# A Configurable ULP Instrumentation Amplifier with Pareto-Optimal Power-Noise Trade-Off Achieving 1.93 NEF in 65nm CMOS

Rémi Dekimpe, *Student Member, IEEE*, and David Bol, *Senior Member, IEEE* *

## Abstract

**Performance trade-offs are central to analog/mixed-signal circuit design as they define the boundaries of the achievable design space. Circuit configurability allows run-time dynamic adaptation of these performance trade-offs to variable operating conditions. In this work, a new design methodology is used to implement an ultra-low-power (ULP) Pareto-optimal biomedical instrumentation amplifier (IA) with configurable power-noise trade-off. A multi-objective genetic algorithm performs the numerical optimization of the parameters at design time. The non-dominated sorting genetic algorithm (NSGA-II) is used along with an efficient simulation framework to limit the computation time. The optimal sizing is then applied to selected devices with digitally-controlled parameters in the amplifier. The configurable IA for biomedical applications has been prototyped in 65nm LP CMOS. It can be digitally set to 4 operating modes with power consumption ranging from 0.56 to 23.8$\mu$W and input-referred noise from 1$\mu$V to 0.17$\mu$V. The minimum noise efficiency factor (NEF) achieved by the amplifier is 1.93. The silicon area is 0.055mm$^2$ excluding the off-chip high-pass filter.**

## 1 Introduction

Analog and mixed-signal circuits are key components in modern integrated circuits, involved in various functions such as signal acquisition, clock generation, voltage regulation or high speed communications. For each analog circuit block, the performance evaluation usually relies on multiple metrics. At design time, several of those metrics can be considered simultaneously as optimization objectives, resulting in a multi-objective optimization problem. Each of these performance objectives is defined as a function $f_i(\mathbf{x})$ of the design parameters $\mathbf{x} = (x_1, x_2, ...)$. These parameters correspond to the sizing variables of a given circuit topology, such as device dimensions, bias voltages or clock frequencies [1].

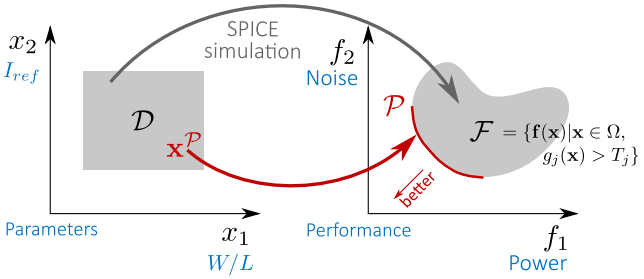


Figure 1: Simulation maps the circuit parameters $\mathbf{x}$ to the circuit performance metrics $\mathbf{f}(\mathbf{x})$. The Pareto curve $\mathcal{P}$ is the boundary of the achievable performance space $\mathcal{F}$ in the optimum direction. The blue text provides an illustration of the optimization task for the IA use-case.

Mathematically, the set of feasible design parameters $\mathcal{D}$ that contains all circuit configurations $\mathbf{x}$ that are physically achievable, maps to a set of feasible performance $\mathcal{F}$, as illustrated in Fig. 1 for a 2-D objective space. Performance constraints can also be defined on the set of achievable performance to account for additional circuit requirements, in which performance metrics are assigned a minimum value $g_j(\mathbf{x}) > T_j$. As the mapping is usually non-linear, an accurate evaluation of the circuit performance is typically obtained using a SPICE circuit simulator.

In the set of achievable performance, the subset that optimizes performance is of interest as it can be inversely mapped to optimal design parameters $\mathbf{x}^{\mathcal{P}}$. When several performance objectives are considered, there is rarely a single combination of parameters that optimizes all metrics simultaneously, as these are often competing against each other, which leads to a performance trade-off. The set of optimal points is known as the Pareto front $\mathcal{P}$ (Fig. 1), i.e. all points for which it is not possible to improve one objective without degrading another [2]. In the case of an instrumentation amplifier (IA), there is a trade-off between e.g. power consumption and intrinsic noise. The Pareto curve is typically used to help the designer select the optimal performance trade-off for his application. In hierarchical system design, Pareto-optimal models of sub-circuits can also be used at system level to take into account low-dimensional low-level design information [1].

