Low-Power Digital Signal Processing via Dynamic Algorithm Transformations (DAT)

Manish Goel and Naresh R. Shanbhag
VLSI Information Processing Systems (VIPS) Group
Coordinated Science Laboratory/ECE Department
University of Illinois at Urbana-Champaign
1308 West Main Street, Urbana, IL 61801.
Phone : (217) 244 0041, Fax : (217) 244 1946.
E-mail : [mgoel,shanbhag]@uivlsi.csl.uiuc.edu

Abstract— **We have proposed dynamic algorithm transformations (DAT) for low power VLSI signal processing. The principle behind DAT is that the conventional signal processing system designed for the worst-case is not optimum (from energy viewpoint) for the nominal and the best cases. Therefore, significant energy savings can be achieved by optimally reconfiguring the hardware in these situations. The reconfiguration strategies are derived by solving an optimization problem with energy as the objective function, and a constraint on the SNR. In this paper, we present reconfiguration strategies for a system comprising of multiple adaptive filters. An example is 155.52 Mb/s ATM-LAN transceiver, where two adaptive filters - namely near-end crosstalk canceller and fractionally spaced equalizer are employed. The input variabilities for this example are due to the different cable lengths at various locations. Simulation results indicate that energy savings range from 0%-85% for cable lengths ranging from 110m to 40m. On an average 69% energy savings are achieved.**

I. Introduction

Power-reduction techniques have been proposed at all levels of VLSI design hierarchy ranging from the circuits to algorithms. Of particular interest in this paper are *algorithm transformation techniques* [1]-[2]. We refer to these algorithm transformations as *static algorithm transformations* (SAT), because these are applied during the algorithm design phase assuming a worst-case scenario and their implementation is time-invariant.

In contrast, we have proposed *dynamic algorithm transforms* (DAT) [3]-[5], which optimize energy dissipation via real-time energy-optimum *reconfiguration* in the presence of input non-stationarities. Traditionally, a signal processing system is designed for the *worst* case, so that it can meet the performance requirement for all cases including the *nominal* and the *best* case. For example, a broadband modem is typically designed for the longest cable length, the maximum cable temperature and the worst case near-end/far-end crosstalk interferers. This worst-case design is an "over-kill" in terms of energy consumption for nominal and best cases. If the worst-case occurs rarely, then much energy savings can be achieved by reconfiguration of the system for the best and the nominal cases.

In past [3]-[5], we have shown that significant energy savings can be achieved via DAT-based reconfiguration of the signal processing systems. For example, in [3], a DAT-based reconfiguration of the near-end crosstalk (NEXT) canceller for 155.52 Mb/s ATM-LAN [3] resulted in energy savings of 21%-62%. Similarly, 88% energy savings were achieved via DAT-based reconfiguration of the fractionally-spaced linear equalizer (FSLE) for 51.84 Mb/s very high-speed digital subscriber loop (VDSL) [5] environments. The reconfigurable parameters are - number of taps, the input and coefficient precisions, and the supply voltage. In this paper, we revisit 155.52 Mb/s ATM-LAN transceiver [6]. Such a transceiver for category-3 copper wiring requires both a NEXT canceller and a FSLE. In contrast to individual reconfiguration of the NEXT canceller [3] or the FSLE [5], we present reconfiguration strategy for *jointly reconfiguring* the NEXT canceller and FSLE. It will be shown that for 155.52 Mb/s ATM-LAN transceiver, joint reconfiguration results in larger overall energy savings.

The outline of the rest of the paper is as follows. In next section, we review dynamic algorithm transformations (DAT) and present a reconfiguration strategy for the problem in this paper. In section III, we describe ATM-LAN transceiver and the simulation set-up for section IV, in which we present the simulation results.

II. Dynamic Algorithm Transformations (DAT)

In this section, we present preliminaries about dynamic algorithm transformations (DAT) [3]-[5]. A DAT-based reconfigurable signal processing system has two components - 1.) signal processing algorithm (**SPA**) block that is a reconfigurable datap-

This work was supported by DARPA contract DABT63-97-C-0025 and NSF CAREER award MIP-9623737.

