



Optimizing EPR pulses for broadband excitation and refocusing

Eric R. Lowe^{a,*}, Stefan Stoll^b, J.P. Kestner^a

^a Department of Physics, University of Maryland Baltimore County, Baltimore, MD 21250, USA

^b Department of Chemistry, University of Washington, Seattle, WA 98195, USA

ARTICLE INFO

Keywords:

Broadband pulses
Hahn echo
Optimal control
EPR spectroscopy

ABSTRACT

In this paper, we numerically optimize broadband pulse shapes that maximize Hahn echo amplitudes. Pulses are parameterized as neural networks (NN), nonlinear amplitude limited Fourier series (FS), and discrete time series (DT). These are compared to an optimized choice of the conventional hyperbolic secant (HS) pulse shape. A power constraint is included, as are realistic shape distortions due to power amplifier nonlinearity and the transfer function of the microwave resonator. We find that the NN, FS, and DT parameterizations perform equivalently, offer improvements over the best HS pulses, and contain a large number of equivalent optimal maxima, implying the flexibility to include further constraints or optimization goals in future designs.

1. Introduction

The use of shaped pulses in electron paramagnetic resonance (EPR) spectroscopy is a topic of recent interest [1–6]. They address the basic challenge that the excitation bandwidth of monochromatic square pulses is much smaller than the spectral line width of samples. This situation can arise in nuclear magnetic resonance (NMR), but it is more broadly relevant in EPR. Typical EPR spectral widths are about 250 MHz for a nitroxide at Q-band (≈ 1.2 T) or ≈ 2 GHz for a Cu(II) complex at X-band (≈ 0.35 T). Using shaped pulses can increase sensitivity and excitation bandwidth.

Broadband pulses have been initially designed for NMR, including the Kunz–Böhlen–Bodenhausen (KBB) approach to generate a Hahn echo [7–9]. This sequence consists of a frequency-swept (chirped) $\pi/2$ pulse of length t_p , followed by a chirped π pulse of length $t_p/2$. The intention of this sequence is to refocus all the spins within a broad excitation window. The Fourier transform of this echo gives the spectral distribution of the excited spin ensemble. The more complete the refocusing, the larger the signal and more accurate the reconstruction of the spectrum. The original KBB scheme used pulses with constant amplitudes and a linear frequency sweep over the designated bandwidth. Performance can be further improved by shaping the pulses to have an adiabatic hyperbolic secant (HS) amplitude with the frequency swept according to a hyperbolic tangent [10–13]. In the last decade, pulse shaping has become possible in EPR as arbitrary waveform generators (AWGs) became fast enough to generate shaped pulses that cover bandwidths larger than those obtainable by hard square pulses. For this, microwave pulse amplitude and phase are modulated with sub-ns timing resolution. Currently, AWGs with sampling rates of 1.25 GS/s or faster are in use.

However, HS pulses are not optimal. In practice, limited available power limits the maximum achievable pulse amplitude which in turn limits the frequency sweep rate and therefore puts a lower bound on the pulse time. Not only that, relaxation can put an upper limit on pulse duration. Also, in some sequences, pulse lengths must be shorter than the evolution periods of the interactions of interest, for instance dipolar couplings in dipolar EPR. [13]. Thus, along with the constraint of limited power, constraining the pulse time can cause the performance of pulses to suffer.

In response, optimal control methods such as composite pulses [14, 15], adiabatic pulses [16,17], optimal control theory (OCT) pulses using different numerical algorithms [18–22], including gradient ascent pulse engineering (GRAPE) [23–28], were developed to accomplish broadband excitation and uniform inversion across a given bandwidth. In practice, optimal shaped pulses are distorted by nonlinearities in the power amplifier and by the resonator transfer function, moving them away from the extremum in the optimization landscape. One can try to compensate for these distortions post-optimization, but the necessary compensation may not be possible while respecting constraints such as limited power at fixed pulse time.

In this paper, we use an optimization method that includes a model of the full experimental transfer chain and resulting shape distortions while limiting both the available power along with the length of pulses. We investigate the use of variously parameterized broadband pulses for a Hahn echo sequence and the effect of varying the pulse length ratio on echo amplitude, refocusing time, and refocusing phase. Due to the freedom of the large parameter space and low number of constraints, we find the individual pulses ($\pi/2$ and π) act cooperatively

* Corresponding author.

E-mail addresses: elowe2@umbc.edu (E.R. Lowe), stst@uw.edu (S. Stoll), jkestner@umbc.edu (J.P. Kestner).

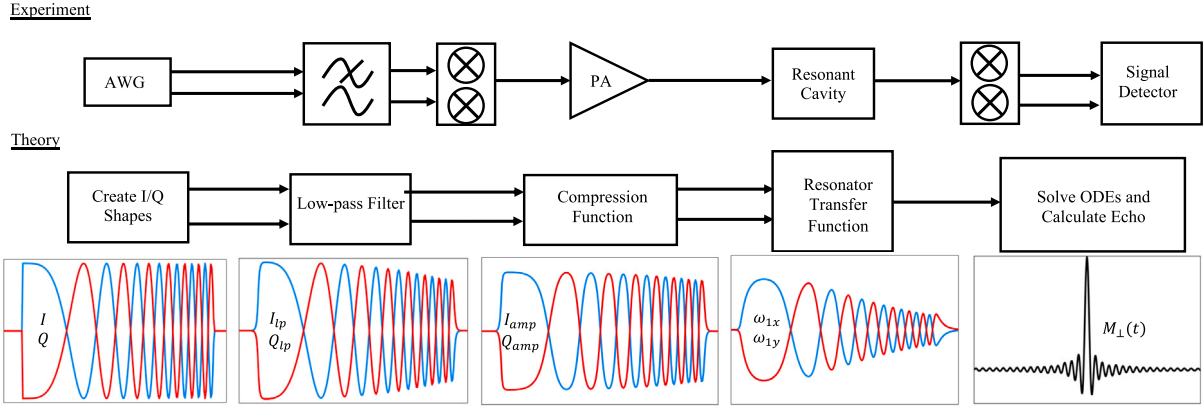


Fig. 1. A schematic overview of the experimental transmitter setup (top), its computational representation (middle) and exemplary signals (bottom). The input I/Q shapes are modeled and sent through a low-pass filter, amplifier compression function, and a resonator transfer function before they are used in the spin quantum dynamics calculation of the echo.

as done previously with NMR COOP pulses [29–32]. This shows that cooperatively performing pulses are still optimal under transmission distortions and in the presence of a power constraint while in limited pulse length time. Performances of individual pulses as well as of the entire pulse sequence are compared. Section 2 summarizes the transmitter model and the spin physics model, Section 3 describes the pulse parameterizations used, Section 4 provides details about the optimization method, and Section 5 discusses the results.

2. Model

In order to model distortions that affect the pulse shapes, we closely follow the transmitter design of a typical EPR spectrometer. This is illustrated in Fig. 1. First, AWG-generated in-phase and quadrature drive functions $I(t)$ and $Q(t)$ with amplitudes in the range $[-1, 1]$ are set up using different parameterizations described in detail in the next section. We model the limited output bandwidth of the AWG by applying a low-pass filter with transfer function

$$H_{lp}(\omega) = \frac{1}{1 + (\omega/\Gamma)^2}. \quad (1)$$

to I and Q via convolution, yielding

$$I_{lp}(t) = V_{DAC} \mathcal{F}^{-1} [H_{lp}(\omega) \cdot \mathcal{F} [I(t)]] \quad (2)$$

where \mathcal{F} represents the Fourier transform, and V_{DAC} represents the conversion factor from the digital-to-analog converter (DAC). Γ represents the 3-dB bandwidth of the filter. The Q channel pulse shape is similarly distorted into $Q_{lp}(t)$. The conversion factor between the dimensionless input and the voltage output is combined with other overall multiplicative factors at the end of the transmission chain model.

