

A SiGe BiCMOS Eight-Channel Multidithering Sub-Microsecond Adaptive Controller

Dimitrios N. Loizos, *Member, IEEE*, Paul P. Sotiriadis, *Senior Member, IEEE*, and Gert Cauwenberghs, *Senior Member, IEEE*

Abstract—A SiGe BiCMOS mixed-signal adaptive controller-on-chip is presented that implements gradient descent of a supplied analog control objective. Eight analog variables controlling the external plant are perturbed in parallel using sinusoidal dithers, and their gradient components are estimated by parallel synchronous detection of the dithers in the control objective. Translinear all-NPN bipolar circuits achieve linear tuning of frequency and amplitude in the oscillators and synchronous detectors, covering a 4-kHz–600-MHz range in dither frequencies with -30 -dB/octave suppression of intermodulation products. Experimental results demonstrate adaptive optimization of a three-variable nonlinear plant within $1 \mu\text{s}$ for dithers in the 100–200-MHz frequency range. The chip measures $3 \text{ mm} \times 3 \text{ mm}$ in $0.5\text{-}\mu\text{m}$ SiGe and consumes 110 mW at 3.3-V supply.

Index Terms—Gradient descent optimization, high-speed adaptive control, model-free adaptation, multidithering, SiGe BiCMOS integrated circuits, translinear circuits, wideband linear tunability.

I. INTRODUCTION

SYNCHRONOUS detection is fundamental to many communications systems performing analog decoding of an amplitude-modulated carrier by measuring the component of the received signal in phase with the carrier. The same principle extends to gradient-descent optimization of an objective “metric” $J = J(u, t)$ of a plant, where the derivative $\partial J/\partial u$, with respect to the control variable u , is needed. This information is retrieved by superimposing a “dither” signal to the control variable u and performing synchronous detection between the received perturbed objective and the dither signal.

For multiple control variables $\mathbf{u} = (u_1, u_2, \dots, u_n)^T$, the gradient ∇J of the metric $J = J(\mathbf{u}, t)$ is estimated by applying mutually orthogonal dithers to the control variables *in parallel* and performing synchronous detection for each of them. In the case of broad-band random dithers, this technique is known as model-free adaptation (MFA) [1] or stochastic parallel gradient descent (SPGD) [2]–[5]. Several analog hardware implementations of SPGD have been presented, with most of them being

based on Bernoulli distributed dithers [6]–[8]. In the case of narrow-band sinusoidal dithers, the technique is known as multidithering (MD) [9] and has been used extensively in adaptive optics for wavefront correction.

One limitation of broad-band excitation is that delays in the plant and in evaluation of the metric distort the gradient estimates and thus limit the speed of gradient-descent adaptation. The circuit presented here circumvents this limitation by applying narrow-band excitation using sinusoidal dithers, for which any delays in the plant and metric reduce to a single parameter, i.e., one phase for each dither. The MD circuit allows for variable phase selection in parallel synchronous detection to compensate for arbitrary phase delay in each control channel. As with SPGD, MD gradient descent implements a model-free form of adaptive control [1], which guarantees convergence to a local optimum (minimum) of the measured control objective independent of model assumptions on the plant, and mismatch in the circuit realization of the analog controller. SiGe BiCMOS circuit implementation supports dither frequencies that are linearly tunable from 4 kHz to 600 MHz, serving a wide range of applications in high-speed adaptive control such as optical wavefront correction [4], [9], multibeam optical communications [10], and aberration correction in two-photon microscopy [11], among others.

A SiGe BiCMOS current-controlled oscillator circuit that provides sinusoidal dither signals with wide tuning range is demonstrated in [12]. Architectural considerations and the circuit design of the MD gradient-descent adaptive control system, as well as the first experimental results from the integrated SiGe BiCMOS implementation, are presented in [13]. Here, we provide a detailed description, analysis, and experimental characterization of the circuits comprising the MD control system, quantify the system performance in terms of normalized energy efficiency, and present experimental results demonstrating closed-loop adaptation settling within $1 \mu\text{s}$ at 110-mW power. The MD adaptation system architecture is briefly reviewed in Section II, and circuits are described in Section III. Performance metrics are defined in Section IV and experimentally evaluated in Section V.

II. SYSTEM ARCHITECTURE

The presented VLSI chip contains eight identical control channels, each serving a variable of the control vector $\mathbf{u} = (u_1, u_2, \dots, u_n)^T$ of the metric $J(\mathbf{u}, t)$. Two or more chips can be connected in parallel and optimize performance metrics with 16 or more variables. Optimization of the externally presented metric $J(\mathbf{u}, t)$ is achieved by realization of the signed version of the gradient-flow algorithm and is performed

Manuscript received September 09, 2008; revised January 14, 2009. First published February 20, 2009; current version published January 20, 2010. This work was supported in part by the Army Research Laboratory, Adelphi, MD and the Defense Advanced Research Projects Agency (DARPA). This paper was recommended by Associate Editor A. J. L. Martin.

D. N. Loizos was with the Division of Biological Sciences, University of California, San Diego, La Jolla, CA 92093 USA. He is now with NetLogic Microsystems Inc., Mountain View, CA 94043 USA (e-mail: dloizos@ucsd.edu).

P. P. Sotiriadis is with Sotekco Electronics LLC, Baltimore, MD 21201 USA.

G. Cauwenberghs is with the Department of Bioengineering, Jacobs School of Engineering, and the Institute for Neural Computation, University of California, San Diego, La Jolla, CA 92093 USA.

Digital Object Identifier 10.1109/TCSI.2009.2015726




Fig. 1. System architecture of the MD adaptive optimization system, with one of the control channels being shown.

in two steps. First, the gradient ∇J is estimated through parallel synchronous detection, and second, the gradient information is quantized and used to update variables u_i .

The simplified architecture of the channels is shown in Fig. 1. The sinusoidal dither of the channel is generated by one of the phases of a three-phase sinusoidal oscillator, added onto the control variable u_i

$$\tilde{u}_i = u_i + \alpha \cos(\omega_i t). \quad (1)$$

The perturbed control variables $\tilde{\mathbf{u}} = (\tilde{u}_1, \tilde{u}_2, \dots, \tilde{u}_n)^T$ are applied to the under optimization plant, which returns a metric signal $J = J(\tilde{\mathbf{u}})$ that is fed back to the inputs of the channels for gradient estimation. Each channel contains a linear multiplier and an adjustable high-order low-pass filter for synchronous detection which result in a gradient estimate¹ [13]

$$\overline{J(\tilde{\mathbf{u}}) \cos(\omega_i t)} = \frac{\alpha}{2} \left. \frac{\partial J}{\partial u_i} \right|_{\mathbf{u}}. \quad (2)$$

The multiplier output is amplified prior to low-pass filtering since both the dither and the perturbed portion of the received metric are low in amplitude.

