

RESIDUAL-DRIVEN ADAPTIVE MULTI-RATE QUADRATIC PROGRAMMING FRAMEWORK FOR NONLINEAR ANALOG AUDIO CIRCUIT EMULATION

Miguel Zea and Luis A. Rivera

Dept. of Electronics, Mechatronics and Biomedical Engineering
Centro de Procesos Industriales (CPI)
Universidad del Valle de Guatemala
Guatemala, GTM
{mezea, larivera}@uvg.edu.gt

ABSTRACT

This work extends our previously proposed Quadratic Programming (QP) approach for the emulation of nonlinear analog audio circuits by formalizing its main numerical ingredients and introducing a residual-driven adaptive multi-rate scheme. Starting from a state-space Differential Algebraic System of Equations (DAE) formulation, the nonlinear algebraic circuit device relations are replaced inside the QP by a first-order surrogate linear constraint, and the post-step nonlinear residual is shown to act as a valid defect indicator for adaptive step-size control. This yields a single-step simulation procedure that avoids the usual combination of nonlinear iterative solves and separate integration updates. The method is evaluated on a diode clipper, a BJT common-emitter amplifier, and a Colpitts oscillator, using SPICE as a baseline reference. The results show that adaptive step sizing considerably improves agreement with the reference solution, that the pseudo-inverse implementation is essentially equivalent to the full equality-constrained QP in the tested cases, and that the proposed formulation remains effective beyond the baseline clipper example, including for a self-oscillating circuit. These results position the proposed method as a promising bridge between SPICE-like interpretability and the efficiency demands of virtual analog (VA) audio applications.

1. INTRODUCTION

In the digital emulation of analog electrical circuitry for audio applications, there exists a tradeoff between accuracy (fidelity with respect to the original analog circuit), efficiency (real-time feasibility), and interpretability (scalability and parameter modulation, in practical terms), for each of the available methods. Data-driven methods, such as [1], can achieve high fidelity with efficient inference, but usually require high-quality training data and may provide limited direct access to physically meaningful circuit parameters. Gray-box approaches, such as [2], mitigate some of these limitations by embedding learned nonlinear components within structured models. However, they still rely on data-driven identification for parts of the nonlinear behavior. Virtual analog (VA) modeling techniques, such as [3], can be efficient and modulable, but often require considerable domain-specific expertise to obtain high fidelity for a given circuit. At the other end of the spectrum, SPICE-based simulators from electrical engineering practice offer

high accuracy and strong interpretability, but at the cost of being too computationally demanding for real-time audio use.

Some methods that try to close the gap between the three aforementioned tradeoff criteria are the so called state-space methods, like the ones presented by Yeh [4, 5], Holters and Zölzer [6]. However, their real-time feasibility is up for debate, as they too, like SPICE, need to solve a Differential Algebraic System of Equations (DAE) at each time-step. This requires an iterative method to solve the nonlinear algebraic equations and then use the result to propagate the system’s state forward in time by using an appropriate discretization of the dynamics. Thus, the solution has two possible sources of approximation error.

In our previous work [7], we proposed a modification of these state-space formulations in which the nonlinear algebraic constraint is replaced by a stabilized linear surrogate inspired by constraint-stabilization ideas from rigid-body dynamics. Furthermore, the simulation problem was reformulated as a convex Quadratic Program (QP), to be able to solve the DAE system with a single iterative method instead of the two originally required. Preliminary results on a diode clipper circuit showed improved agreement with a SPICE baseline and competitive computational behavior with respect to traditional state-space formulations. Nevertheless, several aspects of the method remained underdeveloped, such as the selection of time-step and stabilization parameters, cost function interpretability and feasibility in more complex audio circuits, all of which will be addressed in this work.

2. MATHEMATICAL BACKGROUND

This section summarizes the mathematical preliminaries required later in Section 3, as the proposed method uses notions from nonlinear systems, numerical integration, and constrained optimization that might not be standard in an audio-DSP setting. From this section onward, the notation $\|\cdot\|$ represents the standard Euclidean norm.

2.1. Nonlinear dynamics, linearization, and Taylor remainder

A nonlinear, state-space, dynamical system can be represented by the model

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad (1)$$

where $\mathbf{x} \in \Omega \subset \mathbb{R}^n$ is the state and $\mathbf{f} : \Omega \rightarrow \mathbb{R}^n$ is a vector field referred to as the dynamics of the system. If $\mathbf{f}$ has no explicit dependence on time, the system is called autonomous or time-invariant. Otherwise, the system may be written as $\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x})$

Copyright: © 2026 Miguel Zea et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, adaptation, and reproduction in any medium, provided the original author and source are credited.

and is called non-autonomous or time-varying. A standard sufficient condition for local existence and uniqueness of solutions is local Lipschitz continuity of the dynamics. In particular, $\mathbf{f}$ is locally Lipschitz on Ω if, for every $\mathbf{x}^* \in \Omega$, there exists a neighborhood $V \subset \Omega$ of $\mathbf{x}^*$ and a constant $L > 0$ such that

$$\|\mathbf{f}(\mathbf{x}_1) - \mathbf{f}(\mathbf{x}_2)\| \leq L \|\mathbf{x}_1 - \mathbf{x}_2\|, \quad \forall \mathbf{x}_1, \mathbf{x}_2 \in V.$$

If $\mathbf{f}$ is locally Lipschitz, then for every admissible initial condition, the Initial-Value Problem (IVP) admits a unique local solution [8].