In this work, we propose to use the information of the Pareto curve to design an IA with configurable power-noise



This work is published as [R. Dekimpe and D. Bol. "A Configurable ULP Instrumentation Amplifier with Pareto-Optimal Power-Noise Trade-Off Achieving 1.93 NEF in 65nm CMOS." *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 68, no. 7, pp. 2272-76, 2021.], available at https://doi.org/10.1109/TCSII.2021.3059311.
This work was supported by the National Foundation for Scientific Research (FNRS) of Belgium.
The authors are with the ICTEAM Institute, Université catholique de Louvain, Louvain-la-Neuve 1348, Belgium (e-mail: {remi.dekimpe, david.bol}@uclouvain.be)

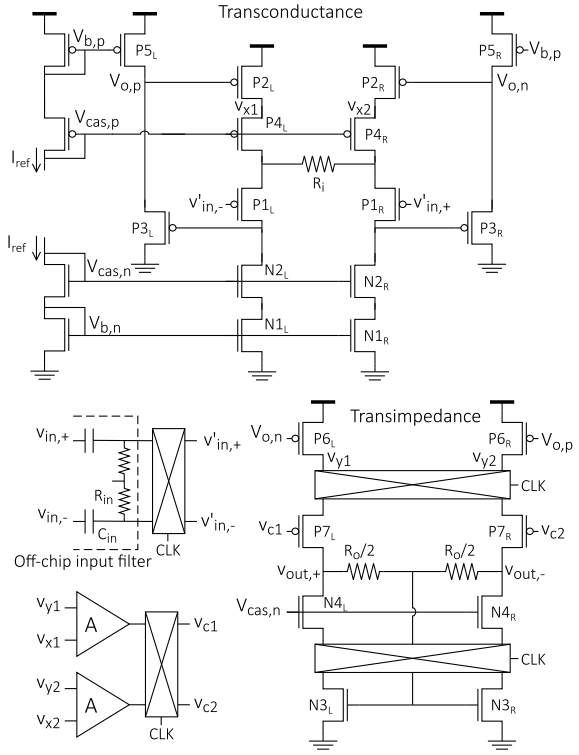


Figure 2: The IA circuit is a fully-differential amplifier based on the CFBIA topology, composed of a transconductance input stage and a transimpedance output stage. The input high-pass filter, chopper clock source and reference current are placed off-chip.

trade-off for biomedical applications. This allows the user to tune the circuit performance along the Pareto curve at runtime, depending on the operating conditions [3]. The proposed design methodology is divided in two stages. First, the Pareto curve in the power-noise performance space is extracted using a genetic algorithm for a baseline IA topology. As speed is usually the bottleneck for such algorithms, two solutions are proposed to reduce the global convergence time. Secondly, the information of the Pareto curve is used to modify the topology and size it optimally, in order to enable performance configurability. Section 2 describes the baseline IA topology and the associated design constraints. In Section 3, the design optimization problem is solved using a genetic algorithm to extract the power-noise Pareto curve. The implementation of the configurability is described in Section 4. Finally, the result is validated by comparing measurements of IAs prototyped in 65nm CMOS with fixed and configurable performance.

# 2 Instrumentation amplifier topology

The IA is a fully-differential current-feedback IA (CFBIA) topology similar to the one proposed by Van Helleputte et al. [4] as shown in Fig. 2. It is composed of an input transconductance (TC) that converts the differential input voltage to a proportional current and a transimpedance (TI) that converts it back to a differential voltage at the output. The total voltage gain ($A_v$) of the amplifier is therefore set by the ratio of the resistances in the TI and the TC and by the current mirror ratio [5] :

$$A_v = \alpha \frac{R_o}{\frac{g_{d,P1}}{g_{m,P1} g_{m,P2}} + R_i} \text{ with } \alpha = \frac{(W/L)_{P6}}{(W/L)_{P2}}. \quad (1)$$

The input TC is based on two flipped voltage followers (P1, P2) in the parallel branches with DC-level shifters (P3) to improve the dynamic range. Chopper stabilization is used to reduce 1/f noise present in the bandwidth of the signal. Operational transconductance amplifiers (OTA) create a feedforward path to regulate the drain voltage of the current mirrors in the TI in order to improve distortion. The IA input is AC-coupled to remove the differential electrode offset. We use an off-chip high-pass filter in this work.

