Statistical Design and Optimization for Adaptive Post-silicon Tuning of MEMS Filters

Fa Wang, Gokce Keskin, Andrew Phelps, Jonathan Rotner, Xin Li, Gary K. Fedder, Tamal Mukherjee and Lawrence T. Pileggi
Carnegie Mellon University, Pittsburgh, PA 15213
{fwang1, gkeskin, rphelps, jrotner, xinli, fedder, tamal, pileggi}@ece.cmu.edu

ABSTRACT

Large-scale process variations can significantly limit the practical utility of microelectro-mechanical systems (MEMS) for RF (radio frequency) applications. In this paper we describe a novel technique of adaptive post-silicon tuning to reliably design MEMS filters that are robust to process variations. Our key idea is to implement a number of redundant MEMS resonators to form an array and then optimally select a subset of these resonators to achieve the desired frequency response. Several new CAD algorithms and methodologies are proposed to optimize and configure the design variables of the proposed MEMS resonator array. A MEMS design example demonstrates that the proposed post-silicon tuning is able to reduce the ripple of the channel filter gain by 7× over other traditional approaches.

Categories and Subject Descriptors

B.7.2 [**Integrated Circuits**]: Design Aids – Verification

General Terms

Algorithms

Keywords

MEMS Filter, Process Variation

1. INTRODUCTION

Advanced radio systems (e.g., software defined radio, cognitive radio, etc.) are of great importance for both commercial and military applications [1]. Designing and manufacturing fully-integrated on-chip radio systems is extremely important to achieve low power consumption and small system size for high-speed wireless communication. This goal, however, is challenging due to the lack of channel filters that can be fully integrated with RF transceivers and their signal processing circuits.

MEMS resonator filters hold great potential for the aforementioned on-chip integration problem [2]-[6], [14]. MEMS resonators made in low-loss materials (e.g., silicon, polysilicon, aluminum nitride, etc.) have been demonstrated to achieve extremely high quality factors (e.g., 3000 ~ 80000). Most importantly, MEMS resonators can be designed with their center frequencies set by layout, thereby enabling on-chip integration of channel filters required by today's radio systems.

While the research and development of MEMS filters has made significant progress in recent years, process variations posed by nanoscale IC technology remain the key bottleneck that limits the practical applications of these MEMS devices. For example, since the quality factor (i.e., Q) of a MEMS resonator is extremely high, a small variation (e.g., less than 1%) of the center frequency can completely change the frequency response of the resonator and make the MEMS filter ineffective for the intended application.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
DAC 2012, June 3-7, 2012, San Francisco, California, USA.
Copyright 2012 ACM 978-1-4503-1199-1/12/06...$10.00

In prior art, mechanically-coupled MEMS resonators have been proposed to create a tunable resonator array and, hence, achieve the desired frequency response [6]. However, mechanical coupling cannot offer a large tuning range that is sufficient to accommodate the large-scale process variations from run-to-run, lot-to-lot, wafer-to-wafer and die-to-die. In addition, mechanically-coupled MEMS resonators often generate spurious modes out of the intended pass-band, thereby distorting the frequency response of the filter. The technical challenge here is how to develop a new design methodology to make a MEMS filter flexibly tunable and, consequently, match the desired frequency characteristics.

In this paper, we borrow the concept of *statistical element selection* [7]-[8] to develop a novel post-silicon tuning methodology for MEMS filters. Namely, we propose to place a large number (i.e., N) of resonators on a die and select a subset (i.e., M where $M < N$) of these resonators to build a MEMS filter. Other resonators that are not selected are simply turned off. The frequency response of the MEMS filter is equal to the summation of the frequency responses of the M selected resonators. Based on the post-silicon measurement results, the optimal set of M resonators is selected to achieve the desired filter response. As such, the impact of large-scale process variations is effectively minimized. Note that the aforementioned post-silicon tuning is completely different from the traditional N-modular redundancy method where only one device out of N candidates is selected [15].

To make the proposed post-silicon tuning efficient, a number of CAD algorithms and methodologies are developed in this paper to optimize the design variables of the redundant resonators (e.g., the total number of resonators, the center frequencies of resonators, etc.). In particular, a linear programming formulation is derived to determine the optimal values of these design variables. Since the linear programming is convex, it can be solved efficiently (i.e., with low computational cost) [13]. Furthermore, an efficient heuristic algorithm is developed to select the optimal subset of resonators based on the post-silicon measurement results. Compared to the traditional branch and bound algorithm [13], the proposed heuristic method achieves more than 1000× runtime speed-up while converging to a similar solution (less than 8% increase in cost function).

The remainder of this paper is organized as follows. In Section 2, we first describe our proposed post-silicon tuning scheme and then develop a number of statistical CAD algorithms and methodologies in Section 3. The efficacy of the proposed post-silicon tuning is demonstrated by a MEMS mixer example in Section 4. Finally, we conclude in Section 5.