The total time delay τ_i in the adaptation loop of channel i can be represented as a phase delay $\varphi_i = \omega_i \tau_i$, which affects gradient estimation by scaling the synchronous detection output with a factor $\cos(\varphi_i)$

$$\overline{J(\tilde{\mathbf{u}})_{\tau\text{-delay}} \cos(\omega_i t)} = \frac{\alpha}{2} \left. \frac{\partial J}{\partial u_i} \right|_{\mathbf{u}} \cos(\varphi_i). \quad (3)$$

Compensation of the delay phase factor $\cos(\varphi_i)$, approximated with $\pi/6$ phase quantization error, can be achieved by choosing one of the three available phases of the oscillator as reference for synchronous detection and by retaining or inverting the polarity of the gradient estimate [13]. With a phase resolution in steps of $\pi/3$, proper selection of one of the six phases (including polarity) leads to a residual net phase error that is less than $\pi/6$.

¹Overline denotes low-pass filtering.

The $\pi/6$ phase quantization error leads to an amplitude error in the gradient estimate that is at most $1 - \cos(\pi/6) \approx 0.134$, producing the correct polarity of the updates and a minor reduction in the rate of convergence. A practical algorithm to dynamically select the phase φ_i during closed-loop adaptation of a varying metric is given in [14]. In what follows, we assume that the phase factor is maintained close to identity, i.e., $\cos(\varphi_i) \approx 1$.

A comparator extracts the signum of the metric's derivative $\partial J / \partial u_i$. The signum controls the direction of the current in the charge pump, which continuously updates the value of the control variable u_i and implies (4), where G is an adjustable gain coefficient that controls the convergence rate of all channels

$$\frac{du_i}{dt} = -G \operatorname{sgn} \left(\left. \frac{\partial J}{\partial u_i} \right|_{\mathbf{u}} \right). \quad (4)$$

Equation (4) is the signed form of the gradient-flow algorithm and ensures convergence of the objective metric J to a (local) extremum as long as J has no saddle points [13]. Expansion of dJ/dt using (4) reveals the L-1 norm convergence of J with time

$$\frac{dJ}{dt} = \sum_i \frac{\partial J}{\partial u_i} \cdot \frac{du_i}{dt} = -G \left\| \left. \frac{\partial J}{\partial \mathbf{u}_i} \right|_1 \right\|. \quad (5)$$

For $G > 0$, (5) converges to a (local) minimum of J , whereas for $G < 0$, a (local) maximum of J is reached.

III. CIRCUIT DESIGN AND IMPLEMENTATION

The circuit has been designed for ultrawide-frequency-range tunability. To this end, a translinear architecture has been chosen and implemented in a $0.5\text{-}\mu\text{m}$ SiGe BiCMOS process providing linear tuning from 4 kHz to 600 MHz. An all-NPN translinear design avoids the use of inferior PNP devices in the particular SiGe process.

A. Oscillator

The oscillator is a differential three-stage G_m - R - C ring oscillator [Fig. 2(a)] with coupled frequency and amplitude control [12]. The Barkhausen criterion implies the following oscillation conditions:

$$G_m = \frac{2}{R} \quad (6)$$

$$\omega = \frac{\sqrt{3}G_m}{2C}. \quad (7)$$

According to (6) and (7), the frequency of oscillation can be linearly controlled through G_m ; however, R needs to scale inversely proportionally to G_m for oscillations to be sustained. To this end, and in order to enhance the frequency tunability range, G_m and R have been implemented as coupled translinear circuits.

Transconductance G_m is implemented by transistors Q_1 - Q_8 [Fig. 2(b)], while R is the differential resistance seen between the emitters of transistors Q_9 and Q_{10} . Linearizing the transconductances of the transistors in Fig. 2(b) and performing a small-signal analysis result in $G_m \cdot R \simeq 3$, a gain that is higher than the minimum required by (6), which guarantees the existence of oscillations. Starting from a zero initial state, the amplitude of oscillation will increase until the nonlinearities in the circuit




Fig. 2. (a) Architecture of the three-phase oscillator. (b) Circuit implementation of the G_m - R blocks.

limit the gain $G_m \cdot R$ to a value of two and the amplitude is stabilized. Assuming, for the moment, that $I_{AMP} = 0$, the values of both G_m and R are linearly proportional to I_{FREQ} , and, through (7), so is the frequency of oscillation.

Amplitude control of the oscillations is achieved by introducing current sources I_{AMP} and, thus, providing a way to adjust the $G_m \cdot R$ gain. Small-signal analysis of the translinear block in that case leads to

$$G_m \cdot R \simeq 3\kappa, \quad \kappa = \frac{I_{FREQ}}{I_{FREQ} + 2I_{AMP}}. \quad (8)$$

The Barkhausen oscillation criterion requires $G_m \cdot R \geq 2$ [see (6)], which implies $(2/3) \leq \kappa \leq 1$ and, therefore, $0 \leq I_{AMP} \leq 0.25I_{FREQ}$. Minimum oscillation amplitude is expected for a zero initial state with $G_m \cdot R = 2$, whereas maximum oscillation amplitude is achieved when the corresponding gain is $G_m \cdot R = 3$. The amplitude of oscillation can therefore be controlled through κ or, equivalently, the ratio $\gamma \triangleq (I_{AMP}/I_{FREQ})$. Oscillations of constant amplitude over a range of frequencies can be achieved by sweeping I_{FREQ} and scaling I_{AMP} proportionally so that γ is kept constant.

It is worth mentioning that, with the introduction of current sources I_{AMP} 's, G_m no longer depends linearly on I_{FREQ} , and therefore, I_{AMP} affects the oscillation frequency to second order. Detailed analysis of the amplitude dependence on I_{AMP} and its secondary effect on oscillation frequency is beyond the scope of this paper and is presented in [12].

B. Multiplier

The multiplier uses a standard Gilbert multiplier topology (Fig. 3) [15]. In order to enhance linearity, the dither input is predistorted. Further linearization, without compromising the translinear nature of the multiplier, is achieved by using multi-tanh “doublets,” i.e., setting a 1:4/4:1 emitter ratio for the input




Fig. 3. Multiplier along with the predistortion circuit for the upper input. Ratios of 1:4 and 4:1 are used for the emitter area of the input transistors to increase linearity without compromising the translinear design.