A reference sizing of this IA was performed manually and will be used as a starting point for the design optimization methodology presented in Section 3. Manually searching for the optimal values of the circuit parameters is a difficult task due to the numerous trade-offs between performance metrics. For example, the noise in this circuit is typically dominated by the TC resistance $R_i$, the input transistor pair P1 and the TC current source N1. To minimize the noise, one would tend to minimize $R_i$. However, choosing $R_i$ too small impacts the gain due to the finite output conductance of P1 as shown in Eq. (1), and also increases the distortion of the signal. As this conductance depends on the bias current, it is directly linked to the power consumption which is the second performance objective. On the other hand, the resistance values $R_i$ and $R_o$ will be limited by the maximum silicon area. Other examples of such trade-offs involve the current ratio between TC and TI ($\alpha$ in Eq. (1)) which helps decreasing the power and the area but reduces the gain, or the chopping frequency that improves the noise and CMRR but impacts the bandwidth requirement and distortion. Similar reasoning applies for most circuit parameters and performance metrics.

# 3 Multi-objective optimization methodology

Automated sizing optimization of analog blocks with several performance objectives requires the use of specific algorithms designed to solve multi-objective optimization problems (MOP). The multi-objective optimization algorithms that have been applied to the task of analog sizing can be classified in different groups. A first group relies on a weight vector to scalarize the multiple objective and solve the problem using a sequence of single-objective problems, such as normal boundary intersection (NBI) [6]. Another group includes evolutionary algorithms, e.g. particle swarm optimization (PSO) [7], multi-objective evolutionary algorithm based on decomposition (MOEA/D) [8] or non-dominated sorting genetic algorithm (NSGA-II) [9], that rely on bio-inspired iterative methods to make a set of candidate solutions converge towards the Pareto curve. More recently, surrogate model-based methods such as multi-objective bayesian optimization (MOBO) have been proposed in an attempt to reduce the simulation time [10]. In this work, the genetic algorithm NSGA-II [9] is used, as it has been often shown to perform well for analog sizing, despite the high simulation time.

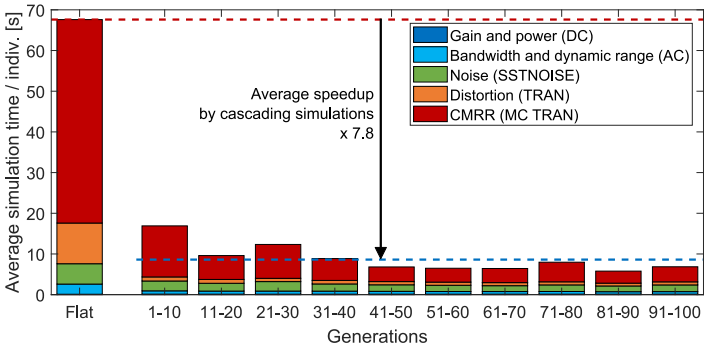


Figure 3: Simulation time to perform SPICE evaluation of a circuit candidate. The first column shows the time for evaluating a single circuit candidate for all analyses. The following columns show the average time spent for each analysis throughout the generations.

Like other evolutionary algorithms, NSGA-II is an iterative process in which a population $\mathbf{X}$, made up of individuals that each correspond to a set of circuit parameters $\mathbf{x}_i$, evolves for several generations. At each generation, the performance objectives and constraints are evaluated using a SPICE simulation. Based on this evaluation, the fittest individuals are selected while the others are discarded. As NSGA-II follows an elitist approach, the fittest individuals are those that meet the performance constraints and improve the performance objectives compared to previous generations. Finally, genetic operators such as mutation and cross-over are used to generate a new population based on the previously selected subset. The process starts over until a stopping criterion is reached.

The drawback of genetic algorithms is that they often require a large number of simulations for solving the MOP. This issue worsens when time-consuming simulations such as long transients or Monte-Carlo simulations are used and when the design space is large, typically for complex analog circuits. Two solutions are proposed in this work to mitigate the simulation time issue. The first solution relies on the fact that when many performance metrics are used as objectives or constraints, several SPICE analyses may be required with respective simulation times that differ in orders of magnitude (e.g. from 0.1s for a DC point calculation to 50s for Monte-Carlo analysis with transient signals). To gain simulation time, the simulations are run from the shortest to the longest and candidates are rejected as soon as requirements are not met. The rejection can be done on the basis of both performance constraints that are violated or performance objectives that are not improving compared to previous generations. This allows to save simulation time on the longest analyses for individuals that are not optimal. Compared to unconditional simulations in which all analyses are performed, we observe a 7.8× reduction in the total simulation time as shown in Fig. 3.