2. ADAPTIVE POST-SILICON TUNING FOR MEMS FILTERS

Figure 1 shows a simplified circuit schematic of a MEMS resonator. The input voltage V_{IN} is placed across electrostatic gap electrodes to create the drive force. As a result, the motional current I_{OUT} is generated at the output electrode. The transfer function from the input V_{IN} to the output I_{OUT} can be approximated as [14]:

$$\frac{I(s)}{V(s)} = \frac{\beta \cdot (\omega_0/Q) \cdot s}{s^2 + (\omega_0/Q) \cdot s + \omega_0^2} \tag{1}$$

where ω_0 represents the center frequency, Q denotes the quality factor and β is a parameter defining the gain of the resonator.

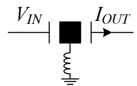

Figure 1. Simplified circuit schematic is shown for a MEMS resonator.

The frequency response of the transfer function in (1) represents a band-pass filter whose center frequency is ω_0 and the bandwidth is ω_0/Q. Since a MEMS resonator is typically made in low-loss materials, its Q value is extremely large (e.g., 3000 ~ 80000). Hence, a MEMS resonator can accurately select the input signals around its center frequency ω_0 and filter out the unwanted signals at other frequencies. More details about MEMS resonator can be found in [14].

In this paper our objective is to design a band-pass MEMS filter instead of a single resonator. In many practical applications (e.g., the channel filter of an RF transceiver) the pass-band of the filter is substantially wider than the pass-band of a single MEMS resonator. Hence, multiple MEMS resonators with different center frequencies should be connected in parallel to form an array. As such, the desired pass-band can be achieved by appropriately designing the resonator array. Due to the high quality factor of MEMS resonators, a sharp transition from the pass-band to the stop-band can be realized and, hence, the resulting MEMS filter can select the signal components within the pass-band and reject the unwanted components outside the pass-band with excellent precision. For illustration purpose, Figure 2 shows an example of the proposed MEMS filter architecture using a resonator array.

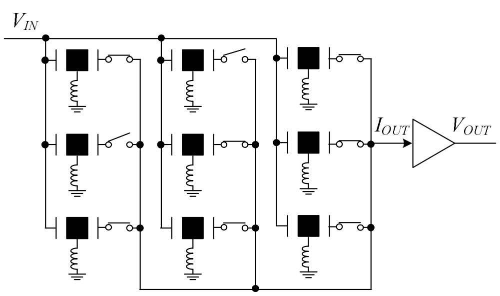

Figure 2. The proposed MEMS filter architecture is based upon a resonator array where the switches are used to turn on/off the resonators for post-silicon selection.

To make the proposed MEMS filter of practical utility, all MEMS resonators in the array must be appropriately designed to meet the following three criteria: (i) the ripple of the filter gain should be minimized within the pass-band to avoid in-band distortion; (ii) the filter gain within the stop-band must be minimized to guarantee sufficient signal attenuation; and (iii) the transition band must be minimized to provide excellent frequency selectivity. Without manufacturing variations, these three criteria can be easily satisfied by appropriately designing the center frequencies and gain values for all of the MEMS resonators.

In practice, due to large-scale process variations, the frequency responses of MEMS resonators can significantly differ from their nominal values. In general, process variations can be classified into two broad categories [9]-[11]: (i) inter-die variations, and (ii) intra-die variations. Inter-die variations refer to the common or average variations across the die. Since inter-die variations identically affect all resonators on the same die, they can be easily compensated by using circuit-level design techniques. On the other hand, intra-die variations represent the individual local variations within the die. These intra-die variations introduce random mismatches among MEMS resonators and must be carefully managed in order to achieve the desired frequency response for the proposed MEMS filter. Therefore, we will mainly focus on intra-die variations in this paper.

From (1) we notice that the transfer function of the MEMS resonator is uniquely determined by three parameters: (i) the center frequency ω_0, (ii) the quality factor Q, and (iii) the gain-related parameter β. All three of these parameters can vary due to intra-die variations. However, the variation of the center frequency is most critical to our filter application because the quality factor of a MEMS resonator is extremely high and, therefore, a small perturbation of its center frequency can completely change the frequency response of the filter [12]. Once the intra-die variation (e.g., random mismatch) of center frequency is successfully handled, the frequency response of the proposed MEMS filter can be accurately controlled.