Fig. 4. Translinear amplifier with a fixed gain of six.

transistors of both the predistorter and the actual multiplier [16], [17].

C. Amplifier

Simulation of the architecture revealed that, for perturbation amplitudes α 's below a few tens of millivolts, amplification of the product between the received metric and the dither by a factor of six was adequate for the correct nonsaturating operation of the comparator. The amplifier is implemented as a cascade of three differential translinear gain stages, biased with the same current as that of the oscillator, as shown in Fig. 4.

D. Low-Pass Filter

Two competing requirements set a tradeoff in selecting the cutoff frequency of the low-pass filter. On the one hand, the cutoff frequency is directly related to the system's closed-loop bandwidth and therefore needs to be set high for fast adaptation speeds. On the other hand, a lower cutoff frequency allows for closer spacing between dither frequency bands and, thus, an increased number of control channels within a specific bandwidth. The tradeoff can be loosened using a high-order low-pass filter.




Fig. 5. (a) Fifth-order Chebyshev low-pass filter using a tunable G_m - C architecture. (b) Circuit diagram of each transconductor. Transistors Q_3 - Q_8 form a differential high-impedance load.

In the proposed architecture, a fifth-order Chebyshev topology with 1-dB ripple is implemented, offering a 30-dB/octave attenuation beyond the cutoff frequency. The low-pass filter is implemented using biquads based on G_m - C integrators [Fig. 5(a)]. The cutoff frequency of the filter is adjusted by controlling the gain of the transconductors. All transconductors have the same topology, and their gains are linearly controlled by replicas of the same current I_{LPF} . The capacitors in the design are scaled according to the fifth-order Chebyshev polynomial.

Fig. 5(b) shows the circuit implementation of transconductors G_m 's. The transconductance gain is provided from the differential input pair (Q_1 and Q_2), while transistors Q_3 - Q_8 form the load circuit. As described in [16], the load circuit provides a high-ohmic impedance for differential currents since the base-emitter voltages and, therefore, the emitter currents of Q_5 and Q_8 are equal. The translinear design of the transconductors ensures a wide tuning range for the cutoff frequency of the filter.

E. Comparator and Charge Pump

The comparator, computing the polarity of the gradient estimate used in the updates (4), is based on the architecture proposed in [18] and shown in Fig. 6. The preamplifier is implemented by a simple differential pair whose bias current I_{bias} sets the response speed of the comparison. The decision circuit is composed of a positive-feedback network (latch). Transistors M_1 and M_4 are designed to have the same dimensions as transistors M_2 and M_3 , respectively, so as to minimize hysteresis. The output buffer converts the differential signal from the decision circuit to a single-ended output. Owing to the fully symmetrical design of the topology, the offset in the comparison is low and affected only by mismatches in the fabrication process.

The charge pump, providing constant-modulus updates of either polarity in the control variable u_i according to (4), is implemented using the design in [7] and shown in Fig. 7(a). The rate G at which capacitor C is charged or discharged (i.e., u_i incremented or decremented) can independently be controlled by separate biasing of the NMOS and PMOS current sources. The corresponding biasing circuits are shown in Fig. 7(b).




Fig. 6. Circuit diagram of the comparator.




Fig. 7. (a) Charge pump with individual biases for the up and down rates. (b) Biasing circuits.

IV. PERFORMANCE ANALYSIS

A. Phase-Noise Immunity

The wide frequency tuning range of the oscillator through a single bias control is bound to result in high phase noise. This is true, first, because the frequency of oscillation is sensitive to any change mainly in I_{FREQ} and secondarily in I_{AMP} . Particularly for lower frequencies, where these bias currents are small, noise in the I_{FREQ} control will add considerable FM noise to the oscillator's output signal, whereas noise in the I_{AMP} control will translate to AM noise. Second, due to the relatively large number of transistors in the design, flicker ($1/f$) noise is non-negligible. Finally, G_m - C architectures have low Q [16] and, therefore, limited capability in filtering the generated AM, FM, and flicker noise.

The closed-loop operation of the proposed architecture, however, is only weakly dependent on phase noise; the benefit of using synchronous detection for retrieving the gradient information is its property of filtering out phase noise. To clarify this point, synchronous detection can be thought of as an FM discriminator and, more specifically, as a delay-line discriminator [19], [20] [Fig. 8(a) and (b)].

For constant delay τ_p , the power spectral density (PSD) of the noise $S_{\phi, out}(f)$ at the output of a delay-line discriminator has been shown [21]-[23] to be

$$S_{\phi, out}(f) = S_d(f) \cdot S_{\phi}(f) = 2(1 - \cos(2\pi f\tau_p)) \cdot S_{\phi}(f)$$

where $S_{\phi}(f)$ is the PSD of the phase noise of the oscillator and f is the frequency offset from the carrier. For small delays τ_p 's and for small frequency offsets f 's, the scaling factor $S_d(f) = 2(1 - \cos(2\pi f\tau_p)) \approx 4\pi^2(f\tau_p)^2$ is close to zero,




Fig. 8. (a) Simplified model of the synchronous detection scheme for the closed-loop operation of the proposed architecture. (b) Simplified model of the delay-line discriminator.

eliminating phase-noise power that is close to the carrier that could downgrade the output from synchronous detection. Intuitively, synchronous detection acts as a high-pass filter on the oscillator's phase noise.

Thus, phase noise is not a detrimental factor in the performance of the overall system, and we have not taken any measures to minimize it. In fact, some amount of phase noise is *desirable* for parallel synchronous detection to avoid coupling between control channels due to possible harmonic relationships between dither frequencies. Phase noise contributes to making the dither signals for each of the control channels orthogonal, so that the gradient estimates are unbiased.²

B. Linearity of the Oscillator

To assess the effect of dither harmonics on the accuracy of gradient estimation, consider the system architecture in Fig. 1. If the dither signals are pure sinusoids of frequencies ω_i , $i = 1, \dots, n$, and of relatively small amplitudes, then the interference between the channels is limited as long as $|\omega_i - \omega_j| \forall i \neq j$ exceeds the cutoff frequencies of the corresponding filters. In this case, the interchannel interference is only due to the nonlinearities of the cost metric J and the multipliers.

