 1$.

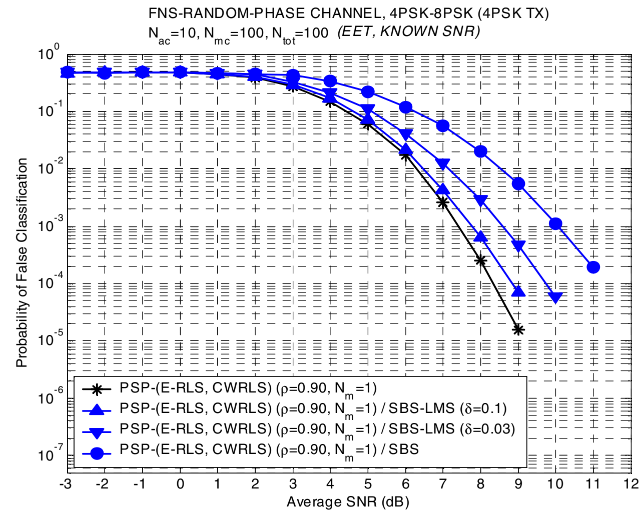

Figure 4. PFC–PSP Cascade FGA, 4PSK-8PSK, N_{ac}=10

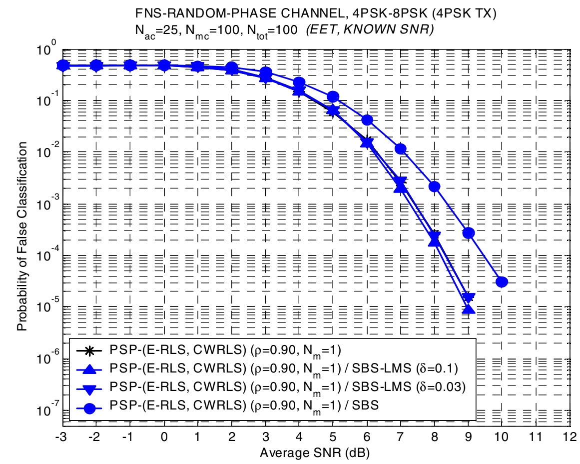

Figure 5. PFC–PSP Cascade FGA, 4PSK-8PSK, N_{ac}=25

For this modulation pair, as well as all other modulation pairs that we have examined, the SBS detector with LMS OCHIR tracking was found to provide better PFC performance than the SBS detector without OCHIR tracking, with only minimal increase of the computational complexity. In the case where only 10 baud-spaced samples are used for OCHIR acquisition, the larger step-size choice provides the best performance. Increasing the number of baud-spaced samples used for OCHIR acquisition to 25, the PFC performance of this FGA for this particular modulation pair is almost identical to that of the Full-Record PSP FGA and independent of the step-size choice.

From the simulation results presented here, as well as from numerous other simulations that we have conducted [5], we have concluded that in FNS environments, the PSP Cascade FGA can provide equivalent MC performance to that of the Full-Record FGA, provided that a sufficient number of baud-spaced samples is used for OCHIR acquisition and the proper step size is used for LMS-based OCHIR tracking.

We must note here that the proper step-size choice may be different for each modulation type of interest. For example, we have found that the step size required for QAM modulations is much larger than the one required for PSK modulations. Furthermore, in channels with large dynamics, it may be necessary to use an adaptive LMS algorithm [6], rather than one with fixed step size.

Finally, in FS environments, we have found that, as is the case with Full-Record PSP FGAs, increasing the trellis memory size beyond the channel memory is quite crucial for obtaining reliable MC performance. This requirement increases the computational complexity of PSP Cascade FGAs significantly, especially when the modulation menu includes many higher-order modulations.

4. THRESHOLD DETERMINATION

The determination of the appropriate thresholds for all required pair-wise comparisons is one of the most critical components of any modulation classifier. Choosing the minimization of the PFC as our performance criterion, the threshold is set at the point where the probability-density functions (pdf's) of the decision statistic, conditioned upon the separate hypotheses, intersect [2], [7]. In most cases, these pdf's are hard to determine analytically and we have to resort to ensemble (histogram-based) techniques for threshold setting.

These techniques require the generation of a large number of decision statistics under each hypothesis. Once these decision statistics are generated, an EET can be determined by finding the point between the sample means of the two decision statistics, which yields the smallest total number of decision errors. As one can appreciate, the need for generation of these decision statistics requires a considerable amount of processing that can only be done off-line. In order to reduce the processing requirements for the determination of the EET, we note that the EET depends on the following parameters:

1. Particular modulation pair
2. FGA and its parameters (such as forgetting factor, step size)
3. Record length
4. SNR and signal power
5. Propagation environment

The dependence of the EET on the first two parameters is obvious: EETs for all pair-wise comparisons required by the particular DMP employed must be generated and stored. Furthermore, the EET value for the FGAs that we described in the previous section depends on the parameters of the particular FGA. Several simulations that we have conducted have shown that the EET dependence on the record length is approximately linear. For example, the EET value when the record length is equal to 1000 baud-spaced samples is approximately equal to 10 times the EET value obtained for a record length of 100 baud-spaced samples. This observation implies that it is possible to generate and store EET values assuming a minimum record length of 100 baud-spaced samples. Then, approximate EET values for larger record lengths are calculated by multiplying the stored EET values with the appropriate scaling factor.

