Delft University of Technology

Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control

Van Dijk, Jeroen P.G.; Patra, Bishnu; Pellerano, Stefano; Charbon, Edoardo; Sebastiano, Fabio; Babaie, Masoud

DOI

[10.1109/TCSI.2020.3019413](https://doi.org/10.1109/TCSI.2020.3019413)

Publication date

2020

Document Version

Final published version

Published in

IEEE Transactions on Circuits and Systems I: Regular Papers

Citation (APA)

Van Dijk, J. P. G., Patra, B., Pellerano, S., Charbon, E., Sebastiano, F., & Babaie, M. (2020). Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 67(12), 5380-5393. [9189938]. <https://doi.org/10.1109/TCSI.2020.3019413>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control

Jeroen P. G. van Dijk¹, *Student Member, IEEE*, Bishnu Patra², *Student Member, IEEE*,
Stefano Pellerano, *Member, IEEE*, Edoardo Charbon³, *Fellow, IEEE*,
Fabio Sebastiano⁴, *Senior Member, IEEE*,
and Masoud Babaie⁵, *Member, IEEE*

Abstract—The design of a large-scale quantum computer requires co-optimization of both the quantum bits (qubits) and their control electronics. This work presents the first systematic design of such a controller to simultaneously and accurately manipulate the states of multiple spin qubits or transmons. By employing both analytical and simulation techniques, the detailed electrical specifications of the controller have been derived for a single-qubit gate fidelity of 99.99% and validated using a qubit Hamiltonian simulator. Trade-offs between several architectures with different levels of digitization are discussed, resulting in the selection of a highly digital DDS-based solution. Initiating from the system specifications, a complete error budget for the various analog and digital circuit blocks is drafted and their detailed electrical specifications, such as signal power, linearity, spurs and noise, are derived to obtain a digital-intensive power-optimized multi-qubit controller. A power consumption estimate demonstrates the feasibility of such a system in a nanometer CMOS technology node. Finally, application examples, including qubit calibration and multi-qubit excitation, are simulated with the proposed controller to demonstrate its efficacy. The proposed methodology, and more specifically, the proposed error budget lay the foundations for the design of a scalable electronic controller enabling large-scale quantum computers with practical applications.

Index Terms—Direct digital synthesis (DDS), quantum computing, qubit control, frequency division multiplexing, specifications, fidelity.

I. INTRODUCTION

QUANTUM computing will provide exponential speed-up over classical computers in several applications.

Manuscript received February 7, 2020; revised June 5, 2020 and July 18, 2020; accepted August 16, 2020. Date of publication September 9, 2020; date of current version December 1, 2020. This work was supported by Intel Corporation. This article was recommended by Associate Editor R. Azarderakhsh. (Jeroen P. G. van Dijk and Bishnu Patra contributed equally to this work. Stefano Pellerano, Edoardo Charbon, Fabio Sebastiano, and Masoud Babaie contributed equally to this work.) (Corresponding author: Jeroen P. G. van Dijk.)

Jeroen P. G. van Dijk, Bishnu Patra, and Masoud Babaie are with the Department of Quantum and Computer Engineering, Delft University of Technology, 2628 Delft, The Netherlands, also with the Kavli Institute of Nanoscience, 2628 Delft, The Netherlands, and also with Qutech, Delft University of Technology, 2628 CJ Delft, The Netherlands (e-mail: jpgvandijk@gmail.com).

Stefano Pellerano is with Intel Corporation, Hillsboro, OR 97124 USA.

Edoardo Charbon is with the Kavli Institute of Nanoscience, 2628 Delft, The Netherlands, also with the Faculty of Engineering, Swiss Federal Institute of Technology Lausanne (EPFL), 1015 Lausanne, Switzerland.

Fabio Sebastiano is with the Department of Quantum and Computer Engineering, Delft University of Technology, 2628 Delft, The Netherlands, and also with Qutech, Delft University of Technology, 2628 CJ Delft, The Netherlands.

Color versions of one or more of the figures in this article are available online at <https://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2020.3019413

In quantum chemistry, the electronic structure of molecular orbitals can be mapped onto a quantum processor to simulate the interaction between molecules. As a result, new molecules and reactions can be designed, and industrial chemical processes can be optimized [1], [2]. For instance, fertilizer production takes up to 1% of the world's energy supply, due to the currently employed high-temperature high-pressure Haber-Bosch industrial process. However, a very similar process, i.e. nitrogen fixation in plants, happens under ambient conditions. While the primary cofactor of the biological nitrogen-fixing enzyme nitrogenase (FeMo cofactor) is not yet fully understood, it could conceivably be simulated using a quantum computer [3].

The computing power of a quantum computer is directly related to the quality and *quantity* of its quantum bits (qubits¹), its fundamental computing unit. Although a quantum computer with only 53 qubits has surpassed the capability of even the most powerful supercomputers [4], large-scale quantum computers with thousands – or even millions – of qubits would be required for any practical computations, thus demanding a scalable qubit and control architecture.

Amongst various qubit topologies, solid-state qubit technologies, such as spin qubits [5] and transmons [4], promise scalability due to their small form factor and fabrication process. However, such qubits typically operate at temperatures below 100 mK inside a dilution refrigerator, while the control electronics is implemented with off-the-shelf equipment operating at room temperature and connected via at least a single RF cable per qubit. Such control setups hinder scalability both because of the excessive complexity of an equipment-based control and because it is impractical to fit thousands of cables inside a dilution refrigerator, while minimizing the heat load in the fridge and ensuring the reliability of the interconnects.

Although several setups have migrated from the use of generic equipment to custom-made electronics, the specifications of such systems are not tailor-made for qubit control, but rather the main goal is to reduce the amount and cost of equipment/interconnect used for qubit control [6]. The ideal solution to build large-scale quantum computers would be to operate both qubits and control electronics at the same cryogenic temperature [5]. Along such direction, custom-made PCBs with commercial off-the-shelf components

¹The term “qubit” refers to physical qubits.

operating at cryogenic temperatures have been designed to interface several control and read-out channels, thus minimizing cabling [7], [8]. Taking a step further, control and readout integrated circuits implemented in standard CMOS technology and operating at cryogenic temperature (cryo-CMOS) have been shown to operate at cryogenic temperatures as low as 4 K and even below, promising a viable solution for scalability [9], [10]. Additionally, such cryo-CMOS electronics can be in principle co-integrated with spin qubits on a single chip, thus providing a compact solution towards the realization of practical quantum computers [11], [12]. To design such systems, the circuit specifications need to be estimated/simulated to produce a power-efficient design. Furthermore, this is indispensable for circuits operating at cryogenic temperatures, due to limited cooling power of the dilution refrigerator.

In this paper, we address the above-mentioned issues by proposing a systematic design technique of the electronic controller for single-qubit operations, employing frequency multiplexing to reduce interconnects and power consumption. This work presents the architecture and specifications of a power-efficient qubit control system, to achieve a single-qubit gate fidelity up to 99.99% by complying with the signal specifications for qubit control, outlined in [13]. That work presents a systematic study of the impact of the classical electrical signals on the qubit fidelity for single-electron spin qubits, considering all operations, i.e., single-qubit rotations, two-qubit gates, and readout, in the presence of errors in the control electronics, such as static, dynamic, systematic, and random errors. Moreover, using case studies, [13] shows how preliminary signal specifications can be derived to achieve a specified gate fidelity. In this work, those results are used as the basis to find the specifications for a DDS-based system for multi-qubit control.

In the following, section II presents the requirements for the qubit control system. Section III discusses the trade-offs between possible transmitter architectures and describes the chosen system architecture. In Section IV the specifications for the different architectural sub-blocks are determined to assess the feasibility of the system. Finally, Section V demonstrates the flexibility of the design for various applications, and a conclusion follows in Section VI.

II. REQUIREMENTS

The main focus of this work is on single-electron spin qubits, since they are very promising both in terms of scaling opportunities and co-integration with CMOS electronics [5], and a co-simulation platform is readily available [13], [14]. Since the control signals required by spin qubits and transmons are very similar, we will also describe the minor changes required to ensure compatibility with transmons.

A. Qubit Control Signal Requirements

In a single-electron spin qubit, information is encoded in the spin of a single electron hosted in a quantum dot. The spin-up

Fig. 1. The qubit state is represented in the Bloch sphere.

and spin-down states encode the $|0\rangle$ and $|1\rangle$, respectively.² Since a large magnetic field is applied, the two states are separated by the Zeeman energy E_z , which is associated to the qubit frequency $f = E_z/h$, with h the Planck constant. Any single-qubit operation can be represented as a rotation of the qubit state in the Bloch sphere, as indicated in Fig. 1. To achieve universal quantum computation, rotations around at least two different axes are required. The accuracy of the implemented qubit rotation can be measured by the *fidelity* of such operation [13]. The fidelity is the most commonly used quantity to benchmark the quality of quantum operations. However, for systematic errors, the actual error rate in an algorithm can deviate from the error predicted by the fidelity $(1 - F)$, and is, in the worst case, bounded by the diamond norm $(\sqrt{1 - F})$. In current experiments, the observed error rate is usually well described by the fidelity and is hence used here.