To address the problem of intra-die center frequency variation, we propose to add redundant resonators into the array and then select an optimal subset of the resonators based on the post-silicon measurement results to accurately match the desired filter response in frequency domain. Other resonators that are not selected are simply turned off so that they do not contribute to the frequency response of the filter, as shown in Figure 2. While the proposed idea of post-silicon tuning looks simple, a number of comprehensive CAD algorithms and methodologies must be developed to optimally design the proposed resonator array with redundancy. In particular, the following two critical problems must be solved. First, a statistical optimization method must be developed to determine the optimal number of redundant resonators and their center frequencies. Second, once the post-silicon measurement results (i.e., the post-manufacturing center frequencies of all resonators) are available, an optimal selection algorithm must be used to find the subset of resonators that should be selected to form the MEMS filter. In what follows, we will discuss our proposed approaches to address these two open problems.

3. STATISTICAL DESIGN METHOD

3.1 Statistical Resonator Array Optimization

Without loss of generality, we assume that M MEMS resonators are used to implement the band-pass filter of interest. In the ideal scenario (i.e., without process variations), these M resonators are designed with the same gain and their center frequencies $\{f_m;\ m = 1,2,\ldots,M\}$ can be optimized by the filter designer to achieve the desired frequency response.

In practice, due to process variations, N ($N > M$) resonators are required to form a resonator array with redundancy for post-silicon tuning. Once process variations are considered, the center frequencies of these N resonators are no longer deterministic. Instead, they must be modeled as N random variables:

$$g_n \sim Gaussian\left(\mu_n, \sigma_n^2\right) \quad (n = 1,2,\cdots,N). \tag{2}$$

Eq. (2) assumes that the center frequency of the nth resonator g_n is modeled as a Gaussian distribution. Its mean value μ_n, referred to as the *nominal center frequency* in this paper, is a design variable that can be controlled by the layout of the MEMS resonator. The standard deviation σ_n models the magnitude of the center frequency variation and is determined by the manufacturing process. Since we focus on intra-die variations as described in Section 2, we further assume that the center frequency variations of different MEMS resonators are dominated by random mismatches and they are statistically independent.

The objective of our proposed post-silicon tuning is to select an optimal subset from $\{g_n; n = 1,2,\ldots,N\}$ to accurately match the desired center frequencies $\{f_m; m = 1,2,\ldots,M\}$. Towards this goal, we define the following interval to quantitatively assess whether a given frequency f_m can be accurately matched:

$$\left[f_m - d, f_m + d\right] \tag{3}$$

where $d \geq 0$ is referred to as the *frequency offset* in this paper. The frequency f_m is *successfully matched* if we can find g_n from the set $\{g_n; n = 1,2,\ldots,N\}$ such that:

$$f_m - d \leq g_n \leq f_m + d\,. \tag{4}$$

Restating in words, there is at least one resonator whose center frequency closely matches the desired value f_m after fabrication. Otherwise, the frequency f_m is *unsuccessfully matched*, if there exists no g_n from the set $\{g_n; n = 1,2,\ldots,N\}$ to satisfy the inequality in (4).

The aforementioned frequency matching plays an important role in minimizing the ripple of the in-band filter gain. If every frequency f_m in the set $\{f_m; m = 1,2,\ldots,M\}$ is successfully matched with a small offset d, it implies that the desired frequency response can be accurately approximated and, hence, a good MEMS filter can be realized even with large-scale process variations. For this reason, the quality of frequency matching, measured by the offset d, is directly correlated to the quality of the filter we design. When designing the proposed MEMS resonator array, we must optimally determine the number of resonators N and their nominal center frequencies $\{\mu_n; n = 1,2,\ldots,N\}$ so that successful matching can be accomplished for every f_m with a small offset d.

Given the Gaussian distribution in (2), we calculate the probability for the inequality (4) to hold:

$$\begin{aligned} P_{mn} &= \mathrm{Prob}\left(f_m - d \leq g_n \leq f_m + d\right) \\ &= \Phi\left(\frac{f_m + d - \mu_n}{\sigma_n}\right) - \Phi\left(\frac{f_m - d - \mu_n}{\sigma_n}\right) \end{aligned} \tag{5}$$

where Prob(•) denotes the probability of a random event and $\Phi(\bullet)$ represents the cumulative distribution function of a standard Gaussian distribution (i.e., zero mean and unit variance). Eq. (5) estimates the probability for the frequency f_m to be successfully matched by g_n. The matching for f_m is successful, if any g_n from the set $\{g_n; n = 1,2,\ldots,N\}$ satisfies the inequality in (4). Hence, the probability of successful matching for f_m is equal to:

$$P_m = 1 - \prod_{n=1}^{N}\left(1 - P_{mn}\right). \tag{6}$$

Studying (5)-(6) reveals two important observations. First, the probability P_m increases, as N increases; namely, the chance of successful matching becomes higher as more redundant resonators are added. Second, the probability P_m decreases as d decreases. This simply implies that if the offset specification becomes tight, the parametric yield decreases.

