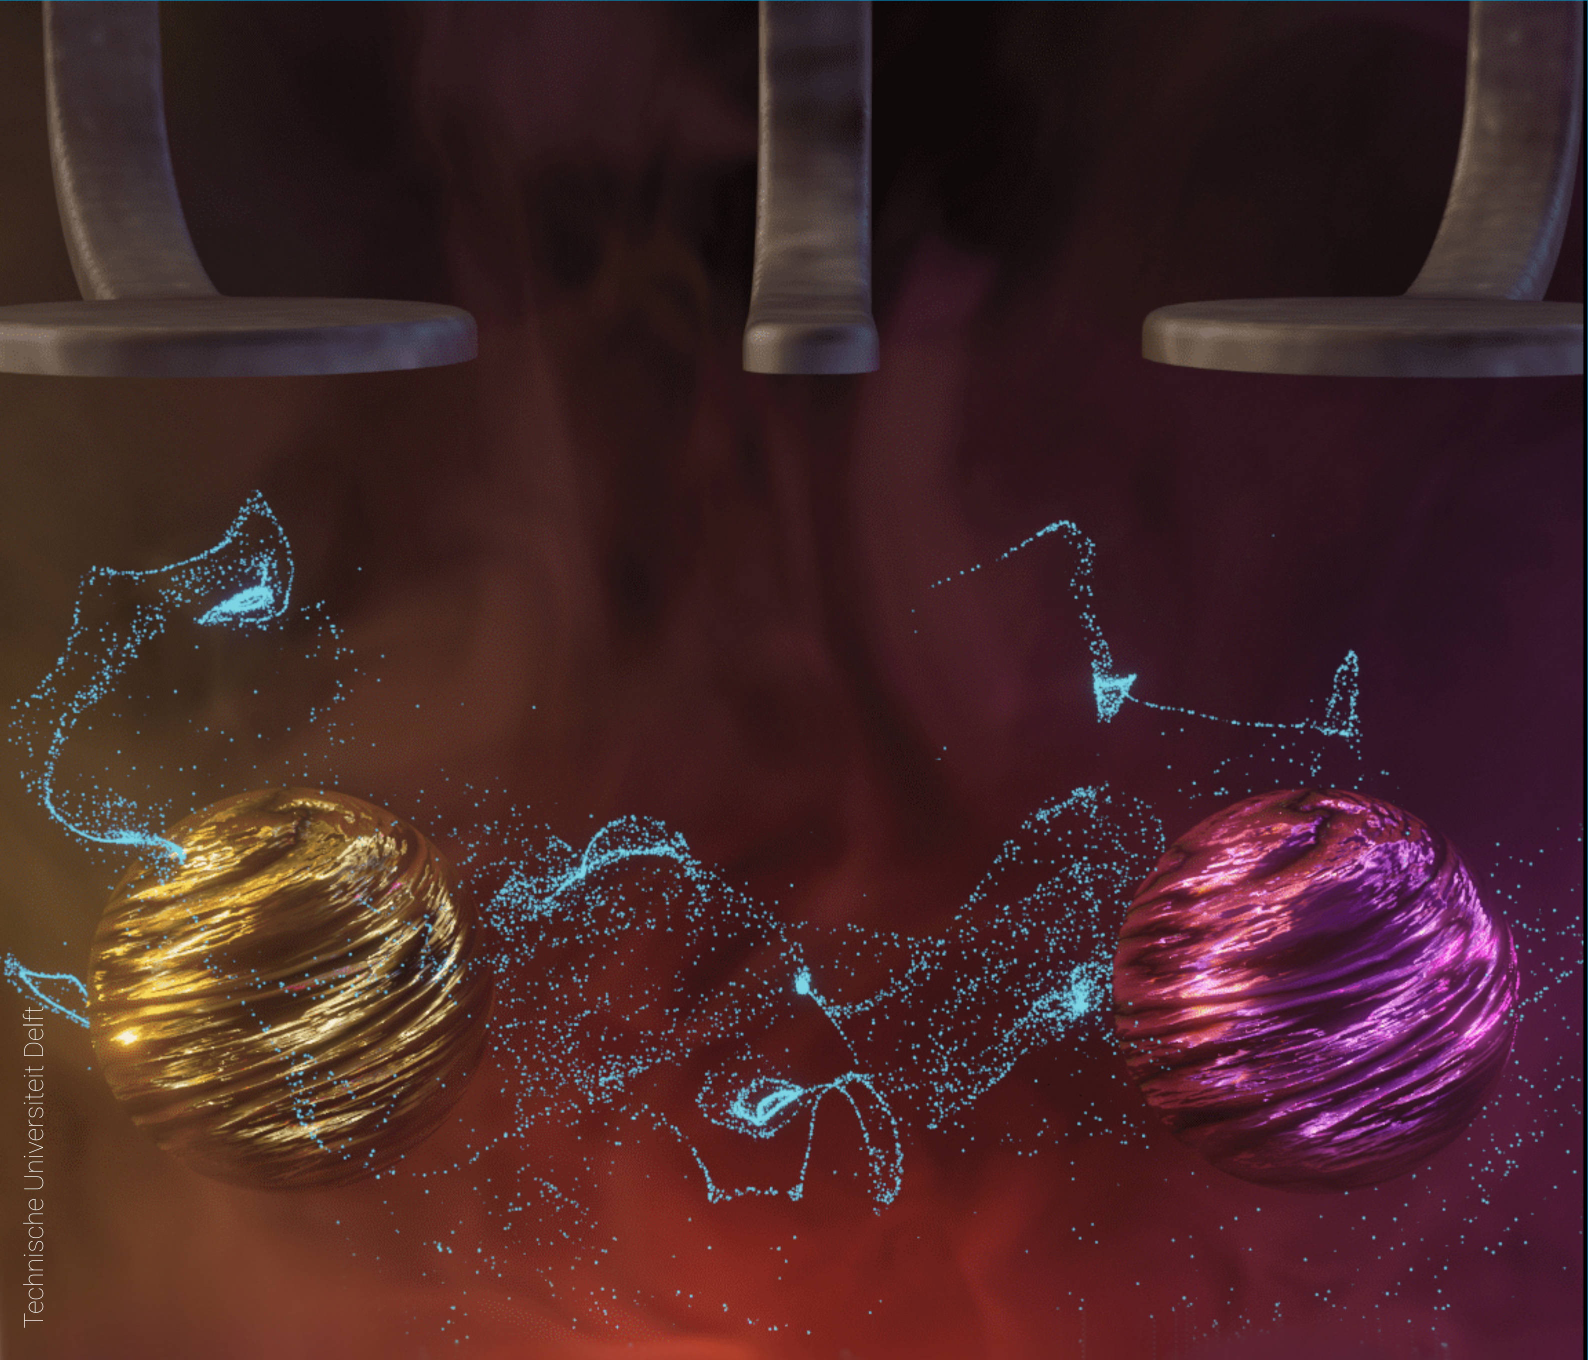

Technische Universiteit Delft

High fidelity single-qubit gates in multi-qubit systems using spin qubits in Si quantum dots

Vincent A. Bejach

High fidelity single-qubit gates in multi-qubit systems using spin qubits in Si quantum dots

by

Vincent A. Bejach

to obtain the degree of Master of Science in Embedded Systems
at the EEMCS faculty of Delft University of Technology,
to be defended publicly on December 14$^{\text{th}}$, 2020.

Student number:	4742796	
Project duration:	March 14, 2019 – December 14, 2020	
Supervisors:	Prof. dr. ir. L.M.K Vandersypen,	TU Delft, supervisor
	Dr. R. Ishihara,	TU Delft, supervisor
	S. Philips,	TU Delft

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

The cover picture was realised by Luca Petit, and has been taken from https://qutech.nl/2020/04/15/hot-dense-and-coherent-scalable-quantum-bits-operate-under-practical-conditions/

TU Delft

Preface

This report marks the end of my studies at Delft University of Technologies, as well as of a journey that lasted more than three years. There have been some highs and lows, and I do have a few regrets about my time spent here, but I am happy to say that these have been very fulfilling and overall happy years, and that I feel richer for having lived them.

My first thank you goes to Lieven, Stephan and Ryoichi, for their guidance and feedbacks through this project. I would also like to thank them for being patient and understanding regarding my multiple delays. I would especially like to thank Lieven for agreeing to take a master student with little experience in quantum computing and for trusting me to still carry out the work. I has sometimes been difficult, but I learned a lot, both about the field and about myself, and I am really glad to have had this opportunity.

Many thanks also go to Max for the theoretical insights he gave me through our discussions, these have helped me a lot.

My greatest appreciation also goes to Jonathan, Theo, Paul and Paul. It has been an immense pleasure to meet or reconnect with you guys during the past three and a half years, and I really hope to see more of you once the current health situation improves a bit.

I would also like to thank my mother, sister and brother, because talking with you has always been a joy, and because despite now living far away from each other, we are no less close than we were when I lived in Paris. Having your support, thoughts and opinions has meant a great deal during the darker periods of my Dutch stay. While I'm talking about family, I also want to thank my mother, sister, brother and grandparents for their valiant efforts to understand my MSc topic.

And last but so very not least, I want to thank Marie. You have been by far the most influential person in my life, and there is no way I can thank you enough for what you have done for me. I'll try nonetheless. Thank you for living through our darkest times with me. Thank you for supporting me when I was so sick I could barely leave my bed. Thank you for the countless hours of encouragement and feedback you gave me throughout my master (if I graduate now, it is in no small part because of you). But above all, thank you for being my best friend, my little sister in heart, *mon kako*.

Vincent A. Bejach
Delft, November 2020

Abstract

Electron spins trapped in quantum dots have recently proven to be a promising technology for the implementation of qubits, already demonstrating high fidelity single- and two-qubits gates. The next step towards fault-tolerant quantum computing is to increase the number of so-called spin qubits on the processor. However, this poses several challenges, one of which being how to implement single-qubit gates in multi-qubit systems, with high fidelity.

This work focused on two aspects. The first one was to identify the physical phenomena and limitations that hinder the realisation of high fidelity single-qubit gates in multi-qubit environments. The second objective was to investigate, implement and optimse methods in order to minimise the impact of these perturbing phenomena,

The identified perturbing effects include unwanted driving between a pulse and the neighbouring qubits, as well as frequency shifts induced by drive-spin coupling and non-ideal pulse in experimental setups. Crosstalk was addressed using pulse shaping (*i.e.* amplitude or phase modulation of the driving pulse to tailor its spectral characteristics and reduce the energy transmitted at unwanted frequencies). The performance of shapes taken from both NMR and signal processing literature was investigated, for various pulse parameters, through an extensive grid search. In addition, a frequency correction algorithm was devised: off-resonant drive frequencies are selected so that the shifted qubits are resonantly driven. It currently accounts for two phenomena: the AC Stark and Bloch-Siegert shifts. The algorithm furthermore tracks the changes caused by the time-dependent characteristics of the shaped pulses.

The proposed frequency correction algorithm was shown to almost entirely negate the effects of the considered shifts, leading to unitary fidelity improvements by up to 55%. In addition, pulse shaping was demonstrated to noticeably improve the fidelity of simultaneous single-qubit rotations compared to unshaped driving. Rotations with fidelities as high as 99.98% were obtained for $\pi/2$ rotations on two-qubits systems. Moreover, shapes whose Fourier transform is narrow and sharp, associated with low Rabi frequencies, were demonstrated to generally provide the highest fidelities of the tested configurations. Lastly, the trends and guidelines highlighted by these results were shown to scale to systems with larger numbers of qubits.

The correction techniques investigated in this work have proven promising for the implementation of high fidelity single-qubit gates in multi-qubit systems. In particular, the guidelines for selecting a well-performing pulse shape should also be useful for the design of optimised driving schemes, regardless of the number of qubits involved. Additionally, the proposed frequency shift correction algorithm is expected to be able to handle arbitrary shifts, and so to be easily adaptable to use in experiments.