0-7803-5148-7/98/$10.00©1998 IEEE

ath; and 2.) signal monitoring algorithm (**SMA**) block that implements reconfiguration strategy/control

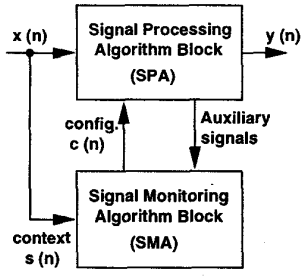

Fig. 1. DAT-based reconfigurable signal processing system.

based upon temporal/spatial variabilities in the input. The input variabilities are described as transition between predefined states. The **SPA** block is characterized by a configuration vector. The readers are referred to [5] for the details of the general framework. Here, we present only the **SPA** and the **SMA** blocks.

A. The **SPA** *block*

A block diagram of the signal processing algorithm **SPA** block is shown in Fig. 2. Assume that K adaptive filters are being employed to generate the output $y(n)$, which should match the desired signal $d(n)$. Let N_i, $0 \leq i \leq K-1$ be the number of taps in the i^{th} filter. The filters can be real-valued, complex-valued or a combination of these two. An example is 155.52 Mb/s ATM-LAN transceiver, where we have two adaptive filters ($K = 2$) - a near-end crosstalk (NEXT) canceller, and a fractionally spaced linear equalizer (FSLE). We will see in section III that the NEXT canceller is a complex adaptive filter and the FSLE is a pair of real adaptive filters.

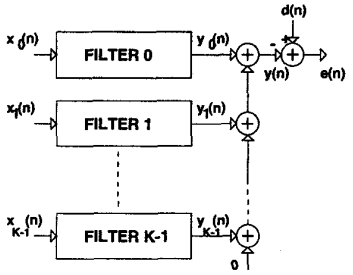

Fig. 2. **SPA** block: K adaptive filters.

Each adaptive filter ($0 \leq i \leq K-1$) is assumed to be reconfigurable with the architecture shown in Fig. 3. The *least mean square (LMS) algorithm* is employed to implement the adaptive filter architecture shown in Fig. 3. The filter (**F**) block implements a finite-impulse response (FIR) filter, whose coefficients are updated by the weight-update block (**WUD**)-block. The multiplication of input by the coefficient and updating of the coefficient constitute a tap as shown in Fig. 3. The configuration signal for the i^{th} adaptive filter is given by,

$$\mathbf{c}_i = [\underline{\alpha_i}, \underline{\beta_i}, B_{x,i}, B_{w,i}, V_{dd,i}], \tag{2.1}$$

where $\underline{\alpha_i}$ is a binary vector of length N_i, whose j^{th} element indicates whether F-block multiplier of j^{th} tap is powered-up ($\alpha_{i,j} = 1$) or powered-down ($\alpha_{i,j} = 1$). This can also be seen from the filter architecture in Fig. 3, where the j^{th} tap can be shut-down by forcing $\alpha_{i,j} = 0$. Similarly, $\underline{\beta_i}$ is the binary vector of length N_i, whose j^{th} element indicates whether **WUD**-block multiplier of j^{th} tap is powered-up ($\beta_{i,j} = 1$) or powered-down ($\beta_{i,j} = 1$). Further, $B_{x,i}$ and $B_{w,i}$ are signals for controlling the input data precision and coefficient precision, respectively of the i^{th} filter. Finally, $V_{dd,i}$ is the supply voltage of the i^{th} filter. The control signals for the K adaptive filters can be combined to form configuration vector $\mathbf{c}(n)$ given as follows:

$$\mathbf{c}(n) = [\mathbf{c}_0(n), \mathbf{c}_1(n), \cdots, \mathbf{c}_{K-1}(n)]. \tag{2.2}$$

The output $y_i(n)$ of the i^{th} filter is computed as:

$$y_i(n) = \sum_{j=0}^{N_i-1} \alpha_{i,j} w_{i,j}^* x_i(n-j), \tag{2.3}$$

where $w_{i,j}$ is coefficient of the j^{th} tap, and $x_i(n)$ is the input to i^{th} filter with tap length N_i. The output $y(n)$ of the combination of the adaptive filter can be computed by adding $y_i(n), 0 \leq i \leq K-1$. The error signal $e(n) = d(n) - y(n)$ is given as follows:

$$e(n) = d(n) - \sum_{i=0}^{K-1} \sum_{j=0}^{N_i-1} \alpha_{i,j} w_{i,j}^* x_i(n-j). \tag{2.4}$$