Working directly with nonlinear systems is generally difficult because no general solution method exists. A simpler and well-understood case is the (homogeneous) Linear Time-Invariant (LTI) system

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}, \quad (2)$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is constant. The matrix $\mathbf{A}$ is said to be Hurwitz if all its eigenvalues have strictly negative real parts. If $\mathbf{A}$ is Hurwitz, then the origin of Eq. (2) is globally exponentially stable, i.e., its state exponentially and asymptotically evolves towards the origin regardless of the initial conditions. Linear systems are also useful because nonlinear systems can be approximated locally by them via linearization. Let $\mathbf{x}^*$ be an equilibrium point of Eq. (1), meaning

$$\mathbf{f}(\mathbf{x}^*) = \mathbf{0}. \quad (3)$$

Then, for $\mathbf{x}$ near $\mathbf{x}^*$ the first-order Taylor expansion gives

$$\mathbf{f}(\mathbf{x}) = \mathbf{f}(\mathbf{x}^*) + \tilde{\mathbf{A}}(\mathbf{x} - \mathbf{x}^*) + \mathbf{r}(\mathbf{x}),$$

where $\tilde{\mathbf{A}} = \frac{d\mathbf{f}}{d\mathbf{x}}(\mathbf{x}^*)$, and $\mathbf{r}(\mathbf{x})$ collects the higher-order nonlinear terms neglected by the linear approximation. Given Eq. (3), the local approximation of Eq. (1) is $\dot{\mathbf{x}} \approx \tilde{\mathbf{A}}(\mathbf{x} - \mathbf{x}^*)$. If the Jacobian of $\mathbf{f}$ is Lipschitz continuous on a compact neighborhood of $\mathbf{x}^*$, then the Taylor remainder is quadratically bounded, i.e., there exists $M > 0$ such that

$$\|\mathbf{r}(\mathbf{x})\| \leq M \|\mathbf{x} - \mathbf{x}^*\|^2 \quad (4)$$

for all $\mathbf{x}$ in that neighborhood [8].

The same local existence, linearization and Taylor remainder ideas extend directly to systems with inputs, with the form $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$, where $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ is an admissible input [8].

2.2. One-step methods, residuals, and adaptive time-stepping

Given that solutions of nonlinear dynamical systems are generally not available in closed form, they must be approximated numerically. A one-step method advances the state using a map

$$\mathbf{x}_{k+1} = \Phi_h(t_k, \mathbf{x}_k), \quad (5)$$

where $h > 0$ is the time-step and Φ_h approximates the solution of the system over one step. Some basic examples are the Forward Euler method,

$$\mathbf{x}_{k+1} = \mathbf{x}_k + h\mathbf{f}(t_k, \mathbf{x}_k), \quad (6)$$

the Backward Euler method,

$$\mathbf{x}_{k+1} = \mathbf{x}_k + h\mathbf{f}(t_{k+1}, \mathbf{x}_{k+1}), \quad (7)$$

and the Trapezoidal Rule,

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{h}{2} [\mathbf{f}(t_k, \mathbf{x}_k) + \mathbf{f}(t_{k+1}, \mathbf{x}_{k+1})]. \quad (8)$$

Forward Euler is explicit, since $\mathbf{x}_{k+1}$ is obtained directly from known quantities. Backward Euler and the trapezoidal rule are implicit, since $\mathbf{x}_{k+1}$ appears on both sides and must be found by solving an algebraic equation [9].

Even though all previous examples are one-step methods, and both Backward Euler and the Trapezoidal Rule are implicit, they differ in their stability properties. These are commonly studied through the scalar test equation $\dot{x} = \lambda x$, $\text{Re}(\lambda) < 0$, which represents a stable linear mode. When a one-step method is applied to this equation, it induces an update $x_{k+1} = R(\xi)x_k$, $\xi = h\lambda$, where R is called the stability function of the method. A method is A-stable if $|R(\xi)| \leq 1$ for all $\text{Re}(\xi) \leq 0$, and L-stable if, in addition, $R(\xi) \rightarrow 0$ as $\xi \rightarrow -\infty$. Backward Euler is both A-stable and L-stable, whereas the Trapezoidal Rule is A-stable but not L-stable [10]. Therefore, although the Trapezoidal Rule is formally more accurate for smooth ODEs, Backward Euler may behave better in stiff or constraint-dominated settings because it damps fast decaying modes more strongly.

The accuracy of a numerical step can be described locally by comparing the numerical update with the exact solution. If $\mathbf{x}(t)$ is the exact solution, the local truncation error (LTE) of a one-step method is $\ell_{k+1} := \mathbf{x}(t_{k+1}) - \Phi_h(t_k, \mathbf{x}(t_k))$. A method has order p when $\ell_{k+1} = \mathcal{O}(h^{p+1})$. In practice, the exact LTE is unavailable, so adaptive methods use a computable local error estimate $\hat{\ell}$. If an implicit method is written as $\mathbf{F}_h(t_k, \mathbf{x}_k, \mathbf{x}_{k+1}) = \mathbf{0}$, then the residual of a candidate step can be defined as

$$\boldsymbol{\rho}_{k+1} := \mathbf{F}_h(t_k, \mathbf{x}_k, \mathbf{x}_{k+1}).$$

Thus, the residual measures how well the candidate step satisfies the discrete numerical equations [9].

For Differential-Algebraic Equations (DAEs), the same idea applies, but the residual must account for both the differential update and the algebraic constraint. For example, a semi-explicit DAE can be written as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{q}), \quad \mathbf{0} = \mathbf{g}(\mathbf{x}, \mathbf{q}),$$

where $\mathbf{x}$ contains the differential variables and $\mathbf{q}$ contains the algebraic variables. A Backward Euler step, for example, satisfies

$$\mathbf{x}_{k+1} = \mathbf{x}_k + h\mathbf{f}(\mathbf{x}_{k+1}, \mathbf{q}_{k+1}), \quad \mathbf{0} = \mathbf{g}(\mathbf{x}_{k+1}, \mathbf{q}_{k+1}).$$

Therefore, a DAE step is itself a coupled algebraic problem, and a valid residual should quantify the defect left in the discrete differential and/or algebraic equations [11].

A residual can be used to adapt the time-step when it has the correct asymptotic scaling. In particular, for a method of order p , a computable defect indicator satisfying

$$\delta_k := \|\boldsymbol{\rho}_k\| = \mathcal{O}(h_k^{p+1}) \quad (9)$$

has the appropriate local order to guide adaptive step-size control schemes. The step is accepted when the residual is below a prescribed tolerance; otherwise, it is rejected and recomputed with a smaller time-step. After an accepted step, the next time-step can be increased or decreased using a standard adaptive controller, such as a deadbeat, PI/PID, or predictive update law [12, 13].