The IQ upconversion of the low-pass filtered $I_{lp}(t)$ and $Q_{lp}(t)$ to the carrier frequency ω_c yields

$$V(t) = I_{lp}(t) \cos(\omega_c t) \pm Q_{lp}(t) \sin(\omega_c t) \quad (3)$$

where the sign depends on whether the carrier reference for the Q channel mixer is phase shifted +90 or −90 degree relative to the carrier for the I channel mixer.

The upconverted signal can be rewritten in terms of an amplitude- and phase-modulated oscillation at the carrier frequency:

$$V(t) = V_0(t) \cos(\omega_c t + \varphi(t)) \quad (4)$$

with the time-varying amplitude

$$V_0(t) = \sqrt{I_{lp}^2(t) + Q_{lp}^2(t)} \quad (5)$$

and phase

$$\varphi(t) = \text{atan2}(Q_{lp}, I_{lp}), \quad (6)$$

where atan2 is the two-argument arctangent. The amplifier amplifies and additionally distorts this signal, leading to the appearance of higher harmonics in the amplifier output, $V_{amp}(t)$. Assuming that the amplifier is memory-less, the higher harmonics are rejected by the narrow-band transmission lines, and assuming separation of timescales, i.e. $V_0(t)$ and $\varphi(t)$ vary much more slowly than the carrier signal $\cos(\omega_c t)$, the amplified signal is represented by

$$V_{amp}(t) = G(V_0(t)) V_0(t) \cos(\omega_c t + \varphi(t)) \quad (7)$$

where G is the gain function, or in terms of I_{lp} and Q_{lp}

$$V_{amp}(t) = I_{amp}(t) \cos(\omega_c t) + Q_{amp}(t) \sin(\omega_c t) \quad (8)$$

with

$$I_{amp}(t) = G(V_0(t)) I_{lp}(t) \quad (9)$$

and a similar expression for $Q_{amp}(t)$.

These equations only model amplitude-to-amplitude modulation (AM/AM) effects of the amplifier and neglect possible amplitude-to-phase modulation (AM/PM) effects. AM/PM effects would alter the term $\varphi(t)$ in Eq. (7) by mixing the I and Q signals.

In order to carry out numerical optimizations we have to specify a gain function, and we will use

$$G(V_0(t)) = g \frac{\tanh(V_0(t)/V_{sat})}{V_0(t)/V_{sat}}, \quad (10)$$

where g is the small-signal gain factor and V_{sat} the input saturation amplitude. V_{sat} parameterizes the nonlinearity: for $V_0(t) \ll V_{sat}$ the amplifier is in the linear regime, while for $V_0(t) \gg V_{sat}$ the amplifier is saturated. We assume V_{sat} is constant over the amplifier bandwidth and that the amplifier bandwidth is wider than the signal bandwidth. In principle, we could use different nonlinear models such as Rapp, Saleh, or polynomial models [33], or a tabulated function.

Next, the amplified pulse is transmitted to the resonator. The resonator transfer function is well described by

$$H_{res}(\omega) = \frac{1}{1 + iQ_L(\frac{\omega}{\omega_{res}} - \frac{\omega_{res}}{\omega})}, \quad (11)$$

where Q_L is the loaded Q-value and ω_{res} is the resonator frequency. This produces the following pulse shape at the sample inside the resonator,

$$B_1(t) = C \text{Re}(\mathcal{F}^{-1}[H_{res}(\omega) \cdot \mathcal{F}[V_{phasor}(t)]]), \quad (12)$$

where

$$V_{phasor}(t) = (I_{amp}(t) + iQ_{amp}(t)) e^{-i\omega_c t} \quad (13)$$

C is the resonator conversion factor. In the sample, spins with gyro-magnetic ratio γ experience the drive function

$$\gamma B_1(t) = \omega_1(t) \cos \phi(t) \quad (14)$$

with the drive amplitude ω_1 and the phase ϕ . With this transmitter chain, the maximum drive amplitude that can be achieved with $|I| = |Q| = 1$ and $\omega = \omega_{\text{res}}$ is $\omega_{1,\text{max}} = \gamma g C V_{\text{sat}} \tanh(\sqrt{2} V_{\text{DAC}} / V_{\text{sat}})$.

Note that there is a subtle difference between imposing separate constraints on I and Q as we do here versus simply imposing an overall amplitude constraint. Constraining I and Q independently restricts the accessible domain in the I/Q plane to a square rather than a circle. Maximum power is only attained at the four corners of the square when $|I| = |Q| = 1$, i.e., phases of $\pm 45^\circ, \pm 135^\circ$. Changing the phase of a maximum-power pulse will generally result in a power loss. However, phase cycling can still be carried out in $\pm 90^\circ$ increments since this domain is symmetric under $\pm 90^\circ$ rotations.

The laboratory frame Hamiltonian for a given spin packet with Larmor frequency ω_{res} is

$$H_{\text{lab}}(\omega_{\text{res}}, t) = \omega_{\text{res}} S_z + \omega_1(t) \cos(\phi(t)) S_x, \quad (15)$$

where S_i ($i = x, y, z$) are spin operators. For EPR experiments, $\omega_{\text{res}}/2\pi$ is typically on the order of 10–100 GHz while $\omega_1/2\pi$ is on the order of tens of MHz, so the rotating-wave approximation is valid when moving to a frame rotating at the carrier frequency, ω_c . The rotating-frame Hamiltonian for a spin packet off-resonant with the carrier frequency by $\Delta\omega = \omega_{\text{res}} - \omega_c$ is

$$H_{\text{rot}}(\Delta\omega, t) = \Delta\omega S_z + \omega_{1x}(t) S_x + \omega_{1y}(t) S_y, \quad (16)$$

where

$$\omega_{1x} = \omega_1(t) \cos(\phi(t) - \omega_c t), \quad (17)$$

$$\omega_{1y} = \omega_1(t) \sin(\phi(t) - \omega_c t). \quad (18)$$

In experiments, the signal passes back through the resonator and, after downconversion, a bandwidth-limited video amplifier. The resulting receiver-side distortions affects all echo signals equally and the optimal echo shape, obtained when all spin packets align on the same axis along the xy -plane, does not change. Incorporating receiver distortions does not affect the pulse optimization, as the optimal echo shape is formed when all spins are aligned, regardless of receiver distortions, and so we do not include it in the model.

We now consider an ensemble of spin packets, each governed by a Hamiltonian with its own ω_{res} , ranging over a spectral distribution of width δ . We choose to place ω_c in the center of this distribution, so that $\Delta\omega$ ranges from $-\delta/2$ to $+\delta/2$. To obtain the effect of a given pulse we must solve an ensemble of Schrödinger equations for the propagators U ,

$$i\dot{U}(\Delta\omega, t) = H_{\text{rot}}(\Delta\omega, t)U(\Delta\omega, t), \quad (19)$$

corresponding to the different resonant frequencies across the relevant spectral width.

With our mapping between the I/Q inputs and the driving functions in the spin Hamiltonian, ω_{1x} and ω_{1y} , we can solve Eq. (19) for both the $\pi/2$ and π pulses for a representative ensemble of Larmor frequencies within the desired band. Using the initial condition of $U(t=0) = \mathbb{1}$, where $\mathbb{1}$ is the identity matrix, we can solve the ensemble of differential equations, giving us a corresponding evolution operator for each Larmor frequency. The evolution operators for the $\pi/2$ and π pulse are denoted as $U_{\pi/2}^i$ and U_{π}^i , where i indexes the Larmor frequencies. We combine these propagators with the free evolution periods in the rotating frame to obtain the total propagators

$$U_{\text{tot}}^i(t) = U_{\text{free}}^i(t) U_{\pi}^i(\tau_1) U_{\pi/2}^i \quad (20)$$

where $t = 0$ now denotes the end of the π pulse and the beginning of the second free evolution period. The signal due to a particular spin packet is

$$M_{\perp,i}(t) = 2 \text{tr} \left(U_{\text{tot}}^i(t) \rho_0 U_{\text{tot}}^{i\dagger}(t) S_+ \right), \quad (21)$$

where ρ_0 is the initial density matrix for each spin, $\rho_0 = \mathbb{1}/2 - S_z$.