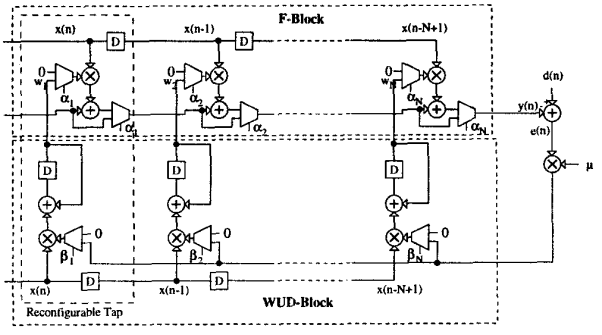

Fig. 3. Reconfigurable Adaptive Filter Architecture.

If we assume that the inputs to all the filters are uncorrelated, then it can be shown that the mean squared

error (MSE) is given by,

$$\mathcal{J} = \sigma_d^2 - \sum_{i=0}^{K-1} \sum_{j=0}^{N_i-1} \alpha_{i,j} |w_{i,j}|^2 \sigma_{x,i}^2, \qquad (2.5)$$

where σ_d^2 is the variance of the desired signal, and $\sigma_{x,i}^2$ is the variance of the input to the i^{th} filter. It can be shown that powering down any tap (i.e. forcing $\alpha_{i,j} = 0$) causes an increase in MSE. The powering down of a tap however results in decrease in energy consumption of the filter, which is given by:

$$\mathcal{E} = \sum_{i=0}^{K-1} \sum_{j=0}^{N_i-1} \alpha_{i,j} \mathcal{E}_m(w_{i,j}), \qquad (2.6)$$

where $\mathcal{E}_m(w_{i,j})$ is the energy consumption of the j^{th} tap in i^{th} filter. The goal of the dynamic algorithm transforms is to find an energy-optimum configuration under the SNR (or equivalently MSE) constraint. This leads to the following optimization problem:

$$\min_{\alpha_{i,j} \in \{0,1\}} \sum_{i=0}^{K-1} \sum_{j=0}^{N_i-1} \alpha_{i,j} \mathcal{E}_m(w_{i,j})$$
$$s.t.\ \sigma_d^2 - \sum_{i=0}^{K-1} \sum_{j=0}^{N_i-1} \alpha_{i,j} |w_{i,j}|^2 \sigma_{x,i}^2 \leq \mathcal{J}_o, \qquad (2.7)$$

where $\mathcal{J}_o$ is the desired MSE. Note that we do not include β_k's in the optimization problem because $\beta_k = 0$ after the adaptive filter has converged, i.e., the **WUD**-block is powered down. The optimization problem in (2.7) is solved via the *Lagrange Multiplier Method* to get a reconfiguration strategy, which is described next.

B. The SMA *block*

It can be shown [5] that the solution to (2.7) is given by:

$$\alpha_{i,j,opt} = \begin{cases} 1, & \frac{\sigma_{x,i}^2 |w_{i,j}|^2}{\mathcal{E}_m(w_{i,j})} \geq \tilde{\lambda}^* \\ 0, & \frac{\sigma_{x,i}^2 |w_{i,j}|^2}{\mathcal{E}_m(w_{i,j})} < \tilde{\lambda}^*, \end{cases} \qquad (2.8)$$

where $\tilde{\lambda}^*$ is a constant. The solution (2.8) indicates that it is better to power down taps with small values of $[\sigma_{x,i}^2 |w_{i,j}|^2 / \mathcal{E}_m(w_{i,j})]$. Intuitively, this makes sense as small values of $[\sigma_{x,i}^2 |w_{i,j}|^2 / \mathcal{E}_m(w_{i,j})]$ imply that the j^{th} tap in the i^{th} filter contributes relatively less to the MSE (as $w_{i,j}$ is small) but consumes more energy ($\mathcal{E}_m(w_{i,j})$ is large). In practice, we don't need to compute the constant $\tilde{\lambda}^*$ in (2.8) if we employ the strategy that $\alpha_{k,opt}$ can be obtained by powering down the taps starting with the smallest value of $[\sigma_{x,i}^2 |w_{i,j}|^2 / \mathcal{E}_m(w_{i,j})]$ until the MSE constraint (see (2.7)) is violated.