Contents

Nomenclature

Abbreviations and acronyms

AWG	Arbitrary wave generators
DFT	Discrete Fourier Transform
FT	Fourier Transform
MC	Monte-Carlo
NMR	Nuclear Magnetic Resonance
TLS	Two-level system
RWA	Rotating-wave approximation

Latin symbols

a	Free parameter of the Gaussian and gaussian-like shapes
	Constant term of polynomial in Hermite shape
b	2^{nd} order coefficient of polynomial in Hermite shape
c_e	Excited state probability amplitude
c_g	Ground state probability amplitude
C	Free parameter of the optimised Blackman shape
$\|e\rangle$	Excited state
f	Frequency
F	Fidelity
$\bar{F}$	Infidelity
$\|g\rangle$	Ground state
$\hbar$	Reduced Planck constant
H	Hamiltonian of the system
$\tilde{H}$	Hamiltonian expressed in the rotating reference frame
I_0	Modified 0^{th}-order Bessel function of the first kind
I^j	Tensor product of the identify operator in dimension 2 with itself j times
J	Exchange rate between two spins
m	Number of driving pulses
n	Number of qubit
P	Population of the state indicated in subscript
s	Shape function
S	Spin momentum
U	Unitary time-evolution operator
	Unitary transformation from non-rotating to rotating frame
t	Time
$\hat{x}$	Discrete Fourier transform

Greek symbols

α	Free parameter of the Kaiser's modified 0^{th}-order Bessel functions
Δ	Detuning between the resonant frequency of a qubit and the pulse frequency
$\Delta\omega_0$	Frequency separation between two adjacent qubits
$\Delta\omega_{BS}$	Block-Siegert shift
$\Delta\omega_S$	AC Stark shift
θ	Stuckelberg angle (*i.e.* angle between the dressed and uncoupled quantization axes)
	Rotation angle
ρ	Density matrix representation of the quantum state
σ	Density matrix of the state produced by an ideal gate
	Width of the curve for a Gaussian, gaussian-like or Hermite shape
σ^x,σ^y,σ^z	Pauli X,Y, and Z matrices
τ	Pulse duration
ϕ	Initial phase of the driving pulse
$\lvert\psi\rangle$	state of the two-level system
ω	Angular frequency of the driving pulse
ω_0	Resonance frequency
Ω	Rabi frequency
$\tilde{\Omega}$	Generalised Rabi frequency

Subscripts

0	Referring to an ideal gate
BS	Referring to the Bloch-Siegert
dressed	Referring to the dressed qubit
drive	Referring to the driving pulse
ex or EX	Referring to the exchange interaction between two adjacent spins
free	Referring to the free evolution of the spin
i, j	Referring to different spins
multi	Multi-qubit system
S	Referring to the AC Stark shift
shape	Referring to the shape whose name is used as subscript
shifted	Referring to the shifted qubit
spin	Referring to the spin
x, y, z	Referring to x, y, and z axes

Superscripts

†	Referring to the complex conjugate

1
Introduction

Interest in quantum computation has been growing over the last years, as it theoretically allows not only significantly more efficient computations compared to classical computers (to the point of being able to tackle classically intractable problems), but also the possibility of improved information security. However, its implementation poses a lot of challenges. Several different technologies are currently being investigated as candidates to implement quantum bits (qubits). Among them, the spin qubit has recently shown encouraging results. Spin qubits consist of the spin of a single electron trapped in a quantum dot (*i.e.* a small conductive region in a semiconductor that can contain an arbitrary but known number of electrons which occupy well-defined, discrete quantum states). The spin state encodes the information and is manipulated using resonant microwave pulses in a transverse magnetic field gradient (produced by a local micromagnet). Over the last years, spin qubits have been shown to perform quite promisingly. So far, there have been demonstrations of high-fidelity control for single- (Kawakami et al., 2014, Veldhorst et al., 2014, Yang et al., 2019, Yoneda et al., 2018) and two-qubit gates (Huang et al., 2019, Watson et al., 2018, Zajac et al., 2018), as well as a combination of both (Veldhorst et al., 2015).

The next step towards the implementation of a fault-tolerant quantum computer based on spin qubits is to expand the number of qubits on the processor. The current aim is to realise three- or five-qubits devices. This scaling of the number of qubits raises several challenges, one of which being how to implement simultaneous single-qubit gates with high fidelity in multi-qubit systems. Indeed, as the microwave drive is shared between all qubits, each pulse creates crosstalk with the other ones. Applying microwaves to one qubit will thus also partially drive the others. Furthermore, couplings between the driving fields and the qubits shift their resonance frequencies, which makes all the driving pulses off-resonant. Both of these phenomena deteriorate the fidelity of the qubit operations and must be corrected for to achieve high fidelity ($> 99.9\%$).

Several solutions already exist to tackle these issues in the case of a single qubit. For instance, schemes for tracking the shifting resonance frequency of a qubit throughout the duration of a driving pulse have been proposed (Steffen et al., 2000). In addition, techniques such as pulse shaping (*i.e.* modulation of the pulse amplitude or phase to modify its spectral properties) have long been used in NMR to fine-tune the frequency selectivity of pulses manipulating electron spins (Freeman, 1998, Kessler et al., 1991, Vandersypen and Chuang, 2005). Pulse engineering methods have also recently been used to optimise the performance of driving pulses and increase coherence times of a single qubit (Yang et al., 2019). This work investigates the applicability and performance of those techniques for simultaneous single-qubit gates in multi-qubit environments. The objective is to identify methods and driving schemes to implement such gates with high fidelity.

1.1. Research questions

As highlighted above, the objective of this work is twofold. Firstly, it is to study the physical as well as experiment-related phenomena that hinder the implementation of high fidelity single-qubit gates in multi-qubit systems. Secondly, it is to investigate and implement techniques to mitigate these perturbations and allow high fidelity driving. The main research question has thus been formulated as follows:

- **How can we design driving strategies to perform high fidelity single-qubit rotations on multiple qubits simultaneously?**

This main question has been subdivided into several smaller ones to structure the research effort:

- Which phenomena prevent the implementation of high fidelity single-qubit rotations in multi-qubit systems, and how do they affect gate fidelity?
- Which techniques can be applied to mitigate the impact of these phenomena?
- How can the aforementioned techniques be fine-tuned to provide the highest possible fidelity?

1.2. Report outline

Chapter 2 describes the modeling of the driving pulses and their interactions with a spin qubit. The model is then extended to multi-qubit systems. The phenomena that are expected to disturb the proper driving of the qubits are extensively discussed. Chapter 3 introduces the methods used to correct for these undesired effects, namely optimisation of the drive frequencies to cancel the shifts and pulse shaping. It subsequently provides theoretical insights into their expected performances. A brief overview of the spectral characteristics of the studied modulation shapes is also given. The software used and developed during this project is described in Chapter 4. Chapter 5 presents and discusses the performance of the aforementioned improvement techniques on simultaneous rotations in two- and three-qubits arrays. It then investigates the correlations between characteristics of the modulation shapes and the rotations fidelities. Afterwards, the performance of the fidelity improvement techniques in a five-qubits array are briefly explored. Detailed conclusions and recommendations for future work are finally given in Chapters 6 and 7, respectively.

To complete the report, additional information is provided in the appendices of this document. Appendix A complements the brief overview of the modulation shapes given in the body of the report by giving a detailed definition of all of them, as well as some discussion on their respective spectral properties. Appendix B describes the main validation and verification steps that were taken during this thesis work. Appendices C and D complete Chapter 5 by regrouping all the results of the fidelity improvement schemes. Finally, the documentation of the software developed during this project is compiled in Appendix F for future reference.

2

Theory and models

In this chapter, the general configuration of the considered system is described, as well as the analytical model that has been used to represent the driven qubit. An analytical expression of the Hamiltonian of such a system is given, and it is extended for the case of multiple qubits driven simultaneously. The ideal time evolution of the system is presented in each case. Lastly, the metrics used to quantify the performance of a driving configuration are defined.

2.1. System of interest

In this section, the qubit and driving configuration considered in this work are presented, and realistic numerical values for the simulation parameters are given.

2.1.1. Qubits layout and driving configuration

In all of the following, the qubits are considered to be arranged in a linear array, *i.e.* the system is composed of several qubits lined up next to each others in a one-dimensional structure. This setup would be the most straightforward experimental realisation of such a device. It also reduces the complexity of the model by reducing the number of exchange and coupling interactions to take into account, as explained in Section 2.3.

The study was conducted using the following general driving configuration: each qubit is driven by a microwave pulse at its resonance frequency, and all of the qubits are driven simultaneously. The pulse frequency is thus determined by that of the spins, but its amplitude and phase are left as free variables. It has furthermore been assumed that each qubit is affected by all the pulses, not only the one targeted at its resonance frequency (see Section 3.2.2 for details). In a n-qubits array, each qubit is then driven by n pulses (one of which being resonant and the others off-resonant). The setup is summarised in Figure 2.1, where the full arrows represent the resonant driving pulses and the dashed ones represent the off-resonant drives.

2.1.2. Numerical values

The qubits are placed in a longitudinal magnetic field gradient, which in experiments is induced by micromagnets patterned above them (*e.g.* in Tokura et al. (2006), Pioro-Ladriere et al. (2008)). Because of this gradient, each qubit has a different resonance frequency. In the present work, the magnetic field gradient as been assumed to be such that the frequency separation between two adjacent qubits is $\Delta\omega_0 = 100$ MHz. This value is similar to what is done in current experiments with micromagnets. Although a larger difference may appear desirable (because it would make frequency selectivity of the qubits easier), in practice a balance needs to be achieved between the decoherence gradient and the addressability. Indeed, a large $\Delta\omega_0$ means a large magnetic field gradient and thus even small noise-induced movements of the qubits would modify their resonance frequencies by a significant amount. The decoherence thus generated would lead to imperfect qubit rotations, which in turn would have a negative effect on the fidelity and must then be avoided. It is worth noting that the exact resonance frequency of the qubit remains unpredictable in a smaller gradient, but the variations can be assumed to be lower.

Furthermore, the magnetic field gradient is assumed to be such that the resonance frequencies of the qubits are of the order of 10 GHz, similarly to what can be found in literature (Watson et al. (2018), Yoneda

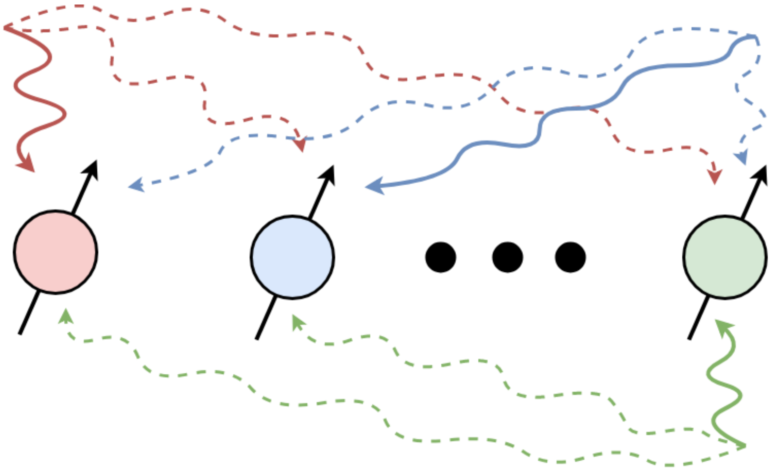

Figure 2.1: Setup of the simulations: all qubits are aligned in a one-dimensional array, simultaneously targeted by resonant driving pulses and are affected by all the pulses at the same time. The full lines represent the pulse resonant with the qubit of the corresponding color, and the dashed lines represent the off-resonant ones.

et al. (2018), Zajac et al. (2018), Kawakami et al. (2014)).

Finally, the amplitudes of the driving pulses are considered to typically be in the range [5.0, 25.0] MHz. It is a fairly common range of values (see for instance Watson et al. (2018), Yoneda et al. (2018), Zajac et al. (2018), Kawakami et al. (2014)), for which parasitic effects can be expected to have an impact small enough to not be too detrimental (see Sections 2.2.1 and 2.3.2 for a detailed description of these parasitic effects). Higher and lower amplitudes will also episodically be considered, in order to study the implications of a very fast (respectively slow) drive on the behaviour of the qubits and on the accuracy of their operations. Fast driving capabilities are especially important when working in noisy environments, where the faster the qubit manipulations are, the less damaged they are by decoherence phenomena.

2.2. Case of a single qubit

2.2.1. Model

The spin qubit is modeled as a two-level system (referred to as TLS in the following), whose ground and excited states are labeled $|g\rangle$ and $|e\rangle$ respectively, and whose resonance frequency is denoted by ω_0. In the remainder of this report, the ground and excited states will occasionally be referred to as spin-up and spin-down respectively.