Since the exact dependence of the EET on the SNR value is unknown, we have to calculate and store EET values for a range of SNR values. Depending on the size of the modulation menu and the desired SNR resolution, such calculation requires a considerable amount of off-line processing. In an effort to establish a functional relationship between the EET and the SNR, we have considered the following approach:

1. For SNR values that yield a PFC in the order of 10^{-3}, we calculate the EET by generating a large number of decision statistics under each hypothesis and then finding the point between the sample means of the two decision statistics, which yields the smallest total number of decision errors.
2. For high SNR values, it is impossible to count any decision errors with the above approach. For these SNR values, the EET is taken to be equal to that corresponding to the highest SNR value in the first step.
3. Polynomials of various degrees are fitted in the EET values calculated in the first step.

Simulation results demonstrating the validity of this approach for the pairs 8PSK-16PSK and 8PSK-16QAM and various degrees of polynomial fitting are presented in Figures 6 and 7. The assumed propagation environment is an FNS, static, Random-Phase channel (i.e., the normalized Doppler spread, f_D, is zero), the record length is $N_{tot} = 100$ baud-spaced samples and the FGA is a PSP Cascade. A PSP-(E-RLS, CWRLS) algorithm with a forgetting factor equal to 0.80 and $N_m = 1$ uses $N_{ac} = 25$ baud-spaced samples for OCHIR acquisition, whereas an SBS detector with LMS OCHIR tracking uses $N_{mc} = 100$ baud-spaced samples to generate the necessary metrics.

Using the same environment and FGA, we have examined the performance impact of the new threshold-approximation method for a modulation menu consisting of 2PSK, 4PSK, 8PSK, 16PSK, and 16QAM. Tables 2-7 present confusion matrices for this menu using the Multi-

Stage DMP of Section 2, and assuming perfect baud-rate estimation, symbol synch and SNR estimation.

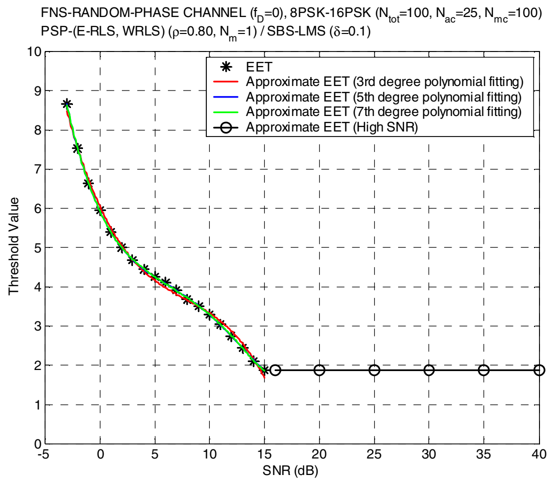

Figure 6. EET Approximation – 8PSK-16PSK

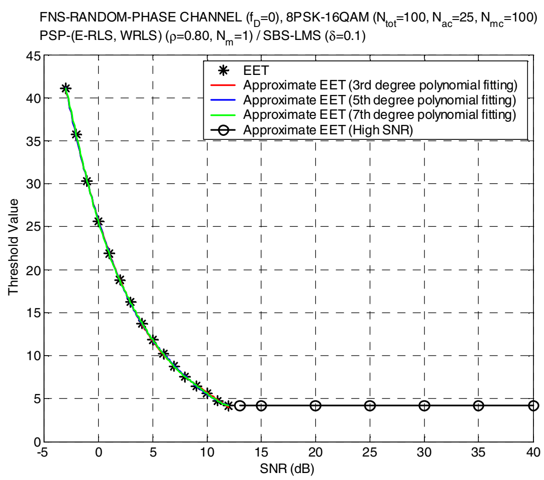

Figure 7. EET Approximation – 8PSK-16QAM

These results demonstrate that a 5th degree polynomial fitting provides very good approximation to the EET, whereas using a 7th degree polynomial fitting leads to performance that is almost identical to that obtained when the EET is used. It is interesting to note that for the particular modulation menu that we have considered and SNR values higher than 15 dB, this threshold-approximation method will lead to the same performance regardless of the polynomial degree used [7].