Secondly, the design parameter space is reduced by relying on designer knowledge. Starting from a manual sizing performed initially, a subset of higher-level parameters are tuned instead of considering all low-level device dimensions individually. This can be done by deriving a few main equations driving the circuit performance to identify the main parameters. Additionally, several parameters can be scaled together to keep the reference relative sizing within functional groups

Table 1: IA design parameters left for optimization.

| **Parameter**, $x_i$ | **Range** |
|---|---|
| Reference current, $I_{ref}$ | 10nA - 1$\mu$A |
| TC/TI ratio, $\alpha$ | 1 - 6 |
| TC resistance, $R_i$ | 1k$\Omega$ - 100k$\Omega$ |
| Resistance ratio, $R_o/R_i$ | 40 - 120 |
| Inversion level of input transistor, $(g_m/I_D)_{P1}$ | 2 - 30 |
| Inversion level of PMOS transistor, $(g_m/I_D)_p$ | 2 - 30 |
| Inversion level of NMOS transistor, $(g_m/I_D)_n$ | 2 - 30 |
| Length NMOS transistor, $L_{nmos}$ | 1$\mu$m - 20$\mu$m |
| Chopper frequency, $f_{CLK}$ | 1kHz - 20kHz |

of devices (e.g. cascode current mirrors). Ideally, the final set of parameters should be as small as possible while still allowing enough flexibility to optimize the objective while meeting the constraints. In the case of the IA, circuit inspection shows that the circuit has at least 28 degrees of freedom. In order to narrow the parameter space, only the main design degrees of freedom are considered. Based on Equation (1), we select the resistances $R_i$ and $R_o$, the inversion level of the input transistor pair P1 and PMOS cascaded current mirror represented by their parameter $g_m/I_D$ [11], and the TI/TC width ratio $\alpha$, as degrees of freedom. The bias current $I_{ref}$ is also included to sweep the power consumption. As the NMOS current sources are major noise contributors, the $g_m/I_D$ of transistor N1 is left as independent parameter, as well as its length so that the area can be modified independently. Finally, due to the low bandwidth of the IA, the chopping frequency $f_{chop}$ has a considerable impact on the noise and is also included. Table 1 summarizes the list of optimization parameters and the range of admissible values. When performing a random space exploration to find candidates that meet the constraints, reducing the number of parameters from 17 to 9 increases the percentage of valid candidates by 40%. This effect is however mitigated during the optimization process as the genetic algorithm only searches in the neighbourhood of valid candidates with mutation or crossover. The improvement in convergence speed is therefore mostly apparent during the early stages of the convergence. Figure 4 shows the convergence based on the hypervolume (HV) metric, which is computed as the volume contained by the Pareto front in the objective space [12]. After convergence of the NSGA, the hypervolumes only differ by 0.8%. The time taken to reach 90% of the final HV is 47% shorter when considering 9 parameters instead of 17.

Fig. 5 illustrates the process of convergence in the case of an IA with input-referred noise (IRN) and power consumption as performance objectives. Performance constraints for a typical biomedical application such as electrocardiogram (ECG) signal acquisition are considered in this work, as summarized in Table 2. The results illustrate power-noise improvement across successive generations, with an even distribution of individuals along the Pareto curve. This algorithm can be directly extended to include more performance objectives, at the cost of an increased convergence time.

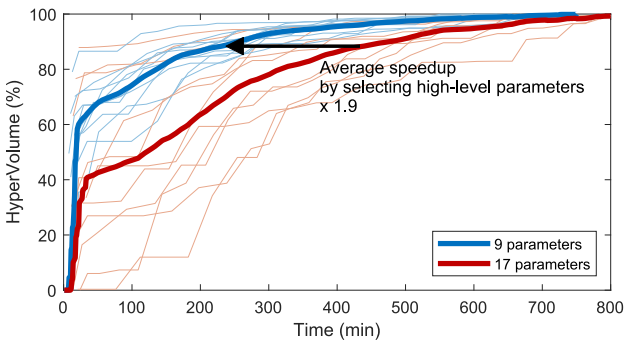


Figure 4: Convergence of the genetic algorithm over 30 generations. Light lines show convergence of individual runs, bold lines show the average convergence.