To perform such a rotation for single-electron spin-qubits, a microwave pulse needs to be applied to the qubit as either an electric (Electric Dipole Spin Resonance, EDSR) or magnetic field (Electron Spin Resonance, ESR), with the frequency accurately matched to the qubit frequency. The amplitude of the microwave pulse sets the rotation speed, called Rabi frequency f_R , and hence, together with the duration of the microwave pulse T , sets the rotation angle θ in Fig. 1. For a rectangular pulse envelope, $\theta = 2\pi f_R T$. The phase of the microwave signal needs to remain coherent with the phase of the qubit, which implies keeping a coherent phase for the whole duration of the quantum algorithm, even over different pulses. Changing the relative phase results in a rotation along a different axis in the Bloch sphere (ϕ in Fig. 1), and can be used to implement both X- and Y-rotations of the qubit.

Frequency multiplexing can be used to drive multiple qubits on the same driveline, with the advantage that the amount of interconnects can be reduced, and the control electronics can potentially be more area and power efficient. However, when applying a short microwave burst to one qubit, the spectrum can contain energy at frequencies corresponding to other qubits. Pulse shaping techniques must then be used to minimize spectral leakage, generally known as cross-talk, to other qubits [15]. However, even with pulse shaping, the frequency

²Introducing the basics of quantum computation is outside the scope of this paper. The interested reader is referred to [13].

of a qubit slightly shifts when applying a signal at a different frequency. This so-called AC-Stark shift causes the qubit to acquire a phase offset. This results in an unintended Z-rotation on the qubit, which needs to be corrected [15].

Two-qubit operations, qubit initialization and qubit readout typically require unmodulated pulses to be applied to the quantum processor [16], and are here assumed to be generated by other control electronics and are therefore outside the scope of this paper.

B. System Specifications

For single-electron spin qubits, the qubit frequency is typically 12-40 GHz with microwave pulse duration in the order of $1\ \mu\text{s}$, and, to achieve a typical Rabi frequency of 1 MHz, a power of $\sim -45\text{ dBm}$ is usually required [17]–[20]. For future systems, it is desirable to operate at lower qubit frequencies and higher Rabi frequencies [12]. Hence, the system presented here will be designed for an output frequency range of 5-20 GHz, and Rabi frequencies in the range of 1-10 MHz, with a maximum rotation angle of π . This sets the nominal duration of a π -rotation to 50-500 ns. The required output power for spin qubits ranges then from -45 dBm to -25 dBm for the selected Rabi frequency range. However, as attenuators (e.g., with 6 dB loss) are typically employed before the qubits to reduce the heat injected into the quantum processor, and the sensitivity of the qubit can easily vary by $\pm 50\%$, the required output power range is extended as -48 dBm to -16 dBm (50 mV_p). Current experiments on single-electron spin-qubits typically do not use frequency multiplexing, and hence, rectangular envelopes for the microwave pulses are allowed [17], [20]. However, for our system, more complex pulse shaping, e.g., Gaussian envelopes, are necessary to support frequency multiplexing. Moreover, for flexibility, it is desirable to program any envelope, with support of I/Q-modulation for the benefit of having X- and Y-rotations.

The fidelity of single-qubit operations is typically above 99% for single-electron spin-qubits [18]. For fault-tolerant quantum computing, a minimum qubit fidelity, typically around 99.9%, is required when using error-correction techniques [21]. In order not to limit the performance of the whole quantum computer, the proposed electronic interface targets a fidelity of 99.99% for a π -rotation performed on a spin-qubit, when taking into account only the errors due to the electronic interface and assuming a perfect qubit. Considering frequency multiplexing, the system will be designed such that both the addressed qubit achieves a 99.99% fidelity for the targeted π -rotation (which generally gives the lowest fidelity [13]) while the idle qubits reach a 99.99% for the identity, or idle, operation.

Since modern CMOS processes allow processing of extremely wide bandwidths, we aim at the maximum feasible bandwidth to maximize the number of qubits that can be served. Fig. 2 shows the number of qubits that can be multiplexed in a 1 GHz bandwidth for different microwave pulse envelopes, when assuming uniformly distributed qubit frequencies, a π -rotation at the maximum supported Rabi frequency of 10 MHz, and Z-corrections to compensate for

Fig. 2. The number of qubits that can be allocated in a 1 GHz band when driving with different envelopes and the required Z-correction, each performing a π -rotation in 50 ns.

the AC-Stark shift [13]. Less than 5 qubits can be served at a 99.9% fidelity with a rectangular envelope. By employing Gaussian pulses, this can be significantly improved, resulting in ~ 40 qubits operating at a 99.99% fidelity in a 2 GHz bandwidth. For the system discussed here, 32 qubits are targeted in a 2 GHz bandwidth, since this number allows for easy addressing of the qubits ($32 = 2^5$), and for a theoretical fidelity $> 99.999\%$.

Even though frequency multiplexing allows for operating on multiple qubits simultaneously, the system will be optimized assuming sequential execution of the operations on the different qubits, as more complicated measures than a simple Z-correction are required when operating on multiple frequency-multiplexed qubits simultaneously [15]. However, as a scalable solution is desired, the chosen system architecture should support the simultaneous excitation of multiple qubits.

C. Extending to Transmons

The control of transmons is very similar to spin qubits, but there are a few key differences that could affect the system specifications. The qubit frequency is typically around 6 GHz for transmons, and microwave pulses as short as 20 ns are used with a signal power of $\sim -60\text{ dBm}$. Hence, the duration and output power specifications are extended to include this. Additionally, pulse shaping (Derivative Removal by Adiabatic Gate, DRAG) is typically used to minimize spectral leakage to higher energy levels of the same qubit. This specific pulse requires I/Q modulation, which is already supported to seamlessly allow X- and Y-rotations. Finally, as state-of-the-art transmons typically achieve fidelities not better than 99.99% [22], the control system will still not limit the achievable fidelity.

A summary of the discussed specifications is given in Table I.

Following the methods presented in [13], preliminary signal specifications can be estimated for performing a π -rotation on the addressed spin qubit with either a Rabi frequency of 1 or 10 MHz and a rectangular envelope, see Table II. Equal error contributions are assumed, and the value given for the amplitude inaccuracy assumes a peak amplitude of 50 mV, which corresponds to the maximum required output power. These preliminary specifications will be used to assess the

TABLE I
THE REQUIREMENTS OF THE MULTI-QUBIT CONTROL SYSTEM

Qubit technology	Focus on single-electron spin-qubits, support transmons
Qubit frequency range	5 GHz to 20 GHz
Qubit Rabi frequency range	1 MHz to 10 MHz
Output power range	-60 dBm to -16 dBm
FDMA, bandwidth	2 GHz
FDMA, number of qubits	32
FDMA, parallel operations	Supported
Operation, maximum angle	π
Operation, duration	Specifications guaranteed from 50 ns to 500 ns
Operation, modulation	I/Q-modulation with any envelope (nominally Gaussian when using frequency multiplexing, rectangular otherwise)
Fidelity, addressed qubit	99.99% for a π -rotation on a spin-qubit
Fidelity, idle qubit	99.99% for identity operation
Power consumption	Minimum

Fig. 3. Possible transmitter architectures: (a) Very high speed DAC, (b) MRZ DAC, (c) low speed DAC with mixer, (d) high speed DAC with mixer (e) SSB modulation with band pass filters for the sidebands, (f) High speed DAC with reconstruction filter and I/Q mixer.

TABLE II
EXAMPLE SPECIFICATIONS FOR ACHIEVING A 99.99%
FIDELITY FOR A π -ROTATION

<i>Rabi frequency:</i>	1.0	10	MHz
Addressed qubit:			
Phase imbalance	0.20	0.20	degrees
Frequency inaccuracy	3.5	35	kHz
Frequency noise	3.5	35	kHz _{rms}
Duration inaccuracy	1.1	0.11	ns
Timing jitter	1.1	0.11	ns _{rms}
Amplitude inaccuracy	0.011	0.11	mV
Amplitude noise	50	50	dB
Wideband additive noise	5.6	18	nV/ $\sqrt{\text{Hz}}$
Idle qubit:			
SFDR	-44	-44	dB

feasibility of different proposed architectures. Most notably, a high spurious-free dynamic range (SFDR) is required, as spurious tones could interfere with the idle qubits in a frequency multiplexing scheme.