Fig. 2. Four different parameterizations of a pulse shape (blue: in-phase, red: out of phase).

The total magnetization signal from all spins is given by averaging over all spin packets,

$$M_{\perp}(t) = \frac{1}{N} \sum_{i=1}^N M_{\perp,i}(t), \quad (22)$$

where in our optimization the Larmor frequencies were sampled from a uniform distribution. Our goal is to determine optimal shapes of $I(t)$ and $Q(t)$ that maximize the total magnetization signal, $M_{\perp}(t)$. As a check on the spin physics, all of the pulses generated were also tested and confirmed using the MATLAB package Easyspin, an open-source software that allows for the simulation and analysis of EPR spectra [34,35]. The optimized pulse shapes and code are publicly available [36].

3. Parameterizations

We start by considering the original KBB broadband pulse sequence using generalized HS shapes for the $\pi/2$ and π pulses. KBB is not the only broadband refocusing sequence we could use. For instance, the CHORUS sequence [37,38] uses linear swept pulses with effectively rectangular amplitude profiles that have more integrated power than hyperbolic secant pulses for a fixed B_1 amplitude and sequence time, and features improved robustness to B_1 field inhomogeneity [21,39]. However, we consider a strongly power-constrained regime with limited amplitude and pulse lengths where B_1 field inhomogeneity is a secondary concern. In this regime, HS pulses are preferable because they require less power off resonance, where it is costly to compensate for the profile of the resonator transfer function. Therefore, in this work we will use the KBB sequence with HS pulses as a baseline against which to compare other shaped pulses (see Fig. 2). Then, while retaining the general bipartite structure of a $\pi/2$ pulse followed by a π pulse, to introduce more shape flexibility, we consider three models with significantly more parameters: a nonlinear amplitude-limited Fourier series (FS), a discrete-time series (DT), and a neural network (NN). We do not constrain the pulse flip angles to be $\pi/2$ and π , but we will still refer to them by those labels in continuity with the KBB design and in anticipation that the optimization process will indeed drive them to be such.

3.1. Generalized HS pulses

For representing generalized HS pulses, we use excitation functions of the form [4,40]

$$\omega_1(t) = A \text{sech}(2^{n-1} \beta |t/T|^n) \quad (23)$$

$$\dot{\phi}_{\text{HS}}(t) = \Delta\omega_{\text{BW}} \left(\frac{\int_{-T/2}^t \text{sech}^2(2^{n-1}\beta|\tau/T|^n) d\tau}{\int_{-T/2}^{T/2} \text{sech}^2(2^{n-1}\beta|\tau/T|^n) d\tau} - \frac{1}{2} \right), \quad (24)$$

where $-T/2 \leq t \leq T/2$, and $\omega_1(t)$ and $\dot{\phi}_{\text{HS}}(t)$ are the amplitude and instantaneous driving frequency at time t in the rotating frame. Although n is usually considered to be a positive integer ($n = 1$ in the original KBB scheme), we extend the definition to include non-integer n by using the absolute value of the time. Similarly, while $\Delta\omega_{\text{BW}}$ is typically set to the desired excitation bandwidth, we allow this to be a free parameter as well.

These functions are converted to ω_{1x} and ω_{1y} used in Eq. (16) by using

$$\omega_{1x} = \omega_1(t) \cos(\phi_{\text{HS}}(t)), \quad (25)$$

$$\omega_{1y} = \omega_1(t) \sin(\phi_{\text{HS}}(t)), \quad (26)$$

with the phase $\phi_{\text{HS}}(t)$ given by

$$\phi_{\text{HS}}(t) = \int_{-T/2}^t \dot{\phi}(\tau) d\tau. \quad (27)$$

The $\pi/2$ and π pulses of this form are chosen such that, following KBB, the two pulse durations have the ratio $T_{\pi/2}/T_{\pi} = 2$. The values of β , n , and $\Delta\omega_{\text{BW}}$, along with ω_{1x} and ω_{1y} for each of the two pulses are free parameters, and we optimize them for refocusing spins across a bandwidth of δ . We do not require the two pulses to produce $\pi/2$ and π rotations for any spin packet, but only restrict $A \leq \omega_{1,\text{max}}$.

To benefit from the intuitive behavior of HS pulses, we have to directly parameterize the pulse at the *output* of the transmitter chain of Fig. 1, then invert the transfer functions in order to obtain the I/Q inputs that should be programmed into the AWG. In contrast, the other parameterizations considered below are for the *input* I/Q pulse shapes themselves and do not require pre-compensation. Also note that the HS pulses satisfy $I(t)^2 + Q(t)^2 \leq 1$, lying within the inscribing circle of the square domain ($|I(t)| \leq 1, |Q(t)| \leq 1$) available to the other parameterizations on the I/Q plane, so the HS parameterization is clearly at a disadvantage to begin with because it cannot access as much power as a parameterization that allows, for example, $I = Q = 1$. However, in order to focus on the less obvious differences between parameterizations, we have allowed the HS pulses to access the *circumscribing* circle, $I(t)^2 + Q(t)^2 \leq 2$.

3.2. Nonlinear amplitude-limited Fourier series

The second pulse shape model we consider consists of a nonlinear amplitude-limited Fourier series (FS) for the two drive functions

$$I(t) = \tanh \left[\sum_{n=1}^N a_{I,n} \cos \left(\pi n \frac{t}{T} \right) \right] \quad (28)$$

$$Q(t) = \tanh \left[\sum_{n=1}^N a_{Q,n} \cos \left(\pi n \frac{t}{T} \right) \right] \quad (29)$$

for each pulse. Here again $-T/2 \leq t \leq T/2$, and $a_{I,n}$ and $a_{Q,n}$ are real-valued coefficients.

The enclosing tanh function limits I and Q to values between -1 and 1 by construction. The reasoning behind imposing the amplitude limit using tanh, as opposed to scaling the coefficients, is that the maximum of the Fourier series signal can only be determined using a numerical search, which renders the cost function itself non-differentiable. Due to the nonlinearity of tanh, the bandwidths of I and Q are not straightforwardly related to the frequencies included in the cosine series, and the modeled shapes are nonlinear and compressed compared to a standard cosine series. We choose N large enough to cover the desired bandwidth, $\pi N/T \approx \delta$. Including higher-order terms in the series does not improve performance, as those terms are severely attenuated by the resonator in the relevant case where δ is comparable to resonator bandwidth ω_{res}/Q_L . We perform an unconstrained optimization over (a_I, a_Q) for each pulse.




Fig. 3. A representation of the neural network utilized where the input layer is a single node consisting of the time t' , 3 hidden layers of 16 nodes each, and the output layer of the 4 drive signal amplitudes at time t' .

3.3. Discrete-time series

The third pulse shape model we consider consists of discrete-time series (DT) for I and Q for each pulse

$$I = \tanh(a_I), \quad Q = \tanh(a_Q), \quad (30)$$

where I and Q are vectors with elements $I_i = I(i\Delta t)$ and $Q_i = Q(i\Delta t)$ for $i = -N$ to N . The time increment Δt is chosen to be on the order of sub-ns, in accordance with the sampling rates of modern AWGs [41,42]. Just as before, we use an element-wise tanh as an enclosing function to constrain the values to between -1 and 1 . The $\pi/2$ and π pulses are parameterized by separate (a_I, a_Q) .