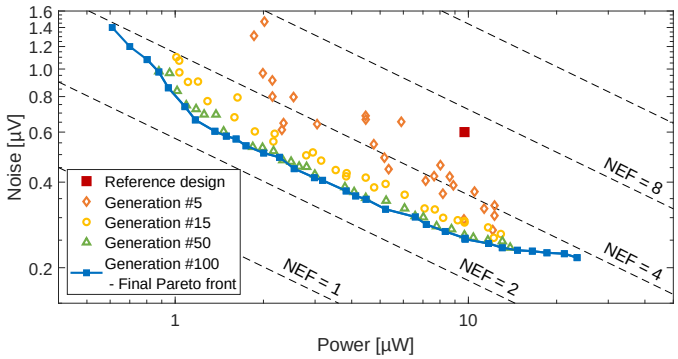


Figure 5: The performance objectives of successive generations converge towards the Pareto-optimal curve. The dashed lines are the constant-NEF contours.

Table 2: Performance constraints.

| **Constraint**, $g(\mathbf{x})$ | **Threshold**, $T$ |
|---|---|
| Gain | $>$ 30 V/V |
| Bandwidth | [0.1 - 250] Hz |
| Total harmonic distortion | $<$ 1% |
| CMRR | $>$ 80 dB |
| Dynamic range | $>$ 0.6 V |
| Active area | $<$ 20 000 $\mu$m$^2$ |

# 4 Configurable power-noise performance

Circuit configurability allows to dynamically change the circuit characteristics based on the current operating conditions. For example, in the case of an IA, input signals with different amplitudes or signal quality requirements may call for different IRN levels, leading to different power consumption as shown in Section 3. A trivial solution to have several noise-power trade-off choices is to implement several amplifiers on the chip and enable them separately. However, this leads to a high area overhead, essentially proportional to the number of configuration points. To avoid this overhead, a single amplifier with configurable characteristics is required. Song et al. [3] tuned the bias current of an amplifier to change its power consumption and IRN. However, simulation shows that this does not ensure optimal performance in all configurations as the bias conditions of the transistors vary. This work proposes to tune simultaneously several design parameters to reach optimal performance as obtained by the genetic algorithm.

Based on the Pareto curve obtained in Fig.5, 4 design points have been selected to cover the noise-power range. Among the parameters considered in building the Pareto curve, transistor length and input transistor P1 width have been set to a fixed value, as they have converged to a near-constant value and they are more difficult to dynamically configure.

The configuration of the circuit parameters by the digital control signals $C_i$ can be separated in three main ideas illustrated in Fig. 6. Firstly, the width of NMOS current sources is changed by modulating the number of transistors in parallel. Each branch is controlled by a separated bias voltage for the cascode transistor $V_{cas,n,i}$, that can be tied to ground to stop the current through that branch (Fig. 6.a). The generation of $V_{cas,n,i}$ is performed in the bias generation circuit where an additional transistor controlled by the digital signal $C_i$ is used to enable the corresponding branch. All branches with $C_i$ at high logic level appear in parallel and the reference current $I_{ref}$ is distributed according to the width ratio. Inserting the additional control transistor in the bias generation circuit saves dynamic range in the signal branches and helps to avoid distortion. Secondly, the width of PMOS current mirrors follows the same behavior in the TC (Fig. 6.b) using $V_{cas,p,i}$ as control voltage. However, unlike NMOS current sources that carry DC current, the bandwidth of PMOS current mirrors is a limiting factor for the performance. In the low-power configuration, the gates of inactive branches in the TC and TI are connected to the $V_o$ voltage, thereby increasing the parasitic capacitance. In order to avoid this, additional switches on the gates of the deactivated branches are used to disconnect the gate from the $V_o$ voltage, which is replaced by a much smaller source capacitance of the control transistor. Finally, changing the resistors $R_i$ and $R_o$ is done using transmission-gate switches to short-circuit parts of the polysilicon resistors (Fig. 6.c). As the resistance of $R_i$ and $R_o$ is quite high, the on-resistance of the switches does not impact the signal quality. Additionally, multiple frequencies of the chopper clock are generated on chip with clock dividers. The variable current reference is provided off-chip.

# 5 Experimental validation

In order to validate the optimization methodology and amplifier configurability, a prototype chip has been built in TSMC 65nm LP CMOS. As shown in Fig. 7, each prototype includes five amplifiers: the first four are stand-alone amplifiers that each correspond to one selected point of the Pareto curve, and the last one is the configurable amplifier that allows modification of the power-noise trade-off using control signals.