III. SYSTEM ARCHITECTURE

Based on the signal requirements for qubit control (Table II), the feasibility of several architectures is discussed and the chosen architecture is presented in this section.

A. Analog/RF Section

To generate the required envelopes for frequency-multiplexed qubits, the simplest architecture would be to design a digital-to-analog converter (DAC) operating at 40 GS/s, as shown in Fig. 3(a). However, the power consumption would be too high due to its large data bandwidth [23]. To reduce the power consumption, a multiple-return-to-zero (MRZ) DAC [Fig. 3(b)] exploiting higher Nyquist zones is capable of synthesizing frequencies up to 20 GHz [24]. However, limited flexibility in choosing the output frequency band (centered around $N \cdot f_s$) and an output spectrum corrupted by DAC replicas does not make this a good candidate. To overcome this, several low-speed DACs along with I/Q mixers [Fig. 3(c)], can be used to generate envelopes at distinct frequencies [10], each covering the bandwidth of one qubit, with the possibility of having an individual RF channel/output per qubit. However, this would require multiple local oscillator (LO) signals, thus making it power/area inefficient for multi-qubit control. Moreover, on-chip implementation of multiple LOs can cause frequency pulling and affect the spectral purity of the synthesizers, thus degrading the transmitter SFDR. Instead, a very high-speed DAC at 4 GS/s and a single mixer can be used for controlling multiple qubits from a single RF cable [Fig. 3(d)].

Single-sideband (SSB) modulation can be implemented instead of double-sideband (DSB) modulation to obtain the same bandwidth with half the DAC sampling frequency, at the cost of increased circuit complexity. This can be achieved by filtering each of the upper/lower sideband (USB/LSB) and combining them at the output, as shown in Fig. 3(e). To achieve an image rejection ratio (IRR) > 44 dB for output frequencies close to the carrier (as required by the SFDR specification), filters with very high order or quality factor are essential. Instead, image rejection can be achieved using a Hartley modulator with I/Q DAC and mixer [Fig. 3(f)] at the cost of requiring an LO with quadrature phases. At the circuit level, instead of cascading the DAC and the mixer, a better solution would be to use a mixing DAC, i.e. combining the DAC and mixer at the circuit level, for power efficiency and linearity [25]. However, the mixing DAC output is corrupted by tones at alias frequencies, which may fall in the upconverted 2 GHz output band when the signal bandwidth is comparable to the carrier frequency, as it will be shown in Section IV (Fig. 6). This would suggest exploring a high-speed DAC followed by a reconstruction filter and a mixer for better spectral purity [Fig. 3(f)].

B. Digital Signal Synthesis

To generate multiple SSB-modulated tones with this front-end design, a digital back-end is required. This work assumes the availability of a reprogrammable on-chip memory that is used to store calibrated waveforms for each of the desired qubit rotations. A qubit algorithm is then executed by playing the various stored waveforms in the desired order. The most straightforward approach is to store all possible combinations of qubit instructions in an SRAM, as shown in Fig. 4(a). The required memory of such an SRAM can be estimated as $SRAM_{mem} = N \times f_s \times t_{pulse} \times m^n$, where N and f_s are the number of bits and the sampling frequency of the DAC, respectively, t_{pulse} is the pulse duration, m the number of possible instructions per qubit and n the number of qubits. Assuming an 8-bit DAC operating at 2.5 GS/s to address 32 qubits and a maximum pulse duration of 500 ns, it would require an impractically large memory of $3.7 \cdot 10^{19}$ bits, considering merely 3 instructions per qubit. Moreover, since qubits require coherent control, intermittent sequential operations on any qubit demand keeping track of the phase of all qubits. Consequently, an individual reference clock would be required for each qubit.

To reduce the required memory, an alternative approach is to store only the amplitude information in the SRAM, which can modulate the amplitude of a sinusoidal waveform with a programmable phase, as shown in Fig. 4(b). Under the mentioned assumptions, this would require less than 1-Mbit SRAM instead (scaling as $m \times n$ instead of m^n), consequently saving area at the cost of a higher power consumption. In order to update the phase for each qubit and ensure coherent control, sine and cosine waveforms scaled by appropriate coefficients can be combined to generate the required phase offset. However, this adds an overhead of 2 multipliers per qubit running at the full sampling speed.

Fig. 4. Possible backends: (a) SRAM for all possible instructions, (b) Reduced memory for on-chip modulation, (c) NCO based modulation. The green blocks are programmable memories.

A power-efficient approach would be to use a Numerically Controlled Oscillator (NCO) for each qubit to generate both the required frequency and the phase offset [26]. A numerically controlled oscillator consists of a phase accumulator running at f_s . An input frequency tuning word (FTW) defines the step size of the phase accumulator to generate the desired output frequency $f_{out} = FTW \times f_{s/2}^N$, where N is the number of bits in the phase accumulator and determines the frequency accuracy ($f_{s/2}^N$). The sine Look-Up Table (LUT) generates a sine wave corresponding to the output phase of the NCO, which is then multiplied with an envelope (stored in the SRAM) to obtain the necessary modulated signal, as shown in Fig. 4(c). This allows for fewer multipliers and the same number of adders compared to Fig. 4(b), thereby saving substantial power, i.e., 2 multipliers per qubit running at 2.5 GHz. Another advantage of such a system is that the NCO can keep track of the phase of individual qubits, thus allowing coherent operation [27].

C. Final Architecture

Considering the above-mentioned trade-offs, a digitally intensive architecture based on direct digital synthesis (DDS) with digital modulation, as shown in Fig. 5, has been selected. Such an architecture benefits from the scaling advantage of advanced CMOS technology nodes in terms of speed and power efficiency and offers the flexibility and robustness of digital signal processing.

Multiple NCOs (one per qubit) are used to keep track of the phase evolution of the qubits. However, the NCO outputs are time-multiplexed to allow operation on one qubit at a time to reduce system complexity, as mentioned in Section II.

Fig. 5. Block diagram of the proposed controller.

The multiplexed output is fed into LUTs to generate the sinusoidal signals, which are then modulated by the envelope memory (ENV_I, ENV_Q) for various gate operations and pulse shaping [28], [29] providing flexibility in qubit control.

Because of the stringent IRR requirement of 44 dB (originating from SFDR) corresponding to a maximum phase and gain imbalance of 0.3° and 0.1 dB, respectively, an I/Q digital correction network is required to compensate for analog I/Q mismatch. Moreover, a DC offset correction is added to cancel the LO feed-through to the output.

Finally, the same transmitter as in Fig. 3(f), comprising I/Q DACs, reconstruction filter, and an I/Q mixer, translate its digital input to the RF band. The only required analog input is then a quadrature LO signal to drive the mixer.

IV. CIRCUIT SPECIFICATIONS

When increasing the signal dynamic range, the rate at which the power consumption increases is much lower in a digital circuit than in its analog counterpart, especially in nanometer CMOS technologies [30]. Therefore, the error budget for the digital section is set an order of magnitude tighter than the target fidelity, i.e. it is set to a 99.999% fidelity, so as to contribute negligibly to the target fidelity of the controller.

To this purpose, a MATLAB simulation model of the entire system is developed, comprising an accurate representation of the digital section (including quantization and rounding effects), an ideal model of every analog block, and a model of the 32-qubit quantum processor. The evolution of each qubit is represented by the Hamiltonian of a single-electron spin-qubit under the excitation of the microwave current $i_{mw}(t)$ generated by the controller (Hamiltonian simulator implementation in [14]):

$$H = \frac{\hbar}{2} \cdot \begin{bmatrix} -\omega_0 & \alpha \cdot i_{mw}(t) \\ \alpha \cdot i_{mw}(t) & \omega_0 \end{bmatrix}, \quad (1)$$

where α and ω_0 represent the sensitivity to the drive signal and qubit frequency respectively of the qubit processor.

The following calculations are based on a rectangular envelope, while the simulations consider both a rectangular and Gaussian envelope. Moreover, as the specifications are typically stricter when operating at a Rabi frequency of 10 MHz, this will be the default assumption, unless otherwise specified.

Besides that, the lowest output frequency band of 5 – 7 GHz will be used in the simulations as this band suffers more from sampling replicas. When simulating the idle qubits, any Z-error is ignored, as these can be corrected in software [31].