In practice, harmonic components always exist, and careful assignment of the dither frequencies is needed. The higher the linearity of the oscillator, the more arbitrarily these frequencies can be chosen. On the other hand, by appropriate selection of ω_i 's, such that no intermodulation product falls within the bandwidth of the filter, any effect of harmonics in the multivariable detection scheme could be avoided. In principle, even square dithers can be used as long as due care is taken in the selection of the dither frequencies, although the dither frequency constraints do not facilitate a large number of control channels and impose severe bandwidth constraints on the available adaptation bandwidth.

²In the limit of very large phase noise, the dithers become broad-band noise signals, and the MD optimization reduces to SPGD [2]–[5]. Practical levels of phase noise retain the narrow-band frequency property of the dither signals, which is necessary for delay-insensitive MFA [14], [24].




Fig. 9. Limit-cycle fluctuations at steady state for single-parameter adaptation with constant-modulus updates.

C. Convergence

The hard-limiting nature of the comparator in the loop implies a steady-state, possibly aperiodic, oscillatory pattern in the control parameters u_i 's after convergence. In the case of a single control parameter, the steady-state oscillations are periodic and take the form of a limit cycle of alternating increments and decrements in u at a frequency determined by the group delay of the adaptation loop τ_{LC} , as shown in Fig. 9. The triangular pattern of the limit cycle in the control variable u induces a near-sinusoidal pattern after synchronous detection at the output of the low-pass filter. In turn, the square form of the quantized output generates charge-pump updates that sustain, after integration onto the capacitor, the triangular pattern of fluctuations u 's around the optimum value of the control parameter. Note that the frequency of the limit cycle, $1/(4\tau_{LC})$, is at most equal to the cutoff frequency of the low-pass filter [and, therefore, much lower than that of the sinusoidal dithers $\delta u_i = \cos(\omega_i t)$] since, otherwise, limit cycles would not be sustained.

The frequency of the limit cycle is directly related to the group delay of the adaptation loop and, therefore, to parameters, such as the delay introduced by the unknown plant, the metric estimation and feedback propagation, as well as the circuit blocks of the controller, such as the low-pass filter. On the other hand, the amplitude A_{LC} of the limit cycle depends on both the frequency of the limit cycle and the gain G in the update rule of the gradient-descent algorithm. Finally, the settling time t_s for the convergence of a control parameter to its optimum value is proportional to G . For a given group delay (directly related to τ_{LC}), the tradeoff between settling time and amplitude of the limit cycle is apparent; higher gain G leads to faster convergence but also higher fluctuation around the optimum value and vice versa.

For the general case of n control parameters, the general observation that amplitude A_{LC} and gain G are directly related through the τ_{LC} delay is still valid. The behavior of the parameters at steady state can be described as coupled bounded-amplitude oscillations around the optimal values. Analysis of the dynamics of the steady-state oscillatory pattern in the multidimensional case is beyond the scope of this paper.

D. Power Efficiency

The translinear bipolar-junction-transistor circuit design throughout the major parts of the architecture also offers, besides wideband operation with current bias control over several decades in frequency, bias-invariant power efficiency

in that power consumption scales with the adaptation speed. Power calculation can be performed by taking into account the G_m - C nature of the circuits and noting that the bandwidth of each stage is proportional to

$$f_{\text{BW}} = \frac{G_m}{C_{\text{eff}}} \propto \frac{I_{\text{bias}}}{C_{\text{eff}}} \quad (9)$$

where C_{eff} is the overall capacitance at the output node of the stage. According to (9), power consumption for each stage, which is proportional to the bias current $P \propto I_{\text{bias}} \cdot V_{\text{supply}}$, scales directly with the bandwidth f_{BW} of each stage.

In a typical scenario of selecting the bandwidth and, therefore, the bias currents for each stage of the system, one starts by determining the bandwidth of the cost function J , $f_{-3 \text{ dB}, J}$. Accordingly the (radial) dither frequencies ω_i 's (where $\omega_i = 2\pi f_i$) for each channel are set below $2\pi f_{-3 \text{ dB}, J}$. The optimal choice of the dither frequencies is a uniform distribution between a user-selectable lowest frequency ω_1 and $\omega_n \lesssim 2\pi f_{-3 \text{ dB}, J}$. Added convenience is provided by setting $\Delta\omega/\omega_{\text{av}} = 1/k = \text{constant}$, where $\Delta\omega = |\omega_i - \omega_j| \forall i \neq j = 1, \dots, n$ and $\omega_{\text{av}} = (1/n) \sum_{i=1}^n \omega_i$ so that the selection of dither frequencies scales in a standard way for any $f_{-3 \text{ dB}, J}$. Under these conditions, the mean power dissipated from the oscillator stage can be written as (see also Section III-A)

$$P_{\text{osc,av}} = m_{\text{osc}} \cdot I_{\text{FREQ,av}} \cdot V_{\text{supply}} \propto \omega_{\text{av}} \cdot V_{\text{supply}} \quad (10)$$

where m_{osc} is a factor representing the multiplicity of current $I_{\text{FREQ,av}}$ in the oscillator stage. More specifically, assuming that I_{AMP} scales proportionally to I_{FREQ} with a factor γ for all channels, and referring to Fig. 2(a) and (b), we conclude that $m_{\text{osc}} = 6(1 + \gamma)$.

The bias current of both the multiplier and the amplifier, for each channel i , is set equal to $I_{\text{FREQ},i}$ so that, according to (9), their bandwidth scales proportionally to frequency ω_i . Therefore, the average power consumed at these stages also scales according to ω_{av} and can be written as

$$P_{\text{mix-amp,av}} = (m_{\text{mix}} + m_{\text{amp}}) \cdot I_{\text{FREQ,av}} \cdot V_{\text{supply}} \quad (11)$$

where $m_{\text{mix}} = 4$ for the design of the multiplier in Fig. 3 and $m_{\text{amp}} = 3$ for the design of the amplifier in Fig. 4.

The cutoff frequency of the low-pass filter needs to be set lower than $\Delta\omega/2$ so that all intermodulation products are attenuated by at least 30 dB. Since the bandwidth of the filter for each channel is set through bias current I_{LPF} (see Section III-B), the power P_{LPF} consumed at this stage is $P_{\text{LPF}} = m_{\text{LPF}} \cdot I_{\text{LPF}} \cdot V_{\text{supply}}$ and scales according to $\Delta\omega$. Assuming that $\omega_{\text{av}} = k \cdot \Delta\omega$, P_{LPF} can be rewritten as

$$P_{\text{LPF}} = \frac{m_{\text{LPF}}}{\eta \cdot k} \cdot I_{\text{FREQ,av}} \cdot V_{\text{supply}} \quad (12)$$

where $I_{\text{FREQ,av}} = k \cdot I_{\text{LPF}}$ can be assumed because of the linear relationship between $\Delta\omega$, ω_{av} and I_{LPF} , $I_{\text{FREQ,av}}$, respectively. Parameter $\eta \geq 2$ represents the ratio between $\Delta\omega$ and the (radial) cutoff frequency of the filter.