3.4. Neural network

Finally, we follow [43] in creating a deep neural network (NN) model to represent the pulse shapes. The model, represented in Fig. 3, is

$$\begin{pmatrix} I_{\pi/2}(t'T_{\pi/2}) \\ Q_{\pi/2}(t'T_{\pi/2}) \\ I_{\pi}(t'T_{\pi}) \\ Q_{\pi}(t'T_{\pi}) \end{pmatrix} = \mathcal{L}_4 \circ \mathcal{L}_3 \circ \mathcal{L}_2 \circ \mathcal{L}_1(t'), \quad (31)$$

where each layer \mathcal{L}_i takes a d_i -dimensional input vector and maps it to a d_{i+1} -dimensional output vector according to the function

$$\mathcal{L}_i(\mathbf{x}) = \tanh(\mathcal{W}_i \mathbf{x} + \mathbf{b}_i) \quad (32)$$

with a $d_{i+1} \times d_i$ weight matrix \mathcal{W}_i , a d_i -dimensional bias vector \mathbf{b}_i , and tanh as an element-wise activation function that ensures the outputs are confined to between -1 and 1 . The first layer contains only a single node (i.e., $d_1 = 1$) and the input to it is the dimensionless time value $t' \in [-1/2, 1/2]$. Layers 2, 3, and 4 have 16 nodes, and the final layer contains four nodes that output the values for the four functions $I_{\pi/2}(t'T_{\pi/2})$, $Q_{\pi/2}(t'T_{\pi/2})$, $I_{\pi}(t'T_{\pi})$, and $Q_{\pi}(t'T_{\pi})$. We optimize the model parameters \mathbf{b}_i and \mathcal{W}_i , 644 in total. In this NN model, all four pulse shapes are controlled by the same set of parameters, which enables the model to represent possible correlations between the two pulses.

4. Optimization

In order to efficiently represent the distortion chain in our cost function, we utilize fast Fourier transforms (FFT), sampling the continuous shapes of the HS, FS and NN models with the same time step, Δt , as in the DT model. To obtain the total echo amplitude, we use a numerical solver to solve the ensemble of Schrödinger equations (19) for 250 frequencies spaced equidistantly across the desired band for a large enough time range to encompass any refocusing point.

For each model, we optimize over the full vector of parameters \mathbf{p} to maximize the echo amplitude irrespective of echo phase. The objective function is

$$\mathcal{J}(\mathbf{p}) = \max_i |M_{\perp}(t)|. \quad (33)$$

In a KBB pulse sequence, the echo occurs at time $t = T_{\pi} + \tau_1$ after the end of the π pulse. However, for the other pulse shape models, the echo can occur at earlier or later times.

The objective function used for the optimization of the hyperbolic secant functions is slightly different. As mentioned earlier, because the HS pulse shapes are the *output* I/Q pulse shapes, we need to ensure that the corresponding *input* I/Q pulse shapes respect the power constraint. For instance, if the HS pulse had an amplitude of $\omega_{1,\max}$ while the instantaneous driving frequency is off-resonant with the resonator, the input pulse required to compensate for the resonator transfer function would exceed $\omega_{1,\max}$. Thus, we include an extra term in the objective function to penalize the pulse for exceeding the power limit at any of the sampling points,

$$\mathcal{J}(\mathbf{p})_{\text{HS}} = \mathcal{J}(\mathbf{p}) - \sum_{i=0}^{T/\Delta t} \max(0, \omega_{\text{in}}^i - \omega_{1,\max}), \quad (34)$$

where ω_{in}^i is the amplitude at time $i\Delta t - T/2$ of the pulse when compensating for the resonator transfer function,

$$\omega_{\text{in}}(t) = \mathcal{F}^{-1} [\mathcal{F}[\omega_1(t)] / H_{\text{res}}(\omega)], \quad (35)$$

with $\omega_1(t)$ as given in Eq. (23). We leave out the factor of Δt in the second term in Eq. (34) in order to more heavily weight this term in the cost function to enforce the amplitude constraint. In practice, the optimization of $\mathcal{J}(\mathbf{p})_{\text{HS}}$ leads to the second term being zero and the optimal value of the objective function is the same as the echo amplitude.

For the FS parameterization, we use 24 terms per pulse shape, totaling 96 free parameters for the sequence. For the DT parameterization, we use a time step of $\Delta t = 0.625$ ns, resulting in 384 parameters for $T_{\pi/2} + T_{\pi} = 120$ ns. The NN parameterization uses 644 free parameters, as described in Section 3.4. We used the Julia package DiffEqFlux.jl [44] to form the NN parameterization and the BFGS optimizer from the Zygote.jl and Optim.jl packages for optimizing the various pulse shape parameterizations [45]. The limiting factor in computational cost is the numerical solution of Schrödinger's equation for each of the 250 Larmor frequencies for both pulses at each optimization step. We solved these in parallel on cluster computing resources with the Bogacki-Shampine method (BS5) as implemented in the DifferentialEquations.jl package. All optimizations were terminated upon the condition that the difference between the current objective function value and the objective function value 20 steps previous was less than 10^{-4} . All three large parameterizations took around 500–1000 optimization steps to converge, whereas the 5-parameter hyperbolic secant model only took tens of steps. On average, these optimizations took around 2–3 hours to complete running with 36 cores in parallel. However, this substantial optimization time is not a problem, as it is a one-time computational cost for a given spectrometer setup. As long as the distortion chain has been properly characterized, the computed pulses should work without (or with minimal) spectrometer-based feedback optimization [46]. The robustness of these pulses to mischaracterization of the distortion chain elements and B_1 field inhomogeneity is discussed in Section 5.

5. Results

We performed the optimizations using the parameter values shown in Table 1, representative of a Q-band EPR spectrometer with a nitroxide sample. The carrier frequency of 33.65 GHz corresponds to a static magnetic field strength of about 1.2 T. The value of Γ corresponds to a 3-dB bandwidth of 450 MHz for the AWG. The loaded Q-value corresponds to a resonator 3-dB bandwidth of 168 MHz. The maximum power limit of $\omega_{1,\max} = 84$ MHz corresponds to an oscillatory magnetic




Fig. 4. The optimized pulse shapes as seen by the spins after passing through the chain of transfer functions shown in Fig. 1, along with the formed echoes, the real part of the magnetization after rephasing the echo such that the maximum is completely real. The optimized HS $\pi/2$ and π pulse lengths are 80 and 40 ns long respectively while the FS, DT, and NN have $\pi/2$ and π pulse lengths of 60 ns each. The phase, ϕ , of each echo is also reported.

field strength of about 3 mT. The choice of a 120 ns total pulse time was made in order to examine a case where the power constraint starts to deteriorate the performance of the hyperbolic secant pulse. The delay time, τ_1 , was chosen to be 100 ns. While this value may affect the path the optimizer takes through the optimization landscape, pulse shapes optimized with one value of τ_1 produce the same echo amplitude for a different value of τ_1 . The nonlinear phase dispersion of the $\pi/2$ pulse will still be cancelled by that of the π pulse, and changing the delay time between them only changes the linear part of the phase dispersion, thus changing the echo time and phase but not the echo amplitude.

The optimization and the resulting analysis does not consider any spins to be coupled. Chirped pulses in particular have been shown to create unwanted artifacts compared to rectangular pulses in situations with coupled electron and nuclear spins [47,48]. A separate analysis can be performed to determine the effect that these optimized pulse shapes have on multi-dimensional spectra.

For each of the parameterizations, we also considered three different ratios $T_{\pi/2}/T_{\pi}$, 2:1, 1:1, and 1:2, keeping the total pulse time fixed at 120 ns. The KBB sequence requires a 2:1 ratio, i.e. 80 and 40 ns for the $\pi/2$ and π pulses, respectively. However, a 40 ns π pulse requires more power to adiabatically flip spins than the available limit, so many of the spins are under-rotated and the performance begins to suffer. Decreasing the time of the $\pi/2$ pulse while increasing the time of the π pulse will alleviate this issue, but it will also cause the phase refocusing aspect of the KBB pulse sequence to suffer [3,49]. Here, we have optimized the HS pulses for 1:1 and 1:2 pulse length ratios to demonstrate this tradeoff inherent to the KBB design.