The design strategy is as follows. First, the sample rate is chosen (Section IV-A), which then allows for the selection of an appropriate reconstruction filter (Section IV-B). Next, the effects of a limited bit length in each digital block on the targeted and idle qubits are individually simulated while keeping the other blocks ideal, i.e. not quantized (Section IV-C). The results of this sensitivity analysis are used to select the number of bits required in each block to achieve the targeted fidelity. The final digital system, including all non-idealities, which are simultaneously accounted for, is simulated in a final verification step (Section IV-D). Finally, in Section IV-E, the specifications of analog blocks can be readily derived from the requirements in Table II.

A. Sample Rate

Due to the chosen I/Q-modulation architecture, there is individual control over the upper and lower sidebands of the upconverted signal. For the required 1-GHz sideband (for a 2-GHz bandwidth), it is sufficient to run the DACs at a sample rate of 2 GS/s to fulfill the Nyquist criteria. However, considering the inherent zero-order hold (ZOH) operation of DACs, the -3-dB bandwidth of a DAC is roughly 40% of the sample rate. Hence, in this design, a sample rate of 2.5 GHz is chosen, thus resulting in a timing resolution of 400 ps for the microwave envelopes.³ The shortest operation of 20 ns is then supported (50 points), while the longest operation (500 ns) sets a minimum memory of e.g. 160 kSa, assuming four instructions for each of the 32 qubits.

B. Reconstruction Filter

The lowest qubit frequency of 5 GHz is achieved using a carrier frequency of 6 GHz and the 1 GHz sideband. A sketch of the output spectrum for this condition is shown in Fig. 6.

³The specified duration inaccuracy cannot be guaranteed with a 400-ps timestep. However, as the total rotation angle is set by both the duration and amplitude, such an under/over-rotation error can be corrected by using the amplitude instead.

Fig. 6. Output spectrum assuming an output tone around 5 GHz, a carrier frequency of 6 GHz, a sample rate of 2.5 GHz, and no reconstruction filter.

The negative frequencies are shown for clarity to illustrate that negative sampling replicas fold back to positive frequencies and eventually fall back close to an in-band qubit. Since the ZOH suppression of the replicas corresponds to a worst-case SFDR of 21 dB, an additional attenuation of at least 33 dB is required at 11 GHz to achieve an SFDR better than 54 dB, as required for 99.999% fidelity (10 dB more than Table II for a $10\times$ smaller error). Since a second-order filter is at least required, a 2nd order Chebyshev-I with 3-dB passband ripple and a 1.8-GHz corner frequency was chosen. The combination of the ZOH and reconstruction filter provides an SFDR better than 58 dB in all cases, resulting in a simulated fidelity of the idle qubit of $>99.9996\%$.

In addition, the chosen filter improves the in-band flatness to 0.14 dB over the full 2-GHz data band. While this is not a strict requirement, this removes the need to predistort the envelopes. As a result, a qubit driven at 5.1 GHz with a rectangular envelope can achieve a fidelity of 99.9995% without any predistortion in an otherwise ideal system. In comparison, a 3rd-order Butterworth filter with a 1.7-GHz corner frequency has an in-band flatness of 2.6 dB, which results in a fidelity of only 99.998% for a non-predistorted rectangular envelope. This is an important result, as it shows that, with proper design, one can use much simpler modulation schemes to achieve the intended performance.

C. Digital Blocks

1) *Number of NCO Accumulator Bits*: The number of bits in the accumulator register b_{acc} (see Fig. 5) sets the frequency resolution f_{res} of the numerically controlled oscillator according to [26]:

$$f_{res} = \frac{f_{clk}}{2^{b_{acc}}}. \quad (2)$$

This results in a maximum frequency error $\Delta f = f_{res}/2$, which results in a theoretical infidelity of

$$1 - F = \left(\frac{\Delta f}{f_R} \right)^2 = \left(\frac{1}{2^{b_{acc}+1}} \frac{f_{clk}}{f_R} \right)^2, \quad (3)$$

when performing a π -rotation using a rectangular envelope [13]. This result, along with the simulated fidelity in the case of both a rectangular and Gaussian envelope is shown

Fig. 7. Infidelity of a π -rotation as a function of the NCO accumulator number of bits. Eq. 3, valid for rectangular envelopes, is plotted as the theoretically expected fidelity.

in Fig. 7. In the simulation, the target qubit frequency is chosen such that the frequency error is maximized. As the Gaussian envelope has a longer duration, a larger frequency error is accumulated. At least 16 accumulator bits are required to achieve a 99.999% fidelity.⁴

2) *Number of LUT Entries*: For a more efficient design, the minimum number of entries ($2^{b_{lut}}$) should be used in the sine/cosine lookup table. However, as this requires the number of bits out of the accumulator (b_{acc}) to be reduced to the number of LUT address bits (b_{lut}), a periodic error would appear, and, as a result, the spectrum will show spurious tones. While the spectrum depends on the generated frequency (see Fig. 8), the spurs are associated with a limited SFDR equal to [26]:

$$SFDR = 6 b_{lut} \text{ dB}. \quad (4)$$

As such a spurious tone can be at the frequency of an idle qubit, its infidelity is expected to increase to [13]:

$$1 - F = \frac{\theta^2}{4} \cdot 10^{-SFDR/10} \approx \frac{\theta^2}{4^{b_{lut}+1}}. \quad (5)$$

The above theoretical bound is compared to simulations in Fig. 9. As the effects of Gaussian and rectangular envelopes are similar, only the results of the Gaussian envelope are presented. Different target frequencies have been simulated, and, in each condition, an idle qubit is considered at the frequency of the largest spur. For the accumulator output bit reduction, both truncation and rounding are considered. Eq. 5 well predicts the fidelity of the idle qubit, both for rounding and truncation. In the case of rounding, the idle-qubit fidelity requires at least 9 bits for a 99.999% fidelity. In the case of truncation, the targeted qubit is affected more and at least 10 bits are required. When targeting a certain fidelity, the required b_{lut} is one bit less when rounding the accumulator output. Note that saving 1 bit is significant as it halves the number of entries required in the LUT.

⁴At very small frequency errors, the simulated infidelity deviates from the expected infidelity, as the practical reconstruction filter limits the achievable fidelity.

Fig. 8. Depending on the choice of f_{out} , the spectrum will either show (a) spurious tones when $f_{out} = f_s/N$ (with an integer N , e.g. $N = 5$ in the plot) due to the repetitive behavior of errors, or (b) a white spectrum when f_{out} is not an integer sub-multiple of f_s as the periodic behavior of the errors is disturbed.

Fig. 9. The simulated infidelity as a function of the number of LUT entries in case of a Gaussian envelope. The top and bottom plots assume truncation and rounding of the accumulator output, respectively. Multiple plotted lines correspond to the different offset frequencies of 450, 495, 499.5, and 500 MHz, while the theoretically expected fidelity refers to Eq. 5.

3) *Number of LUT Data Bits*: A finite number of data bits in the sine/cosine lookup table (b_{data}) results in a quantization error. Generally, such a quantization error can be modeled as white noise spread over the full Nyquist bandwidth $f_s/2$ with associated Signal-to-Quantization-Noise Ratio of

$$SQNR = 4^{b_{data}} \cdot \frac{3}{2}. \quad (6)$$

Since the qubit is only sensitive to noise in a bandwidth $ENBW = f_R \cdot \frac{\pi}{\theta}$ due to the intrinsic noise filtering of the qubit [13], the expected infidelity for the driven qubit is given by:

$$1 - F = \frac{\theta^2}{4} \cdot \frac{1}{SQNR} \cdot \frac{ENBW}{BW} = \frac{\pi\theta}{3} \cdot \frac{f_R}{f_s} \cdot \frac{1}{4^{b_{data}}}. \quad (7)$$

Fig. 10. The simulated infidelity when reducing the number of data bits in the LUT when using a Gaussian envelope for an offset frequency of 450 MHz. The theoretically expected fidelity due to the spur and noise are given by Eqs. 5 and 7, respectively.

This noise affects both the targeted and idle qubits. For certain output frequencies, however, quantization noise is more tonal (similar to Fig. 8), and the spur could be at the frequency of an idle qubit. To capture these different cases, again, different offset frequencies are used when determining the number of LUT entries, and a victim qubit is simulated at the frequency of the highest spur. In Fig. 10, only the simulated fidelity for an offset frequency of 450 MHz is shown for clarity, as the spectrum shows many spurious tones resulting in significant tones affecting the qubit more than expected from the white-noise model. The simulations with the various offset frequencies show that at least 8 data bits are required for a 99.999% fidelity.