The value of such approximations cannot be overstated. There is a range of SNR values for which obtaining the exact EET is almost impossible. Using these approximations, we can obtain an approximate EET for any desired SNR value. The savings in storage requirements are significant also, since only the polynomial coefficients for each modulation pair of interest need to be stored.

TABLE 2
EET, SNR=5dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***75.763***	3.175	3.484	17.578
8PSK	0.002	3.690	***46.426***	41.101	8.781
16PSK	0.001	3.684	45.313	***42.111***	8.891
16QAM	0.004	2.130	6.287	1.890	***89.689***

TABLE 3
APPROXIMATE EET (5TH DEGR. POLYNOM.), SNR=5dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***75.426***	3.217	3.513	17.844
8PSK	0.002	3.615	***45.928***	41.488	8.967
16PSK	0.001	3.603	44.864	***42.478***	9.054
16QAM	0.004	2.066	6.218	1.888	***89.824***

TABLE 4
APPROXIMATE EET (7TH DEGR. POLYNOM.), SNR=5dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***75.989***	3.244	3.544	17.223
8PSK	0.002	3.620	***46.277***	41.538	8.563
16PSK	0.001	3.601	45.204	***42.535***	8.659
16QAM	0.004	2.158	6.427	1.950	***89.461***

TABLE 5
EET, SNR=10dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***99.642***	0	0	0.358
8PSK	0	0	***79.201***	20.765	0.034
16PSK	0.001	0.003	13.304	***86.606***	0.086
16QAM	0.007	0	0.024	0.032	***99.937***

TABLE 6
APPROXIMATE EET (5TH DEGR. POLYNOM.), SNR=10dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***99.556***	0	0	0.444
8PSK	0	0	***78.966***	20.988	0.046
16PSK	0.001	0.003	13.100	***86.765***	0.131
16QAM	0.007	0	0.021	0.029	***99.943***

TABLE 7
APPROXIMATE EET (7TH DEGR. POLYNOM.), SNR=10dB

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***99.636***	0	0	0.364
8PSK	0	0	***79.225***	20.739	0.036
16PSK	0.001	0.003	13.316	***86.593***	0.087
16QAM	0.007	0	0.024	0.032	***99.937***

As we mentioned above, the EET depends also on the signal power. We have found that this dependence is a linear one. Since it is impossible to know the transmitted signal power, it is important that all EET values used for fitting are calculated assuming unit signal power. The signal power estimate provided by the SNR estimator can then be used to scale the approximate EET prior to each pair-wise comparison.

It is intuitive to expect that the EET depends also on the particular propagation environment. In order to examine this dependence, we have conducted several simulations in four time varying, FNS, Random-Phase channels [5]. We have found that for normalized Doppler spreads lower than 10^{-3}, the EET is almost identical to that obtained for a static channel. This implies that we can use the EET of a static channel even for slowly changing, time-varying channels. However, when the temporal variations are significantly large, the EET values for static and time-varying channels are vastly different and using the EET of a static channel yields very inaccurate results.

5. SNR ESTIMATION

The implementation of PSP-based MC algorithms requires the estimation of the SNR in order to select the appropriate threshold. In [2], we suggested that the PSP-derived data and OCHIR estimates could be used to obtain an estimate of the SNR for each of the hypothesized modulation types in the menu of interest. Extensive performance simulation results [5] have revealed that this SNR estimation method can cause significant performance imbalance when the operating SNR is low. This performance imbalance has been attributed to the poor estimation performance of this method in the low SNR region.

An alternative SNR estimation method, which can be used to alleviate this problem, is based on the Moment-Matching technique [9]. Using this technique and assuming an FNS channel, the estimates of the signal and noise powers for the k^{th} modulation type are given by

$$\begin{aligned} \hat{S}^{(k)} &= \sqrt{\frac{2\cdot\mu_{z,2}^2-\mu_{z,4}}{2\cdot\mu_{a,2}^{(k)2}-\mu_{a,4}^{(k)}}}, \quad k=1,\ldots,K \\ \hat{N}_0^{(k)} &= \mu_{z,2}-\mu_{a,2}^{(k)}\cdot\hat{S}^{(k)} \end{aligned}$$

where $\mu_{z,m}$ and $\mu_{a,m}^{(k)}$ are the m^{th} moments of the received signal and the k^{th} hypothesis' constellation, respectively. The latter moments can be easily calculated. Calculation of the received signal's moments, however requires knowledge of the channel statistics. Since, in most cases, these statistics are unknown, the received signal's ensemble averages are approximated by their time (sample) averages, i.e.,

$$\mu_{z,m} \approx \frac{1}{N_{snre}}\sum_{n=0}^{N_{snre}-1}|z_n|^m$$

A drawback of this method is that, in several cases, it fails to produce an SNR estimate, because either the numerator for the signal power estimate or the denominator for the noise power estimate yields negative values. In most cases, however we have found that this method provides very good SNR estimation performance, even for low SNR values and short records, when the hypothesized modulation matches the transmitted one.