For systems with inputs, the same one-step notation is used with the input evaluated or approximated over the current step, e.g., $\mathbf{x}_{k+1} = \Phi_{h_k}(t_k, \mathbf{x}_k, \mathbf{u}_k)$.

2.3. Equality-constrained strictly convex quadratic programs

Consider the equality-constrained QP (ECQP)

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{c}^\top \mathbf{x} \quad \text{s.t.} \quad \mathbf{A}_{\text{eq}} \mathbf{x} = \mathbf{b}_{\text{eq}},$$

with $\mathbf{Q}$ being positive definite. If the equality constraints are feasible, then the problem has a unique solution [14]. In the special case $\mathbf{c} = \mathbf{0}$, the change of variables $\tilde{\mathbf{x}} = \mathbf{Q}^{1/2} \mathbf{x}$, $\tilde{\mathbf{A}} = \mathbf{A} \mathbf{Q}^{-1/2}$, transforms the problem into

$$\min_{\tilde{\mathbf{x}}} \frac{1}{2} \|\tilde{\mathbf{x}}\|_2^2 \quad \text{s.t.} \quad \tilde{\mathbf{A}}_{\text{eq}} \tilde{\mathbf{x}} = \mathbf{b}_{\text{eq}}.$$

Hence, an equality-constrained strictly convex QP is equivalent to a Euclidean minimum-norm problem in transformed coordinates, and to a weighted minimum-norm problem in the original coordinates. If $\mathbf{A}_{\text{eq}}$ has full row rank, then

$$\mathbf{x}^* = \mathbf{Q}^{-1} \mathbf{A}_{\text{eq}}^\top (\mathbf{A}_{\text{eq}} \mathbf{Q}^{-1} \mathbf{A}_{\text{eq}}^\top)^{-1} \mathbf{b}_{\text{eq}}.$$

When $\mathbf{Q} = \mathbf{I}$ this reduces to the ordinary Moore–Penrose pseudo-inverse $\mathbf{x}^* = \mathbf{A}_{\text{eq}}^\dagger \mathbf{b}_{\text{eq}}$ [14].

2.4. QP-based state-space method recap

We now recall the circuit formulation used in our previous work [7], using the notation introduced in the preceding subsections. Following the state-space formulation of nonlinear audio circuits in [4, 5], we consider circuits described by

$$\dot{\mathbf{x}} = \mathbf{A} \mathbf{x} + \mathbf{B} \mathbf{u} + \mathbf{G} \mathbf{q}, \quad (10a)$$

$$\mathbf{H} \mathbf{q} = \mathbf{D} \mathbf{x} + \mathbf{E} \mathbf{u}, \quad (10b)$$

$$\boldsymbol{\eta}(\mathbf{q}) = \mathbf{0}. \quad (10c)$$

Here, $\mathbf{x} \in \mathbb{R}^n$ contains the circuit state variables, $\mathbf{u} \in \mathbb{R}^m$ is the input, and $\mathbf{q} \in \mathbb{R}^r$ collects the voltages and currents associated with the nonlinear devices. The map $\boldsymbol{\eta} : \mathbb{R}^r \rightarrow \mathbb{R}^d$ represents the nonlinear current-voltage relations, while $\mathbf{A}, \mathbf{B}, \mathbf{G}, \mathbf{H}, \mathbf{D}$ and $\mathbf{E}$ are determined by the linear part of the circuit and by the selected state-space coordinates.

The key idea in our previous formulation was to avoid solving $\boldsymbol{\eta}(\mathbf{q}_k) = \mathbf{0}$ as a nonlinear algebraic equation at every step. Instead, the nonlinear constraint was replaced inside the numerical step by a first-order surrogate. More precisely, using Backward Euler (Eq. (7)) for the linear state dynamics and Forward Euler (Eq. (6)) for $\dot{\boldsymbol{\eta}}(\mathbf{q}) = \mathbf{J}(\mathbf{q}) \dot{\mathbf{q}} = \mathbf{K} \boldsymbol{\eta}(\mathbf{q})$:

$$(\mathbf{I} - h_k \mathbf{A}) \mathbf{x}_k - h_k \mathbf{G} \mathbf{q}_k = \mathbf{x}_{k-1} + h_k \mathbf{B} \mathbf{u}_k, \quad (11a)$$

$$-\mathbf{D} \mathbf{x}_k + \mathbf{H} \mathbf{q}_k = \mathbf{E} \mathbf{u}_k, \quad (11b)$$

$$\mathbf{J}_{k-1} \mathbf{q}_k = h_k \mathbf{K}_k \boldsymbol{\eta}(\mathbf{q}_{k-1}) + \mathbf{J}_{k-1} \mathbf{q}_{k-1}, \quad (11c)$$

where $\mathbf{J}_{k-1} := \frac{d\boldsymbol{\eta}(\mathbf{q}_{k-1})}{d\mathbf{q}}$. The first two equations are the discretized circuit dynamics and the linear algebraic relation. The third equation is the linearized surrogate for the nonlinear device constraint, with $\mathbf{K}_k$ selected to be Hurwitz.

By defining $\mathbf{z}_k = [\mathbf{x}_k^\top \quad \mathbf{q}_k^\top]^\top$, the three equations can be written as a single linear equality system

$$\mathbf{A}_{\text{eq}}^{(k)} \mathbf{z}_k = \mathbf{b}_{\text{eq}}^{(k)}, \quad (12)$$

with

$$\mathbf{A}_{\text{eq}}^{(k)} = \begin{bmatrix} \mathbf{I} - h_k \mathbf{A} & -h_k \mathbf{G} \\ -\mathbf{D} & \mathbf{H} \\ \mathbf{0} & \mathbf{J}_{k-1} \end{bmatrix}, \quad (13)$$

and

$$\mathbf{b}_{\text{eq}}^{(k)} = \begin{bmatrix} \mathbf{x}_{k-1} + h_k \mathbf{B} \mathbf{u}_k \\ \mathbf{E} \mathbf{u}_k \\ h_k \mathbf{K}_k \boldsymbol{\eta}(\mathbf{q}_{k-1}) + \mathbf{J}_{k-1} \mathbf{q}_{k-1} \end{bmatrix}. \quad (14)$$

The corresponding one-step update was posed as the ECQP

$$\mathbf{z}_k = \arg \min_{\mathbf{z}} \frac{1}{2} \mathbf{z}^\top \mathbf{Q} \mathbf{z} + \mathbf{c}^\top \mathbf{z} \quad \text{s.t.} \quad \mathbf{A}_{\text{eq}}^{(k)} \mathbf{z} = \mathbf{b}_{\text{eq}}^{(k)}, \quad (15)$$

with optional inequality constraints when state or device-variable bounds are required. This formulation is useful because it separates the general optimization framework from the particular solver used in simple cases. When no inequality constraints are active and $\mathbf{c} = \mathbf{0}$, the QP reduced to the minimum-norm equality solve described in Section 2.3.