The $\beta_{i,j}$ can similarly be obtained by choosing $\beta_{i,j} = 0$ whenever $\alpha_{i,j} = 0$ or when the filter has converged. The optimum input precision $B_{x,i,opt}$ for the i^{th} filter can be obtained by computing the input peak-to-average ratio (PAR_i) (defined as the ratio of the maximum to the root mean squared value of the input signal) and then determining the minimum precision required for the desired signal-to-quantization ratio (SQNR) at the input. In general, the input precision of i^{th} filter can be reduced by one bit for each $6dB$ reduction in PAR of $x_i(n)$. The optimum coefficient precision $B_{w,i,opt}$ for the i^{th} filter is chosen by deriving expression for the round-off error in terms of the coefficient precision and then determining the precision to achieve the desired SQNR at the output. It can be shown that SQNR at the output is a function of number of powered-up taps, and the coefficient precision can be reduced by 1 bit for each four-fold reduction in number of powered-up taps. As some of the taps are powered down, the critical path delay of the filter decreases. Thus the power supply can be reduced for further power savings. The reader is referred to [5] for a closed-form expression for the optimum supply voltage selection.

The reconfiguration strategy in (2.8) requires computation of energy values $\mathcal{E}_m(w_{i,j})$. The energy $\mathcal{E}_m(w_{i,j})$ is computed by evaluating the energy consumed by the j^{th} tap of the i^{th} filter. The energy models based on zero-delay simulations are presented in [5] for estimating $\mathcal{E}_m(w_{i,j})$. For the sake of simplicity, we employ these zero-delay energy models in the reconfiguration strategy. However, the energy savings are computed with the real-delay energy models. In the next section, we employ a DAT-based adaptive system for 155.52 Mb/s ATM-LAN.

III. The 155.52 Mb/s ATM-LAN

In this section, we will present a brief overview of the 155.52 Mb/s ATM-LAN transceiver [6] followed by the simulation set-up.

A. The Transceiver

The block diagram of a digital carrierless amplitude phase (CAP) transceiver is shown in Fig. 4. At the transmitter, the scrambled bit-stream is fed into an encoder, which maps blocks of $m = 6$ bits onto one of $k = 2^m = 64$ different complex symbols for a 64-CAP line code. The complex symbols are processed by digital shaping filters operated at a sampling frequency f_s of $77.78MHz$. The outputs of the filters are subtracted and the result is passed through a digital-to-analog (D/A) converter, which is followed by an interpolating low-pass filter (LPF).

At the receiver (see Fig. 4), the received signal is distorted further due to the superimposition of the near-end crosstalk (NEXT) signal. This composite signal is

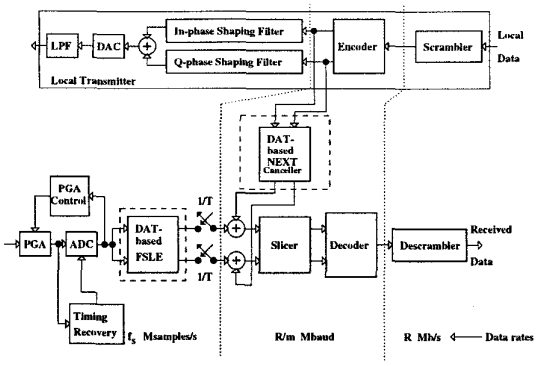

Fig. 4. 155.52 Mb/s ATM-LAN transceiver.

processed by a $T/3$ fractionally-spaced linear equalizer (FSLE), which is a pair of adaptive filters. In addition, the local transmitted symbols are passed through a complex adaptive NEXT canceller, which tries to cancel the effect of the NEXT in the received signal. In past, we have presented strategies for reconfiguring the NEXT canceller [3] and the FSLE [5]. In this paper, we employ the reconfiguration strategy in section II-B for *joint reconfiguration* of the NEXT canceller and FSLE. For 64-CAP, a slicer SNR of $29.45dB$ is sufficient to obtain a probability of error less than 10^{-10}.

B. Simulation Setup

We assume a spatial variation in the length of the UTP-3 cable from $110m$ to $40m$. An estimate of the probability distribution of the cable lengths can be obtained by estimating typical distances from the switch to the desktops. In this paper, we consider eight cable lengths and assume them to be Gaussian distributed with a mean of $75m$. The changes in the input state can be detected by monitoring slicer SNR. We assume that $SNR_o = 31dB$ (this is $1.55dB$ more than the minimum of $29.45dB$) is the desired performance level. We choose $\delta = 3dB$ to remove undesired glitching in steady state.