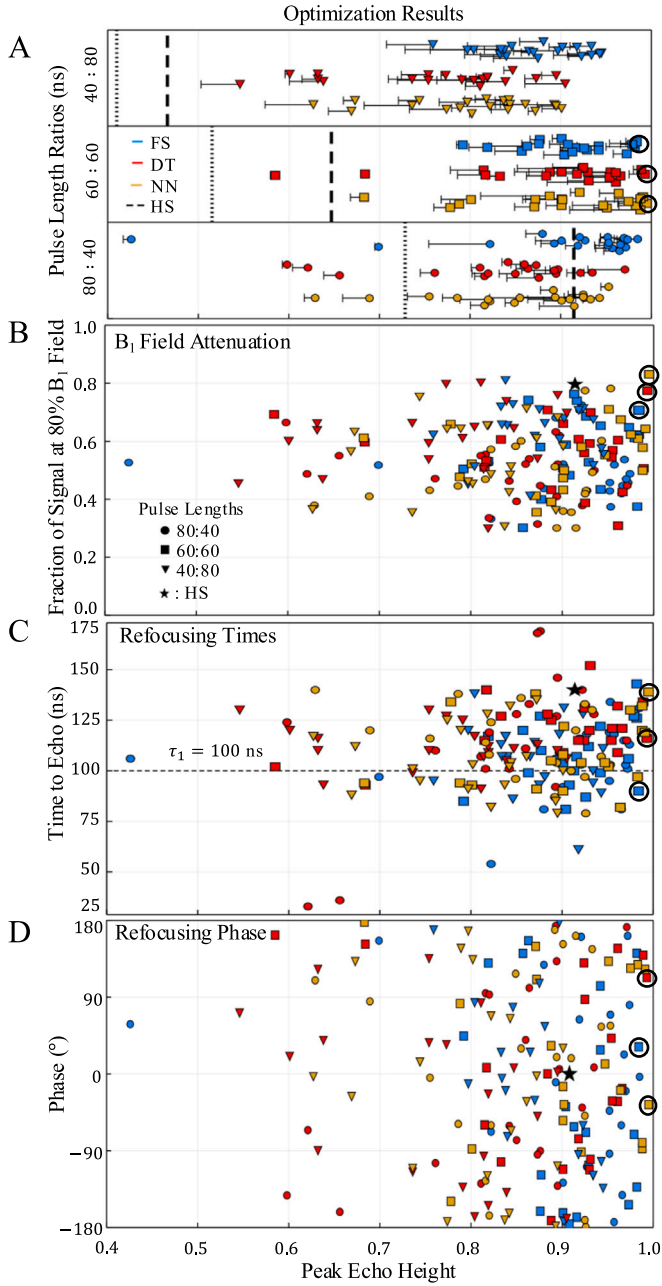


Fig. 5. Results of 20 optimizations each for the three large parameterizations of Fourier series (FS), discrete time (DT), and neural network (NN) with 80/40 ns, 60/60 ns, and 40/80 ns pulse lengths from random starting points. The top panel shows the echo amplitudes for the different pulse types, where the dashed lines are the echo amplitudes of the best performing hyperbolic secant (HS) pulses for each respective pulse length ratio. The 3 circled points represent the 3 best performing pulse shapes for each parameterization that are plotted in Fig. 4. The error bars and dotted lines show the effect of altering the amplifier compression function as described in the text. The bottom three panels show the relative echo amplitude reduction for the case where B_1 field is reduced to 80% of its original value, the time between the end of the π pulse and the echo maximum, and the phase with which the spins are refocusing, all versus the echo amplitudes.

The best HS result overall was for a 2:1 pulse length ratio, with $A_{\pi/2}/2\pi = 23.42$ MHz and $A_{\pi}/2\pi = 61.66$ MHz for the amplitudes of the $\pi/2$ and π pulses, $\beta = 7.09$, $n = 1.57$, and $\Delta\omega_{BW}/2\pi = 138.51$ MHz, which produces an echo amplitude of 0.9070 (see Fig. 4 top). To quantify the effect of artificially allowing the HS to exceed the power limit as discussed in Section 4, we also optimized strictly within the inscribing circle of the allowed domain in the I/Q plane, $I(t)^2 + Q(t)^2 \leq$

1. This more restricted optimization resulted in an echo amplitude of 0.8621 for our example parameters. Of course, the total extra power gained from accessing the corners of the square domain diminishes as the value V_{sat} decreases, since the amplifier saturates more readily. Having quantified this effect, below we set it aside and focus on the performance of other parameterizations compared to the HS with equivalent maximum power whose echo amplitude is 0.9070. As an alternative point of comparison, the time length required for an HS pulse sequence to perform as well as the other optimized pulse shapes is 180 ns, 50% longer than the other pulses.

For those other pulse parameterizations, all three pulse length ratios do allow for both full rotation across the bandwidth and refocusing, with a 1:1 pulse length ratio performing the best. The maximized echo amplitudes are, in increasing order, 0.9852 for the FS, 0.9928 for the DT, and 0.9951 for the NN. These pulses plotted in Fig. 4 are the drive functions seen by the spins, i.e. after passing through the low-pass filter, the amplifier, and the resonator transfer function (thus the pulses are several ns longer than the nominal 60 ns due to the finite ring-down time). The pulses show irregular shapes; there are no apparent interpretable features. Clearly, the FS, DT, and NN pulses offer performance gains compared to the HS for this scenario of short pulse time and limited power. Among these three parameterizations, though, there is not one that stands out as particularly advantageous. As mentioned in Section 3, the reasoning behind using the NN was to efficiently represent possible correlations between the two pulses. However, both the FS and DT show equivalent cooperativity in compensating for phase accumulation between the $\pi/2$ and π shapes, and the NN parameterization offered no extra advantage in that regard. The pulse shapes before passing through the distortion chain are plotted in Fig. S1 of the Supplementary Material [50].

All three parameterizations are equivalently efficient computationally, taking around the same number of optimization steps. The reason for the wide range of performance with many equivalent maxima is that the landscape contains many local maxima. Any large parameterization flexible enough to access a large area of that landscape in an unbiased way will lead to many maxima of varying quality. The outcome of a particular optimization depends on the initial point, so it is clearly useful to try many random initial seeds as we have done.

Fig. 5A shows the optimized echo amplitudes for the three parameterizations and three pulse length ratios. For each combination of shape and duration, the results of 20 separate optimization runs with random starting points are shown. The plots show that for any ratio the best FS, DT, and NN pulse shapes all outperform even the best 2 : 1 HS pulse (indicated by vertical dashed lines). For the FS, DT, and NN parameterizations, the best performing pulses have 1 : 1 pulse length ratios.

A similar observation regarding pulse length ratios was made by Kallies and Glaser [32], where they found an optimal ratio of 1 : 1.3 for their set of parameters. They used a different set of constraints, e.g., $\omega_{1,\text{max}}/\delta = 0.2$ and $\omega_{1,\text{max}} \cdot (T_{\pi/2} + T_{\pi}) = 6$, compared to our $\omega_{1,\text{max}}/\delta = 0.35$ and $\omega_{1,\text{max}} \cdot (T_{\pi/2} + T_{\pi}) = 10$. However, the two scenarios are more similar than these numbers would suggest because Ref. [32] did not account for the effect of a resonator and their pulses allowed a maximum pulse amplitude that is independent of the instantaneous driving frequency. Scaling their pulses to our desired bandwidth and compensating for low-pass filter and resonator, the required pulse amplitude is about twice as large in order to recover their pulse design when driving near the edges of the bandwidth. So, accounting for the transfer chain effectively makes their constraints to be $\omega_{1,\text{max}}/\delta \approx 0.4$ and $\omega_{1,\text{max}} \cdot (T_{\pi/2} + T_{\pi}) \approx 12$, similar to ours.