The measured performance of the stand-alone and configurable amplifiers match the results of the genetic algorithms obtained with pre-layout simulations. The performance of the configurable IA ranges from $0.17\mu V_{rms}$ to $1\mu V_{rms}$ of IRN and $23.8\mu W$ to $0.56\mu W$ of power consumption. The minimum NEF is 1.93. All constraints from Table 2 are met in all configurations. No performance penalty is introduced by the

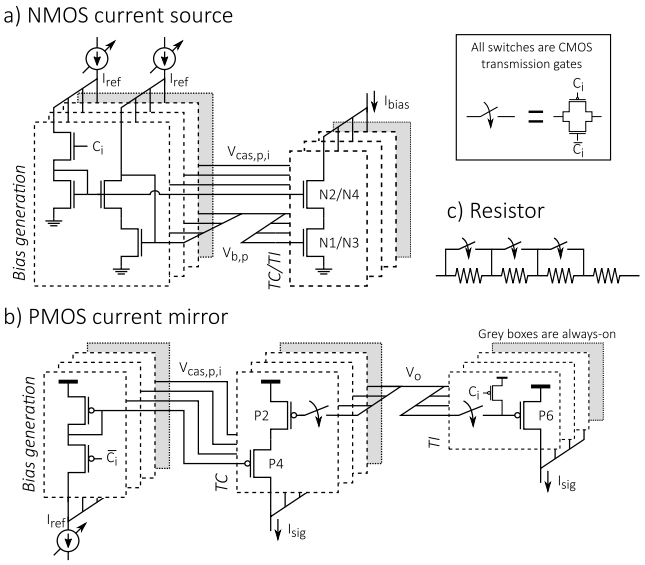


Figure 6: Digital control signals $C_i$ allow configuring the width of transistors in current sources and current mirrors, and the resistance of polysilicon resistors. The branches in grey do not include digital control as they are always-on. Variable current source is provided off-chip.

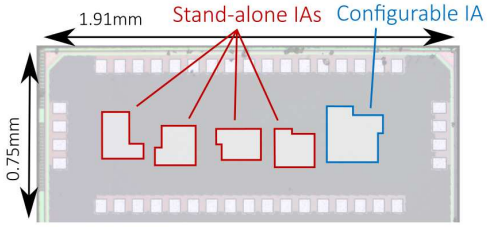


Figure 7: The chip photograph shows the location of four stand-alone IAs in red and one configurable IA in blue. The area is approximately 0.035 and 0.055mm$^2$ for stand-alone and configurable IAs respectively.

configurability.

The proposed IA sized using a genetic algorithm performs better in terms of noise-power trade-off compared to most state-of-the-art references. It can be noted that the works from Yaul et al. [13] and Song et al. [3] that achieve lower power consumption both use an extremely low supply voltage below 0.3V on the input stage. This reduces the input dynamic range, which is a limiting factor in most biomedical applications.

The chip measures 1.91×0.75mm. The configurable IA area is 0.055mm$^2$, compared to around 0.035mm$^2$ for each stand-alone amplifier. The difference is explained by the inverse relationship between resistors and transistors sizes in the different configurations according to the genetic algorithm. Indeed, while low-power configurations require small transistors, the resistance value is higher, and inversely for low-noise configurations. The configurability logic and switches incurs a negligible area overhead.

# 6 Conclusion

In this paper, an original design approach is proposed for building circuits with optimal configurable performance, illustrated with a biomedical IA. The genetic algorithm identifies the parameter sizing that provides optimal performance. This sizing information is then used to include configurability in the circuit. The measurement results show that this methodology improves the power and noise performance compared to the state of the art. The power consumption ranges from $0.56\mu W$ to $23.8\mu W$ with respective IRN from $1\mu V_{rms}$ to $0.17\mu V_{rms}$. No performance penalty is induced by configurability while area is significantly reduced compared to stand-alone amplifiers.