After deriving the probability in (6) we are ready to formulate the following optimization problem:

$$\begin{array}{ll} \min\limits_{N,\mu_1,\cdots,\mu_N} & N \\ \text{S.T.} & P_m \geq 1 - \varepsilon \quad \left(m = 1,2,\cdots,M\right) \end{array} \tag{7}$$

where ε is a given specification defining the maximum probability of unsuccessful matching for each frequency f_m. The value of ε should be sufficiently small (e.g., $\varepsilon = 0.01$) in order to guarantee a high probability for successful matching. The objective of (7) is to find: (i) the minimum number of resonators (i.e., N), and (ii) the nominal center frequencies $\{\mu_n; n = 1,2,\ldots,N\}$ for all resonators. Note that directly solving the unknowns $\{\mu_n; n = 1,2,\ldots,N\}$ and N from (7) is not trivial, since the constraint set in (7) is not convex. For this reason, we will propose an alternative optimization formulation in order to solve these problem unknowns efficiently (i.e., with low computational cost).

We discretize the pass-band of the filter and represent it by K discrete frequency values $\{h_k; k = 1,2,\ldots,K\}$. In other words, instead of solving the unknowns $\{\mu_n; n = 1,2,\ldots,N\}$ as continuous variables, we constrain the solution space to a set of discrete values defined by $\{h_k; k = 1,2,\ldots,K\}$. If the step size of the discretization is sufficiently small (i.e., the value of K is sufficiently large), the discrete solution space accurately approximates the original continuous solution space and the approximation error is negligible.

Once the discretization step is complete, we know that the nominal center frequency of the nth resonator (i.e., μ_n) must belong to the finite set $\{h_k; k = 1,2,\ldots,K\}$. Hence, we define a new symbol n_k ($n_k \geq 0$) as the number of resonators whose nominal center frequencies equal h_k. Since we have N resonators in total, the summation of $\{n_k; k = 1,2,\ldots,K\}$ is equal to N:

$$n_1 + n_2 + \cdots + n_K = N \quad \left(n_1 \geq 0 \quad n_2 \geq 0 \quad \cdots \quad n_K \geq 0\right). \tag{8}$$

Based on this notation, $\{n_k; k = 1,2,\ldots,K\}$ can be considered as a set of new optimization variables that should be solved. Namely, instead of solving $\{\mu_n; n = 1,2,\ldots,N\}$ and N, we propose to derive an alternative-yet-equivalent optimization formulation to find the optimal values of $\{n_k; k = 1,2,\ldots,K\}$.

Given the aforementioned definition of $\{h_k; k = 1,2,\ldots,K\}$ and $\{n_k; k = 1,2,\ldots,K\}$, it is easy to verify that the probability of successfully matching in (6) can be re-written as:

$$P_m = 1 - \prod_{k=1}^{K}\left(1 - Q_{mk}\right)^{n_k} \tag{9}$$

where

$$Q_{mk} = \Phi\left(\frac{f_m + d - h_k}{\sigma_k}\right) - \Phi\left(\frac{f_m - d - h_k}{\sigma_k}\right). \tag{10}$$

In (10), σ_k represents the standard deviation of the center frequency variations for the resonators whose nominal center frequencies equal h_k. Note that once $\{f_m; m = 1,2,\ldots,M\}$ (i.e., the desired center frequencies), $\{h_k; k = 1,2,\ldots,K\}$ (i.e., the discrete solution space), and d (i.e., the frequency offset) are known, Q_{mk} in (10) is determined. The probability P_m in (9) is a function of the unknown variables $\{n_k; k = 1,2,\ldots,K\}$ that we need to solve.

Substituting (9) into (7), the constraints in (7) become:

$$1 - \prod_{k=1}^{K}\left(1 - Q_{mk}\right)^{n_k} \geq 1 - \varepsilon \quad \left(m = 1,2,\cdots,M\right) \tag{11}$$

or equivalently:

$$\prod_{k=1}^{K}\left(1 - Q_{mk}\right)^{n_k} \leq \varepsilon \quad \left(m = 1,2,\cdots,M\right). \tag{12}$$

Taking the logarithm for both sides of (12) yields:

$$\sum_{k=1}^{K} n_k \cdot \log(1 - Q_{mk}) \le \log(\varepsilon) \quad (m = 1,2,\cdots,M). \tag{13}$$

Combining (7) and (13) results in the following optimization:

$$\begin{aligned} &\min_{n_1,\cdots,n_K} \quad n_1 + n_2 + \cdots + n_K \\ &\text{S.T.} \quad \sum_{k=1}^{K} n_k \cdot \log(1 - Q_{mk}) \le \log(\varepsilon) \quad (m = 1,2,\cdots,M) \\ &\qquad\quad n_k \ge 0 \qquad\qquad\qquad\qquad\quad (k = 1,2,\cdots,K) \end{aligned} \tag{14}$$

In (14), both the cost function and the constraints are linear. Hence, it is a linear programming problem and the optimal solution can be efficiently found by an interior point algorithm [13].