The derivations of the Hamiltonians presented in this Section are given in Steck (2007) and Shahriar et al. (2004).

In the case of a single qubit, the state of the TLS can be represented by an element of $\mathbb{C}^2$ denoted $|\psi\rangle = \begin{pmatrix} c_g \\ c_e \end{pmatrix}$, where c_g and c_e are the ground and excited states probability amplitudes. Another possible representation of the state is the so-called density matrix, defined by $\rho = |\psi\rangle\langle\psi|$, if $|\psi\rangle$ is a pure state. This form is more commonly used in quantum information theory, and is a more general representation of the state. It is indeed able to represent incoherent superpositions of states, and hence can be used to describe open quantum systems.
The time evolution of such a quantum system is governed by the Schrödinger equation (Eq 2.1).

$$i\hbar\frac{\partial}{\partial t}|\psi\rangle = H|\psi\rangle, \tag{2.1}$$

where H is the Hamiltonian of the system and $\hbar$ is the reduced Planck constant.

This equation transforms into the Schrödinger-von Neumann equation (Eq 2.2) when considering density matrices. $[.,.]$ is the commutator of two operators acting upon the tensor product $\left(\mathbb{C}^2\right)^{\otimes N}$, N being the number of spins.

$$\frac{\partial}{\partial t}\rho = -\frac{i}{\hbar}\left[H,\rho\right], \tag{2.2}$$

The driving pulse is modelled as a monochromatic electromagnetic field with amplitude Ω, angular frequency ω and initial phase ϕ. For the moment, the amplitude is considered constant (corresponding to a so-called block or rectangular pulse), but the consequences of it being time-dependent will be studied in Section 3. In the following, the detuning between the resonant frequency of the qubit and the pulse frequency will be noted $\Delta = \omega - \omega_0$.

In this work, we will consider the so-called dressed atom model, which consists in the TLS representing the qubit coupled to the driving pulse. As described in Steck (2007), Shahriar et al. (2004) and Henley (2008), the Hamiltonian of such a system can be expressed as the sum of two parts: one representing the free evolution of the spin and the other due to its interaction with the electromagnetic field of the pulse (Eq 2.3).

$$H_{spin} = H_{free} + H_{drive} \tag{2.3}$$

The free evolution Hamiltonian H_{free} simply corresponds to the spin precessing around the $\hat{z}$ axis at the Larmor frequency ω_0. If the ground state energy is chosen to be 0, it can be expressed as follows:

$$H_{free} = \hbar\omega_0 \left|e\right\rangle\left\langle e\right|. \tag{2.4}$$

We express the interaction Hamiltonian H_{drive} using the dipole approximation, which allows to neglect any variation in the field over the size of the electron. This is equivalent to considering that the wavelength of the pulse is much larger than the size of the electron, which in the case of microwave is a perfectly fair assumption. The expression is then:

$$H_{drive} = \frac{\hbar\Omega}{2}\left[\sigma^x e^{i\omega t} + \sigma^x e^{-i\omega t}\right], \tag{2.5}$$

where σ^x is the Pauli X matrix and ω is the pulse frequency. Ω is the so-called Rabi frequency of the pulse, which characterises the strength of its coupling with the spin and is directly related to its amplitude.

In addition to the rotation around the $\hat{x}$- or $\hat{y}$-axis that we aim for, the energy splitting caused by placing the qubit in a static magnetic field generates a precession of the spin around the $\hat{z}$-axis at frequency ω. It is then convenient in the simulation to eliminate this precession by switching to a co-rotating reference frame. The Hamiltonian is transformed using Equation 2.6 (see Steck (2007) and Zeuch et al. (2018) for the details).

$$\tilde{H} = UHU^\dagger + i\hbar\frac{\partial U}{\partial t}U^\dagger, \tag{2.6}$$

where H is the Hamiltonian in the initial reference frame, $\tilde{H}$ is the one expressed in the rotating reference frame and U is the unitary transformation representing the conversion from the non-rotating laboratory frame to the rotating one. It is defined in equation 2.7.

$$U = e^{i\omega\left|e\right\rangle\left\langle e\right|} \tag{2.7}$$

The Hamiltonian in the rotating frame of the pulse is then given by Equation 2.8.

$$\tilde{H} = \frac{\hbar\Omega}{2}\left[e^{i2(\omega t+\phi)}+1\right]\left|g\right\rangle\left\langle e\right| + \frac{\hbar\Omega}{2}\left[e^{-i2(\omega t+\phi)}+1\right]\left|e\right\rangle\left\langle g\right| - \hbar\Delta\left|e\right\rangle\left\langle e\right| \tag{2.8}$$

This transformation introduces terms which are oscillating at twice the pulse frequency. It is at this point relevant to investigate the applicability of the so-called rotating-wave approximation (that will be commonly referred to as RWA in the following). The RWA focuses on slow dynamics by replacing the rapidly oscillating terms by their zero average value. This is equivalent to neglecting the component of the drive that rotates in the opposite direction compared to the spin. Making such an approximation would greatly simplify the expression of the Hamiltonian by removing any explicit time dependency, making the Schrödinger equation much easier to solve.

Three conditions need to be verified for the RWA to be applicable (Zeuch et al., 2018):

- The drive is weak: $\Omega << \omega$
- The drive is near resonant: $\omega \approx \omega_0$
- The drive varies slowly on the time-scale of the inverse qubit frequency $\frac{1}{\omega}$

According to Zeuch et al. (2018), the RWA is already not applicable anymore in the $\Omega/\omega \geqslant 0.01$ regime. Taking into account the values considered in Section 2.1.2, we have $\Omega/\omega \in [0.0005, 0.0025]$), which means that we fulfill the first condition of applicability of the RWA. The third requirement is also easily met. However, the second condition is more problematic. Indeed, in our working configuration, where multiple qubits are driven at the same time, each qubit will receive off-resonant pulses in addition to its resonant drive. Consequently, the RWA does not seem to be applicable to our problem.

To validate this reasoning, a numerical simulation has also been performed: the time evolution of a one-qubit system submitted to one resonant and several off-resonant driving pulses has been computed both when using and when discarding the RWA. This test represents the most common use case of the considered setup. Figure 2.2 shows part of the evolution of the ground and excited states populations during the pulse.

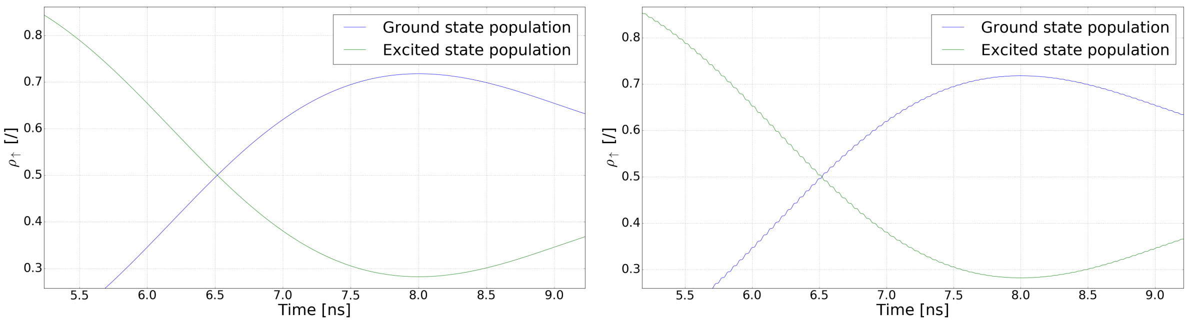

(a) Using the RWA

(b) Without the RWA

Figure 2.2: Evolution of a qubit submitted to several driving pulses of amplitude $\Omega = 30.0$ MHz. One is resonant and the others are detuned.

We see when comparing the two panels that the evolution computed outside of the RWA displays fast oscillations that are not present in the RWA evolution (they are investigated in more details in Section 2.2.2). Similar features, albeit less clearly visible, can also be observed when reconducting similar simulations with lower pulse amplitudes. Although these fast oscillations are admittedly of small amplitude and could potentially be neglected, it has been decided to not make the rotating-wave approximation, as making it still introduces sensible modification in the system's dynamics. Consequently, only the non-RWA expression of the Hamiltonian, given in Equation 2.8, will be used in the following.

2.2.2. Time evolution

Injecting Hamiltonian 2.8 in Schrödinger's equation, we can solve for the ground and excited states populations. This is done using the method of adiabatic elimination to handle the explicit time dependency introduced by not making the RWA (see Shahriar et al. (2004) for the details). This method is valid to first order in $\eta = \dfrac{\Omega}{4\omega}$, which is a reasonable approximation to make in our range of values (*cf.* Section 2.1.2). Assuming a resonant pulse, as well as an initial state where the qubit is in the ground state $|g\rangle$, it yields the following expressions:

$$\begin{aligned} P_{|e\rangle}(t) &= \sin^2\left(\frac{\Omega t}{2}\right) + \eta \sin(\Omega t)\sin\left(2\phi_t\right) \\ P_{|g\rangle}(t) &= 1 - P_{|e\rangle}(t), \end{aligned} \tag{2.9}$$

where $\phi_t = \omega t + \phi$ is the phase of the driving field at time t. We thus see two components in the time evolution of the populations. The first one makes the TLS oscillate from the ground to the excited state at the angular frequency Ω, and corresponds to the traditional Rabi oscillations that one would observe when making the RWA. This component is the one responsible for the driving of the qubit.

The second component produces faster phase-dependent oscillations that superimpose on the Rabi oscillations. These so-called Bloch-Siegert oscillations can be observed on Figure 2.2b, and are due to the interference between the excitations caused by the co- and counter-rotating driving fields. This effect directly comes from discarding the RWA and may constitute an important error source in precise driving of the qubits.

Finally, the time evolution of the system is modified when there is a detuning between the drive and the qubit's resonance frequency. Figure 2.3 illustrates how the excited state population of a single qubit varies when driven by a pulse detuned by from 0 to 2.5% ($\Delta = \omega - \omega_0$ ranging from 0 to 250.00 MHz) compared to the qubit's resonance frequency. The exact expression of the populations is not of primary importance here, but two features are important to keep in mind.

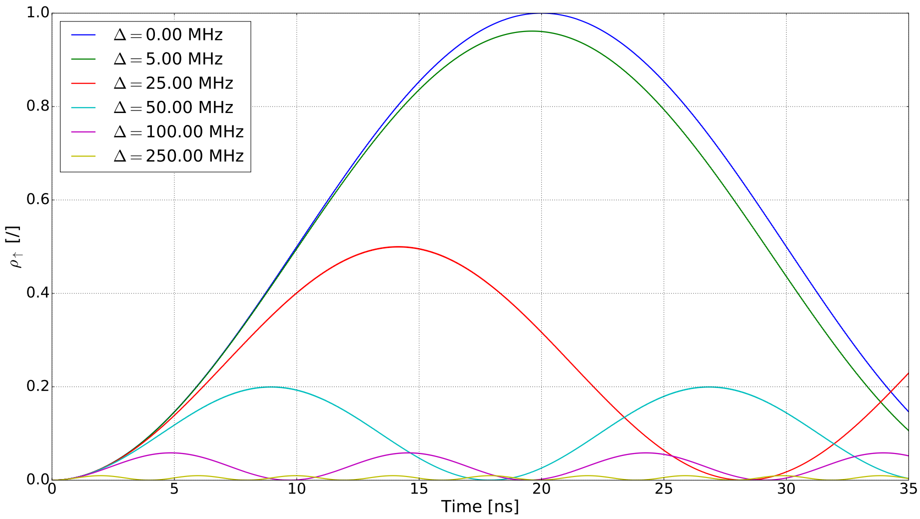

Figure 2.3: Excited state population of a qubit driven by an off-resonant pulse, modeled without making the RWA. The detuning Δ ranges from 0 to 250 MHz.

Similarly to the RWA case, the population oscillates at the generalised Rabi frequency $\tilde{\Omega} = \sqrt{\Omega^2 + \Delta^2}$ and with a lower amplitude $\frac{\Omega^2}{\tilde{\Omega}^2}$, preventing complete population inversion. In addition to these oscillations, the smaller and faster Bloch-Siegert oscillations are still present, albeit at a slightly offset frequency. Figure 2.3 clearly highlights the need for the driving pulse to be exactly on resonance for the qubit operation to be accurate: a detuning as small as 0.25% of the resonance frequency already leads to roughly half of the population not being inversed.