The error bars in Fig. 5A show the effect on the echo amplitude of reducing the amplifier compression V_{sat} by 60%, i.e. increasing amplifier compression without changing maximum output power. The FS, DT, and NN pulses are all relatively robust against this. This is because many of the numerically shaped pulses use as much power as possible, so only the maximum power of the amplifier matters rather

Table 1

A list of all the values used in the optimizations.

| Quantity | Value |
|---------------------------------|-----------|
| $\omega_c/2\pi$ | 33.65 GHz |
| $\delta/2\pi$ | 240 MHz |
| $\Gamma/2\pi$ | 450 MHz |
| $V_{\text{sat}}/V_{\text{DAC}}$ | 1.131 |
| Q_L | 200 |
| τ_1 | 100 ns |
| $\omega_{\text{res}}/2\pi$ | 33.65 GHz |
| $T_{\pi/2} + T_x$ | 120 ns |
| $\omega_{1,\text{max}}/2\pi$ | 84 MHz |

than the shape of the compression function at intermediate power. This is not entirely clear from the plots of the output pulse shapes in Fig. 4, but the input pulse shapes of I and Q plotted in Fig. S1 of the Supplementary Material [50] are typically toggling between their maximum values of -1 and 1 , and the values plotted are only diminished from maximum amplitude due to the effects of the low-pass filter and resonator. Thus, changing the general shape of the amplifying function does not affect the pulse shapes much. The HS pulse shapes, on the other hand, do use the full range of the amplifier, and the dotted lines show how their performance is substantially diminished under the same V_{sat} reduction.

In the B_1 field attenuation plot in Fig. 5B, we plot the fractional reduction of the echo amplitude when the B_1 field is reduced to 80%. This probes how robust the echo amplitude is to B_1 field inhomogeneity. The best HS pulse is one of the more robust pulses against B_1 field inhomogeneity, since it is constructed as an adiabatic frequency sweep. Other pulse types include equivalently robust pulses, but generally there is no correlation between echo amplitude and robustness to inhomogeneity because the pulses were not optimized for robustness. In Fig. S3 of the Supplementary Material [50], we also plot echo amplitude versus resonator quality factor and versus B_1 field inhomogeneity.

In the refocusing times plot of Fig. 5C, we plot the time between the end of the π pulse and the peak of the echo signal for $\tau_1 = 100$ ns. Most refocusing times fall in the range between 80 and 140 ns. Similarly to the results of Ref. [32], the middle of this optimal range is slightly longer than the waiting time τ_1 . There is no correlation of echo refocusing time with pulse parameterization or pulse length ratio.

The refocusing phase plot in Fig. 5D shows no evidence for a predominant echo phase. This is surprising, since one could imagine echo phase preferences in the presence of the $|I| \leq 1$, $|Q| \leq 1$ power constraint, where the maximum drive amplitude can only be achieved with driving phases $\pm 45^\circ$, $\pm 135^\circ$.

This ensemble of optimizations from different random seeds shows that there is no unique pulse shape that is optimal. Of course, there is always a trivial degeneracy of rotating both I and Q by $\pm 90^\circ$, $\pm 180^\circ$ resulting only in a change of echo phase by the same angle (the degeneracy is four-fold rather than continuous due to the power constraint being a square rather than a circle in the I/Q plane), retaining the same refocusing time and echo amplitude. But beyond that, Fig. 5 shows that there are many pulse shapes with equivalent echo amplitudes that are not simply related by a phase transformation, as is clear from the different echo refocusing times. The fact that the landscape harbors many comparable maxima suggests that there remains a significant amount of flexibility in the parameterized pulses that could be used to satisfy additional constraints, such as favoring a particular refocusing time or phase (as in [32]), robustness to B_1 inhomogeneity, echo phase independence from τ_1 , or some other desirable property.

In Fig. 6 we examine the action of the individual pulses by plotting the z -projection of a spin packet after it is rotated from the ground state by the shaped $\pi/2$ pulse,

$$2\langle S_z(\omega) \rangle_{\pi/2} = 2\text{tr}(U_{\pi/2} \rho_0 U_{\pi/2}^\dagger S_z), \quad (36)$$

where ρ_0 is the ground state density matrix, and similarly for the π pulse. We plot over a frequency range extending slightly outside the




Fig. 6. The excitation profiles of the optimized pulses for the $\pi/2$ pulse (left), and the π pulse (right), calculated from Eq. (36).

δ bandwidth. Recall that we have not constrained the first and second pulse to be a $\pi/2$ and a π pulse. Yet, Fig. 6 shows that inside the desired band the optimization always produces nearly perfect $\pi/2$ and π pulses, suggesting this basic structure is optimal for refocusing. This is not an artifact of a particular initialization of the parameters — we initialize randomly and the initial pulse shapes are not $\pi/2$ or π rotations. We also see that for the FS, DT, and NN parameterizations, the effect of the pulses on spins outside the bandwidth varies wildly with frequency compared to the HS pulses which do not excite these spins.

Fig. 6 strongly suggests that the improvement of the optimal FS, DT, and NN pulses compared to the HS pulse comes from (i) improved performance near the band edges and (ii) improved intra-band π rotation performance in the presence of the power constraint. For a more comprehensive visualization of the spin dynamics, in Fig. S2 of the Supplementary Material [50] we plot the total magnetization in x , y , and z during each of the optimized pulse sequences.

In Fig. 7 we further characterize the action of the sequence as a whole on any given spin packet. The top left panel in Fig. 7 is a plot of the phase dispersion at the refocusing time. This is calculated by first computing the individual spin packet phases as a function of offset frequency,

$$\phi(\omega) = \arg\left(\text{tr}(\rho_{\text{refocus}} S_+^\dagger)\right), \quad (37)$$

where

$$\rho_{\text{refocus}} = U_{\text{tot}}(t_{\text{refocus}}) \rho_0 U_{\text{tot}}^\dagger(t_{\text{refocus}}), \quad (38)$$

and t_{refocus} denotes the time at which the peak of the echo occurs. The average phase for spin packets within the band, which corresponds to the phase of the echo is

$$\phi_{\text{avg}} = \frac{1}{\delta} \int_{-\delta/2}^{\delta/2} \phi(\omega') d\omega', \quad (39)$$

and the phase dispersion plotted is

$$\Delta\phi(\omega) = \phi(\omega) - \phi_{\text{avg}}. \quad (40)$$

In other words, this is the azimuthal angle between the spin packet and the refocusing axis. The top right panel in Fig. 7 is a plot of the polar angle each spin packet forms with the z -axis at the refocusing time,

$$\theta(\omega) = \arccos\left(2\text{tr}(\rho_{\text{refocus}} S_z)\right). \quad (41)$$

From these top two panels we see that the phase dispersion and the polar angle with the FS, DT, and NN pulses are closer to ideal than with the HS pulse. The highly oscillatory polar angle compared to individual polar angles obtained in Fig. 6 is because the effect of a π pulse on a spin packet depends upon the phase of the spin packet, which varies rapidly as a function of offset frequency due to the waiting time between the pulses.

In the bottom left panel of Fig. 7 we plot the projection of each spin packet onto the refocusing axis in the xy -plane,

$$\langle R \rangle = 2\text{tr}\left(\rho_{\text{refocus}} (\cos(\phi_{\text{avg}}) S_x + \sin(\phi_{\text{avg}}) S_y)\right). \quad (42)$$

The FS, DT, and NN pulses are clearly more consistent in cooperatively producing an echo across the bandwidth.