It is worth mentioning that two practical issues must be carefully considered when applying the optimization formulation in (14). First, the variables $\{n_k; k = 1,2,\ldots,K\}$ solved by the linear programming in (14) are non-negative and real-valued, while they should be non-negative integers by definition. Therefore, a post-processing step must be applied to "legalize" the solution of (14). In this paper we simply round these variables to the nearest integers during the legalization step.

Second, the frequency offset d must be appropriately determined in order to calculate $\{Q_{mk}; m = 1,2,\ldots,M; k = 1,2,\ldots,K\}$ in (10) that are required by the optimization in (14). As previously discussed, the value of d is directly correlated to the ripple of the in-band filter gain. If d is large, the optimization in (14) will result in a small number of resonators to implement the band-pass filter (i.e., N is small). The ripple of the filter gain, however, is expected to be large. Hence, varying the value of d allows us to explore the design trade-offs between the number of required resonators and the ripple of the in-band filter gain. In practice, if the ripple is given as a performance specification, we can repeatedly solve the optimization in (14) with different values of d and then find the appropriate d value that meets the ripple specification.

3.2 Optimal Post-silicon Resonator Selection

In the previous section we formulate the optimization in (14) that enables us to optimally determine the nominal center frequencies $\{\mu_n; n = 1,2,\ldots,N\}$ for all resonators during the design stage. Once these resonators are manufactured, their actual center frequencies differ substantially from the designed values due to process variations. We therefore must measure the actual center frequencies after fabrication and then optimally select a subset of these resonators for frequency matching. Namely, given the post-manufacturing center frequencies of N resonators $\{\tilde{g}_n; n = 1,2,\ldots,N\}$, we aim to select M resonators out of these N ($N > M$) candidates such that the M selected center frequencies match the desired values $\{f_m; m = 1,2,\ldots,M\}$ as closely as possible. Here, we use a new symbol $\tilde{g}_n$, instead of g_n shown in (2), to emphasize that $\{\tilde{g}_n; n = 1,2,\ldots,N\}$ represent the post-manufacturing center frequencies with deterministic values (i.e., they are not random variables any more). The aforementioned resonator selection, however, is not trivial, because both M and N are typically large (e.g., 100 ~ 1000), as will be demonstrated by our design example in Section 4. As a result, there are numerous possible choices to select M resonators out of N candidates. Instead of enumerating all possible options, we propose a heuristic algorithm to solve this combinatorial optimization problem efficiently.

To derive our proposed resonator selection algorithm, we first mathematically define the cost function that we aim to minimize. Remember that our objective is to match the desired center frequencies $\{f_m; m = 1,2,\ldots,M\}$ as closely as possible. Therefore, we use the following cost function to quantitatively measure the quality of frequency matching:

$$F = \sum_{m=1}^{M} \left| f_m - \tilde{g}_{sm} \right| \tag{15}$$

where $\tilde{g}_{sm}$ denotes the selected center frequency to match f_m. Ideally, if all frequencies $\{f_m; m = 1,2,\ldots,M\}$ are exactly matched, the cost function in (15) reaches its minimum (i.e., zero).

Next, we define the *distance matrix* $V \in R^{M \times N}$. The (m, n)-th element of the M-by-N matrix V is the distance between the desired center frequency f_m and the post-manufacturing center frequency $\tilde{g}_n$:

$$V_{mn} = \left| f_m - \tilde{g}_n \right|. \tag{16}$$

Our proposed heuristic algorithm iteratively selects the optimal $\tilde{g}_n$ to match f_m based on a greedy search method. At the first iteration step, it finds the minimum element (say, V_{mn}) in the distance matrix V and assign $\tilde{g}_n$ to match f_m. Since the matching between f_m and $\tilde{g}_n$ is complete, the mth row and nth column of the distance matrix V is removed. It simply means that the frequencies f_m and $\tilde{g}_n$ will not be considered any more during the following iteration steps. The dimension of the distance matrix V is reduced to $(M-1)$-by-$(N-1)$. Next, during the second iteration step, we apply the same heuristic to the distance matrix V with reduced dimension and find a new pair of f_m and $\tilde{g}_n$ with minimum distance. The aforementioned procedure is repeated until every desired center frequency f_m is matched to an optimal post-manufacturing center frequency $\tilde{g}_n$.