Strictly speaking, the linear relation between power, bandwidth, and oscillator bias current is not valid for the entire MD control architecture since some of the building blocks, in particular the comparators and the output buffers, are implemented

using MOSFETS biased in the above threshold regime, where the MOSFET transconductance exhibits a square-root dependence on bias current. Nonetheless, overall, the bandwidth scales approximately linearly with the oscillator bias, and thus, the adaptation speed scales roughly linearly with power.

E. FOM

In order to compare the performance of the presented architecture with adaptive systems previously presented in the literature, the following figure of merit (FOM) is proposed here:

$$FOM = \frac{BW \cdot SNR \cdot n}{P} \quad (13)$$

where BW is the adaptation bandwidth that is equivalent to the inverse of settling time t_s (Fig. 9), $SNR = A_{\text{max}}/A_{\text{LC}}$ is the signal-to-noise ratio between the desired transition voltage A_{max} and the amplitude A_{LC} of the steady-state (limit-cycle or aperiodic) oscillations due to the signed constant-modulus updates (Fig. 9), n is the number of channels for which adaptation is demonstrated, and P is the total dissipated power.

From the aforementioned definition of bandwidth and referring to Fig. 9, parameter BW can be rewritten as $BW = G/A_{\text{max}}$, in which case the product $BW \cdot SNR$ collapses to

$$BW \cdot SNR = \frac{G}{A_{\text{max}}} \cdot \frac{A_{\text{max}}}{A_{\text{LC}}} = \frac{1}{\tau_{\text{LC}}}$$

For comparison purposes, it is worthwhile to contrast the constant-modulus updates (4) to adaptive circuits and systems that implement a linear unthresholded form of gradient updates. For the case of linear updates, where convergence is inversely exponential and no oscillatory behavior is observed, the FOM definition of (13) can directly be evaluated by expressing SNR explicitly and by expressing BW as the inverse of settling time. Thus, at $SNR = 1$, the FOM for systems with linear and constant-modulus updates can directly be compared by comparing the settling time t_s of the linear system with the delay τ_{LC} for the constant-modulus system [or a quarter of the period of the limit cycles (see Fig. 9)]. For higher values of SNR, values of t_s should be accordingly scaled in comparison with τ_{LC} .

V. EXPERIMENTAL MEASUREMENTS

The eight-channel MD controller was implemented on a 3 mm \times 3 mm chip fabricated in 0.5- μm SiGe BiCMOS technology. A micrograph of the chip and a detailed view of one of the eight channels are shown in Fig. 10(a) and (b), respectively.

Fig. 11(a) shows the linear operating frequency range of the oscillator, extending to more than five decades (from below 4 kHz to above 600 MHz) and the linear dependence of frequency on biasing current I_{FREQ} . The frequency and amplitude of the signal were measured for three different values of ratio γ . As can be seen in Fig. 11(a), for a given value of I_{FREQ} , the amplitude of oscillation can be controlled by the value of γ . For $\gamma = 0$, a maximum differential amplitude of 60 mV_{pp} (corresponding to a single-ended measurement of 30 mV_{pp} or ~ -30 dBm on a 50- Ω load) is reached for most of the frequency range. The dependence of oscillation frequency with




Fig. 10. (a) Chip microphotograph and (b) single-channel detail. The dimensions of the chip are 3 mm \times 3 mm, and the cell measures 500 μ m \times 900 μ m in 0.5- μ m SiGe BiCMOS technology.




Fig. 11. (a) Oscillation frequency and amplitude with respect to current I_{FREQ} recorded for three different values of γ . The operating frequencies range from below 4 kHz to above 600 MHz. (b) Spectrum of oscillator's output measured for $I_{\text{FREQ}} \approx 40 \mu\text{A}$ and $\gamma = 0$. The amplitude of oscillation from the single-ended output is $\approx 20 \text{ mV}_{\text{pp}}$.

γ (and, therefore, I_{AMP}), as discussed in Section III-A, can also be observed.

Fig. 11(b) shows the spectral content of the oscillator output signal for $I_{\text{FREQ}} \approx 40 \mu\text{A}$ and $\gamma = 0$. The oscillator signal is observed at the control output u_i for one of the channels. This output conveys an ac-coupled single-ended version of the internally differential oscillator signal. The effect of harmonic distortion on adaptation performance is negligible as long as the dithers of different channels are not harmonically related




Fig. 12. Transfer function of the filter measured for four different values of bias I_{LPF} (a) 66.2 nA. (b) 814 nA. (c) 10.6 μA . (d) 128 μA . The achieved cutoff frequencies are 200 kHz, 2 MHz, 20 MHz, and 200 MHz, respectively.

(Section IV-B). Since synchronous detection makes use of differential signals (for one of the oscillator phases), only the odd-order harmonics are relevant to assess linearity in synchronous detection. For other applications of the very wide frequency range tunable oscillator where harmonic distortion may be an issue, differential measurements directly from an isolated oscillator circuit are included in [12].

The transfer function of the low-pass filter was measured for four different values of I_{LPF} , and the results are shown in Fig. 12. For biases of 66.2 nA, 814 nA, 10.6 μA , and 128 μA , the measured cutoff frequencies correspond to 200 kHz, 2 MHz, 20 MHz, and 200 MHz, respectively, spanning a tuning range of at least four decades. Measurements of the filter transfer function for lower values of I_{LPF} were limited by the operating range of our measuring equipment. Moreover, the seemingly lower stop-band rejection at higher cutoff frequencies is due to higher resolution and video bandwidths in the spectrum analyzer used to acquire the data, which, in turn, increased the noise floor in the measurements. Note, finally, the linear relation between the control bias and the corresponding bandwidth of the filter.

To demonstrate synchronous detection performance, four channels were perturbed with oscillation frequencies of 97, 122, 139, and 160 MHz. Their outputs were summed, using an RF power combiner, resulting in a multitone signal whose spectrum is shown in Fig. 13(a). The combined signal was fed back, single-ended, to the metric input of the chip. Fig. 13(b) shows the spectrum after multiplication with the 139-MHz signal but before any filtering. The expected products around 20 and 40 MHz are present. Fig. 13(c) shows the spectrum at the output of the low-pass filter following the multiplier. The cutoff frequency has been set to approximately 10 MHz. Components above that frequency are significantly attenuated down to the noise floor. Comparison of Fig. 13(b) and (c) shows also the expected 30-dB/octave attenuation of the signal at 20 MHz due to the fifth-order filter. Note also the sharp peak of the fundamental synchronous detection component at dc frequency, in contrast to the significantly wider spurs. This is consistent with the observation on the immunity of synchronous detection to phase noise, as noted in Section IV.