Finally, the circuit output is obtained as

$$\mathbf{y}_k = \mathbf{M} \mathbf{x}_k + \mathbf{L} \mathbf{u}_k + \mathbf{N} \mathbf{q}_k, \quad (16)$$

as in standard state-space formulations of nonlinear audio circuits [4].

3. METHOD EXTENSION AND FORMALIZATION

This section formalizes the adaptive extension of the QP-based formulation recalled in Section 2.4. The goals are twofold. Firstly, we justify the parameter choices that were mostly left heuristic in our previous work, particularly the step-dependent selection of the stabilization matrix $\mathbf{K}_k$ and the interpretation of the QP matrices. Secondly, we show that the nonlinear residual left after the QP solve can be interpreted as a local defect indicator, which allows the method to be equipped with an adaptive multi-rate time-stepping scheme.

For a previous state $\mathbf{z}_{k-1}$, input $\mathbf{u}_k$, and trial time-step h_k , the QP solve defines a one-step map $\mathbf{z}_k = \hat{\boldsymbol{\Phi}}_{h_k}(\mathbf{z}_{k-1}, \mathbf{u}_k)$. Thus, the QP is treated as a numerical one-step method for the circuit DAE. It enforces the discretized linear dynamics, the linear algebraic circuit relation, and a local linear surrogate of the nonlinear devices' constraints.

3.1. Operating region and regularity assumptions

We assume that the simulation remains inside an operating region $D \subset \mathbb{R}^{n+r}$, where the circuit model (Eqs. (10a)–(10c)), the nonlinear map $\boldsymbol{\eta}$, and its Jacobian are well defined. Thus, every accepted iterate satisfies $\mathbf{z}_k \in D$. We also assume that the relevant signals remain bounded on the time interval of interest. That is, there exist constants $X_{\max}, U_{\max}, Q_{\max} > 0$ such that

$$\|\mathbf{x}_k\| \leq X_{\max}, \quad \|\mathbf{u}_k\| \leq U_{\max}, \quad \|\mathbf{q}_k\| \leq Q_{\max}.$$

Additionally, the input is assumed to vary at a bounded discrete rate,

$$\|\mathbf{u}_k - \mathbf{u}_{k-1}\| \leq U_{BL} h_k,$$

for some $U_{BL} > 0$, i.e., the input is assumed to be band-limited. If these bounds are not guaranteed a priori, they may be imposed as inequality constraints in the QP. Finally, we assume that the Jacobian of $\boldsymbol{\eta}$ is Lipschitz continuous on the compact subset of D

visited by the simulation. By the Taylor-remainder result of Section 2.1, the linearization error of $\boldsymbol{\eta}$ is then quadratically bounded.

A technical regularity condition is also required for the equality system solved at each step. We assume that the matrix $\mathbf{A}_{\text{eq}}^{(k)}$, as defined in Eq. (13), is uniformly regular on the operating region. Equivalently, there exists $\Gamma_{\text{eq}} > 0$, independent of k , such that

$$\|\delta \mathbf{z}\| \leq \Gamma_{\text{eq}} \|\mathbf{A}_{\text{eq}}^{(k)} \delta \mathbf{z}\|, \quad \forall \delta \mathbf{z} \in \mathbb{R}^{n+r}.$$

This means that a small mismatch in the equality system induces a proportionally small step increment. In practical terms, the equality solve is locally well conditioned.

Finally, we also assume that the time-step's behavior remains bounded and does not change abruptly. In particular,

$$0 < h_{\min} \leq h_k \leq h_{\max} < 1, \quad \frac{1}{\gamma_h} \leq \frac{h_k}{h_{k-1}} \leq \gamma_h, \quad (17)$$

for some $\gamma_h \geq 1$.

3.2. QP and equality-solve interpretations

Our method's formulation is written as a QP because this is the most general version of the proposed one-step update. It allows the inclusion of cost metrics, regularization terms, and inequality constraints on the state or nonlinear device variables. However, in most of the experiments to be reported later in Section 4, no inequality constraints are active and the selected metric is $\mathbf{Q} = \mathbf{I}$, $\mathbf{c} = \mathbf{0}$. Under these conditions, the QP reduces to

$$\min_{\mathbf{z}} \frac{1}{2} \|\mathbf{z}\|^2 \quad \text{s.t.} \quad \mathbf{A}_{\text{eq}}^{(k)} \mathbf{z} = \mathbf{b}_{\text{eq}}^{(k)}.$$

Therefore, the solution is the minimum-norm solution of the equality system, $\mathbf{z}_k = \left(\mathbf{A}_{\text{eq}}^{(k)}\right)^\dagger \mathbf{b}_{\text{eq}}^{(k)}$. Thus, one of the goals for the numerical results in this paper is to validate the equality-solve implementation obtained from the more general QP formulation. The QP form provides the natural framework for adding metrics and bounds, while the equality-solve form is the simpler implementation used when those additional terms are inactive, which leads to a performance boost.