2.2.3. Resonance frequency shifts

In addition to driving the spin rotations, the microwave pulses couple to the system, effectively dressing it. This dressing tilts the qubit's precession axis by an angle θ and modifies the energy of the qubit's eigenstates, thus modifying its resonance frequency. Figure 2.4 illustrates the evolution of the spin's precession axis between its uncoupled and dressed configurations. This shift in resonance frequency is a major error source as it makes the otherwise well targeted drives off-resonant, hence preventing complete Rabi oscillations.

When considering the driving due to both the co-rotating and counter-rotating components of the microwave field (as it is done when discarding the RWA), two different phenomena are to be considered:

- the AC Stark shift, resulting from the interaction between the qubit and the co-rotating component of the driving pulse (*i.e.* the component that would be taken into account in the RWA, as discussed in Section 2.2.1),
- the Bloch-Siegert shift, that is caused by the coupling of the qubit with the counter-rotating component of the driving pulse (which would have been discarded when making the RWA).

The AC Stark shift is extensively described in Bakos (1977), Steck (2007) and Steffen et al. (2000). In particular, Steffen et al. (2000) give that, away from resonance, the shift can be approximated by the simple expression given in Eq. 2.10.

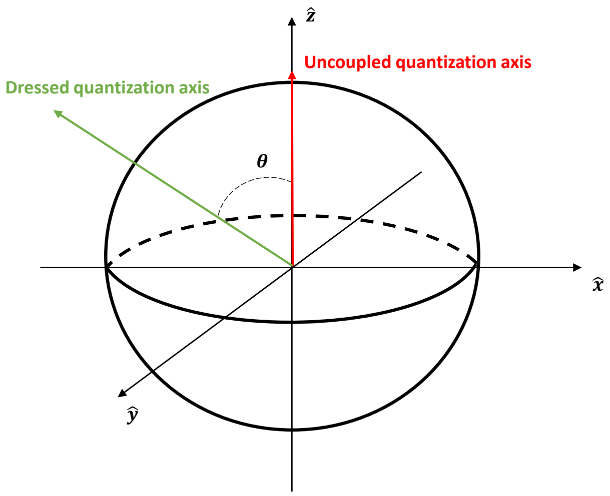

Figure 2.4: Evolution of the spin's quantization axis between its bare (uncoupled) and dressed configurations.

$$\Delta\omega_S = \frac{\Omega^2}{2\,(\omega_0 - \omega)} \tag{2.10}$$

However, this study requires an analytical expression valid very close to and on resonance. The complete expression thus needs to be derived.
In the rotating reference frame, the detuning of the dressed qubit from the electromagnetic field is:

$$\omega_{0,\text{shifted}} - \omega = (\omega_0 + \Delta\omega_S) - \omega = \begin{cases} \tilde{\Omega} & \text{, if } \Delta = \omega - \omega_0 < 0 \\ -\tilde{\Omega} & \text{, if } \Delta = \omega - \omega_0 > 0 \\ 0 & \text{, if } \Delta = \omega - \omega_0 = 0 \end{cases} \tag{2.11}$$

where ω_0 is the qubit's unshifted resonance frequency, ω is the drive frequency and $\Delta\omega_S$ is the Stark shift. Equation 2.11 translates into:

$$\Delta\omega_S = \begin{cases} \tilde{\Omega} + \Delta & \text{, if } \Delta = \omega - \omega_0 < 0 \\ -\tilde{\Omega} + \Delta & \text{, if } \Delta = \omega - \omega_0 > 0 \\ 0 & \text{, if } \Delta = \omega - \omega_0 = 0 \end{cases} \tag{2.12}$$

Equation 2.12 gives the value of the energy splitting along the dressed precession axis of the TLS, which is tilted compared to the uncoupled one. If used as-is, it gives an overestimation of the value of the shift. To recover the exact value, the equation is projected onto the uncoupled quantization axis as follows:

$$\Delta\omega_S = \Delta\omega_{S,\text{dressed}} \times \cos\theta \tag{2.13}$$

where $\Delta\omega_{S,\text{dressed}}$ is the energy splitting along the dressed quantization axis (Eq 2.12), and θ is the Stückelberg angle, defined in Steck (2007) as the angle between the dressed and uncoupled quantization axes (see Figure 2.4) by

$$\cos\theta = \frac{|\Delta|}{\sqrt{\Omega^2 + \Delta^2}}. \tag{2.14}$$

From here we have the final expression for the AC Stark shift of the resonance frequency:

$$\Delta\omega_S = \begin{cases} \dfrac{|\Delta|}{\sqrt{\Omega^2 + \Delta^2}} \times \left(\tilde{\Omega} + \Delta\right) & \text{, if } \Delta = \omega - \omega_0 < 0 \\ \dfrac{|\Delta|}{\sqrt{\Omega^2 + \Delta^2}} \times \left(-\tilde{\Omega} + \Delta\right) & \text{, if } \Delta = \omega - \omega_0 > 0 \\ 0 & \text{, if } \Delta = \omega - \omega_0 = 0. \end{cases} \tag{2.15}$$

Note that this shift pushes the resonance frequency “away" from the driving frequency: the shifted resonance frequency will be higher if $\omega < \omega_0$ and *vice-versa*. This is especially clear in Equation 2.10, but remains valid close to resonance in Equation 2.15. Furthermore, the value of the shift rapidly tends to zero when getting close to $\Delta = 0$, which is expected as the AC Stark shift is an off-resonant phenomenon.

Figure 2.5 presents the evolution of the AC Stark shift induced by a pulse on a qubit of resonance frequency $\omega_0 = 10.0$ GHz, for different representative amplitude values. In complement, Table 2.1 gives some shift values, expressed in megahertz, at specific detunings for various pulse amplitudes.

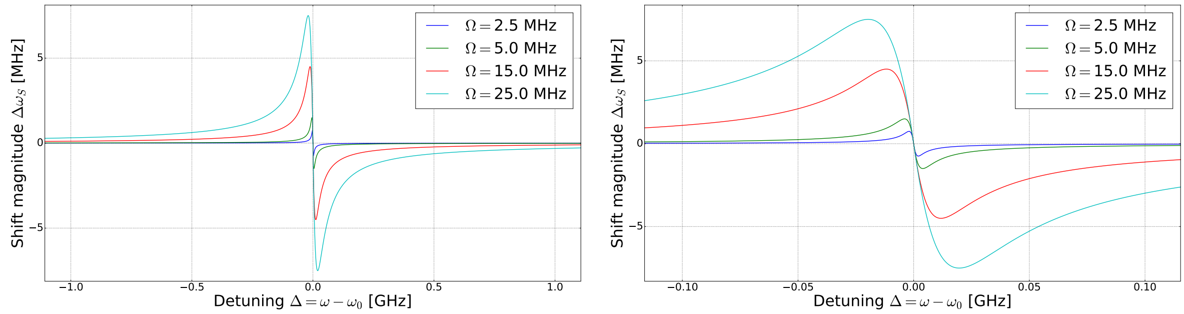

(a) Stark shift over a large detuning range

(b) Zoom on the resonance region

Figure 2.5: AC Stark shift induced by a single driving pulse as a function of its detuning from the spin's resonance frequency, for various pulse amplitudes. The resonance frequency of the qubit has been fixed to $\omega_0 = 10.0$ GHz.

First of all, Figure 2.5a shows that, regardless of the pulse amplitude, the Stark shift vanishes for large detunings. Secondly, as expected, higher pulse amplitudes clearly result in higher shift values: the stronger you are driving, the stronger the coupling between the spin and the pulse will be. Additionally, near resonance all the considered pulse amplitudes lead to similarly low shift values, until vanishing at resonance.

	Δ = 0.01 MHz	**Δ = 2.5 MHz**	**Δ = 5 MHz**	**Δ = 15 MHz**	**Δ = 25 MHz**	**Δ = 50 MHz**	**Δ = 100 MHz**
Ω = 2.5 MHz	-0.00996	-0.73223	-0.52786	-0.20409	-0.12407	-0.06238	-0.03124
Ω = 5.0 MHz	-0.00998	-1.38197	-1.46447	-0.76975	-0.48548	-0.24814	-0.12477
Ω = 15 MHz	-0.00999	-2.08900	-3.41886	-4.39340	-3.56268	-2.10869	-1.10636
Ω = 30 MHz	-0.01000	-2.29239	-4.17801	-8.29180	-8.99539	-7.12535	-4.21737

Table 2.1: AC Stark shift induced by several typical drive configurations. As the Stark shift is symmetric with respect to resonance, only positive detunings have been considered. The only difference that would be induced by considering negative detunings is that the shift values would be positive. The shifts are expressed in MHz.

Looking at the largest values in Table 2.1 and comparing them with the off-resonant driving illustrated in Figure 2.3 gives a preliminary indication of the extent to which the AC Stark shift can damage the performance of the qubit operations.

Similarly to the AC Stark shift, the Bloch-Siegert shift is caused by the coupling of the spin to the drive. However, in this phenomenon, the spin couples with the counter-rotating component of the pulse instead of the co-rotating one. Shirley (1965) derives an expression for this shift up to 6^{th} order in Ω, valid in the weak driving regime ($\Omega/\omega < 2$), using perturbation theory. This expression is confirmed to be accurate using a purely quantum mechanical treatment in Cohen-Tannoudji et al. (1973). Such an approximated expression is sufficient for the present work, since all the envisioned driving configurations clearly stay in this weak driving regime. The expression for the shift is given in Equation 2.16.

$$\Delta\omega_{BS} = \frac{(\Omega/4)^2}{\omega_0} + \frac{(\Omega/4)^4}{4\omega_0^3} - \frac{35\,(\Omega/4)^6}{32\omega_0^5} \tag{2.16}$$

where Ω is the Rabi frequency and ω_0 is the resonance frequency of the qubit.

Contrarily to the AC Stark shift, the Bloch-Siegert shift only depends on the amplitude of the pulse and the qubit frequency, and not directly on the detuning between the drive and the qubit. Furthermore, it always leads to an increase of the qubit's resonance frequency. This leads to asymmetric values of total shift between the $\Delta > 0$ and $\Delta < 0$ regimes.

Bloch-Siegert shift values corresponding to typical parameters of our system are summarised in Table 2.2 for comparison purposes.

Rabi frequency Ω	[MHz]	2.5	5.0	15.0	30.0
Bloch-Siegert shift $\Delta\omega_{BS}$	[kHz]	0.039	0.156	1.406	5.625

Table 2.2: Bloch-Siegert shift induced by typical driving pulse amplitudes. The shift values are given in kilohertz.

In our driving regime, it appears that the amplitude of the Bloch-Siegert shift is on average lower than that of the AC Stark shift by several orders of magnitude, which would suggest that the Bloch-Siegert shift may be neglected. However, that is not possible due to the fact that the Stark shift sharply drops to zero when the detuning is low enough (see Figure 2.5). Consequently, the two shifts are then of comparable magnitudes when close to resonance. Since our analysis requires a precise evaluation of the qubits' resonance frequencies very close to resonance, neglecting the Bloch-Siegert is ruled out.

Seeing the wide range of values the shifts can take, it is insightful to examine from what point a shift in resonance frequency actually become relevant. To asses this, the performance of a driving pulse slightly detuned from its target qubit have been computed, for various detunings, both in the case of a π- and a $\pi/2$-rotation (see Section 2.4 for details on the performance measurements). The outcome of this test is presented in Figure 2.6.