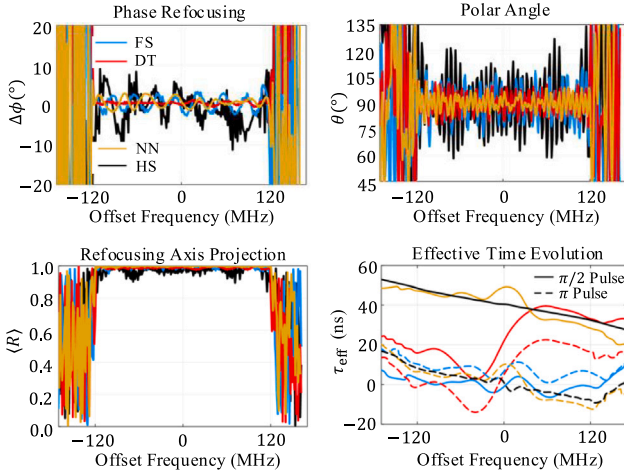


Fig. 7. Behavior of spin packets as a function of offset frequency. Phase difference $\Delta\phi$ relative to average phase at refocusing time (top left), polar angle θ at refocusing time (top right), projection onto the refocusing axis in the xy -plane at refocusing time (bottom left), and effective time evolution for each of the $\pi/2$ and π pulses (bottom right) calculated from Eqs. (40)–(42) and (48) respectively.

In the bottom right panel of Fig. 7 we explore the cooperativity of the shaped $\pi/2$ and π rotations using the effective evolution time defined in Ref. [32]. Consider the density matrix as a function of offset frequency at three intermediate times: immediately after the $\pi/2$ pulse, immediately before the π pulse, and immediately after the π pulse,

$$\rho_1(\omega) = U_{\pi/2} \rho_0 U_{\pi/2}^\dagger \quad (43)$$

$$\rho_2(\omega) = U_{\text{free}}(\tau_1) \rho_1(\omega) U_{\text{free}}^\dagger(\tau_1) \quad (44)$$

$$\rho_3(\omega) = U_{\pi} \rho_2(\omega) U_{\pi}^\dagger. \quad (45)$$

The phase accumulated from the $\pi/2$ pulse is

$$\phi_{\pi/2}(\omega) = \arg\left(\text{tr}(\rho_1(\omega) S_+^\dagger)\right) \quad (46)$$

and that accumulated from the π pulse is

$$\phi_{\pi}(\omega) = \arg\left(\text{tr}(\rho_3(\omega) S_+^\dagger)\right) - \arg\left(\text{tr}(\rho_2(\omega) S_+^\dagger)\right). \quad (47)$$

Unwrapping these phases about the center frequency, the effective evolution times are defined as [32]

$$\tau_{\text{eff},\pi/2}(\omega) = \frac{\phi_{\pi/2}(\omega) - \phi_{\pi/2}(0)}{\omega} \quad (48)$$

$$\tau_{\text{eff},\pi}(\omega) = \frac{\phi_{\pi}(\omega) - \phi_{\pi}(0)}{\omega}. \quad (49)$$

(The values of $\tau_{\text{eff},\pi/2}(0)$ and $\tau_{\text{eff},\pi}(0)$ are obtained via interpolation.) This is how long each spin packet would have to freely evolve following an ideal, instantaneous rotation in order to obtain the same final dispersion as produced by the actual rotation. For an echo to form, one must have a linear total phase dispersion (with negative slope) after the π rotation. This means any nonlinear phase dispersion acquired from the $\pi/2$ rotation must be cancelled by the nonlinear phase dispersion acquired from the π rotation.

In other words, since all the spin packet phases are flipped by the π rotation, echo formation requires $\tau_{\text{eff},\pi/2}(\omega) - \tau_{\text{eff},\pi}(\omega)$ to be a constant [32]. Fig. 7 shows that this is indeed the case for the FS, DT, and NN pulses, which have nonlinear phase dispersions for the two pulses but cooperate such that their nonlinear parts mutually cancel. Note that the difference $\tau_{\text{eff},\pi/2}(\omega) - \tau_{\text{eff},\pi}(\omega) = t_{\text{refocus}} - \tau_1$. For example, in this case, the FS pulse has $\tau_{\text{eff},\pi/2}(\omega) - \tau_{\text{eff},\pi}(\omega) = -10$ ns and so refocuses at 90 ns while the NN pulse has $\tau_{\text{eff},\pi/2}(\omega) - \tau_{\text{eff},\pi}(\omega) = 39$ ns and refocuses at 139 ns. The difference $\tau_{\text{eff},\pi/2}(\omega) - \tau_{\text{eff},\pi}(\omega)$ is shown in Fig. S4 of the Supplementary Material [50].

As a specific example, in Fig. 8 we use the spectral distribution of a solid-state dilute disordered sample of a nitroxide radical with




Fig. 8. Left: an exemplary nitroxide EPR spectrum. Right: the difference between the exemplary nitroxide spectrum and the Fourier transform of the resulting echoes formed by the Fourier Series (FS), discrete time (DT), neural network (NN), and hyperbolic secant (HS) optimized pulse sequences.

a bandwidth of about 240 MHz, shown in the left panel. Because the bandwidth of this distribution is slightly more than the optimized bandwidth of $\delta/2\pi = 240$ MHz, the chosen value of the pulse carrier frequency for each parameterization was also optimized in order to achieve maximum echo height. In the right panel we plot the differences between the actual spectrum and the spectra one would recover from the Fourier transforms of the echoes (note the different scale compared to the left panel). The FS, DT, and NN parameterizations result in less error, particularly around the spectral maximum. All parameterizations have some unrecovered spectral density towards the lower edge of the spectrum due to the bandwidth of the exemplary nitroxide spectrum being slightly larger than the optimized bandwidth.

These optimized pulse shapes have not yet been experimentally implemented. To record the echo generated by these pulses with sufficient fidelity, a receiver of sufficient bandwidth (> 240 MHz) is required. Higher-bandwidth detection systems with up to 1 GHz of bandwidth have become available commercially recently (Bruker SpecJet 3 and VideoAmp 3), so the presented approach is timely and feasible. Also, non-commercial wideband receiver systems have been built [4].

6. Conclusion

In this paper, we have shown that it is possible to obtain about a 10% improvement of the Hahn echo amplitude over an optimal, generalized hyperbolic secant KBB sequence by optimizing nonlinear amplitude-limited Fourier series, discrete time series, or neural network parameterized pulses. With these parameterized pulses, we find a 1:1 pulse length ratio is favorable in the presence of power constraints because it allows more energy to be allotted to the π pulse while the pulse shaping is still able to maintain the refocusing. Interestingly, the optimization landscape for this type of problem has many equivalent maxima, all of which involve the $\pi/2$ and π pulses cooperatively working together to compensate phase dispersions in each other. We find no marked differences among the three parameterizations. Furthermore, we have demonstrated that nonlinear effects due to amplifier compression and resonator transfer can be included in the optimization workflow, allowing for the usage of the full power of an amplifier, including its nonlinear region. These results demonstrate a pathway towards optimal broadband spectral acquisition with constrained power.

CRedit authorship contribution statement

Eric R. Lowe: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation. **Stefan Stoll:** Writing – review & editing, Supervision, Project administration, Investigation, Funding acquisition, Conceptualization. **J.P. Kestner:** Writing – review & editing, Supervision, Project administration, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

I have nothing to declare.

Acknowledgments

This research was in part supported by The Under Secretary of Defense for Research and Engineering (OUSD/R&E) National Defense Education Program (NDEP)/BA-1, Basic Research (E. R. L.), NSF CHE-2154302 (S. S.), NSF PHY-1915064 (J. P. K.), and is based upon work supported by the National Science Foundation, USA under the STC IMOD Grant No. 2019444 (S. S. and J. P. K.).

This work used the Advanced Research Computing at Hopkins (ARCH) core facility (rockfish.jhu.edu) through allocation XSEDE-PHY230011 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program [51], which is supported by National Science Foundation, USA grants #2138259, #2138286, #2138307, #2137603, and #2138296. This work also used the UMBC High Performance Computing Facility (HPCF), which is supported by the NSF, USA through the MRI program (Grant Nos. CNS-0821258, CNS-1228778, OAC-1726023, and CNS-1920079) and the SCREMS program (Grant No. DMS-0821311), with additional substantial support from the University of Maryland, Baltimore County (UMBC), USA.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jmr.2024.107807>.