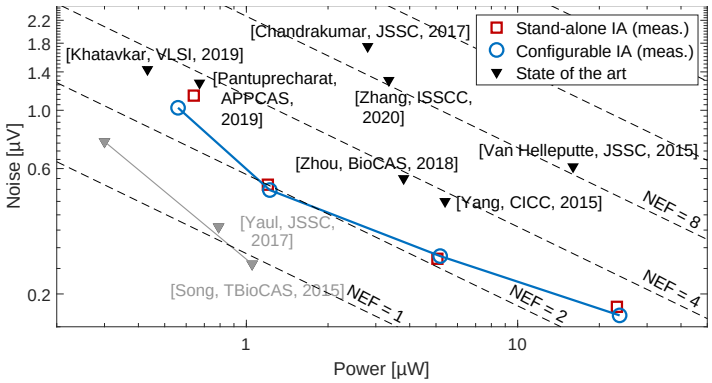


Figure 8: Measurement results are compared for stand-alone and configurable amplifiers, as well as state-of-the-art references. The IRN from the state of the art has been scaled to a [0.5 - 100] Hz bandwidth.

# Acknowledgment

The authors would like to thank Dr. Nick Van Helleputte from IMEC for his help to design the reference IA.

# References

[1] R. A. Rutenbar, G. G. Gielen, and J. Roychowdhury, "Hierarchical Modeling, Optimization, and Synthesis for System-Level Analog and RF designs," *Proc IEEE*, vol. 95, no. 3, pp. 640–669, 2007.

[2] S. Kundu and P. Mandal, "An Efficient Method of Pareto-Optimal Front Generation for Analog Circuits," *Analog Integr Circuits Signal Process*, vol. 94, no. 2, pp. 289–316, 2018.

[3] S. Song, M. Rooijakkers, P. Harpe, C. Rabotti, M. Mischi, A. H. van Roermund, and E. Cantatore, "A Low-Voltage Chopper-Stabilized Amplifier for Fetal ECG Monitoring with a 1.41 Power Efficiency Factor," *IEEE Trans Biomed Circuits Syst*, vol. 9, no. 2, pp. 237–247, 2015.

[4] N. Van Helleputte, M. Konijnenburg, J. Pettine, D.-W. Jee, H. Kim, A. Morgado, R. Van Wegberg, T. Torfs, R. Mohan, A. Breeschoten, H. de Groot, C. Van Hoof, and R. F. Yazicioglu, "A 345 $\mu$W Multi-Sensor Biomedical SoC With Bio-Impedance, 3-Channel ECG, Motion Artifact Reduction, and Integrated DSP," *IEEE J Solid-State Circuits*, vol. 50, no. 1, pp. 230–244, 2015.

[5] N. Van Helleputte, S. Kim, H. Kim, J. P. Kim, C. Van Hoof, and R. F. Yazicioglu, "A 160 $\mu$A Biopotential Acquisition IC with Fully Integrated IA and Motion Artifact Suppression," *IEEE Trans Biomed Circuits Syst*, vol. 6, no. 6, pp. 552–561, 2012.

[6] G. Stehr, H. Graeb, and K. Antreich, "Performance Trade-Off Analysis of Analog Circuits by Normal-Boundary Intersection," in *Proceedings of the 40th Design Automation Conference*. IEEE, 2003, pp. 958–963.

[7] S. Mallick, R. Kar, D. Mandal, and S. P. Ghoshal, "Optimal Sizing of CMOS Analog Circuits Using Gravitational Search Algorithm with Particle Swarm Optimization," *Int J Mach Learn Cybern*, vol. 8, no. 1, pp. 309–331, 2017.

[8] Q. Zhang and H. Li, "MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition," *IEEE Trans Evol Comput*, vol. 11, no. 6, pp. 712–731, 2007.

[9] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II," in *International Conference on Parallel Problem Solving from Nature*, 2000, pp. 849–858.

[10] W. Lyu, F. Yang, C. Yan, D. Zhou, and X. Zeng, "Multi-Objective Bayesian Optimization for Analog/RF Circuit Synthesis," in *Proceedings of the 55th Annual Design Automation Conference*, 2018, pp. 1–6.

[11] F. Silveira, D. Flandre, and P. G. Jespers, "A $g_m/I_D$-based Methodology for the Design of CMOS Analog Circuits and its Application to the Synthesis of a Silicon-On-Insulator Micropower OTA," *IEEE J Solid-state circuits*, vol. 31, no. 9, pp. 1314–1319, 1996.

[12] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca, "Performance Assessment of Multiobjective Optimizers: an Analysis and Review," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 117–132, 2003.

[13] F. M. Yaul and A. P. Chandrakasan, "A Noise-Efficient 36 nV/$\sqrt{Hz}$ Chopper Amplifier Using an Inverter-Based 0.2-V Supply Input Stage," *IEEE J Solid-State Circuits*, vol. 52, no. 11, pp. 3032–3042, 2017.