4) *Number of Envelope Bits*: A limited number of bits used for the envelope in the I/Q-modulation (b_{env} , signed) causes an error in the pulse amplitude. For a rectangular envelope, the maximum amplitude inaccuracy is

$$\frac{\Delta A}{A} = \frac{1}{2^{b_{env}}}, \quad (8)$$

leading to an infidelity of [13]:

$$1 - F = \frac{\theta^2}{4} \cdot \left(\frac{\Delta A}{A} \right)^2 = \frac{\theta^2}{4^{b_{env}+1}}. \quad (9)$$

While the amplitude could be different for a rectangular envelope due to quantization noise, the shape of the envelope is unaffected. This is not the case for e.g., a Gaussian envelope, where quantization leads to distortion of the envelope, affecting the signal spectrum.

A simulation is performed by setting the qubit properties such that the ideal driving amplitude for a rectangular envelope is in-between two quantization levels. Although the effect of the quantization noise on another qubit may be relevant for a Gaussian envelope, and it is hence simulated as well, the results in Fig. 11 indicate that such effect is negligible. The simulated fidelity follows the prediction of Eq. 9, resulting in a minimum of 9 bits for a fidelity of 99.999%.

Fig. 11. The simulated infidelity versus the number of envelope bits when using a rectangular or Gaussian envelope for an offset frequency of 500 MHz.

5) *Number of Bits in the Correction Network:* The tolerable phase imbalance (ϕ) and gain imbalance (ϵ) follow from the required image rejection ratio (IRR) (Section 4.2.4 of [32]):

$$IRR \approx \frac{4}{\epsilon^2 + \phi^2}. \quad (10)$$

For an SFDR of 54 dB, to achieve a fidelity of 99.999% due to the image spur (see Section IV B in [13]), the maximum gain imbalance and the maximum phase imbalance are 0.4% (0.035 dB) and 0.32° , respectively.

A correction network is added to compensate for inaccuracies in the analog blocks (see Fig. 5). In this correction network, the coefficients α_I , α_Q , β_I and β_Q are unsigned fractions of b_{frac} bits.⁵

A gain imbalance can be compensated for by lowering either α_I or α_Q , and for a maximum error of 0.4% at least 7 bits are required ($\frac{\Delta A}{A} = \frac{1}{2^{b_{frac}+1}}$). However, since the relation is non-linear for phase imbalance, both the α and β coefficients need to be adapted. As it is difficult to predict the worst-case scenario, a system-level simulation is performed where any phase imbalance from -25° to $+25^\circ$ is introduced and subsequently corrected using a finite number of bits. The situation of the worst-case IRR is further considered when simulating the system along with the quantum processor. The results of this simulation, when using a Gaussian envelope, are shown in Fig. 12.

It can be clearly seen that the fidelity of the qubit at the image frequency equals the fidelity as expected from the spur power (Eq. 5). Besides the victim qubit, the targeted qubit seems affected in the same way. From this simulation, it follows that at least 9 fractional bits in the fixed-point number are required to achieve a fidelity of 99.999%.

D. Total Digital System

To summarize, for a 99.999% fidelity, it was found that at least a 16-bit accumulator is required, of which the

⁵In case different corrections are required at different frequencies, these coefficients could be selected based on the selected NCO.

Fig. 12. The simulated infidelity when reducing the number of fractional bits in the fixed-point number in the I/Q-correction network. Each simulation is performed at the worst-case phase imbalance, and uses a Gaussian envelope to drive a qubit at an offset frequency of 500 MHz. The victim qubit is placed at the image frequency of -500 MHz, and its fidelity is estimated from the simulated SFDR following Eq. 5.

9 most-significant bits, after rounding, are used to index the LUT holding 8-bit values. Moreover, both the envelope and I/Q-correction network require a 9-bit resolution. In the sensitivity analysis, only part of the digital datapath under investigation was quantized, and hence all multiplier outputs were not quantized. As an initial estimate for the entire digital system, these minimum specifications were used and all multiplier outputs were truncated to 9 bits, as at least 9 bits were found necessary for the envelope. Reducing the number of multiplier output bits is critical to save power and to find the minimum number of bits required for the DAC.

A full system simulation was done, where on each qubit a Gaussian-shaped microwave pulse was applied to perform a π -rotation at a 10 MHz Rabi frequency. The operating frequencies of the 32 qubits are evenly spaced over the available 2 GHz band. Furthermore, the system was again simulated with the worst DC and I/Q errors. The fidelity of the performed rotation is recorded, as well as the fidelity of all unaddressed qubits, including an additional one placed at the highest spectral spur.

The fidelity of the resulting system was limited to $\sim 99.996\%$ by the unaddressed qubit at the image frequency when truncating the multiplier outputs. After implementing rounding in the multipliers of the I/Q correction network, the fidelity improved to $\sim 99.998\%$, limited by the unaddressed qubit at the highest spectral spur. When increasing the number of LUT entry bits (b_{lut}) by 1, we are at the edge of achieving the desired fidelity. The result of this simulation is shown in Fig. 13. Finally, the number of accumulator bits (b_{acc}) is increased to 19 to ensure the required frequency accuracy when operating at the lowest Rabi frequency of 1 MHz. A summary of the specifications is given in Table III.

E. Analog Blocks

The coarse specifications for a 99.99% fidelity (100 ppm infidelity) in Table II assume an equal contribution from the

Fig. 13. The simulated infidelity of the digital system with specifications in Table III.

TABLE III
SPECIFICATIONS FOR THE DIGITAL SYSTEM

Number of accumulator bits	19 bits
Number of LUT entries	10 bits (after rounding)
Number of LUT data bits	8 bits
Number of envelope bits	9 bits
Result of envelope multiplication	9 bits (after truncation)
Number of bits in the correction network	9 bits
Result of I/Q correction multiplication	9 bits (after rounding)
Envelope memory	~100 kSa

different errors (~ 10 ppm each), with the previously discussed digital system contributing another ~ 10 ppm to the infidelity. While the assumption of equal error contribution is useful for drafting initial specifications, the trade-offs between these specifications are analyzed in this section in order to budget the different errors for feasibility.

As long as the digital clock frequency and analog gain are stable enough, the pulse amplitude, generated frequency, I/Q phase imbalance, and duration can be guaranteed by the digital section. Following Table II, a variable gain of 44 dB with stability of 0.22% is required from the analog circuit. The frequency accuracy of 3.5 kHz (for a 1-MHz Rabi frequency and a 20-GHz output) requires a 0.18 ppm frequency stability. Such stability can be achieved by a crystal oscillator [33], and easily satisfies the required duration accuracy of $0.11 \text{ ns}/50 \text{ ns} = 0.22\%$. Hence, the duration inaccuracy will hardly contribute to the infidelity.

Assuming that the same frequency generator is used to derive the clock and the LO, the tolerable frequency noise (σ_f) can be translated to the required clock jitter (σ_t) as

$$\sigma_t = \frac{1}{2\pi} \frac{1}{f_0} \sqrt{\frac{f_b}{f_a} \frac{\sigma_f}{f_b}}, \quad (11)$$

where f_0 is the clock frequency, and a phase noise profile of a narrowband PLL with $\sim 1/f^2$ over the frequency range of interest from f_a to f_b ($f_a \ll f_b$) is assumed. A qubit is only sensitive to noise in a bandwidth of $f_b = f_R \cdot \frac{\pi^2}{4}$ for a π -rotation at a Rabi frequency of f_R [13]. For the case of a 1-MHz Rabi oscillation ($\sigma_f = 3.5 \text{ kHz}_{\text{rms}}$) and a 2.5-GHz

clock, this requires an absolute jitter of $\sigma_t < 0.9 \text{ ps}_{\text{rms}}$ ($f_a = f_b/100$ for a total duration of ~ 100 quantum operations). Consequently, the timing jitter requirement of $0.11 \text{ ns}_{\text{rms}}$ is well satisfied, and this error source will hardly contribute to the infidelity. Achieving such a frequency noise is however not trivial; assuming the same phase noise profile, a single-sideband phase noise of -116 dBc/Hz is required at a 1 MHz offset from the carrier.

As the maximum output swing of -16 dBm (50 mV_p) can be directly generated by the DAC, no gain is assumed in the following stages, thus each stage contributes equally to the noise⁶ and distortion. As a representation of those blocks, a single stage CMOS class-A resistive-loaded common-source amplifier, that can serve as the $50\text{-}\Omega$ output driver,⁷ is analyzed in the following.