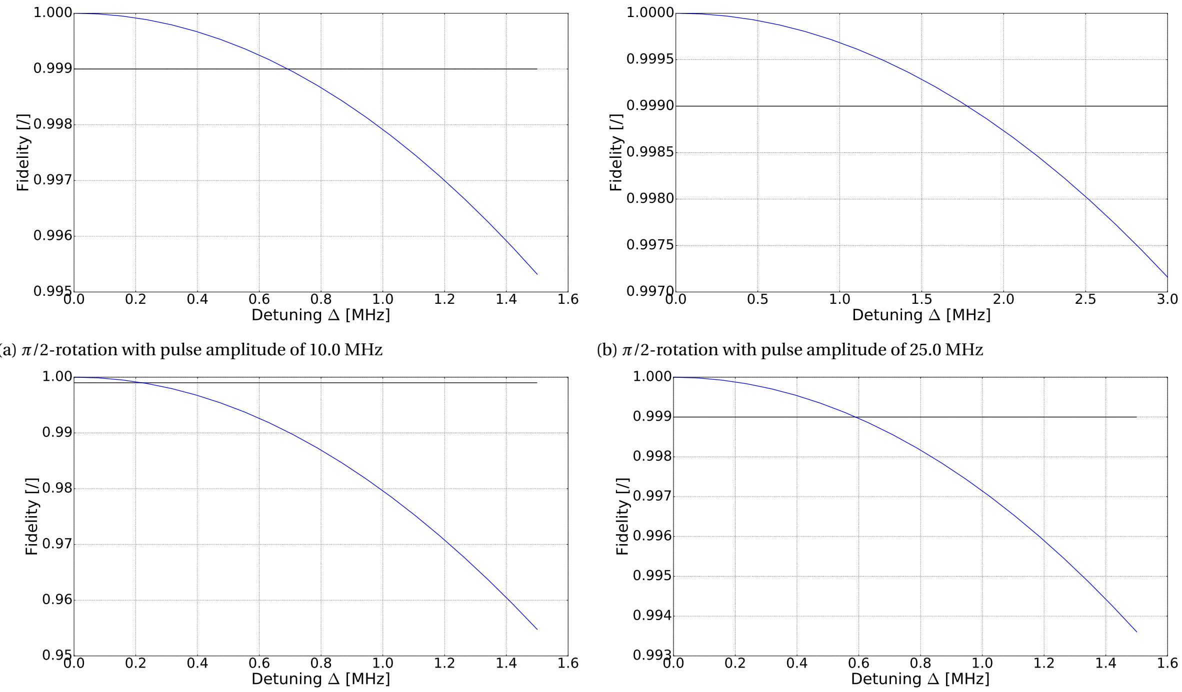

(a) $\pi/2$-rotation with pulse amplitude of 10.0 MHz

(b) $\pi/2$-rotation with pulse amplitude of 25.0 MHz

(c) π-rotation with pulse amplitude of 10.0 MHz

(d) π-rotation with pulse amplitude of 25.0 MHz

Figure 2.6: Evaluation of the fidelity of a driving pulse with regards to its detuning from the qubit's resonance frequency

It clearly shows that, depending on the driving configuration, a frequency shift as small as 300 kHz can bring the fidelity lower than 99.9%. The actual threshold depends on the precise driving configuration (target rotation angle, duration of the pulse, *etc.*), but from these tests it appears that a shift of a couple of megahertz is usually enough to lower the fidelity below 99.9%. Shifts of this order of magnitude can clearly be produced by the phenomena considered in this Section (*c.f.* Table 2.1 and 2.2). This further proves the negative effect the frequency shifts can have on the fidelity of the qubit operations, and highlights the need to correct for it in the driving scheme.

2.3. Extension to multi-qubits systems

2.3.1. Model

Now that an expression for the Hamiltonian of a single qubit has been derived and that the time evolution of the driven system has been described, the model can easily be extended to systems of multiple qubits. The multi-qubit Hamiltonian is presented in Equation 2.17.

$$H_{\text{multi}} = \sum_{i=1}^{n} I^{i-1} \otimes H \otimes I^{n-i}, \tag{2.17}$$

where I^j is the tensor product of the identity operator in dimension 2 with itself j times and H is the single-qubit Hamiltonian defined in Equation 2.8.

The exchange interaction between the adjacent spins must also be taken into account. It has been modeled using the Heisenberg exchange Hamiltonian, given in Eq. 2.18 between two spins labeled i and j (Henley, 2008).

$$\begin{aligned} H_{ex,i,j} &= J \times \vec{S}_i \cdot \vec{S}_j \\ &= -\frac{1}{4} \times \sum_{i<j} J_x \sigma_i^x \sigma_j^x + J_y \sigma_i^y \sigma_j^y + J_z \sigma_i^z \sigma_j^z, \end{aligned} \tag{2.18}$$

where J is the exchange rate between spins i and j, $\vec{S}_i$ and $\vec{S}_j$ are i^{th} and j^{th} spin momenta respectively, and σ_i^x, σ_i^y and σ_i^z are the X, Y and Z Pauli matrices of the i^{th} qubit.

This expression can be simplified by assuming that the exchange rate is the same in each direction: $J_x = J_y = J_z = J$ (Meunier et al., 2011, Watson et al., 2018). The exchange Hamiltonian then becomes (Henley, 2008):

$$H_{ex,i,j} = -\frac{1}{4} J \times \sum_{i<j} \sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z, \tag{2.19}$$

with the same notations as Eq 2.18.

As discussed in Section 2.1.1, only one-dimensional qubit arrays have been considered in this work. Consequently, the exchange interaction has only been modeled between adjacent spins. Furthermore, coupling interactions between more than two spins have been shown to be very small and rarely taken into account in spin Hamiltonians (Henley, 2008). They have thus been neglected.

The exchange Hamiltonian for a n-qubits system is then:

$$H_{EX} = \sum_{i=0}^{n-2} I^i \otimes H_{ex,i,i+1} \otimes I^{n-i-2}, \tag{2.20}$$

where again I^j denotes the tensor product of the identity operator in dimension 2 with itself j times, and $H_{ex,i,j}$ is the two-qubits exchange Hamiltonian defined in equation 2.19.

Finally, the full Hamiltonian of a n-qubits system is given in equation 2.21 below. The unitary evolution of the system is then computed by solving Von Neumann equation (Eq 2.2).

$$H = H_{\text{multi}} + H_{EX} \tag{2.21}$$

2.3.2. Unwanted off-resonant interactions

In the physical implementation of multi-qubit quantum chips, a driving pulse applied to drive a specific qubit is also seen by the other qubits, with varying intensity. A gate applied to one qubit thus results in off-resonant driving for all the other qubits, with an strength depending on the magnitude of the detuning between the resonance frequency and that of the drive. As discussed in Section 2.2.2, such an off-resonant pulse imperfectly drives the spin rotation (Figure 2.3). These additional unaccounted-for rotations negatively affect the achievable precision of the qubit operations, and must hence be minimised or corrected for.

Furthermore, these additional pulses also couple to each of the qubits, contributing to the resonance frequency shifts described in Section 2.2.3. Table 2.3 illustrates the impact of coupling with multiple pulses with different detunings on a qubit's resonance frequency.

Configuration number	1	2	3	4	5	6	7
Set of pulses	$\Delta = 0.0$	$\Delta = 0.0$	$\Delta = 0.0$	$\Delta = 0.0$	$\Delta = 0.0$	$\Delta = 0.0$	$\Delta = 0.0$
		$\Delta = 0.1$	$\Delta = 0.1$	$\Delta = 0.1$	$\Delta = 0.1$	$\Delta = 0.1$	$\Delta = 0.1$
				$\Delta = -0.1$	$\Delta = -0.1$	$\Delta = -0.1$	$\Delta = -0.1$
			$\Delta = 0.2$		$\Delta = 0.2$	$\Delta = 0.2$	$\Delta = 0.2$
						$\Delta = -0.2$	$\Delta = -0.2$
							$\Delta = 0.3$
Frequency shift [MHz] $\boldsymbol{\Delta\omega = \Delta\omega_S + \Delta\omega_{BS}}$	**0.00391**	**-2.97794**	**-4.51846**	**0.01172**	**-1.52880**	**0.01953**	**-1.01284**

Table 2.3: Resonance frequency shift of a qubit (initially at $\omega_0 = 10.0$ GHz) submitted to different sets of resonant and off-resonant pulses. Each line corresponds to an additional pulse of similar amplitude ($\Omega = 25.0$ MHz in this example) with a different detuning from the qubit. The pulses are assumed to have constant amplitude over space (see Section 3.2.2 for discussions on this point).

This table highlights the existence of large differences between what could be called symmetric configurations (where the number of positively and negatively detuned driving pulses are the same, with similar detuning values on each side) and asymmetric ones (where the number of either positively or negatively detuned pulses is superior to the other, or where one type of pulse has smaller detunings that the other). Indeed, in symmetric configurations, the Stark shifts caused by the positively and negatively detuned pulses compensate each other, leaving only the much smaller Bloch-Siegert shift. On the other hand, in asymmetric pulse sets, the Stark shifts accumulate and lead to very high shift values. The more asymmetric the pulse set is, the more accentuated this effect will be, as can be seen by comparing configuration sets 2 and 3 in Table 2.3, between which the asymmetry is increased. The fact that the presented asymmetric configurations (pulse sets 2, 3, 5 and 7) induce negative shifts is due to the choice of the detuning of the pulses. Indeed, the "odd" pulses have each been chosen with a positive detuning compared to the qubit, which induces a negative Stark shift. Should the detuning of the "odd" pulses have been taken negative, very similar positive shifts would have been found (although slightly higher due to the addition of the always positive Bloch-Siegert shift).

Additionally, as expected and as illustrated by pulse sets 2, 5 and 7, the more detuned the "odd" pulse in an asymmetric set is, the lower the total frequency shift is. This is due to the Stark shift reducing to zero for large detunings.

Lastly, as illustrated by pulse sets 1, 4 and 6, the frequency shift in symmetric pulse configurations increases with the number of pulses. This stems from the Bloch-Siegert shift (the only one not compensated in these configurations) always being positive, and so increasing with each new pulse acting on the system.

In our case of a linear array of simultaneously driven qubits, this results in the following. Assuming the qubits are arranged by order of increasing resonance frequency (most straightforward implementation in a magnetic field gradient), the qubits at both ends of the array experience the largest frequency shifts, while they are the lowest for the centermost spins. Similarly to the unwanted driving, these accumulated resonance frequency shifts are expected to severely impact the achievable performance of our qubit operations by making the driving pulses off-resonant with their targeted spins.

2.4. Rotation performance metrics

Now that the object of this study has been defined, as well as the phenomena that could influence the precision of the operations performed by the qubits, the metrics quantifying this precision need to be properly discussed. The two metrics that are considered, the state and the average gate fidelity, are described below.

Although the state fidelity is not a metric in the mathematical sense (since $F(\rho_1, \rho_2) = 0$ does not mean that $\rho_1 = \rho_2$), it does have the intuitive interpretation of measuring how close two quantum states are from each other. It takes values between 0 and 1, where 1 is obtained for identical density matrices and 0 corresponds to states that are orthogonal in their Hilbert space (*i.e.* as different from each other as possible, mathematically speaking). It is defined in Nielsen and Chuang (2011) as:

$$F(\rho,\sigma) = \mathrm{Tr}\left(\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\right), \tag{2.22}$$

where ρ is the density matrix of the quantum state produced by the driving and σ is the density matrix of the state an ideal gate would take the system to.

State fidelity is a good metric to evaluate the performance of a specific driving configuration, but is has the important drawback of being initial state dependent. This means that in order to obtain results applicable to so-called universal rotations, it would be necessary to average the state fidelity over a very large, ideally infinite, number of initial states, which is not desirable.

In order to circumvent this issue, the average gate fidelity is used (Pedersen et al., 2007). It operates on the unitary time-evolution operators U that evolve a system from its initial to final state: $|\psi(t)\rangle = U(t,t_0)|\psi(t_0)\rangle$, and is consequently independent of the initial state of the evolution. It measures the closeness between the unitary operator U representing the driving and the one that corresponds to an ideal gate U_0. Pedersen et al. (2007) defines it as follows (Equation 2.23).

$$\begin{aligned} F(U,U_0) &= \int_{S^{2n-1}} |\langle\psi| M |\psi\rangle|^2 dV \\ &= \frac{1}{n(n+1)}\left[\mathrm{Tr}\left(MM^\dagger\right) + |\mathrm{Tr}(M)|^2\right], \end{aligned} \tag{2.23}$$

where $M = U_0^\dagger U$, and $|\psi\rangle$ is a state vector belonging to the unit sphere S^{2n-1} in a n-dimensional Hilbert space, acting as initial state of the unitary evolution. Similarly to the state fidelity, it takes values in the $[0,1]$ range, with 1 corresponding to indistinguishable unitary matrices and 0 to orthogonal ones. It is important to note that because this metric operates on unitary matrices, it cannot be used as defined here to quantify the performance of systems including decoherence.