3.3. Post-step nonlinear residual as a defect indicator

For the one-step map $\mathbf{z}_k = \hat{\boldsymbol{\Phi}}_{h_k}(\mathbf{z}_{k-1}, \mathbf{u}_k)$, the exact nonlinear algebraic constraint at the new step would be $\boldsymbol{\eta}(\mathbf{q}_k) = \mathbf{0}$. The QP step instead enforces the first-order approximation

$$\hat{\boldsymbol{\eta}}(\mathbf{q}_k) := \boldsymbol{\eta}(\mathbf{q}_{k-1}) + \mathbf{J}_{k-1}(\mathbf{q}_k - \mathbf{q}_{k-1}).$$

The last equality constraint of the QP is

$$\mathbf{J}_{k-1}(\mathbf{q}_k - \mathbf{q}_{k-1}) = h_k \mathbf{K}_k \boldsymbol{\eta}(\mathbf{q}_{k-1}).$$

With

$$\mathbf{K}_k = -\frac{1}{h_k} \mathbf{I}, \quad (18)$$

we obtain $\hat{\boldsymbol{\eta}}(\mathbf{q}_k) = \mathbf{0}$. Therefore, $\mathbf{K}_k$ is chosen so that the linearized nonlinear constraint is annihilated over the current one-step update.

The post-step nonlinear residual is defined as

$$\boldsymbol{\nu}_k := \boldsymbol{\eta}(\mathbf{q}_k) - \hat{\boldsymbol{\eta}}(\mathbf{q}_k). \quad (19)$$

Since $\hat{\boldsymbol{\eta}}(\mathbf{q}_k) = \mathbf{0}$ by construction, $\boldsymbol{\nu}_k = \boldsymbol{\eta}(\mathbf{q}_k)$, up to solver tolerances. More importantly, $\boldsymbol{\nu}_k$ is the Taylor remainder left by replacing $\boldsymbol{\eta}$ with its first-order approximation. It therefore measures the part of the nonlinear algebraic constraint not enforced by the linearized QP step. We now show that this residual has the correct local order for adaptive time-stepping. Since the QP enforces $\mathbf{A}_{\text{eq}}^{(k)} \mathbf{z}_k = \mathbf{b}_{\text{eq}}^{(k)}$, evaluating the same equality system at $\mathbf{z}_{k-1}$ gives

$$\mathbf{A}_{\text{eq}}^{(k)} \Delta \mathbf{z}_k = \mathbf{e}_k, \quad \Delta \mathbf{z}_k := \mathbf{z}_k - \mathbf{z}_{k-1},$$

where $\mathbf{e}_k$ is the induced mismatch. Under $\mathbf{K}_k = -\mathbf{I}/h_k$, this mismatch is

$$\mathbf{e}_k = \begin{bmatrix} h_k(\mathbf{A}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{u}_k + \mathbf{G}\mathbf{q}_{k-1}) \\ \mathbf{E}(\mathbf{u}_k - \mathbf{u}_{k-1}) \\ -\boldsymbol{\eta}(\mathbf{q}_{k-1}) \end{bmatrix}. \quad (20)$$

The first two blocks are $\mathcal{O}(h_k)$ by boundedness of the signals and by the bounded input-rate assumption. The third block is the previous nonlinear residual. By consistent initialization and the induction implied by accepted steps, $\|\boldsymbol{\eta}(\mathbf{q}_{k-1})\| = \|\boldsymbol{\nu}_{k-1}\| = \mathcal{O}(h_{k-1}^2)$. Since the accepted time-step ratio is bounded (Eq. (17)), this also implies $\|\boldsymbol{\eta}(\mathbf{q}_{k-1})\| = \mathcal{O}(h_k^2)$. Thus, the third block is of higher order than the first two, and since $|h_k| < 1$,

$$\|\mathbf{e}_k\| = \mathcal{O}(h_k).$$

By the uniform regularity assumption,

$$\|\Delta \mathbf{z}_k\| \leq \Gamma_{\text{eq}} \|\mathbf{e}_k\| = \mathcal{O}(h_k),$$

and therefore

$$\|\mathbf{q}_k - \mathbf{q}_{k-1}\| = \mathcal{O}(h_k).$$

Since the Jacobian of $\boldsymbol{\eta}$ is Lipschitz continuous, the Taylor-remainder bound gives

$$\|\boldsymbol{\nu}_k\| \leq M \|\mathbf{q}_k - \mathbf{q}_{k-1}\|^2 = \mathcal{O}(h_k^2), \quad (21)$$

for some $M > 0$. Thus, the post-step nonlinear residual has the expected local scaling of a first-order one-step method. We use the scalar quantity

$$\delta_k^{\text{QP}} := \|\boldsymbol{\nu}_k\| \quad (22)$$

as the adaptive defect indicator of the QP one-step map (QPDI).

3.4. Adaptive QP one-step method

Putting the previous elements together, the proposed adaptive QP-based simulation method is as follows:

1. Represent each nonlinear device through its current-voltage relation and collect the corresponding nonlinear voltages and currents in $\mathbf{q}$. This defines the nonlinear constraint $\boldsymbol{\eta}(\mathbf{q}) = \mathbf{0}$.
2. Derive the circuit matrices $\mathbf{A}, \mathbf{B}, \mathbf{G}, \mathbf{H}, \mathbf{D}, \mathbf{E}$, and the output matrices $\mathbf{M}, \mathbf{L}, \mathbf{N}$, so that the circuit is written in the form of Eqs. (10a)–(10c).
3. Choose the state discretization used in the first block row of the equality system (Eqs. (12)–(14)). In this work, we consider Backward Euler (Eq. (7)) and the Trapezoidal Rule (Eq. (8)).
4. Initialize $\mathbf{x}_0, \mathbf{q}_0$, and the initial step size h_0 . The initial condition must be consistent, i.e., $\mathbf{H}\mathbf{q}_0 = \mathbf{D}\mathbf{x}_0 + \mathbf{E}\mathbf{u}_0$ and $\boldsymbol{\eta}(\mathbf{q}_0) = \mathbf{0}$.