Fig. 13. (a) Spectrum of the combination of four channel outputs with channel dither frequencies at 97, 122, 139, and 160 MHz. The combined signal is fed to the input of the chip. (b) Recorded spectrum after multiplying the multitone signal with the 139-MHz dither. (c) Recorded spectrum after low-pass filtering of the multipliers' output (cutoff at 10 MHz).




Fig. 14. Synchronous detection output measured with different phase selections for a metric containing a single delayed dither.

The influence of metric delay on gradient estimation has been studied, and the effect of phase selection on the phase error of the estimate has been validated based on the following experiment. The loop was closed by connecting the output of a single channel to the metric input using a transmission line of considerable time delay τ . In this case, the metric is $f(x(t)) = x(t - \tau)$, and the adaptation rule dictates a continuous increase in the control voltage of the channel until saturation. Fig. 14 shows the change in filter output voltage versus dither frequency due to variable phase delay, introduced by the delay line, for each of the three oscillator phases selected as inputs to synchronous detection. The three curves consistently show the 120° separation in phase according to the three oscillator phases. For any given phase selection, the delay of the line causes the sign of the output to change with frequency. However, for any given frequency, at least one of the three phases produces the correct sign of the derivative estimate. Correct phase selection is critical for convergence at high dither-frequency \times time-delay products [24], and an algorithm for adaptively selecting the optimal phase is given in [14].

The relation between the control biases V_{bp} and V_{bn} of the charge pump and the corresponding update rates was characterized by the measurements shown in Fig. 15. The uprate corresponds to how fast the capacitor of the charge pump is charged through the PMOS current source and is controlled by V_{bp} . The downrate corresponds to the discharge of the capacitor through




Fig. 15. Update rates G 's by the charge pump, measured as a function of voltage biases V_{down} (V_{bn}) and V_{up} (V_{bp}).

the current sink of the NMOS and is controlled by V_{bn} . Through appropriate biasing, the update rates can take values from below $1 \text{ V}/\mu\text{s}$ to above $10^6 \text{ V}/\mu\text{s}$.

The closed-loop performance of the system was evaluated using an external circuit that serves as an analog plant with a characteristically nonlinear min-max cost metric. The simplified schematic of the external plant, implemented with resistors and diodes, is shown in Fig. 16. The differential output $V_{max} - V_{min}$ realizes approximately the metric $f(V_1, \dots, V_n, V_{ref}) = \max(V_1, \dots, V_n, V_{ref}) - \min(V_1, \dots, V_n, V_{ref}) - 2V_F$, where $V_i > 0$, $i = 1, \dots, n$, are the voltage outputs from the n channels of the system, $V_{ref} > 0$ is a reference voltage provided by a function generator, and V_F is the forward voltage drop of the used diodes. The metric f has a global minimum that is equal to $-2V_F$, which is reached when $V_1 = V_2 = \dots = V_n = V_{ref}$. Fig. 17 shows how the outputs from three channels (V_1 , V_2 , and V_3) adapt to a 100-kHz triangular reference voltage V_{ref} of 500-mV_{pp} amplitude, minimizing the output of the metric. The dithers of the three channels were set at 90, 120, and 150 MHz.

To test the adaptation speed of the system, the same diode-based external circuit supplying the min-max metric was used, but the reference was set to a 500-mV_{pp} square wave. The gain of the charge pump was adjusted so as to achieve fast adaptation with tolerable limit cycles. Fig. 18 shows results for the case of a single-channel control voltage perturbed at 90 MHz. At each transition of the square wave, tracking is momentarily lost (peaks in the metric) before the channel output signal reaches




Fig. 16. Experimental setup for the characterization of closed-loop controller dynamics. A metric function $f(V_1, \dots, V_n, V_{ref}) = \max(V_1, \dots, V_n, V_{ref}) - \min(V_1, \dots, V_n, V_{ref}) - 2V_F$ is provided by an external analog plant, which is implemented using diodes and resistors.




Fig. 17. Experimental adaptation of three channels to a 100-kHz triangular reference waveform. Exact tracking of the reference indicates minimization of the metric function. Dithers are set at 90, 120, and 150 MHz.

the new level of the reference voltage. The minimum value of the metric is -650 mV and very close to $-2V_F$ ($\simeq -700$ mV). Adaptation is achieved in less than $1 \mu s$.

The experimental results of closed-loop adaptation shown here are typical and are consistent across several channels across several chips. Because of differences in dither frequencies and corresponding phase differences in the control loop, it was necessary to individually adjust the phase parameters for each channel. A practical procedure for selecting the phase parameters is given and demonstrated in [14]. No other parameters required tuning, although we observed that for certain control metrics it was necessary to amplify the external metric signal, so that the gains $\partial J / \partial u_i$ exceed a given threshold, in order for all channels to lock into closed-loop adaptation. We traced this sensitivity in the amplitude of the control metric to analog mismatch in the circuit realization of the comparator in the synchronous detector (Fig. 6). An offset-independent




Fig. 18. Single-channel control tracking a 100-kHz square waveform. With dither set at 90 MHz, adaptation is achieved in less than $1 \mu s$.

TABLE I
CHIP DESIGN PARAMETERS AND MEASURED PERFORMANCE

Process	SiGe BiCMOS 0.5 μ m
Area	3mm \times 3mm
Channels	8
Linear Dither Frequency Range	4kHz - 600MHz
Update Rate	0.4 - $4 \cdot 10^6$ V/s
Adaptation Speed	< $1 \mu s$
Power for dithers at 12MHz-20MHz	50mW
Power for dithers at 120MHz-200MHz	110mW

autoadaptive comparator that alleviates this sensitivity is described in [32], whereas an alternative approach could be a wideband comparator with adaptable offset, as presented in [33].

The measured power dissipation for simultaneous operation of all eight channels ranges between 50 mW for dither frequencies in the range of 12–20 MHz and 110 mW for dither frequencies in the range of 120–200 MHz. The design parameters and measured performance of the chip are summarized in Table I.