The maximum RMS output voltage of such an amplifier is given by

$$V_{out} = \frac{I_d \cdot R_L}{\sqrt{2}}, \quad (12)$$

and the RMS output noise voltage by

$$v_n = \sqrt{4 \cdot k_B \cdot T \cdot \gamma \cdot g_m \cdot BW \cdot R_L}, \quad (13)$$

where I_d is the bias current, R_L is the load resistance, k_B is Boltzmann's constant, T the temperature, $\gamma \sim 2$ is the excess noise factor for sub-micron devices and g_m the device transconductance. The Signal-to-Noise Ratio follows as

$$SNR = \frac{V_{out}^2}{v_n^2} = \frac{I_d}{8 \cdot k_B \cdot T \cdot \gamma \cdot \left(\frac{g_m}{I_d}\right) \cdot BW}. \quad (14)$$

Assuming $T = 300 \text{ K}$, a transistor overdrive voltage where $\left(\frac{g_m}{I_d}\right) \sim 10 \text{ V}^{-1}$, and the bandwidth for which the qubit is sensitive to amplitude noise $BW = f_R$ (for a π -rotation) [13], it is found that a bias current $I_d > 0.66 \mu\text{A}$ is required to achieve the 50-dB SNR requirement with a 10-MHz Rabi frequency. Note that this is easily satisfied as $I_d > 1 \text{ mA}$ is required to obtain the desired output voltage swing over a $50\text{-}\Omega$ load.

Assuming a CMOS single-ended amplifier, and an ideal square law device, the 2nd-order distortion is given by (Section 5.7 of [32])

$$HD2 = \frac{1}{4} \cdot \frac{V_{in}}{V_{gs} - V_T}, \quad (15)$$

where V_{in} is the input voltage, V_{gs} is the gate-to-source voltage and V_T is the device threshold voltage. Given the requirement of $HD2 < -44 \text{ dB}$, and assuming no gain ($V_{in,max} = 50 \text{ mV}_p$), an unrealistic overdrive voltage $V_{gs} - V_T > 2 \text{ V}$ is required. In order not to be limited by HD2, a differential circuit topology can be considered with a 3rd-order distortion of (Section 5.7 of [32])

$$HD3 = \frac{1}{18} \left[\frac{V_{out,p}}{(V_{gs} - V_T) \cdot g_m \cdot R_L} \right]^2, \quad (16)$$

⁶The DAC quantization noise is already accounted for in the digital specifications.

⁷The same noise analysis is also valid for e.g. a current-steering DAC.

TABLE IV
SPECIFICATIONS FOR THE ANALOG SYSTEM

Clock sample rate (f_s)	2.5 GHz
Reference clock stability	0.18 ppm
Clock jitter	0.9 ps _{rms}
LO frequency range	6-19 GHz
LO phase noise	-116 dBc/Hz at 1 MHz offset
DAC Resolution	9 bit
DAC ENOB	7 bit
Reconstruction filter	2 nd -order Chebyshev-I at 1.8 GHz
Analog gain stability	0.22%
Analog gain control	44 dB
SNR	50 dB ($I_d > 0.66 \mu\text{A}$)
HD3	-44 dB ($I_d > 1.5 \text{ mA}$)

where $V_{out,p}$ is the peak amplitude. Achieving HD3 < -44 dB requires an overdrive $V_{gs} - V_T > 0.15 \text{ V}$ (assuming the gain $g_m \cdot R_L = 1$). For a device in saturation, $V_{gs} - V_T = 2 \cdot \left(\frac{g_m}{I_d}\right)^{-1}$. With a g_m of $\frac{1}{50 \Omega}$, a bias current larger than 1.5 mA is required.⁸ Finally, an SFDR < -44 dB requires a DAC with an effective number of bits (ENOB) of 7.

To summarize, the proposed design requirements are specified in Table IV. Of these requirements, the reference clock stability and LO frequency noise requirements appear most stringent. As the duration accuracy, timing jitter and amplitude noise specifications of Table II are most easily satisfied, their error contribution can be reduced to relax the specifications on the more stringent ones to save power. However, the SFDR as specified in Table II is not part of any error budgeting as it is the only error source considered affecting idle qubits, and hence this specification cannot be relaxed.

F. Power Consumption Estimate

While an accurate estimation of the power consumption requires knowledge of the exact digital and analog circuit implementation, in this section an estimate is given based on the previously found specifications and implementation examples found in literature.

A direct digital synthesizer with similar specifications (9-bit amplitude, 2-GHz clock and 55-dB SFDR), has been implemented in 55-nm CMOS while consuming 25 mW in the 32-bit NCO and 37 mW in the phase-to-amplitude conversion [34]. Considering our system with 32 19-bit NCOs operating at 2.5 GHz and a single phase-to-amplitude conversion block, a power consumption of 640 mW is expected. Similarly, in 65-nm CMOS, a 10-bit multiplier operating at 2.5 GHz consumes 14 mW [35], and hence an additional 112 mW is expected in our digital modulation and I/Q correction network (8 multipliers), bringing the total digital power consumption to ~750 mW. Based on the study presented in [36], a power consumption of ~160 mW is expected in a 22-nm CMOS node, with 80% of the power consumed in the NCOs, i.e. 4 mW/NCO.

As found from the analog specifications, a single-transistor bias current of 1.5 mA is required to meet the linearity

⁸In case of simultaneous excitation of multiple qubits (Section V-B), an IM3 = 3 HD3 < -44 dB is required and consequently a bias current larger than 2.6 mA.

requirement if the entire circuit consists of a single stage. However, a more realistic implementation consists of at least 2 stages contributing to the distortion, e.g. current-steering DACs driving a 50- Ω passive reconstruction filter which in turn drives a double-balanced I/Q mixer driving the 50- Ω output load. Considering a 2-stage implementation, a single-transistor bias current of $1.5 \text{ mA} \cdot \sqrt{2} = 2.1 \text{ mA}$ is required.⁹ As there are 2 stages, each differential, with I/Q, a total current of at least 17 mA is required (17 mW with a 1-V supply).

As about 36 mW is expected for the digital section in case of a single NCO and further reduction in digital power is promised going to a more advanced CMOS node, the power consumption is well-balanced between the analog and digital section, with another 4 mW required for every NCO, i.e. qubit, that is added.

V. APPLICATION EXAMPLE

Compared to state-of-the-art controllers based on general-purpose instruments or tailor-made controllers employing FPGAs [16], the presented solution offers the highest number of frequency multiplexed control channels and is maximally tailored to the quantum processor requirements allowing for a reduced power consumption. Implementing the proposed controller as a CMOS SoC will reduce its form factor, potentially enabling operating this power-efficient controller physically close to the qubits. The advantages of such a digital-intensive microwave signal generator can be observed by considering application examples for qubit control, as illustrated in this section.

A. Qubit Tune-Up

Besides the intended application of performing single-qubit operations, the control architecture can, for example, be used to tune-up the qubit processor. Part of this tune-up protocol is to find the qubit resonance frequency. The adiabatic fast passage technique uses a chirp pulse to sweep the microwave frequency across the spin resonance frequencies of multiple qubits in an FDMA setup, thereby smoothly rotating all spins whose resonant frequencies lie within the range [37]. Generating such a chirp pulse using the system architecture presented in the paper can be readily implemented using the following waveform for the in-phase (I) and quadrature-phase (Q) part of the envelope:

$$\phi[n] = 2\pi \cdot \frac{f_{\max} - f_{\min}}{f_s} \cdot \sum_{i=0}^{n-1} \left(\frac{i}{N} - \frac{1}{2} \right) \quad (17)$$

$$I[n] = A \cdot \sin(\phi[n]) \quad (18)$$

$$Q[n] = A \cdot \cos(\phi[n]), \quad (19)$$

for envelope samples $n = 1$ to N , resulting in a chirp from frequency $f_{\min}$ to $f_{\max}$ using N samples (total chirp time $T_{\text{chirp}} = N/f_s$) and amplitude A .

As an example, a binary search for the qubit frequency is shown in Fig. 14 using the designed system. Multiple frequency chirps are used over a frequency band that is halved every cycle of the search, narrowing down on the actual qubit

⁹Each stage requires a 6 dB stricter HD3 of -50 dB (Section 2.2.5 of [32]).

frequency indicated by the black dashed line. In the first chirp, the frequency is swept over the lower sideband (LSB) from 5 to 6 GHz, and it is observed whether the qubit rotates or not. In case the qubit rotates, the qubit resonance frequency is in the LSB and the search continues there, otherwise the search continues in the upper sideband (USB). In order to keep the power spectral density the same when the frequency band is halved, the signal amplitude is gradually reduced with each step.

Besides the qubit resonance frequency, the required pulse duration and amplitude should be determined during tune-up to calibrate the rotation angle. This is typically done by performing a Rabi oscillation where either the pulse duration or amplitude is incremented in small steps and the resulting rotation angle is measured. Due to the option to program any pulse envelope, the pulse duration and/or amplitude can easily be varied to perform such a Rabi oscillation and finalize the calibration of the qubit operation.