- At each step, given $\mathbf{z}_{k-1}$, $\mathbf{u}_k$, and a trial step size h_k , compute

$$\mathbf{J}_{k-1} = \frac{d\boldsymbol{\eta}(\mathbf{q}_{k-1})}{d\mathbf{q}}, \quad \mathbf{K}_k = -\frac{1}{h_k} \mathbf{I}.$$

- Build $\mathbf{A}_{\text{eq}}^{(k)}$ and $\mathbf{b}_{\text{eq}}^{(k)}$ as in Section 2.4. For Backward Euler, the equality system given by Eqs. (13)–(14) is used directly. For the Trapezoidal Rule, the first block row is replaced by

$$\mathbf{A}_1 = \left[\mathbf{I} - \frac{h_k}{2} \mathbf{A} \quad -h_k \mathbf{G} \right],$$

$$\mathbf{b}_1 = \left(\mathbf{I} + \frac{h_k}{2} \mathbf{A} \right) \mathbf{x}_{k-1} + h_k \mathbf{B} \mathbf{u}_k.$$

- Solve the resulting one-step QP, or the equality-solve form of Section 3.2 when no additional costs or inequality constraints are active. This gives the trial state $\mathbf{z}_k^* = \hat{\Phi}_{h_k}(\mathbf{z}_{k-1}, \mathbf{u}_k)$.

- Compute the post-step nonlinear residual

$$\hat{\boldsymbol{\eta}}(\mathbf{q}_k^*) = \boldsymbol{\eta}(\mathbf{q}_{k-1}) + \mathbf{J}_{k-1}(\mathbf{q}_k^* - \mathbf{q}_{k-1}),$$

$$\boldsymbol{\nu}_k = \boldsymbol{\eta}(\mathbf{q}_k^*) - \hat{\boldsymbol{\eta}}(\mathbf{q}_k^*), \quad \delta_k^{\text{QP}} = \|\boldsymbol{\nu}_k\|.$$

- Accept the trial step if $\delta_k^{\text{QP}} \leq \text{tol}$. If the step is rejected, reduce h_k and recompute the same interval from $\mathbf{z}_{k-1}$. The reduction of h_k is done using one of the residual-control rules considered in this work: the deadbeat (classical) update, the PID H321 update, or the predictive H0211 update as defined by Söderlind in [13]. Since $\delta_k^{\text{QP}} = \mathcal{O}(h_k^2)$, the controller order is selected as $\kappa = 2$. In all cases, the step size is clipped to $h_{\min} \leq h_k \leq h_{\max}$. Once the step is accepted, set $\mathbf{z}_k = \mathbf{z}_k^*$.

- For an accepted step, compute the circuit output $\mathbf{y}_k$ (Eq. (16)).

- Reset the time-step to a higher value, $h_k = \bar{h}_0 \leq h_0$, to avoid unnecessarily small values that may cause the simulation to take too long, and repeat from Step 5.

An example MATLAB implementation can be found in the Code section on the accompanying website ¹.

4. EXPERIMENTS AND RESULTS

To evaluate the performance of the proposed method, we analyzed the three circuits depicted in Figure 1, where the diode clipper (a) is considered a baseline comparison with respect to our original formulation in [7] before the modifications. In addition, the common-emitter amplifier (b) is proposed as a way to evaluate BJT transistor-based circuits, and it is the same as the one employed in [5]. Finally, mostly as a stress test for the numerical method, we proposed a Colpitts-type [15] oscillator (c), as it is well known that self-oscillating systems are particularly sensitive to numerical errors and instabilities. The respective derivations and required matrices for each circuit can be found in the Supplementary Material section on the accompanying website ².

As a baseline DAE defect indicator to compare the proposed δ_k^{QP} (Eq. (22)) to, we used a basic step-doubling scheme in which the defect indicator (SDDI) is computed as

$$\delta_k^{\text{SD}} = \left\| \hat{\Phi}_{h_k}(\mathbf{z}_k, \mathbf{u}_k) - \hat{\Phi}_{h_{k/2}} \left(\hat{\Phi}_{h_{k/2}}(\mathbf{z}_k, \mathbf{u}_k), \mathbf{u}_k \right) \right\|,$$

¹<https://mezea-uvg.github.io/RAMA/#code>

²<https://mezea-uvg.github.io/RAMA/supplementary/>

and the time-step keeps halving until acceptance. This indicator is well known [9, 10] to work in general, but it is seldom used given its considerable computational cost. This is especially relevant when a QP needs to be solved. Additionally, we considered two discretization schemes for the LTI state dynamics in the form of Backward Euler and the Trapezoidal Rule, and the three time-step size update rules mentioned in Step 9 of Section 3.4. Table 1 presents the full list of experiments carried out, along with the corresponding average execution times obtained. All computations were performed in a Windows-based personal computer with an AMD Ryzen 7 5800H CPU.

Figures 2a–2d show the results for the clipper; Figures 3a–3c show the results for the amplifier; and Figures 4a–4c show the results for the oscillator. For all cases, we start with the standard $h_0 = 1/44100$ s and consistent initial conditions equal to zero.

Table 1: List of experiments. a.0 – a.8 correspond to the diode clipper circuit, with a 2.205 kHz, 5V sinusoidal input (2.5 periods). b.0 – b.6 correspond to the common-emitter amplifier, with a 220 Hz, 0.5V sinusoidal input (22 periods). c.0 – c.3 correspond to the Colpitts oscillator, with a 12V step input (5 ms). The average execution times are computed over 10 repetitions for each case. As a reminder: QPDI is the adaptive defect indicator of the QP one-step map, SDDI stands for the step-doubling scheme’s defect, ECQP stands for the equality-constrained QP, PINV corresponds to using the Moore-Penrose pseudo-inverse, BE stands for Backward Euler and TRAP for the Trapezoidal Rule, pid stands for the PID H321 update and predictive for the predictive H0211 update.