In situations where it is more convenient, the infidelity (defined in Equation 2.24) may also be used. Its definition is the same whether considering the state or the average gate fidelity.

$$\bar{F} = 1 - F \tag{2.24}$$

3

Fidelity improvement techniques

Chapter 2 presented the different phenomena that can impact the achievable fidelity of a rotation. This chapter discusses the methods and techniques to minimise their impact. First, a systematic correction algorithm to compensate for the shifts of the qubits's resonance frequencies is presented. Then, the pulse shaping technique, which aims at reducing the crosstalk caused by off-resonant pulses through amplitude modulation, is introduced. Following that, the pulse shapes that were used in this study are briefly described and analysed. Their definition, as well as their complete characteristics, are given in Appendix A.
Please note that in the following, we will refer interchangeably to shaped and amplitude-modulated pulses.

3.1. Frequency shift compensation scheme

The shifts in resonance frequency caused by the coupling of the spins with the driving fields have been shown to reduce the fidelity of qubit operations. This reduction is caused by the shifts making the driving pulses off-resonant with their target. The qubit rotations they produce are then incomplete (see Figure 2.3). This section presents a correction algorithm that was devised to counter this effect.

3.1.1. Optimisation of the drive frequencies to nullify the shift

A first possibility to reduce the impact of the frequency shifts described in Section 2.2.3 is to select the frequencies of the driving pulses in such a way that the shifts they induce make the qubits' resonance frequencies match that of the initially off-resonant pulses.

Considering a n-qubit system, this is equivalent to solving the following non-linear system of equations for the set of drive frequencies $(\omega_1, \omega_2, ..., \omega_n)$:

$$\begin{cases} \omega_1 & = \omega_{0,1} + \sum_i \left[\text{Stark}\left(\omega_{0,1}, \Omega_i, \omega_i\right) + \text{BS}\left(\omega_{0,i}, \Omega_i\right)\right] \\ \omega_2 & = \omega_{0,2} + \sum_i \left[\text{Stark}\left(\omega_{0,2}, \Omega_i, \omega_i\right) + \text{BS}\left(\omega_{0,i}, \Omega_i\right)\right] \\ \vdots & \\ \omega_n & = \omega_{0,n} + \sum_i \left[\text{Stark}\left(\omega_{0,n}, \Omega_i, \omega_i\right) + \text{BS}\left(\omega_{0,i}, \Omega_i\right)\right] \end{cases}, \tag{3.1}$$

where Ω_i and ω_i are respectively the Rabi and drive frequencies of the i^{th} pulse, and where $\omega_{0,i}$ is the unshifted resonance frequency of the i^{th} qubit.

This non-linear system can easily be solved using numerical integration. It yields a set $(\omega_1, \omega_2, ..., \omega_n)$ of initially off-resonant pulse frequencies that effectively dress the qubits to make them resonant.

As an illustration, let's consider the case of a two-qubits system of resonance frequencies $\omega_{0,1} = 10.0$ GHz and $\omega_{0,2} = 10.1$ GHz. This system is driven by two pulses of amplitude $\Omega_1 = \Omega_2 = 15.0$ MHz. As discussed in Section 2.2.3 and 2.3.2, if the frequency of these pulses were chosen to match that of the qubits, the values of $\omega_{0,1}$ and $\omega_{0,2}$ would shift to 9.998896 and 10.101109 GHz respectively. This would result in off resonant driving and thus in incomplete qubit rotations. Solving the system of Equation 3.1 yields $(\omega_1, \omega_2) = (9.999441, 10.100562)$ GHz. In that configuration, the shifted resonance frequencies of the qubits now exactly match the drive frequencies. The rotation of the qubits is now expected to be complete, and the fidelity of the operation is increased.

This method is very powerful because it eliminates all the detuning that would be induced by the Stark and Bloch-Siegert shifts. However, it has the major drawback that it is an instantaneous correction. As a consequence, it can only be used when the driving pulses are of constant amplitudes and identical durations. Indeed, if at some point during the driving the amplitude varies, the pulse will then induce a different frequency shift than the one that was compensated for at the start, and all the drives will again be off-resonant. Similarly, if at least one of the considered pulses has a different duration than the others, once this pulse is turned off the dressing of the qubits will change and hence the initial correction will not be appropriate anymore. This characteristic severely limits the applicability of this method to the present study. Indeed, as will be discussed in Section 3.2, the main technique used to limit the crosstalk due to off-resonant drives is to exploit the different spectral properties of pulses with time-varying amplitudes (the so-called pulse shaping). This scheme of selecting the appropriate drive frequencies to cancel the shifts thus needs to be adapted.

3.1.2. Frequency shift correction algorithm

The method discussed in Section 3.1.1 can be adapted to pulses with time-varying amplitudes with a simple alteration, inspired by the way Steffen et al. (2000) handle such pulses. They model a shaped pulse as a sequence of short time slices during which the amplitude and phase are constant. This is in practice equivalent to considering that the qubit is driven by successive very short pulses of constant amplitude and phase. The method described in Section 3.1.1 can then be applied to each of those, *i.e.* to each time slice. No memory of the shifted resonance frequencies of the previous slices is necessary, since at the end of each slice the qubits “return" to their bare states before being dressed again by the pulse of the next slice. The previous method yields, for each element of the discretised pulse, a set of driving frequencies that will be resonant with the dressed qubits, effectively suppressing the impact of frequency shifts on the fidelity of the operations.

Furthermore, it is worth noting that this modeling of the driving pulses actually corresponds to the way they would be implemented in an experimental setup. In this study, the length of the time slice has been set to 1 nanosecond, in accordance with the capabilities of waveform generators used by the lab in which this work was done (Keysight3201A, with a sampling rate of 1 GHz).

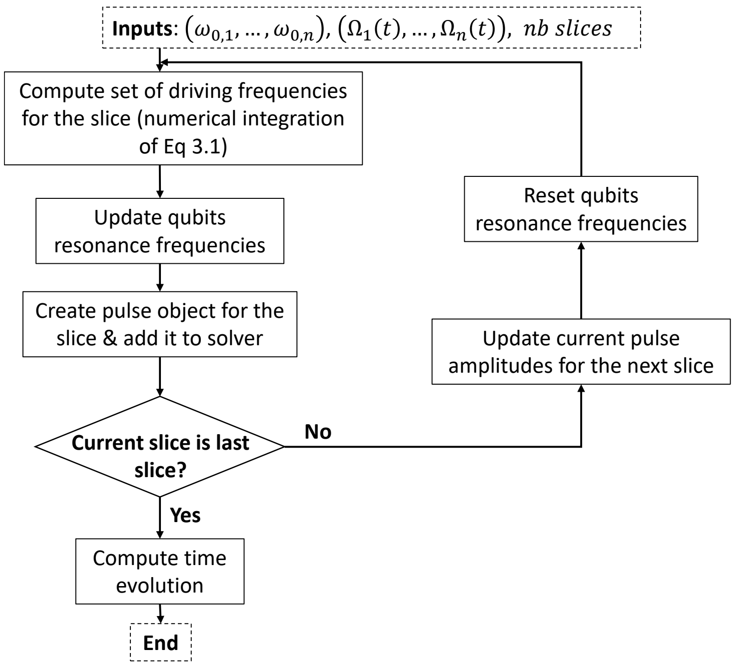

Figure 3.1: Flowchart representation of the frequency shift correction algorithm.

The case of driving configurations with pulses of different durations remains to be discussed. Because each pulse is discretised at the same points in time and each slice is independent from the others, it is in theory possible to handle having pulses of different durations. Indeed, having a pulse end before another would simply result in a different non-linear system (with one less equation and one less variable) to solve. This is however not necessarily relevant, since a driving configuration where qubits have idle time is not desirable.

The algorithm is summarised in Figure 3.1. It processes the time slices sequentially and, for each of them, creates a pulse object (see Chapter 4 for details about it) that is added to the solver. When this is done, the algorithm moves to the next slice by updating the pulse amplitude (the value the continuous pulse would have at the middle of the slice has been chosen as the constant value) and resetting the qubits' resonance frequencies.

3.1.3. Construction of the time-domain pulse

The driving pulse s can be constructed in two ways, defined in Equations 3.2 and 3.3, respectively.

$$s(t) = cos\big(\omega(t) * t + \phi\big) \tag{3.2}$$

$$s(t) = cos\big(\omega * t + \phi(t)\big) \tag{3.3}$$

In Equation 3.2, the drive has a time dependent frequency (determined by the tracking algorithm described before) and a constant phase, while in Equation 3.3 it is defined by a constant frequency (equal to the non-shifted qubit frequency) and a time-varying phase that compensates the shifting of the resonance frequencies.

In this work, the first construction option (Eq 3.2) has been used.

In this case, additional care must be taken when constructing these successive pulses with varying frequency, that their respective phases match. Indeed, if the pulse for each time slice is constructed independently from the previous ones, a phase offset will accumulate through the slices, leading to discontinuities in the pulse. It in turns introduces errors in the driving of the qubits submitted to this pulse. The phase discontinuity is illustrated in Figure 3.2a, where pulses with different frequencies are concatenated (as they would be in the simulation). The discontinuity between the segments is clearly visible.

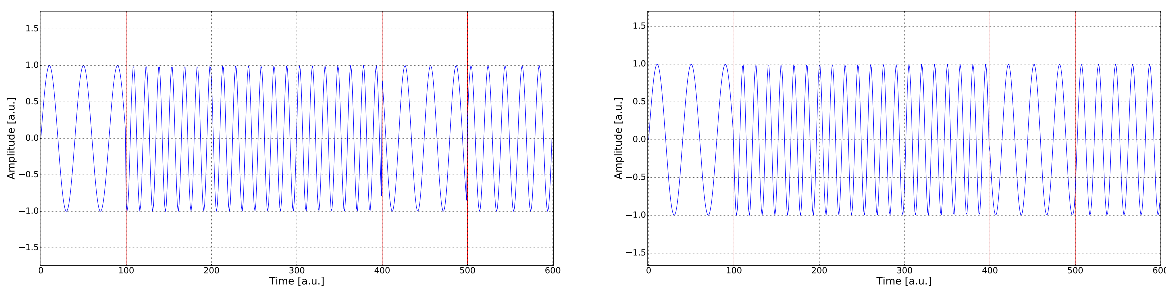

(a) Without phase offset correction

(b) With phase offset correction

Figure 3.2: Pulse with varying frequency with and without phase offset correction. The vertical lines highlight the points at which the frequency is modified.

The phase offset is expressed in Equation 3.4, where $\omega(t)$ is the pulse frequency at time t and t_0 is the end of the time slice being considered. Subtracting this offset to the phase of the driving pulse gives the kind of evolution illustrated in Figure 3.2b, where no discontinuity remain.

$$\phi = \int_0^{t_0} [\omega(t_0) - \omega(x)]\, dx \tag{3.4}$$

3.2. Shaped pulses

This section introduces amplitude modulation of the driving pulses (the so-called pulse shaping) as a powerful tool to mitigate the crosstalk between a pulse targeted at one resonance frequency and the other qubits.

3.2.1. Principle

Pulse shaping is the use of a non-constant time domain evolution (*i.e.* a time-varying amplitude and/or phase) in order to modify the spectral characteristics of the driving pulses. Intuitively, a pulse that has spectral components of low amplitude at the frequencies of the qubits it is not targeting will induce lower crosstalk than one with very high such components. The goal is then to find and/or design such a shape to have the lowest possible components at the frequencies of interest. It has been decided that, as the setup of the qubits is not fully settled, only existing shapes would be used in this study. This choice will furthermore allow to draw trends by comparing the characteristics and performance of various shapes.

3.2.2. Assumptions and model

Before describing the modulation shapes that have been studied, it is important to consider the assumptions and modeling decisions that were taken. First of all, it was assumed that the amplitude of a pulse is constant over space. It means that a pulse aiming at driving a qubit at one end of the array will be “seen” with exactly the same amplitude by the qubit at the other end of the array. Depending on the type of experimental setup, this may be a realistic assumption (in Takeda et al. (2020) for instance). However, in some other types of devices, the driving field is not constant over space. In any case, the aforementioned assumption can serve at least as a sensible worst-case scenario. Indeed, attenuated pulses will always induce less crosstalk and lower frequency shifts that what is obtained here. Consequently, if the techniques presented in this document are shown to be efficient, they would perform even better in an experimental setup with location-dependent drive amplitude.