As stated previously, the system is optimized by assuming sequential execution of operations on different qubits. However, the chosen system architecture supports the excitation of multiple qubits simultaneously when having a digital modulator and correction network for each channel (Fig. 15). The required DAC resolution increases to:

$$b_{dac} = 9 + \log_2(N_{ch}), \quad (20)$$

Fig. 15. Simplified block diagram of the system for the case of 2 simultaneous excitation channels.

Fig. 16. Fidelity for the simultaneous excitation of 2 qubits spaced by a frequency Δf when using uncompensated and compensated Gaussian envelopes.

Fig. 17. The compensated Gaussian envelope for the simultaneous excitation of 2 qubits spaced by a frequency $\Delta f = 5 f_R$. Top: amplitude and phase components; bottom: the in-phase and quadrature-phase components.

where N_{ch} is the number of simultaneous channels. For the following example, the simulation model is adapted to allow for the simultaneous excitation of 2 qubits with a 10-bit DAC.

When simultaneously exciting 2 qubits using standard Gaussian envelopes (Fig. 17, amplitude modulation only) the fidelity is limited when the qubits are close in frequency due to the AC-Stark shift [15]¹⁰ (Fig. 16). This effect shifts the resonance frequency of the qubit when an off-resonance pulse is applied. To account for this frequency shift and to compensate for its effect, phase modulation must be added

¹⁰Moreover, additional spurious tones can be present due to intermodulation distortion affecting the performance of other qubits.

besides the Gaussian amplitude modulation [15], as shown in Fig. 17 (top). The resulting in-phase and quadrature-phase components that are used for the digital modulation are shown in Fig. 17 (bottom). With these compensated Gaussian envelopes, a high fidelity can be achieved for 2 qubits spaced closely in frequency while being driven simultaneously in the presented control system (Fig. 16).

Thanks to the digital-intensive back-end that allows individual I/Q modulation for each channel, simultaneous excitation of multiple qubits is easily implemented in the presented system.

VI. CONCLUSION

Deriving the system specifications of the classical electronic controller for qubits and determining the optimal error budget are crucial in designing power efficient circuits. To meet these specifications, design trade-offs between several system architectures have been compared in this paper, resulting in the proposal of an efficient architecture exploiting frequency multiplexing for multi-qubit control. Co-simulation of the proposed electronic system and the qubits was used to assess the effect of non-idealities of each circuit block on qubit fidelity. Based on such analysis, the design specifications of each block have been determined to achieve the required gate fidelity while optimizing power consumption. Finally, the effectiveness and flexibility of such a system has been shown by demonstrating relevant practical applications, such as qubit tune-up and simultaneous qubit excitation. As a result of the proposed design methodology, we have obtained the blueprint for a power-efficient integrated electronic controller to realize single-qubit operations for practical large-scale quantum computers.

REFERENCES

- [1] A. Kandala *et al.*, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, p. 242, 2017.
- [2] M. Ganzhorn *et al.*, "Gate-efficient simulation of molecular eigenstates on a quantum computer," *Phys. Rev. A, Gen. Phys.*, vol. 11, no. 4, Apr. 2019, Art. no. 044092.
- [3] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, "Elucidating reaction mechanisms on quantum computers," *Proc. Nat. Acad. Sci. USA*, vol. 114, no. 29, pp. 7555–7560, Jul. 2017. [Online]. Available: <https://www.pnas.org/content/114/29/7555>
- [4] F. Arute *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [5] L. Vandersypen *et al.*, "Interfacing spin qubits in quantum dots and donors—Hot, dense, and coherent," *NPJ Quantum Inf.*, vol. 3, no. 1, p. 34, 2017.
- [6] C. A. Ryan, B. R. Johnson, D. Ristè, B. Donovan, and T. A. Ohki, "Hardware for dynamic quantum computing," *Rev. Sci. Instrum.*, vol. 88, no. 10, Oct. 2017, Art. no. 104703.
- [7] H. Homulle *et al.*, "A reconfigurable cryogenic platform for the classical control of quantum processors," *Rev. Sci. Instrum.*, vol. 88, no. 4, Apr. 2017, Art. no. 045103.
- [8] D. J. Reilly, "Engineering the quantum-classical interface of solid-state qubits," *NPJ Quantum Inf.*, vol. 1, no. 1, p. 15011, Dec. 2015.
- [9] B. Patra *et al.*, "Cryo-CMOS circuits and systems for quantum computing applications," *IEEE J. Solid-State Circuits*, vol. 53, no. 1, pp. 309–321, Jan. 2018.
- [10] J. C. Bardin *et al.*, "29.1 A 28nm bulk-CMOS 4-to-8GHz 12 mW cryogenic pulse modulator for scalable quantum computing," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 456–458.
- [11] L. Petit *et al.*, "Spin lifetime and charge noise in hot silicon quantum dot qubits," *Phys. Rev. Lett.*, vol. 121, no. 7, Aug. 2018, Art. no. 076801.
- [12] J. P. G. van Dijk *et al.*, "The impact of classical control electronics on qubit fidelity," 2018, *arXiv:1803.06176*. [Online]. Available: <http://arxiv.org/abs/1803.06176>
- [13] J. P. G. van Dijk *et al.*, "Impact of classical control electronics on qubit fidelity," *Phys. Rev. A, Gen. Phys.*, vol. 12, no. 4, Oct. 2019, Art. no. 044054, doi: [10.1103/PhysRevApplied.12.044054](https://doi.org/10.1103/PhysRevApplied.12.044054).
- [14] J. van Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano, "A co-design methodology for scalable quantum processors and their classical electronic interface," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2018, pp. 573–576.
- [15] M. Steffen, L. M. K. Vandersypen, and I. L. Chuang, "Simultaneous soft pulses applied at nearby frequencies," *J. Magn. Reson.*, vol. 146, no. 2, pp. 369–374, Oct. 2000. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1090780700921785>
- [16] J. P. G. van Dijk, E. Charbon, and F. Sebastiano, "The electronic interface for quantum processors," *Microprocessors Microsyst.*, vol. 66, pp. 90–101, Apr. 2019.
- [17] T. F. Watson *et al.*, "A programmable two-qubit quantum processor in silicon," *Nature*, vol. 555, no. 7698, pp. 633–637, Mar. 2018, doi: [10.1038/nature25766](https://doi.org/10.1038/nature25766).
- [18] J. Yoneda *et al.*, "A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%," *Nature Nanotechnol.*, vol. 13, no. 2, p. 102, 2018.
- [19] E. Kawakami *et al.*, "Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot," *Nature Nanotechnol.*, vol. 9, no. 9, p. 666, Sep. 2014, doi: [10.1038/nnano.2014.153](https://doi.org/10.1038/nnano.2014.153).
- [20] M. Veldhorst *et al.*, "An addressable quantum dot qubit with fault-tolerant control-fidelity," *Nature Nanotechnol.*, vol. 9, pp. 981–985, Oct. 2014, doi: [10.1038/nnano.2014.216](https://doi.org/10.1038/nnano.2014.216).
- [21] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Phys. Rev. A, Gen. Phys.*, vol. 86, no. 3, Sep. 2012, Art. no. 032324.
- [22] R. Barends *et al.*, "Superconducting quantum circuits at the surface code threshold for fault tolerance," *Nature*, vol. 508, no. 7497, pp. 500–503, Apr. 2014.
- [23] Y. M. Greshishchev *et al.*, "A 56GS/S 6b DAC in 65nm CMOS with 256× 6b memory," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2011, pp. 194–196.
- [24] L. Duncan *et al.*, "A 10-bit DC-20-GHz multiple-return-to-zero DAC with >48-dB SFDR," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3262–3275, Dec. 2017.
- [25] E. Bechthum, G. I. Radulov, J. Briare, G. J. G. M. Geelen, and A. H. M. van Roermund, "A wideband RF mixing-DAC achieving IMD < -82 dBc up to 1.9 GHz," *IEEE J. Solid-State Circuits*, vol. 51, no. 6, pp. 1374–1384, Jun. 2016.
- [26] "A technical tutorial on digital signal synthesis," Analog Devices, Inc., Appl. Note, 1999, pp. 1–122.
- [27] C. A. Ryan, B. R. Johnson, D. Ristè, B. Donovan, and T. A. Ohki, "Hardware for dynamic quantum computing," *Rev. Sci. Instrum.*, vol. 88, no. 10, Oct. 2017, Art. no. 104703, doi: [10.1063/1.5006525](https://doi.org/10.1063/1.5006525).
- [28] Z. Chen *et al.*, "Measuring and suppressing quantum state leakage in a superconducting qubit," *Phys. Rev. Lett.*, vol. 116, no. 2, Jan. 2016, Art. no. 020501, doi: [10.1103/PhysRevLett.116.020501](https://doi.org/10.1103/PhysRevLett.116.020501).
- [29] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient Z gates for quantum computing," *Phys. Rev. A, Gen. Phys.*, vol. 96, no. 2, Aug. 2017, Art. no. 022330, doi: [10.1103/PhysRevA.96.022330](https://doi.org/10.1103/PhysRevA.96.022330).
- [30] E. A. Vittoz, "Low-power design: Ways to approach the limits," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 1994, pp. 14–18.
- [31] L. M. K. Vandersypen, "Experimental quantum computation with nuclear spins in liquid solution," Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, USA, 2001.
- [32] B. Razavi, "RF microelectronics," in *Prentice-Hall Communications Engineering and Emerging Technologies Series*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2011.
- [33] M. H. Roshan *et al.*, "11.1 dual-MEMS-resonator temperature-to-digital converter with 40k resolution and FOM of 0.12pJK²," in *IEEE ISSCC Dig. Tech. Papers*, Jan. 2016, pp. 200–201.
- [34] T. Yoo *et al.*, "A 2 GHz 130 mW direct-digital frequency synthesizer with a nonlinear DAC in 55 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2976–2989, Dec. 2014.
- [35] A. Martinez Alonso, M. Miyahara, and A. Matsuzawa, "A 12.8-ns-latency DDFS MMIC with frequency, phase, and amplitude modulations in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 53, no. 10, pp. 2840–2849, Oct. 2018.