Case	Experiment parameters	$\bar{h}_0$	Avg. exec. times (ms)
a.0	SPICE	N/A	80.70
a.1	QPDI; BE; ECQP; fixed h	N/A	43.14
a.2	QPDI; BE; PINV; fixed h	N/A	1.38
a.3	QPDI; BE; PINV; deadbeat	$h_0/4$	5.26
a.4	QPDI; BE; PINV; pid	$h_0/4$	5.55
a.5	QPDI; BE; PINV; predictive	$h_0/4$	8.50
a.6	QPDI; TRAP; ECQP; fixed h	N/A	43.87
a.7	QPDI; TRAP; PINV; deadbeat	$h_0/4$	5.89
a.8	SDDI; BE; PINV; fixed h	N/A	3.05
b.0	SPICE	N/A	210.80
b.1	QPDI; BE; PINV; deadbeat	h_0	100.90
b.2	QPDI; BE; PINV; pid	$h_0/2$	238.56
b.3	QPDI; BE; PINV; predictive	h_0	99.01
b.4	QPDI; TRAP; PINV; deadbeat	h_0	102.60
b.5	QPDI; TRAP; PINV; pid	$h_0/2$	248.32
b.6	QPDI; TRAP; PINV; predictive	h_0	99.39
c.0	SPICE	N/A	215.60
c.1	QPDI; TRAP; PINV; deadbeat	$h_0/8$	99.98
c.2	QPDI; TRAP; PINV; pid	$h_0/8$	186.19
c.3	QPDI; TRAP; PINV; predictive	$h_0/8$	85.65

5. DISCUSSION

Figure 2a experimentally validates the equivalence between the equality-constrained QP and pseudo-inverse formulations of the method, as stated in Section 3.2. This is confirmed by the plots corresponding to cases a.1 and a.2, which are indistinguishable. Additionally, the figure shows that a Backward Euler scheme has better performance than a Trapezoidal Rule-based scheme under




Figure 1: Circuits considered in the experiments: (a) Baseline diode clipper circuit from [6, 7]. (b) BJT transistor-based common-emitter amplifier from [5]. (c) Simplified common-emitter Colpitts oscillator topology adapted from standard BJT Colpitts forms [15].




Figure 2: SPICE baseline and numerical results for the diode clipper circuit for several cases of the proposed method (a.0 – a.8). (a) Time comparison of 2.5 periods of the outputs. (b) Squared error in dB between both fixed and adaptive step-size results of the proposed method against the SPICE baseline. (c) Defect indicator δ_k . (d) Time-step size for a fixed scheme and adaptive schemes against the step-doubling with fixed step-size baseline.




Figure 3: SPICE baseline and numerical results for the common emitter amplifier for several cases of the proposed method (b.0 – b.6). (a) Time comparison of 22 periods of two cases of the proposed method's output. (b) Zoomed-in segment of the time comparison, for all cases. (c) Squared error in dB between the SPICE baseline and all cases.

a fixed time-step, proving that the former's stability properties are better suited to filter out the high frequency artifacts introduced by the clipper's nonlinearities. This difference in discretization, however, becomes negligible when using an adaptive time-step scheme like the deadbeat, as can be seen in the plots corresponding to cases a.3 and a.7, which are also indistinguishable. The use of an adaptive time-step scheme allows the proposed method's output to track the SPICE baseline more accurately, even with an input frequency of $1/20$ of the standard sampling frequency of 44.1 kHz. This can be verified by comparing the plots corresponding to cases a.3 and a.7 with the one corresponding to a.0.




Figure 4: SPICE baseline and numerical results for the Colpitts oscillator using the proposed method (case c.1, input extended to 20 ms). (a) Time comparison of the output for the entire 20 ms after the step input. (b) Zoomed-in segments of the output signals in (a): on the left, the first 1 ms; on the right, the last 0.5 ms. (c) Comparison of the single-sided magnitude spectrum of the outputs, in dB.

The squared-error metric in Figure 2b experimentally shows that the discrepancy between the method’s output and the SPICE baseline significantly diminishes when it is allowed to use an adaptive step-size. Overall, the squared-error is considerably lower for cases a.3 and a.7. The discrepancy is largest when the input begins to be affected by the nonlinearities of the circuit, i.e., when the diodes effectively start clipping the input. The adaptive time-stepping scheme allows the method to mitigate this increase in nonlinearities by reducing the step-size accordingly. This adaptive behavior is captured in Figures 2c–2d, where the proposed residual indicator spikes when most of the nonlinearities are present. Correspondingly, the step-size reduces in those instances. This behavior is considerably clearer than the one obtained with the baseline step-doubling indicator in case a.8. While the proposed indicator localizes the start of nonlinear events through narrow residual peaks, the step-doubling indicator remains comparatively flat and exhibits wider drops before and after the clipping regions, contrary to what intuitively would be expected. Thus, for this benchmark, the proposed indicator provides a more interpretable measure of the nonlinear DAE defect and leads to a more physically meaningful adaptive step-size profile. On the other hand, in the absence of an adaptive time-stepping scheme, the residual remains high until the circuit returns to a mostly linear region of operation, experimentally validating that it effectively captures the strong nonlinearities of the circuit. We also observe that the specific adaptive update law is not critical for the simple clipper example, since the deadbeat, PID, and predictive cases (a.3, a.4, and a.5) show almost the same behavior and accepted time-step values.

Figures 3a–3c show that the proposed method also generalizes well to the common-emitter amplifier, which is a more representative BJT-based circuit than the baseline clipper. In this case, all tested adaptive cases remain in close agreement with the SPICE baseline, both in the time-domain and in the squared error waveforms. This suggests that the proposed residual indicator and adaptive update rules remain effective beyond diode-only examples, and that the method preserves its practical accuracy when applied

to transistor-based amplifier topologies of clearer relevance in virtual analog audio.

Figures 4a–4c show encouraging results for the Colpitts oscillator, which was included mainly as a numerical stress test due to the well-known sensitivity of self-oscillating circuits to discretization errors and long-term instability. Even in this more demanding case, the proposed method reproduces the oscillation build-up, steady-state waveform, and output spectrum with very good agreement with respect to the SPICE baseline. The spectrum shows an almost perfect match between the resonant frequencies obtained with SPICE and our method (11.3 kHz vs. 11.35 kHz in this circuit). The slight difference explains what is shown in Figure 4b. At first, our method’s output slightly lags behind SPICE’s output (left image), but as time goes by, our method’s output slightly leads ahead (right image). Nevertheless, the waveforms are essentially the same, and the oscillating nature of the circuit is well reproduced by our method. However, this is at the cost of ending up with a minimum time-step of around 200 ns, which would correspond to an oversampling factor of 128 with respect to the initial time-step of 1/44100 s.