Secondly, it is necessary to account for the non-ideality of the arbitrary waveform generators (AWG) that produce the pulses. In particular, the pulses have a finite rise time: it takes a non-zero time for the output voltage to go from zero to its target value. This transition time affects the drive and so must be accounted for. This has been done by adding an effective bandwidth to the pulse signal. More specifically, an IIR lowpass filter is applied to the previously computed ideal pulse voltages. An elliptic filter with a cutoff frequency of 400 MHz above the pulse frequency has been chosen. These characteristics have been calibrated by comparison with experimental pulse data.

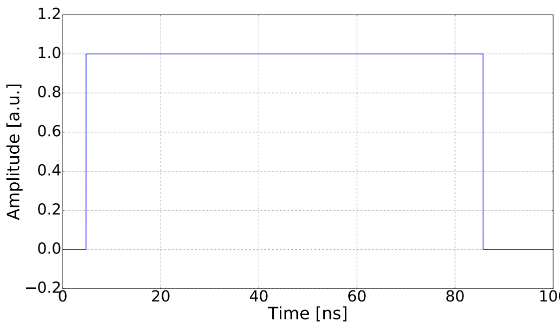

(a) Rectangle pulse, without rise time

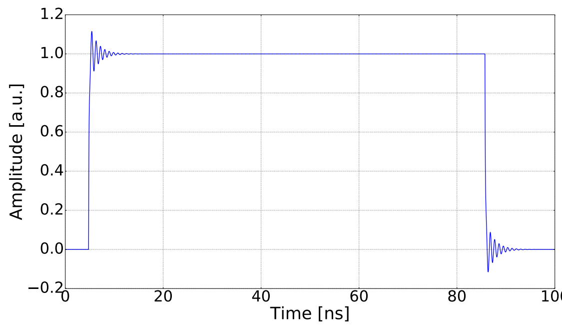

(b) Rectangle pulse, with rise time

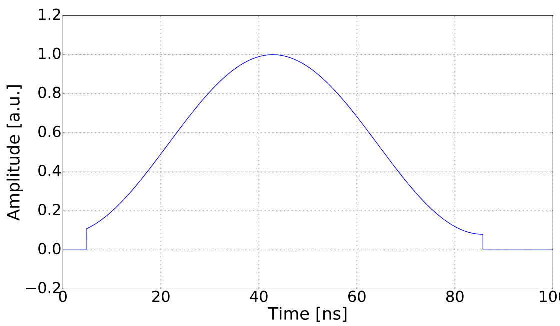

(c) Hamming pulse, without rise time

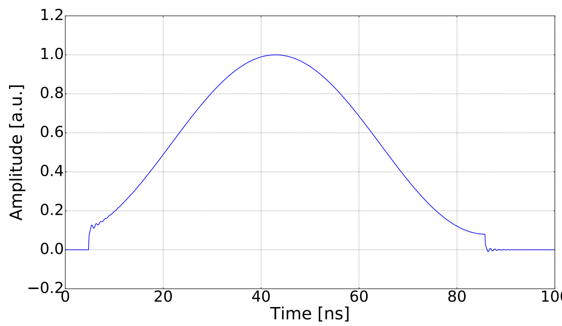

(d) Hamming pulse, with rise time

Figure 3.3: Comparison of the time domain profile of rectangular and Hamming pulse shapes with and without rise time. Each pulse has the same duration, and their amplitudes have been normalised for ease of comparison.

It is worth noting that the addition of rise time will only affect pulse shapes that have sharp-edged features (*e.g.* rectangle pulse which in theory transitions from 0 to its maximal amplitude instantaneously). Since most of the pulse shapes described in Section 3.2.5 and Appendix A actually begin and end at an amplitude of 0 (and thus do not present sharp-edged features), their high-fidelity driving capabilities should not be affected by the addition of finite pulse rise time. Figure 3.3 illustrates the effect of the addition of rise time to the rectangular and the so-called Hamming shapes (see Appendix A for a description of those shapes).

The main effect of introducing rise time to the pulse is the appearance of ringing on the sharp features of their time domain evolution. The comparison between the rectangular (which presents the sharpest feature of the shapes considered in this study) and Hamming shapes (which has a much smaller but still present discontinuity) is insightful. Indeed, Figures 3.3b and 3.3d show that the largest the discontinuity in the time evolution of the shape is, the more important the aforementioned ringing will be.

3.2.3. Pulse duration

In order to perform a rotation of a specific angle, the duration for which the driving field must be turned on must be varied. In the case of a pulse of constant amplitude Ω, the duration required to perform a rotation of angle θ is given by:

$$\tau = \frac{\theta}{\Omega}. \tag{3.5}$$

In the case of a time-dependent amplitude, a pulse will perform a rotation of angle θ when its "area" equals θ. This can be expressed more formally as follows: the duration τ of the pulse must fulfill (Steck, 2007):

$$\Omega \times \int_0^\tau s(t)dt = \theta, \tag{3.6}$$

where s is the shape function (normalized with respect to its maximum value Ω), Ω is the Rabi frequency of the pulse (*i.e.* its peak amplitude) and θ is the desired rotation angle.

The deviations of the non-ideal pulse shapes from their analytical expressions, described in Section 3.2.2, can easily be taken into account by replacing the analytical integration of Equation 3.6 by the numerical one of the filtered time-domain values of the pulse amplitude. This will usually result, at constant peak amplitude, in only slightly modified required durations. This makes sense, considering the ringing introduced by the filtering is approximately symmetrical at the beginning and the end of the pulse.

Another interesting possibility is to define pulses by their duration instead of their maximal amplitude. In that case, the appropriate peak amplitude can also be derived by solving for Ω in Equation 3.6. This method of defining pulses seems beneficial in the context of this analysis because it allows to compare the performance of the different shapes in similar conditions regarding decoherence: as all the durations are equal, decoherence would impact each driving to a similar degree. Furthermore, the pulses that would have required a shorter duration at fixed Rabi frequency will benefit in that case from a lower maximum amplitude, most likely resulting in a reduced amount of crosstalk regardless of their spectral properties. Lastly, this approach to designing driving pulses is similar to what would be expected in a "clocked" quantum computer, where a qubit operation would have to be completed in a fixed clock cycle (regardless of which operation it is, and of which type of pulse it requires). In that case, comparing pulses of a fixed duration would allow to better select the optimal pulse for each type of operation at a fixed clock rate.

On the other hand, the very fact that comparing the pulse shapes from the fixed-duration perspective may partially mask the influence of their spectral characteristics hinders the goal of this analysis. Gaining insight into which shapes induce the less crosstalk by design, independently of their amplitude, is indeed valuable. It would for instance allow to select the best initial candidate pulse, whose duration and amplitude can later be optimised to further increase performance.

Those two approaches of defining driving pulses thus seem to lead to complementary analyses, and are hence both conducted in this document.

3.2.4. Analytical estimation of the performance of a shaped pulse

In order to properly investigate the performance of a pulse shape, it is necessary to be able to relate its spectral characteristics to the spin dynamics it produces. The immediate "brute force" solution would be to compute the response of a spin to excitation of varying frequency using numerical integration methods. However, a tool less computationally intensive would be better suited to practical uses, as it would allow to very quickly get an idea of how well a specific pulse is likely to perform.

The first candidate for this fast analytical tool is the Fourier transform. However, the spin dynamics close to resonance are in general appreciably non-linear, a feature not shared by the Fourier transform. It is then expected to provide inaccurate results in this spectral region of particular interest to this study. It is nevertheless important to note that in some cases (low rotation angle, far off-resonance), the Fourier transform is still a good approximation for the spin dynamics (Freeman, 1998, Kessler et al., 1991).

To better account for this non-linearity, several solutions have been envisioned. The most promising one was an adaptation of the filter function formalism (see Biercuk et al. (2011) and Cywiński et al. (2008)). Filter functions are generally used to describe the way a pulse or pulse sequence suppresses noise-induced dephasing in a quantum system. This suggests that they would qualify as a good descriptor of the non-linear dynamics of a spin. However, it has proved very difficult to find a precise adaptation of the filter function that decently approximated the amplitude of the spin dynamics in all cases. Estimate functions have been defined that would fit spin dynamics pretty well on some driving configurations, but would be a gross approximation on others. Overall, finding a decent approximation of the spin dynamics that covered all the driving scenarios encountered in this study proved difficult. This task has thus been set aside.

In order to still be able to draw some conclusions regarding the spectral properties of the driving shapes, the Fourier transform was used. Although it is far from an ideal approximation, extensive testing highlighted the fact that it simulates some characteristics of the spin dynamics well enough. In particular, after taking into account the frequency shifts in the x-axis of the Fourier transform, the position of the zeros of the FT correspond well to those of the numerically-computed spin dynamics. Furthermore, excluding very close to resonance, the amplitude of the Fourier transform is always superior to that of the spin dynamic. This allows to make qualitative and conservative predictions of the amount of crosstalk a pulse will induce at a specific frequency. A comparison between the Fourier transform of a rectangular pulse and the inversion probability of the spin as a function of the drive frequency is given in Figure 3.4.

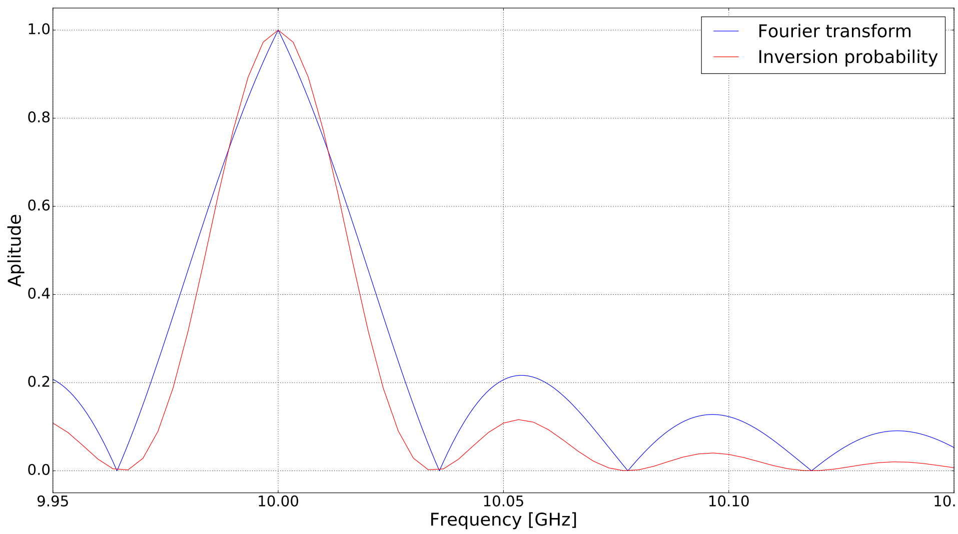

Figure 3.4: Comparison between the Fourier transform of a rectangular pulse and the inversion probability of a spin (ω_0 = 10.0 GHz, initially in the $|\downarrow\rangle$ state) subjected to a π-pulse, as a function of the frequency of the driving pulse. The x-axis of the Fourier transform has been adjusted to include the frequency shifts described in Section 2.2.3.

The pointy shape of the FT very close to resonance is caused by the shifting of the abscissa values of the data points. Due to the nature of the Stark shift, frequencies above resonance are pushed down and *vice versa* for those below, leading to this surprising shape. These values are clearly not a good estimation of the spin inversion probability, but it can be noted that they remain above those of the spin evolution (except on the immediate surrounding of resonance, where having other qubits is not desired anyway). Thus, it is still an acceptable worst-case estimation.

A decent estimation of the shape of the central peak is however useful to get insights into how the pulse would be affected by decoherence. A conservative one can be obtained by looking at the unshifted FT, which retains the proper shape of the main lobe.

Additionally, the energy spectral density can be a useful other tool. It is defined as (Stoica et al., 2005):

$$s(\omega) = \hat{x}(\omega) \times \hat{x}^{\dagger}(\omega) \tag{3.7}$$

where $\hat{x}(\omega)$ is the discrete Fourier transform of the signal and $\hat{x}^{\dagger}(\omega)$ its complex conjugate.