Data availability

The code used in the paper has a reference to Zenodo with all of the available code.

References

- [1] T. Prisner, M. Rohrer, F. MacMillan, *Annu. Rev. Phys. Chem.* 52 (1) (2001) 279–313.
- [2] A. Doll, S. Pribitzer, R. Tschaggelar, G. Jeschke, *J. Magn. Reson.* 230 (2013) 27–39.
- [3] P.E. Spindler, P. Schöps, W. Kallies, S.J. Glaser, T.F. Prisner, *J. Magn. Reson.* 280 (2017) 30–45.
- [4] A. Doll, G. Jeschke, *J. Magn. Reson.* 280 (2017) 46–62.
- [5] C.E. Tait, S. Stoll, *J. Magn. Reson.* 277 (2017) 36–44.
- [6] T.F. Prisner, *J. Magn. Reson.* 306 (2019) 98–101.
- [7] D. Kunz, *Magn. Reson. Med.* 3 (3) (1986) 377–384.
- [8] D. Kunz, *Magn. Reson. Med.* 4 (2) (1987) 129–136.
- [9] J.-M. Bohlen, M. Rey, G. Bodenhausen, *J. Magn. Reson.* (1969) 84 (1) (1989) 191–197.
- [10] M.S. Silver, R.I. Joseph, D.I. Hoult, *J. Magn. Reson.* 59 (2) (1984) 347–351.
- [11] R. Fu, G. Bodenhausen, *J. Magn. Reson. A* 119 (1) (1996) 129–133.
- [12] J.-Y. Park, M. Garwood, *Magn. Reson. Med.* 61 (1) (2009) 175–187.
- [13] B. Endeward, M. Bretschneider, P. Trenkler, T.F. Prisner, *Prog. Nucl. Magn. Reson. Spectrosc.* 136–137 (2023) 61–82.
- [14] C. Counsell, M.H. Levitt, R.R. Ernst, *J. Magn. Reson.* (1969) 63 (1) (1985) 133–141.
- [15] M.H. Levitt, *Prog. Nucl. Magn. Reson. Spectrosc.* 18 (2) (1986) 61–122.
- [16] J. Baum, R. Tycko, A. Pines, *Phys. Rev. A* 32 (6) (1985) 3435–3447.
- [17] M. Garwood, L. DelaBarre, *J. Magn. Reson.* 153 (2) (2001) 155–177.
- [18] K.E. Cano, M.A. Smith, A.J. Shaka, *J. Magn. Reson.* 155 (1) (2002) 131–139.
- [19] T.E. Skinner, T.O. Reiss, B. Luy, N. Khaneja, S.J. Glaser, *J. Magn. Reson.* 163 (1) (2003) 8–15.
- [20] K. Kobzar, T.E. Skinner, N. Khaneja, S.J. Glaser, B. Luy, *J. Magn. Reson.* 170 (2) (2004) 236–243.
- [21] M. Foroozandeh, M. Nilsson, G.A. Morris, *J. Magn. Reson.* 302 (2019) 28–33.
- [22] J.D. Haller, D.L. Goodwin, B. Luy, *Magn. Reson.* 3 (1) (2022) 53–63.
- [23] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, S.J. Glaser, *J. Magn. Reson.* 172 (2) (2005) 296–305.
- [24] N.C. Nielsen, C. Kehlet, S.J. Glaser, N. Khaneja, *Optimal Control Methods in NMR Spectroscopy*, John Wiley & Sons, Ltd, 2010.
- [25] Z. Tošner, T. Vosegaard, C. Kehlet, N. Khaneja, S.J. Glaser, N.C. Nielsen, *J. Magn. Reson.* 197 (2) (2009) 120–134.
- [26] T.W. Borneman, D.G. Cory, *J. Magn. Reson.* 225 (2012) 120–129.
- [27] D.L. Goodwin, I. Kuprov, *J. Chem. Phys.* 144 (20) (2016) 204107.
- [28] P.E. Spindler, Y. Zhang, B. Endeward, N. Gershernzon, T.E. Skinner, S.J. Glaser, T.F. Prisner, *J. Magn. Reson.* 218 (2012) 49–58.
- [29] M. Braun, S.J. Glaser, *J. Magn. Reson.* 207 (1) (2010) 114–123.
- [30] D. Wei, Y. Chang, S.J. Glaser, X. Yang, *Appl. Phys. Lett.* 104 (24) (2014).
- [31] S. Asami, W. Kallies, J.C. Günther, M. Stavropoulou, S.J. Glaser, M. Sattler, *Angew. Chem.* 130 (44) (2018) 14706–14710.
- [32] W. Kallies, S.J. Glaser, *J. Magn. Reson.* 286 (2018) 115–137.
- [33] J. He, S. Huang, Y. Huang, S. Chang, S. Ying, B.-Z. Shen, Z. Feng, *IEEE Trans. Commun.* (2023) 1.
- [34] S. Stoll, A. Schweiger, *J. Magn. Reson.* 178 (1) (2006) 42–55.
- [35] S. Stoll, R.D. Britt, *Phys. Chem. Chem. Phys.* 11 (31) (2009) 6614.
- [36] Data available at <https://zenodo.org/records/11153641>.
- [37] J. Power, M. Foroozandeh, R. Adams, M. Nilsson, S. Coombes, A. Phillips, G. Morris, *Chem. Commun.* 52 (14) (2016) 2916–2919.
- [38] J.E. Power, M. Foroozandeh, P. Moutzouri, R.W. Adams, M. Nilsson, S.R. Coombes, A.R. Phillips, G.A. Morris, *Chem. Commun.* 52 (42) (2016) 6892–6894.
- [39] J.-B. Verstraete, W.K. Myers, M. Foroozandeh, *J. Chem. Phys.* 154 (9) (2021).
- [40] A. Tannús, M. Garwood, *J. Magn. Reson.* 120 (1) (1996) 133–137.
- [41] G.G. Brown, B.C. Dian, K.O. Douglass, S.M. Geyer, S.T. Shipman, B.H. Pate, *Rev. Sci. Instrum.* 79 (5) (2008) 053103.
- [42] M. Tseitlin, R.W. Quine, G.A. Rinard, S.S. Eaton, G.R. Eaton, *J. Magn. Reson.* 213 (1) (2011) 119–125.
- [43] U. Güngördü, J.P. Kestner, *Phys. Rev. Res.* 4 (2) (2022) 023155.
- [44] C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, A. Ramadhan, A. Edelman, *Universal differential equations for scientific machine learning*, 2021, <http://dx.doi.org/10.48550/arXiv.2001.04385>.
- [45] M. Innes, Don't unroll adjoint: Differentiating SSA-form programs, 2019, <http://dx.doi.org/10.48550/arXiv.1810.07951>.
- [46] D.L. Goodwin, W.K. Myers, C.R. Timmel, I. Kuprov, *J. Magn. Reson.* 297 (2018) 9–16.
- [47] A. Doll, G. Jeschke, *Phys. Chem. Chem. Phys.* 18 (33) (2016) 23111–23120.
- [48] S. Pribitzer, T.F. Segawa, A. Doll, G. Jeschke, *J. Magn. Reson.* 272 (2016) 37–45.
- [49] A. Bowen, N. Erlenbach, P. Van Os, L. Stelzl, S.T. Sigurdsson, T. Prisner, *Appl. Magn. Reson.* 49 (2018) 1355–1368.
- [50] Supplementary Material available at link to be provided by publisher.
- [51] T.J. Boerner, S. Deems, T.R. Furlani, S.L. Knuth, J. Towns, *Practice and Experience in Advanced Research Computing*, PEARC '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 173–176.