Algorithm 1: Optimal Post-silicon Resonator Selection

1. Start from a set of desired center frequencies $\{f_m; m = 1,2,\ldots,M\}$ and a set of post-manufacturing center frequencies $\{\tilde{g}_n; n = 1,2,\ldots,N\}$.
2. Initialize the distance matrix V based on (16). Set the iteration index $i = 1$.
3. Find the minimum element V_{mn} in the distance matrix V. Assign $\tilde{g}_n$ to match f_m.
4. Remove the mth row and nth column from the distance matrix V.
5. If i is equal to M, stop iteration. Otherwise, set $i = i + 1$ and go to Step 3.

Algorithm 1 summarizes the major steps of the proposed optimal resonator selection method that is based on simple greedy search. The time complexity of this algorithm is M^2N. Compared to the traditional branch and bound algorithm [13], the proposed heuristic method achieves more than 1000× runtime speed-up and converges to a similar solution (less than 8% increase in cost function), as will be demonstrated by our design example in Section 4.

4. DESIGN EXAMPLE

In this section, we demonstrate the efficiency of the proposed adaptive post-silicon tuning by using a MEMS mixer example with channel filters. All numerical experiments are run on a 2.9 GHz Linux server with 4 GB memory.

4.1 MEMS Mixer and Channel Filter Structure

Shown in Figure 3 is the geometrical structure of a CMOS MEMS square frame resonator [3]. An RF signal V_{RF} and a local oscillator (LO) signal V_{LO} are placed across electrostatic gap electrodes to control the motional current I_{OUT}. The polarizing voltage V_P is used to turn on and off the MEMS resonator for post-silicon selection.

In this example, the LO frequency is set to 840 MHz and the RF frequency band is from 850 MHz to 850.1 MHz. The RF signal is down-converted to the intermediate frequency (IF) band from 10 MHz to 10.1 MHz. The gain of the MEMS resonator, measured by the transfer function in (1), is −160 dBS. Its quality factor is 2000 and, hence, the bandwidth of a single resonator is equal to: 10 MHz / 2000 = 5 kHz. As previously described in Section 2, multiple resonators with different center frequencies are connected in parallel to form an array to implement the band-pass channel filter that is required for the given IF band. Here, the in-band gain of the channel filter must be no less than −135 dBS and its bandwidth is equal to: 10.1 MHz – 10 MHz = 100 kHz.

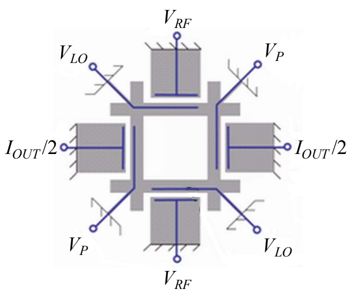

Figure 3. Geometrical structure is shown for a CMOS MEMS square frame resonator [3].

The MEMS resonators were designed based on a commercial 0.35 µm CMOS process. From the device-level variation models in the design kit, it was estimated that the random mismatch of the resonator center frequency has a standard deviation of $\sigma = 21$ kHz. Note that the value of σ (i.e., 21 kHz) is significantly larger than the bandwidth of a single resonator (i.e., 5 kHz). Table 1 summarizes the important parameters for the aforementioned MEMS mixer example. In what follows, we will apply the proposed post-silicon tuning to design and optimize the MEMS mixer to achieve the desired frequency response.

Table 1. Important parameters of the MEMS mixer example

Technology Node	0.35 µm CMOS
RF Frequency Band	[850 MHz, 850.1 MHz]
LO Frequency	840 MHz
IF Frequency Band	[10 MHz, 10.1 MHz]
IF Bandwidth	100 kHz
Channel Filter Gain (I/V)	≥ −135 dBS
Resonator Quality Factor	2000
Resonator Bandwidth	5 kHz
Resonator Center Frequency Variation	$\sigma = 21$ kHz
Resonator Gain (I/V)	−160 dBS

4.2 Resonator Array Optimization

Given the MEMS resonator in Figure 3, we first manually design the desired center frequencies $\{f_m;\ m = 1,2,\ldots,M\}$ for the ideal scenario (i.e., without process variations). It is determined that 228 resonators in total (i.e., $M = 228$) are needed to implement the specified channel filter. The targeted center frequencies of these 228 resonators are distributed over the pass-band of the channel filter, i.e., [10 MHz, 10.1 MHz].

We next apply the linear programming in (14) to optimally determine the nominal center frequencies of the redundant resonators $\{\mu_n;\ n = 1,2,\ldots,N\}$. Figure 4 shows the optimization results where the pass-band of the filter is partitioned into 100 discrete frequencies and the frequency offset d is set to different values. Studying Figure 4 we observe that the resonators are allocated at six nominal center frequencies only: 10.018 MHz, 10.019 MHz, 10.049 MHz, 10.050 MHz, 10.080 MHz and 10.081 MHz. At other frequencies the number of resonators is set to zero by the proposed optimization. This implies that we do not need to design many different resonators with different nominal center frequencies during the design stage, thereby reducing the design complexity. Due to random device mismatches, the actual center frequencies of these resonators will naturally spread over a wide frequency range. In the next sub-section we will demonstrate the proposed resonator selection algorithm to optimally find a subset of the redundant resonators and construct the channel filter with the desired frequency response.