In order to examine the performance impact of this method, we conducted several simulations using the same environment and FGA assumptions as the ones used in Section 4. Note that for the modulation menu under consideration (2PSK, 4PSK, 8PSK, 16PSK, 16QAM), this method produces two SNR estimates; one for all PSK modulations and one for the 16QAM modulation. For all simulation results presented in Tables 8 and 9, we have chosen to use the smaller of these two estimates for threshold selection. Furthermore, based on the results of the previous section, we have chosen to approximate the EET using a 7^{th} degree polynomial fitting.

TABLE 8
APPROXIMATE EET (7^{TH} DEGR. POLYNOM.), SNR=5dB
MOMENT-MATCHING SNR ESTIMATION

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***76.908***	5.197	4.313	13.582
8PSK	0.001	2.579	***47.213***	45.069	5.138
16PSK	0.001	2.594	46.300	***46.015***	5.090
16QAM	0.111	16.771	18.571	10.725	***53.757***

TABLE 9
APPROXIMATE EET (7^{TH} DEGR. POLYNOM.), SNR=10dB
MOMENT-MATCHING SNR ESTIMATION

	SELECTED MODULATION (%)				
TX MOD	*2PSK*	*4PSK*	*8PSK*	*16PSK*	*16QAM*
2PSK	***100***	0	0	0	0
4PSK	0	***99.763***	0	0	0.237
8PSK	0	0	***76.846***	23.129	0.025
16PSK	0.001	0	11.991	***87.971***	0.037
16QAM	0	3.052	5.820	2.398	***88.730***

Comparing the confusion matrices in Tables 8 and 9, with those in Tables 4 and 7, we see that the Moment-Matching SNR estimation method does not introduce any performance degradation when any of the PSK modulations is transmitted. When a 16QAM is transmitted, however there is performance degradation, which can be significant for low SNR values (e.g., 5 dB). This is due to the fact that the minimum of the two SNR estimates obtained by this method is usually the one produced under the PSK hypothesis, which is typically much lower than the actual SNR value when 16QAM is transmitted [8]. Our future research efforts will focus on methods remedying this problem and alleviating its effects on threshold setting. Potential methods that will be considered include the addition of a bias term in the derived SNR estimates and the establishment of thresholds as a function of the estimated SNR rather than the actual transmitted one.

6. CONCLUSIONS

In this paper, we presented an efficient, PSP-based digital MC engine, which employs:

a. An efficient decision-making procedure that reaches a final decision in a tree-like search.

b. A reduced-complexity feature-generation algorithm that blindly and rapidly acquires the overall channel and then employs low-complexity tracking algorithms for generating the necessary discriminating metrics.

c. A threshold-setting mechanism that employs polynomial-fitting techniques so as to establish a functional relationship between the required threshold(s) and the prevailing signal-to-noise ratio.

Selected performance simulation results have demonstrated its performance, benchmarked against the exhaustive version.

REFERENCES

[1] R. Rubenstein, “Radios Get Smart,” *IEEE Spectr.*, pp. 46-50, Feb. 2007.

[2] P. Panagiotou, A. Polydoros, and C. Oh, “Per-Survivor-Processing (PSP) methods for digital modulation classification in stochastic channels,” in *Proc. MILCOM*, Boston, MA, 2003, *Classified Section*.

[3] P. Panagiotou, A. Anastasopoulos, and A. Polydoros, “Likelihood ratio tests for modulation classification,” in *Proc. MILCOM*, Los Angeles, CA, 2000, vol. 2, pp. 670-674.

[4] R. Raheli, A. Polydoros, and C-K Tzou, “Per-Survivor Processing: a general approach to MLSE in uncertain environments, ” *IEEE Trans. Commun.*, vol. 43, pp. 354–364, Feb-Mar-Apr 1995.

[5] A. Polydoros et al, *Weak Signal Modulation Classifier (MC) and Specific Emitter Information (SEI) in Time-Varying Multipath*, TrellisWare Technologies Technical Report, 2003.

[6] S. Haykin. *Adaptive Filter Theory*. 3rd Ed., Prentice Hall: Upper Saddle River, NJ, 1996

[7] N. E. Lay, *Modulation Classification and Data Detection of Signals in Unknown ISI Environments*, Ph.D. thesis, Communication Sciences Institute, University of Southern California, Los Angeles, CA, May 1996.

[8] A. Polydoros et al, *Weak Signal Modulation Classifier (MC) and Specific Emitter Information (SEI) in Time-Varying Multipath*, TrellisWare Technologies Technical Report, 2004.

[9] S. M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall: Upper Saddle River, NJ, 1993.