As far as VLSI parameters are concerned, it is assumed that the standard cells based $0.18\mu m$, $2.5V$ CMOS technology are being employed. The energy models are obtained by real-delay simulations via gate-level simulator MED [7]. The maximum number of taps in the NEXT canceller and FSLE are 30 (complex) and 252 (real), respectively. The data precision B_x for the NEXT canceller is kept constant at 4 bits because the inputs to the NEXT-canceller belong to the 64-CAP signal set $\{-7,-5,-3,-1,1,3,5,7\}$, which can be represented in 4 bits. The input data precision for the FSLE is assumed to be 8 bits. Furthermore, for NEXT canceller, we employ strength-reduced **SR** architecture [8] with three real filters. We follow the reconfiguration strategy presented in section II-B to obtain the energy-optimum configuration. In the next section, we present the simulation results for the DAT-based receiver.

IV. Simulation Results

In this section, we study the performance of the DAT-based adaptive receiver (section II) for 155.52 Mb/s ATM-LAN (section III). As discussed in section III, the blocks of interest are - NEXT canceller and FSLE. The three scenarios of interest are: 1.) NEXT canceller reconfigurable and FSLE hardwired; 2.) NEXT canceller hardwired and FSLE reconfigurable; and 3.) both NEXT canceller and FSLE reconfigurable. All energy savings are computed with respect to hardwired NEXT canceller and FSLE with maximum complexity.

A. Reconfigurable NEXT canceller and hardwired FSLE

In this experiment, we assume that only the NEXT canceller is reconfigurable. The FSLE is hardwired with the maximum number of taps, maximum coefficient precision, maximum data precision and the maximum supply voltage. The results for this scenario are plotted in Fig. 5. The number of powered up taps in

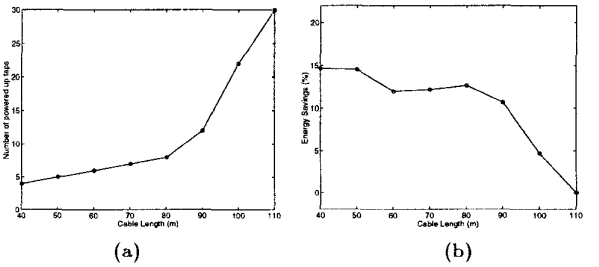

(a) (b)

Fig. 5. Results for reconfigurable NEXT canceller: (a) number of powered-up taps and (b) energy savings.

NEXT canceller vary from 4 to 30, as the cable length varies from 40m to 110m, respectively. The coefficient precision varies from 9 bits to 10 bits for the corresponding cases. The supply voltage varies from 2.1 V to 2.5 V for cable lengths ranging from 40m to 110m. The energy savings for this case are plotted in Fig. 5(b). The energy savings range from 14% to 0% for cable lengths ranging from 40m to 110m.

B. Hardwired NEXT canceller and reconfigurable FSLE

In contrast to the previous experiment, we assume that the NEXT canceller is hardwired and FSLE is reconfigurable. Recall that there are a total of 252

real taps in FSLE. The NEXT canceller is hardwired to the maximum number of taps, maximum coefficient precision, maximum data precision and the maximum supply voltage. The simulation results are plotted in Fig. 6. The number of powered up taps in FSLE vary

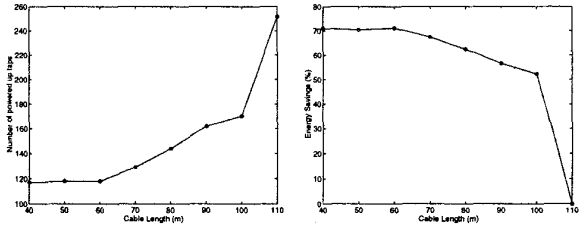

(a) (b)

Fig. 6. Results for reconfigurable FSLE: (a) number of powered-up taps and (b) energy savings.

from 118 to 252, as the cable length varies from 40m to 110m, respectively. Since the coefficient precision can be reduced only if there is a 4 times reduction in number of powered-up taps, the coefficient precision for this case was constant at 10 bits for all cable lengths. varies from 9 bits to 10 bits for the corresponding cases. The supply voltage varies from 2.1 V to 2.5 V for cable lengths ranging from 40m to 110m. The energy savings for this case are plotted in Fig. 6(b). The energy savings range from 71% to 0% as the cable length change from 40m to 110m. It is worth-noting that the the reconfiguration of FSLE has a bigger impact than the reconfiguration of NEXT canceller. This is because of the larger complexity of the FSLE. Next, we consider the scenario where both the NEXT canceller and FSLE are reconfigurable.