Finally, a comparison between the proposed architecture and other adaptive controller implementations reported in the literature was performed. The comparison included recently presented adaptive systems, both analog and digital, for which chip data were provided, without any constraints on the application for which they were used for. The lack of SNR information for most references prohibited the full evaluation of the proposed FOM in Section IV-E, and therefore, comparison was limited to the three other design variables comprising the FOM: the number of controlled parameters n , the power consumed by the controller P , and the lowest reported adaptation time t_s . The collected data are shown in Table II. A visual representation of the Table II data is shown in Fig. 19, where the horizontal axis represents power per control channel P/n and the vertical axis denotes the minimum reported adaptation time t_s . On the logarithmic display, the dashed lines in the graph indicate equal values of FOM for a given SNR.

VI. CONCLUSION

A VLSI implementation of a model-free architecture for adaptive control, using narrow-band MD gradient descent, has been presented. A fully translinear implementation using

TABLE II
COMPARISON BETWEEN HIGH-SPEED VLSI ADAPTIVE CONTROLLERS

Reference	[25]	[26]	[27]	[28]	[29]	[30]	[31]	This work
Process	2 μ m CMOS	0.4 μ m CMOS	0.13 μ m CMOS	0.13 μ m CMOS	0.25 μ m CMOS	0.25 μ m CMOS	IBM SiGe	0.5 μ m SiGe BiC-MOS
Parameters	4	4	4	4	2	1	4	8
Power (mW)	236	43	100	N/A	43.2	5.7	N/A	110
Adaptation time (μ s)	20	10	25	20	25	80	100	1




Fig. 19. Comparison of the presented work with previous reported adaptive controller implementations in terms of settling time and consumed power per channel. The lines shown in the graph indicate equal values of FOM for a given SNR.

a SiGe BiCMOS process has provided linear tunability over five decades of frequency. Characterization and closed-loop experiments have demonstrated synchronous detection and metric adaptation with up to 200-MHz dither frequencies at 110-mW power consumption.

ACKNOWLEDGMENT

The authors would like to thank R. Philipp and A. Çelik for the valuable help with the layout and M. Vorontsov for the support and for stimulating discussions. Chip fabrication was provided through MOSIS.

REFERENCES

- [1] A. Dembo and T. Kailath, "Model-free distributed learning," *IEEE Trans. Neural Netw.*, vol. 1, no. 1, pp. 58–70, Mar. 1990.
- [2] J. Alspector, R. Meir, B. Yuhas, A. Jayakumar, and D. Lippe, "A parallel gradient descent method for learning in analog VLSI neural networks," in *Proc. Adv. NIPS*, San Mateo, CA, 1993, vol. 5, pp. 836–844.
- [3] G. Cauwenberghs, "A fast stochastic error-descent algorithm for supervised learning and optimization," in *Proc. Adv. NIPS*, San Mateo, CA, 1993, vol. 5, pp. 244–251.
- [4] M. A. Vorontsov, G. W. Carhart, and J. C. Ricklin, "Adaptive phase-distortion correction based on parallel gradient-descent optimization," *Opt. Lett.*, vol. 22, no. 12, pp. 907–909, Jun. 1997.
- [5] A. C. Carusone and D. A. Johns, "Obtaining digital gradient signals for analog adaptive filters," in *Proc. ISCAS*, May 1999, pp. 54–57.
- [6] D. B. Kirk, D. Kerns, K. Fleischer, and A. H. Barr, "Analog VLSI implementation of multi-dimensional gradient descent," in *Proc. Adv. NIPS*, San Mateo, CA, 1993, vol. 5, pp. 789–796.
- [7] G. Cauwenberghs, "Analog VLSI stochastic perturbative learning architectures," *Int. J. Analog Integr. Circuits Signal Process.*, vol. 13, no. 1/2, pp. 195–209, May/Jun. 1997.
- [8] A. C. Carusone and D. A. Johns, "Digital LMS adaptation of analog filters without gradient information," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 9, pp. 539–552, Sep. 2003.
- [9] T. R. O'Meara, "The multidither principle in adaptive optics," *J. Opt. Soc. Amer.*, vol. 67, no. 3, pp. 306–315, Mar. 1977.
- [10] M. A. Vorontsov and S. L. Lachinova, "Laser beam projection with adaptive array of fiber collimators. I. Basic considerations for analysis," *J. Opt. Soc. Amer. A, Opt. Image Sci.*, vol. 25, no. 8, pp. 1949–1959, Jul. 2008.
- [11] M. A. A. Neil *et al.*, "Adaptive aberration correction in a two-photon microscope," *J. Microsc.*, vol. 200, no. 2, pp. 105–108, Nov. 2000.
- [12] D. N. Loizos, P. P. Sotiriadis, and G. Cauwenberghs, "A translinear SiGe BiCMOS current-controlled oscillator with 80 Hz–800 MHz tuning range," *Int. J. Analog Integr. Circuits Signal Process.*, vol. 57, no. 1/2, pp. 107–115, Nov. 2008.
- [13] D. N. Loizos, P. Sotiriadis, and G. Cauwenberghs, "High-speed, model-free adaptive control using parallel synchronous detection," in *Proc. IEEE 20th SBCCI*, 2007, pp. 224–229.
- [14] D. N. Loizos, P. P. Sotiriadis, and G. Cauwenberghs, "Adaptive delay compensation in multi-dithering adaptive control," in *Proc. ISCAS*, May 2008, pp. 2182–2185.
- [15] B. Gilbert, "A precise four-quadrant multiplier with subnanosecond response," *IEEE J. Solid-State Circuits*, vol. SSC-3, no. 4, pp. 365–373, Dec. 1968.
- [16] H. Voorman and H. Veenstra, "Tunable high-frequency G_m-C filters," *IEEE J. Solid-State Circuits*, vol. 35, no. 8, pp. 1097–1108, Aug. 2000.
- [17] B. Gilbert, "The multi-tanh principle: A tutorial overview," *IEEE J. Solid-State Circuits*, vol. 33, no. 1, pp. 2–17, Jan. 1998.
- [18] R. J. Baker, H. W. Li, and D. E. Boyce, *CMOS Circuit Design, Layout and Simulation*. Englewood Cliffs, NJ: Prentice-Hall, 1998, pp. 687–699.
- [19] "Understanding and Measuring Phase Noise in the Frequency Domain," Hewlett-Packard, Austin, TX, Oct. 1976, Application Note 207.
- [20] A. L. Lance, W. D. Seal, and F. Labaar, "Phase noise and AM noise measurements in the frequency domain," in *Infrared and Millimeter Waves*, K. J. Button, Ed. New York: Academic, 1984, vol. 11, pp. 239–289.
- [21] D. Halford, The delay line discriminator Nat. Bureau Standards, Gaithersburg, MD, Technical Note 10, 1976, pp. 19–38.
- [22] J. E. Reese and W. S. Ishak, "Automatic phase noise measurements using MSW-delay-line-based discriminators," in *Proc. IEEE Ultrason. Symp.*, 1985, pp. 169–173.
- [23] D. N. Loizos, "A multi-dithering controller with sub-microsecond response time: Analysis and SiGe BiCMOS implementation," Ph.D. dissertation, The Johns Hopkins Univ. Press, Baltimore, MD, 2007.
- [24] D. N. Loizos, P. P. Sotiriadis, and G. Cauwenberghs, "Multi-channel coherent detection for delay-insensitive model-free adaptive control," in *Proc. ISCAS*, May 2007, pp. 1775–1778.
- [25] F. J. Kub and E. W. Justh, "Analog CMOS implementation of high frequency least-mean square error learning circuit," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1391–1398, Dec. 1995.
- [26] T.-C. Lee and B. Razavi, "A 125-MHz mixed-signal echo canceller for gigabit ethernet on copper wire," *IEEE J. Solid-State Circuits*, vol. 36, no. 3, pp. 366–373, Mar. 2001.
- [27] J. E. Jaussi *et al.*, "8-Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation, and clock de-skew," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 80–88, Jan. 2005.
- [28] R. Payne *et al.*, "A 6.25 Gb/s binary adaptive DFE with first post-cursor tap cancellation for serial backplane communications," in *Proc. IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 68–585.