- [36] G. G. Shahidi, "Chip power scaling in recent CMOS technology nodes," *IEEE Access*, vol. 7, pp. 851–856, 2019.
- [37] C. J. Hardy, W. A. Edelstein, and D. Vatis, "Efficient adiabatic fast passage for NMR population inversion in the presence of radiofrequency field inhomogeneity and frequency offsets," *J. Magn. Reson.*, vol. 66, no. 3, pp. 470–482, Feb. 1986. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0022236486901903>

Jeroen P. G. van Dijk (Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 2013 and 2016, respectively, where he is currently pursuing the Ph.D. degree in applied quantum architectures under supervision of Prof. Dr. E. Charbon. During his Ph.D., he worked as an RF & Mixed Signal IC Design Intern with Intel Labs, Hillsboro, OR, USA, from 2017 to 2018. His research interests include analog and mixed-signal circuit design, low-power

techniques, and single-electron spin-qubits. He was a recipient of the Best IP Award at DATE in 2018.

Bishnu Patra (Student Member, IEEE) received the M.Sc. degree in electrical engineering (Microelectronics) from the Delft University of Technology, Delft, The Netherlands, in 2016, where he is currently pursuing the Ph.D. degree, focusing on cryogenic RF Integrated circuits for quantum computing applications.

From 2013 to 2014, he was a B.Sc. Researcher with the ELCA Department, Delft University of Technology. During his Ph.D., he worked as an RF & Mixed Signal IC Design Intern with Intel Labs, Hillsboro, OR, USA, from 2017 to 2018. His current research interests include RF transmitters for qubit control, CMOS frequency synthesizers, and design and modeling of microwave passive components at cryogenic temperatures. He was a recipient of the 2014 M.Sc. Faculty Scholarship by the Microelectronics Research Department, Delft University of Technology, and the IEEE SSCS Predoctoral Achievement Award (2019–2020).

Stefano Pellerano (Member, IEEE) was born in Bari, Italy, in 1977. He received the Laurea (*summa cum laude*) and Ph.D. degrees in electronics engineering from the Politecnico di Milano, Milan, Italy, in 2000 and 2004, respectively. During his Ph.D., his activity was focused on the design of fully integrated low-power frequency synthesizers for WLAN applications. In 2003, he has been a Consultant with Agere Systems (former Bell Labs), Allentown, PA, USA. Since 2004, he has been with Intel Labs, Hillsboro, OR, USA. He is currently a Principal

Engineer leading the Next Generation Radio Integration Lab, where he drives several research activities focused at enabling radio circuit integration in deeply scaled CMOS technologies. His main research contributions include MIMO transceivers for WiFi, digital PLLs, high-efficient digital architectures for polar and outphasing transmitters, mm-wave radio transceiver and phased-array systems, and low-power radios. Recently, he is also exploring cryogenic CMOS integrated electronics for qubit control in fault tolerant scalable quantum computers. He has authored or coauthored more than 40 IEEE conference and journal papers, one book chapter, and holds more than 15 issued patents. He served as the Technical Program Chair and General Chair for the IEEE Radio Frequency Integrated Circuit (RFIC) Symposium in 2018 and 2019, and he is currently a part of the RFIC Executive Committee. He is currently serving as the Wireless Subcommittee Chair of the IEEE International Solid-State Circuit Conference (ISSCC).

Edoardo Charbon (Fellow, IEEE) received the Diploma degree from ETH Zürich in 1988, the M.Sc. degree from the University of California at San Diego in 1991, and the Ph.D. degree from the University of California at Berkeley in 1995, all in electrical engineering and EECS. He has consulted with numerous organizations, including Bosch, X-Fabs, Texas Instruments, Maxim, Sony, Agilent, and the Carlyle Group. He was with Cadence Design Systems from 1995 to 2000, where he was an Architect of the company's initiative on information hiding for intellectual property protection. In 2000, he joined Canesta Inc., as the Chief Architect, where he led the development of wireless 3-D CMOS image sensors. Since 2002, he has been a member of the Faculty of EPFL, where has been a Full Professor since 2015. From 2008 to 2016, he was a Full Professor with the Delft University of Technology and the Chair of VLSI design. He has been the driving force behind the creation of deep-submicron CMOS SPAD technology, which is mass-produced since 2015, and is present in telemeters, proximity sensors, and medical diagnostics tools. Since 2014, he has advocated cryo-CMOS for the classical control of qubits in scalable quantum computers. He is a Distinguished Visiting Scholar with the W. M. Keck Institute for Space at Caltech and a Fellow of the Kavli Institute of Nanoscience Delft. He has authored or coauthored over 250 articles and two books and holds 22 patents. His interests span from 3-D vision, FLIM, FCS, NIROT to super-resolution microscopy, time-resolved Raman spectroscopy, and cryo-CMOS circuits and systems for quantum computing.

Fabio Sebastiano (Senior Member, IEEE) received the B.Sc. (*cum laude*) and M.Sc. (*cum laude*) degrees in electrical engineering from the University of Pisa, Italy, in 2003 and 2005, respectively, the M.Sc. degree (*cum laude*) from the Sant'Anna School of Advanced Studies, Pisa, Italy, in 2006, and the Ph.D. degree from the Delft University of Technology, The Netherlands, in 2011.

From 2006 to 2013, he was with NXP Semiconductors Research, Eindhoven, The Netherlands, where he conducted research on fully integrated CMOS frequency references, deep-submicron temperature sensors, and area-efficient interfaces for magnetic sensors. In 2013, he joined the Delft University of Technology, where he is currently an Associate Professor. He has authored or coauthored one book and over 70 technical publications and holds 11 patents. His main research interests include cryogenic electronics for quantum computing, quantum computing, sensor read-outs, and fully integrated frequency references.

Dr. Sebastiano is on the technical program committee of the IEEE RFIC Symposium. He is currently serving as an Associate Editor of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS. He was a co-recipient of the 2008 ISCAS Best Student Paper Award and the 2017 DATE Best IP Award, and has served as a Distinguished Lecturer of the IEEE Solid-State Circuit Society.

Masoud Babaie (Member, IEEE) received the Ph.D. degree (*cum laude*) in electrical engineering from the Delft University of Technology, Delft, The Netherlands, in 2016.

In 2006, he joined the Kavoshcom Research and Development Group, Tehran, where he was involved in designing wireless communication systems. From 2009 to 2011, he was a CTO of that company. From 2014 to 2015, he was a Visiting Scholar Researcher with the Berkeley Wireless Research Center, Berkeley, CA, USA. In 2016, he joined the Delft University of Technology, where he is currently a tenured Assistant Professor. His current research interests include RF/millimeter-wave integrated circuits and systems for wireless communications, and cryogenic electronics for quantum computation.

Dr. Babaie has been a Committee Member of the Student Research Preview (SRP) of the IEEE International Solid-State Circuits Conference (ISSCC) since 2017. He is currently serving on the technical program committee of the IEEE European Solid-State Circuits Conference (ESSCIRC). He was a co-recipient of the 2015–2016 IEEE Solid-State Circuits Society Pre-Doctoral Achievement Award and the 2019 IEEE ISSCC Best Demo Award. In 2019, he also received the Veni Award from the Netherlands Organization for Scientific Research (NWO).