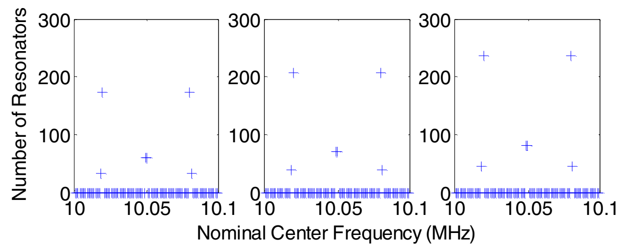

Figure 4. Nominal center frequencies are optimally determined by the linear programming in (14) where the pass-band of the filter is partitioned into 100 discrete frequencies and the frequency offset d is set to different values: (Left) 530 resonators in total, (Middle) 628 resonators in total, and (Right) 724 resonators in total.

4.3 Optimal Resonator Selection

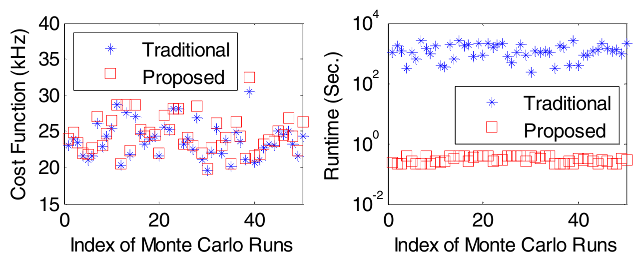

Figure 5. Comparison is shown for the proposed optimal resonator selection algorithm and the traditional branch and bound algorithm [13] based on 50 Monte Carlo runs with 628 resonators in the array: (Left) both algorithms minimize the cost function in (15) and the difference of the minimized cost function values is less than 8%; (Right) the proposed algorithm achieves more than 1000× runtime speed-up over the traditional approach.

Once the MEMS resonator array is manufactured, the center frequencies of all resonators are measured and the proposed heuristic algorithm (i.e., Algorithm 1) is used to select an optimal subset of resonators to minimize the cost function in (15). As such, the response of the selected resonators matches the desired frequency response. For comparison purpose, a traditional branch and bound algorithm [13] is implemented to minimize the same cost function in (15). Figure 5 compares the results arising from these two algorithms when applied to 50 Monte Carlo runs of an array of 628 MEMS resonators.

The data in Figure 5 leads to two important observations. First, the minimum cost function found by Algorithm 1 is close to that found by the branch and bound method (less than 8% difference).

This implies that the proposed heuristic algorithm can reliably find a good sub-optimal solution for the combinatorial optimization problem in (15). Second, but more importantly, the proposed heuristic algorithm is computationally cheaper than the traditional branch and bound method. The branch and bound method often takes more than 15 minutes to converge, while our proposed algorithm can reach convergence within 1 second (more than 1000× runtime speedup). Since Algorithm 1 has substantially lower computational cost, it can be more easily integrated into a practical post-silicon testing and configuration flow to improve the manufacturing yield of MEMS devices.

Table 2. Estimated ripple (dBS) of in-band filter gain for different design methods based on 1000 Monte Carlo runs

	# of Resonators	Mean	Sigma
No Redundancy	228	5.65	1.08
TMR [15]	684	4.56	0.81
Proposed Post-silicon Tuning	530	0.62	0.17
	628	0.51	0.08
	724	0.46	0.06

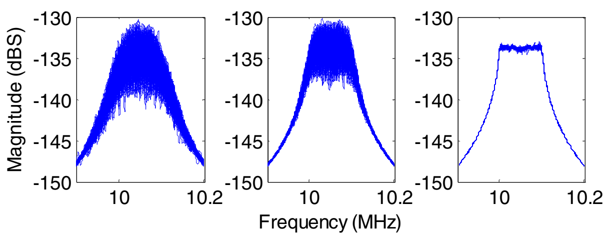

Figure 6. Frequency response (magnitude) is calculated by 1000 Monte Carlo runs: (Left) 228 resonators in total without including redundancy, (Middle) 684 resonators in total with the traditional triple modular redundancy (TMR) [15], and (Right) 628 resonators in total with the proposed post-silicon tuning.