The preliminary average execution times shown in Table 1 are also encouraging and important. In addition to comparing the outputs produced by our method with the SPICE baseline, we wanted to compare the execution times to further validate the claims made in [7] regarding real-time implementation feasibility. In the case of the clipper circuit, all cases outperform the SPICE execution, including those using the equality-constrained QP. As expected, the execution times were the lowest when using the pseudo-inverse approach. The decrease was by an order of magnitude with respect to SPICE, even though the implementation was done in MATLAB and not in a more efficient language such as C/C++. It is worth noting that the best trade-off between output accuracy and execution time was achieved with $\bar{h}_0 = h_0/4$ for all time-step update rules. In the case of the amplifier circuit, the average execution times were less than half the average of SPICE, except when using the PID step-update rule. The best results were obtained with

$\bar{h}_0 = h_0$ for the deadbeat and predictive rules, and with $\bar{h}_0 = h_0/2$ for the PID rule. For the oscillator circuit, once again the average execution times were smaller using our method compared to those obtained with SPICE. The times were less than half when using the deadbeat and predictive rules. In these cases, the best results were obtained using $\bar{h}_0 = h_0/8$.

Finally, as a preliminary real-time implementation test, we generated VST audio plugin demonstrations of the diode clipper and common-emitter amplifier benchmarks using MATLAB’s Audio Toolbox. In both cases, the Toolbox was used to export the implementations as .dll VST plugins, which were then loaded and tested in REAPER. The plugins ran in real-time without perceptible dropouts, even when using the computer’s integrated sound card drivers, providing encouraging preliminary evidence of practical real-time feasibility. More importantly, the plugin implementations allowed circuit and device-level parameters to be exposed as real-time controls. The generated VST plugins are available for download from the accompanying website³.

6. CONCLUSIONS AND FUTURE WORK

This work formalized and extended our QP-based formulation for nonlinear analog circuit emulation presented in [7]. In particular, the method was reframed in a more explicit numerical-analysis framework, the post-step nonlinear residual was justified as a valid defect indicator for adaptive integration, and the resulting adaptive multi-rate scheme was tested on three representative circuits of increasing difficulty. Across the experiments, the proposed method achieved very low discrepancy with respect to SPICE while preserving an interpretable state-space structure. The results also support two practical observations: firstly, that adaptive step sizing is more important than using a fixed high-rate discretization alone; and secondly, that in the tested cases, solving the equality-constrained problem through the Moore–Penrose pseudo-inverse provides essentially the same output as the QP formulation at a lower practical complexity and computational cost.

Future work will address the main limitations of the current implementation. This includes a more systematic study of real-time behavior, with latency-oriented benchmarks and sparse or embedded-oriented solvers. The choice of $\bar{h}_0$ and the QP metric should also be further investigated. The latter, especially from energy- or passivity-motivated perspectives for circuits. The method should then be evaluated on a broader set of nonlinear and modulated audio circuits, using longer simulations and realistic musical inputs. Finally, an automated pipeline from circuit netlists or schematics to the matrices and nonlinear constraints used by the method would make the approach easier to apply beyond hand-derived benchmark examples.

7. ACKNOWLEDGMENTS

We would like to thank Carlos Esquit and the Department of Electronics, Mechatronics and Biomedical Engineering at UVG, for the financial support during the development of this work.

8. REFERENCES

- [1] E. Moliner, M. Švento, A. Wright, L. Juvela, P. Rajmic, and V. Välimäki, “Unsupervised estimation of nonlinear audio effects: Comparing diffusion-based and adversarial approaches,” in *Proceedings of the 28th International Conference on Digital Audio Effects (DAFx25)*, Ancona, Italy, Sep. 2025, pp. 366–373.
- [2] O. Massi, A. Ilic Mezza, R. Giampiccolo, and A. Bernardini, “Training neural models of nonlinear multi-port elements within wave digital structures through discrete-time simulation,” in *Proceedings of the 28th International Conference on Digital Audio Effects (DAFx25)*, Ancona, Italy, Sep. 2025, pp. 94–101.
- [3] F. Bernardo, M. Davison, and A. McPherson, “Impedance synthesis for hybrid analog-digital audio effects,” in *Proceedings of the 28th International Conference on Digital Audio Effects (DAFx25)*, Ancona, Italy, Sep. 2025, pp. 31–38.
- [4] D. T. Yeh, J. S. Abel, and J. O. Smith, “Automated physical modeling of nonlinear audio circuits for real-time audio effects—part I: Theoretical development,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 4, pp. 728–737, 2010.
- [5] D. T. Yeh, “Automated physical modeling of nonlinear audio circuits for real-time audio effects—part II: BJT and vacuum tube examples,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 4, pp. 1207–1216, 2012.
- [6] M. Holters and U. Zölzer, “A generalized method for the derivation of non-linear state-space models from circuit schematics,” in *2015 23rd European Signal Processing Conference (EUSIPCO)*, Nice, France, Aug.-Sept. 2015, pp. 1073–1077.
- [7] M. Zea and L. A. Rivera, “Towards efficient emulation of nonlinear analog circuits for audio using constraint stabilization and convex quadratic programming,” in *Proceedings of the 28th International Conference on Digital Audio Effects (DAFx25)*, Ancona, Italy, Sep. 2025, pp. 1–6.
- [8] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2002.
- [9] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd ed., ser. Springer Series in Computational Mathematics. Berlin, Heidelberg: Springer, 1993, vol. 8.
- [10] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, 2nd ed., ser. Springer Series in Computational Mathematics. Berlin, Heidelberg: Springer, 1996, vol. 14.
- [11] K. E. Brenan, S. L. Campbell, and L. R. Petzold, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*. Philadelphia, PA: SIAM, 1996.
- [12] G. Söderlind, “Automatic control and adaptive time-stepping,” *Numerical Algorithms*, vol. 31, no. 1–4, pp. 281–310, 2002.
- [13] —, “Digital filters in adaptive time-stepping,” *ACM Transactions on Mathematical Software*, vol. 29, no. 1, pp. 1–26, 2003.
- [14] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York: Springer, 2006.
- [15] D. A. Bell, *Electronic Devices and Circuits*, 5th ed. Oxford University Press, 2008.

³<https://mezea-uvg.github.io/RAMA/#demo>