- [29] V. Aparin, G. J. Ballantyne, C. J. Persico, and A. Cicalini, "An integrated LMS adaptive filter of TX leakage for CDMA receiver front ends," *IEEE J. Solid-State Circuits*, vol. 41, no. 5, pp. 1171–1182, May 2006.
- [30] J. Kim and M. A. Horowitz, "An efficient digital sliding controller for adaptive power-supply regulation," *IEEE J. Solid-State Circuits*, vol. 37, no. 5, pp. 639–647, May 2002.
- [31] B. Wedding, A. Chiarotto, W. Kuebart, and H. Bülow, "Fast adaptive control for electronic equalization of PMD," in *Proc. Opt. Fiber Commun. Conf. Exhib.*, 2001, vol. 2, pp. TuP4-1–TuP4-3.
- [32] N. Kumar, G. Cauwenberghs, and A. Andreou, "Auditory feature extraction using self-timed, continuous-time discrete-signal processing circuits," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 44, no. 9, pp. 723–728, Sep. 1997.
- [33] Y. L. Wong, M. H. Cohen, and P. A. Abshire, "A 1.2-GHz comparator with adaptable offset in 0.35- μm CMOS," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 9, pp. 2584–2594, Oct. 2008.



Dimitrios N. Loizos (S'06–M'08) received the Diploma in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2003 and the M.Sc.E. and Ph.D. degrees in electrical and computer engineering from the Johns Hopkins University, Baltimore, MD, in 2005 and 2007, respectively.

He then joined the Division of Biological Sciences, University of California at San Diego, La Jolla, as a Postdoctoral Fellow. Since 2008, he has been with NetLogic Microsystems Inc., Mountain

View, CA. His research interests include model-free optimization techniques and their VLSI implementation, analog and RF IC design for high-speed wire-line transceivers, RF and microwave design, as well as monolithic-microwave integrated-circuit design.

Dr. Loizos was the recipient of the Best Paper Award in the IEEE Symposium on Integrated Circuits and Systems Design 2007, as well as the third place for the Best Student Paper in the IEEE International Symposium on Circuits and Systems 2007.



Paul P. Sotiriadis (S'99–M'02–SM'09) received the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2002, the M.S. degree in electrical engineering from Stanford University, Stanford, CA, in 1996, and the Diploma in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1994.

In 2002, he joined the Johns Hopkins University, Baltimore, MD, as an Assistant Professor of electrical and computer engineering. In 2007, he joined Apex/

Eclipse Inc. as the Chief Technology Officer, and shortly after that, he started Sotekco Electronics LLC, Baltimore, MD, an electronics research company. His research interests include the design, optimization, and mathematical modeling of analog- and mixed-signal circuits, RF and microwave circuits, advanced frequency synthesis, biomedical instrumentation, and interconnect networks in deep-submicrometer technologies. He has led several projects in these fields funded by U.S. organizations and has collaborations with industry and national laboratories. He has authored and coauthored over 70 technical papers in IEEE journals and conferences, most of them as the leading or the single author, has contributed chapters to technical books, is a holder of one patent and several patents pending.

Dr. Sotiriadis serves as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II. He is a member of several technical and conference committees. He regularly reviews for many IEEE transactions and conferences. He also regularly serves on proposal review panels at the National Science Foundation.



Gert Cauwenberghs (S'89–M'94–SM'04) received the M.Eng. degree in applied physics from the University of Brussels, Brussels, Belgium, in 1988 and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1989 and 1994, respectively.

He is currently a Professor of bioengineering with the University of California at San Diego, La Jolla, where he co-directs the Institute for Neural Computation. Previously, he held positions as a Professor of electrical and computer engineering with The Johns

Hopkins University, Baltimore, MD, and as a Visiting Professor of brain and cognitive science with the Massachusetts Institute of Technology, Cambridge. His research interests are in advancing silicon adaptive microsystems to the understanding of biological neural systems and to the development of sensory and neural prostheses and brain–machine interfaces. He pioneered the design and implementation of highly energy-efficient massively parallel microchips that emulate the function and structure of adaptive neural circuits in silicon. His previous contributions include micropower parallel silicon support vector machines for real-time adaptive pattern recognition, and acoustic microarrays for auditory separation and localization.

Dr. Cauwenberghs is a Francqui Fellow of the Belgian American Educational Foundation and was the recipient of the National Science Foundation Career Award in 1997, the Office of Naval Research Young Investigator Award in 1999, and the Presidential Early Career Award for Scientists and Engineers in 2000. He was a Distinguished Lecturer of the IEEE Circuits and Systems Society in 2003–2004 and was the Chair of the Analog Signal Processing Technical Committee in 2001–2002. He is currently an Associate Editor for the IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS and the IEEE TRANSACTIONS ON NEURAL SYSTEMS AND REHABILITATION ENGINEERING. He is also a Senior Editor for the IEEE SENSORS JOURNAL.