To demonstrate the efficacy of the proposed post-silicon tuning, we further implement two traditional design methodologies for the MEMS channel filter. First, a simple MEMS resonator array is implemented without any redundancy (228 resonators in total). Second, the traditional triple modular redundancy (TMR [15]) approach is implemented where three resonators are designed for each desired center frequency $\{f_m;\ m = 1,2,\ldots,228\}$ and one of these three resonators is selected to optimally match the desired value of f_m. Table 2 summarizes the ripple of the in-band filter gain estimated from 1000 Monte Carlo runs for these different methods. Note that the ripple of the two traditional methods is extremely large. Their mean values are greater than 4.5 dBS. On the other hand, the proposed post-silicon tuning reduces the mean value of the ripple from 4.56 dBS to 0.62 dBS (more than 7× reduction). It, therefore, leads to significantly reduced in-band signal distortion when the proposed MEMS resonator array is used as a channel filter.

Figure 6 plots the frequency response (magnitude) of 1000 Monte Carlo runs for three different methods. Note that the two traditional methods result in non-flat in-band filter gain with significant variations. The proposed post-silicon tuning, however, yields a robust filter implementation that is not sensitive to large-scale process variations. These observations are consistent with the results previously shown in Table 2.

5. CONCLUSIONS

In this paper an efficient technique of adaptive post-silicon tuning is proposed to design robust MEMS filters with consideration of large-scale process variations. The proposed post-silicon tuning applies statistical element selection [7]-[8] to optimally select a group of resonators from a large array so that the desired frequency response can be achieved. Several CAD algorithms and methodologies are developed to optimize the design variables of the proposed MEMS resonator array. In particular, a new linear programming formulation is derived for statistical resonator array optimization and a heuristic algorithm is developed for optimal post-silicon resonator selection. Our design example of a MEMS mixer demonstrates that the proposed post-silicon tuning can efficiently reduce the ripple of the channel filter gain by 7× over other traditional approaches. The design methodologies and CAD algorithms proposed in this paper can be further integrated into a practical post-silicon testing and configuration flow to improve the manufacturing yield of MEMS devices.

6. ACKNOWLEDGEMENTS

The authors acknowledge the support of the C2S2 Focus Center, one of six research centers funded under the Focus Center Research Program (FCRP), a Semiconductor Research Corporation entity. This material is based upon work supported by DARPA and SPAWAR under Award No. N66001-09-1-2084. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of DARPA and SPAWAR.

7. REFERENCES

[1] B. Razavi, "RF CMOS transceivers for cellular telephony," *IEEE Communications Magazine*, pp.144-149, 2003.
[2] F. Chen, J. Brotz, U. Arslan, C. Lo, T. Mukherjee and G. Fedder, "CMOS-MEMS resonant RF mixer-filters," *IEEE MEMS*, pp. 24-27, 2005.
[3] C. Lo, F. Chen and G. Fedder, "Integrated HF CMOS-MEMS square-frame resonators with on-chip electronics and electrothermal narrow gap mechanism," *IEEE Transducers*, pp. 2074-2077, 2005.
[4] E. Quevy, A. Paulo, E. Basol, R. Howe, T. King and J. Bokor, "Back-end-of-line poly-SiGe disk resonators," *IEEE MEMS*, pp. 234-237, 2006.
[5] S. Li, Y. Lin. Z. Ren and C. Nguyen, "Disk-array design for suppression of unwanted modes in micromechanical composite-array filters," *IEEE MEMS*, pp. 866-869, 2006.
[6] H. Chandrahalim and S. Bhave, "Digitally-tunable MEMS filter using mechanically-coupled resonator array," *IEEE MEMS*, pp. 1020-1023, 2008.
[7] G. Keskin, J. Proesel and L. Pileggi, "Statistical modeling and post manufacturing configuration for scaled analog CMOS," *IEEE CICC*, pp. Sep. 2010.
[8] J. Proesel, G. Keskin, J. Plouchart and L. Pileggi, "An 8-bit 1.5GS/s flash ADC using post-manufacturing statistical selection," *IEEE CICC*, Sep. 2010.
[9] Semiconductor Industry Associate, *International Technology Roadmap for Semiconductors*, 2010.
[10] P. Drennan and C. McAndrew, "Understanding MOSFET mismatch for analog design," *IEEE JSSC*, vol. 38, no. 3, pp. 450-456, Mar. 2003.
[11] P. Kinget, "Device mismatch and tradeoffs in the design of analog circuits," *IEEE JSSC*, vol. 40, no. 6, pp. 1212-1224, Jun. 2005.
[12] J. Wang, Y. Xie and C. Nguyen, "Frequency tolerance of RF micromechanical disk resonators in polysilicon and nanocrystalline diamond structural materials," *IEEE IEDM*, pp. 285-289, 2005.
[13] D. Bertsekas, *Nonlinear Programming*, Athena Scientific, 1999.
[14] S. Senturia, *Microsystem Design*, Springer, 2000.
[15] M. Stanisavljevi, A. Schmid and Y. Leblebici, *Reliability of Nanoscale Circuits and Systems: Methodologies and Circuit Architectures*, Springer, 2010.