C. Both NEXT canceller and FSLE reconfigurable

The simulation results for this scenario are shown in Fig. 7. The number of powered up taps for the NEXT canceller and FSLE are plotted in Fig. 7(a). The number of powered-up taps in NEXT canceller and FSLE range from 7 to 30 and 113 to 252 for the cable length changes from 40m to 110m, respectively. The coefficient precision of the NEXT canceller vary from 9 to 10 bits, while that of the FSLE stays fixed at 10 bits. The supply voltage vary from 2.1 V to 2.5 V for the cable lengths ranging from 40m to 110m, respectively. The energy savings are plotted in Fig. 7(b). The energy savings ranging from 85% to 0% are obtained for cable lengths variations from 40m to 110m, respectively. A comparison of the energy savings in subsection A,B and C indicate that the joint reconfiguration of the NEXT canceller and FSLE has the biggest impact on the energy savings. The average energy savings can be computed by multiplying the energy savings for each cable

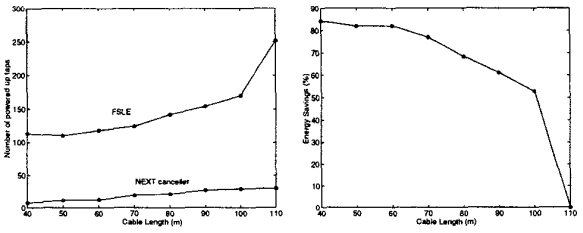

(a) (b)

Fig. 7. Results for reconfigurable NEXT canceller and FSLE: (a) number of powered-up taps and (b) energy savings.

length with their probabilities and then taking the summation. It was found that on an average, 69% energy savings are achieved by the DAT-receiver as compared to a worst case design.

Therefore, we conclude that DAT-based receiver is a good alternative for ATM-LAN transceivers. Future work involves the application of DAT to the forward error correction (FEC) blocks, and to wireless systems.

REFERENCES

[1] A. Chandrakasan *et al.*, "Minimizing power using transformations," *IEEE Trans. Computer-Aided Design*, vol. 14, no. 1, pp. 12–31, Jan. 1995.

[2] K. K. Parhi, "Algorithm transformation techniques for concurrent processors," *Proceedings of the IEEE*, vol. 77, no. 12, pp. 1879–1895, Dec. 1989.

[3] M. Goel and N. R. Shanbhag, "Dynamic algorithm transformations (DAT) for low-power adaptive signal processing," in *International Symposium on Low Power Electronics and Design (ISLPED)*, (Monterey, CA), pp. 161–166, Aug. 1997.

[4] M. Goel and N. R. Shanbhag, "Low-power reconfigurable signal processing via dynamic algorithm transformations (DAT)," in *Proc. ICASSP*, vol. 4, (Seattle, WA), pp. 3081–3084, May 1998.

[5] M. Goel and N. R. Shanbhag, "Low-power equalizers for 51.84 Mb/s very high-speed digital subscriber loop (VDSL) modems," in *IEEE Workshop on Signal Processing Systems (SiPS): Design and Implementation*, (Boston, MA), pp. 317–326, Oct. 1998.

[6] G. H. Im and J. J. Werner, "Bandwidth-efficient digital transmission up to 155 Mb/s over unshielded twisted-pair wiring," *IEEE J. on Selected Areas in Comm.*, vol. 13, no. 9, pp. 1643–1655, Dec. 1995.

[7] M. G. Xakellis and F. N. Najm, "Statistical estimation of the switching activity in digital circuits," in *Design Automation Conference*, pp. 728–733, June 1994.

[8] N. R. Shanbhag and M. Goel, "Low-power adaptive filter architectures and their application to 51.84 Mb/s ATM-LAN," *IEEE Trans. Signal Processing*, vol. 45, no. 5, pp. 1276–1290, May 1997.