The energy spectral density does not offer a good approximate of the spin dynamics. However, as it measures the energy carried by the pulse as a function of the frequency, it can be useful to estimate for instance the amount of frequency shift that will be induced.

3.2.5. Modulation shapes

A wide range of shapes, taken from different general families, have been chosen for this study. The choice was motivated by their theoretical spectral characteristics as well as their traditional use in literature. An extensive description of those shapes is given in Appendix A. In this Section, we will discuss general characteristics of the shapes, as well as the performances that can be expected based on them.

Shape	Main lobe 3dB width [bins]	Peak side-lobe level [dB]	Sidelobe decrease rate [dB/oct]	Spectral flatness [dB]
Rectangle	0.8845	-13.26	-6 (f^{-1})	-4.2712e-08
Triangular	1.2736	-26.52	-12 (f^{-2})	-1.0532e-08
Sine	1.2031	-23.00	-12 (f^{-2})	-1.3606e-08
Hann	1.4382	-31.47	-18 (f^{-3})	-6.4755e-09
Hamming	1.3008	-42.65	-6 (f^{-1})	-9.8053e-09
Blackman	1.6518	-58.11	-18 (f^{-3})	-3.9244e-09
Opt Blackman, $C = 0.11$	2.6172	-33.12	-18 (f^{-3})	-4.8177e-11
Opt Blackman, $C = 0.12$	2.7914	-31.25	-18 (f^{-3})	-3.4910e-11
Opt Blackman, $C = 0.14$	3.1024	-27.75	-18 (f^{-3})	-1.4921e-09
Papoulis	1.7086	-46.00	-30 (f^{-5})	-3.4001e-09
Gaussian	1.7562	-62.39	-6 (f^{-1})	-3.2891e-09
Half-Gaussian	1.0637	-24.19	-6 (f^{-1})	-7.9694e-10
Hermite	2.1995	-28.60	-6 (f^{-1})	-1.6657e-08
SFT3F	3.1502	-31.73	-18 (f^{-3})	-5.4215e-12
SFT4F	3.7618	-44.73	-30 (f^{-5})	-1.2980e-12
SFT5F	4.2910	-57.26	-42 (f^{-7})	-4.6770e-13
SFT3M	2.9183	-44.20	-6 (f^{-1})	-2.3677e-11
SFT4M	3.3451	-66.49	-6 (f^{-1})	-3.9663e-12
SFT5M	3.8340	-89.91	-6 (f^{-1})	-1.2064e-12
HFT90D	3.8340	-90.22	-18 (f^{-3})	-1.2266e-12
HFT116D	4.1579	-116.80	-18 (f^{-3})	-6.9335e-13
HFT169D	4.7588	-167.88	-18 (f^{-3})	-2.4880e-13
Kaiser 0, $\alpha = 2.0$	1.4270	-45.85	-6 (f^{-1})	-6.7902e-09
Kaiser 0, $\alpha = 4.0$	2.0533	-94.41	-6 (f^{-1})	-2.0338e-09
Kaiser 0, $\alpha = 6.0$	2.3499	-145.53	-6 (f^{-1})	-9.5673e-10
Kaiser 0, $\alpha = 7.0$	2.5297	-171.53	-6 (f^{-1})	-7.1424e-10

Table 3.1: Main spectral characteristics of the presented shapes. The Fourier transform is computed on a pulse sampled at eight times the driving frequency. For better comparison, the shapes are normalised in amplitudes before those metrics are computed. For each metric, the best and worst values are highlighted in blue and red, respectively.

The shapes used in this work are described in more details in Heinzel et al. (2002), Prabhu (2013), Salvatore and Trotta (1988), Bauer et al. (1984), Warren (1984) and Geen and Freeman (1991).

First of all, Table 3.1 summarises the main spectral characteristics of the shapes. Those parameters give a first indication of their performance. The width of the main lobe can for instance directly be related to the frequency selectivity of the pulse. Indeed, a narrow main lobe indicates a very selective pulse (as it is the case for the rectangle, Hamming and half-Gaussian shapes, among others), but in return such a pulse will be very sensitive to noise-induced variations of the qubit's resonance frequency. On the other hand, a pulse with a wide main lobe is likely to be more robust against decoherence, but also to induce more crosstalk (*e.g.* as is the case for the SFT5F and HFT169D shapes).

The sidelobe decrease rate pretty straightforwardly shows how fast the magnitude of the sidelobes will vanish when getting away from the main peak. Intuitively, a high decrease rate will be very beneficial to our problem because it will mean on average a good frequency selectivity. Indeed, assuming the main lobe is not excessively wide and the first sidelobe not excessively high, a high decrease rate will make the sidelobes negligible closer to the resonance. However, a low decrease rate is not necessarily indicative of bad performance. For instance, if the peak sidelobe level is already low, a slow decrease rate might not be a problem at all. The decrease rate is expressed in decibels per octave, which is a usual unit for signal processing applications. An octave is a logarithmic unit for ratios between frequencies, with one octave corresponding to a doubling of frequency. For good readability, the decrease rate is also expressed in terms of asymptotic behaviour: a decrease in f^{-n} means that, asymptotically, the sidelobe level decreases as the inverse of the n^{th} power of the frequency.

Spectral flatness (also called Wiener entropy) is a measure commonly used in digital signal processing. It is defined as the ratio between the geometric and arithmetic mean of the energy spectral density (Dubnov, 2004, Johnston, 1988). It is habitually used to characterise an audio spectrum and evaluate how close a signal is to pure noise. This is not directly relevant to quantum computing but, as white noise has a flat power spectrum, this measure can be used to evaluate the flatness of the drive's FT main lobe. The spectral flatness presented in Table 3.1 is computed over a range of 400 kHz centered on the drive frequency. This range has been selected to focus on the main lobe of the Fourier transform but also to encompass the expected linewidth of a qubit in a real experimental setup. It gives a good indication of how robust a pulse can be against decoherence: the closer the value is to 0 dB, the flatter the main lobe is, and thus the more robust it is. Indeed, a pulse with a flat FT main lobe can efficiently drive a qubit which is beginning to loose coherence and whose resonance frequency starts to shift.

Complementary information is given by Table 3.2, which presents the magnitude of the Fourier transform at detuning values that would correspond to the resonance frequencies of neighbouring qubits in our setup. To be as representative as possible, several typical values of Rabi frequency have been covered. Moreover, in order to limit the impact of some configurations coincidentally falling very close to local extrema of the Fourier transform, the values presented are averaged over a frequency range of 30 MHz centered on the considered detuning. These magnitudes are theoretically a qualitative indicator of the level of crosstalk that can be expected on qubits at the corresponding frequencies.

From these two tables, it is already possible to make educated guesses on which shapes might perform well or not. In particular, shapes with high to medium frequency selectivity and small peak sidelobe levels are good candidates. Such shapes include the Hamming, Blackman, Papoulis, Gaussian and Kaiser shapes. Furthermore, if added to the other characteristics, a high sidelobe decrease rate suggests an even better performance, as the already low highest sidelobe will decrease very fast. According to this reasoning, modulations such as the Papoulis, Blackman and Hann shapes should provide among the best results.

Shapes with a larger main lobe could still be efficient solutions in experimental conditions. Admittedly, a wide main lobe means a lower selectivity and thus a higher risk of inducing crosstalk. However, it also means a higher robustness against decoherence (which is also indicated by a value of passband flatness close to 0 dB). Furthermore, shapes with a wider main lobe are often designed to have very high sidelobe decrease rate. This is for instance the case of the SFT4M, SFT5M and HFT90D modulations shapes. Kaiser shapes with high α values could also work, because despite their low sidelobe decrease rate, they show the lowest peak sidelobe level of all the shapes considered, as well as a decent spectral flatness.

	$\Delta = \pm 50$ MHz			$\Delta = \pm 100$ MHz			$\Delta = \pm 200$ MHz		
Ω	5 MHz	15 MHz	25 MHz	5 MHz	15 MHz	25 MHz	5 MHz	15 MHz	25 MHz
Rectangle	1.45e4	1.34e5	2.18e5	7.11e3	6.18e4	1.17e5	3.53e3	3.24e4	5.99e4
Triangular	3.85e2	1.13e4	1.74e4	9.10e1	2.20e3	5.11e3	2.24e1	6.20e2	1.36e3
Sine	6.39e2	1.69e4	5.96e4	1.44e2	3.98e3	1.89e4	3.52e1	1.10e3	4.50e3
Hann	4.30e1	3.40e3	2.45e4	4.60e0	3.53e2	2.95e3	5.54e-1	4.60e1	3.75e2
Hamming	5.69e2	1.71e3	1.96e4	2.91e2	2.32e3	4.23e3	1.57e2	1.29e3	3.13e3
Blackman	8.78e0	2.80e2	1.70e4	9.95e-1	6.45e1	4.71e2	1.18e-1	8.68e0	6.03e1
Opt Black ($c = 0.10$)	5.54e0	5.89e2	5.30e3	6.08e-1	5.06e1	4.01e2	7.24e-2	5.85e0	4.70e1
Opt Black ($c = 0.12$)	5.55e0	4.46e2	4.65e3	5.89e-1	5.04e1	3.54e2	7.35e-2	5.76e1	4.88e1
Opt Black ($c = 0.14$)	4.52e0	3.42e2	2.55e3	4.84e-1	4.02e1	3.10e2	5.97e-2	4.94e0	3.70e1
Papoulis	3.69e0	8.87e2	1.51e4	2.08e-1	4.69e1	4.05e2	1.27e-2	3.06e0	4.79e1
Gaussian	2.74e1	3.88e2	2.40e4	1.38e1	1.20e2	3.08e2	6.90e0	6.19e1	1.46e2
Half-Gauss.	3.25e3	2.67e4	1.25e5	1.66e3	1.51e4	3.44e4	7.94e2	6.65e3	2.63e4
Hermite	1.45e4	5.03e3	1.09e3		2.50e3	6.02e2	3.52e3	1.24e3	3.13e2
SFT3F	5.80e0	4.57e2	4.67e3	6.10e-1	4.81e1	3.85e2	7.20e-2	6.05e0	4.61e1
SFT4F	4.18e-2	3.02e1	8.75e2	8.89e-4	6.89e-1	1.57e1	2.53e-5	1.83e-2	4.23e-1
SFT5F	4.31e-4	3.47e0	7.69e2	1.83e-6	1.29e-2	8.35e-1	1.20e-8	7.58e-5	4.50e-3
SFT3M	1.68e2	1.13e3	1.49e3	8.57e1	7.61e2	1.79e3	4.18e1	3.66e2	1.03e3
SFT4M	1.73e1	4.60e1	8.84e2	9.30e0	6.62e1	1.26e2	4.61e0	4.01e2	1.06e2
SFT5M	1.60e-1	3.19e0	1.10e3	1.89e-2	1.26e0	6.10e0	2.12e-3	1.81e-1	1.28e0
HFT90D	3.21e3	8.07e5	1.66e6	5.87e0	1.31e5	1.12e6	1.14e0	5.21e0	3.84e4
HFT116D	8.40e3	8.59e5	1.68e6	1.85e-1	2.14e5	1.25e6	1.56e-1	6.12e-1	1.06e5
HFT169D	2.44e4	9.24e5	1.71e6	5.83e-4	3.69e5	1.41e6	5.62e-4	4.06e2	2.87e5
Kaiser 0 ($\alpha = 0.643$)	5.01e3	5.00e4	1.15e5	2.37e3	4.00e0	7.46e4	1.21e3	9.65e3	2.54e4
Kaiser 0 ($\alpha = 1.929$)	1.02e2	1.06e3	1.79e4	4.94e1	4.79e2	1.26e3	2.48e1	2.29e2	6.86e2
Kaiser 0 ($\alpha = 3.857$)	0.26e0	4.19e1	1.72e4	0.12e0	1.17e0	4.04e0	5.79e-2	5.48e-1	1.50e0
Kaiser 0 ($\alpha = 5.786$)	6.18e-4	9.71e1	1.56e4	2.83e-4	9.00e0	1.31e-2	1.41e-4	1.35e-3	4.06e-3

Table 3.2: Magnitude of Fourier transform at frequencies corresponding to neighbouring qubits, for various representative values of maximum Rabi frequency Ω. The Fourier transforms have been computed in the case of a π-pulse. Furthermore, the values have been averaged over a 30 MHz interval centered on the indicated Δ in order to eliminate coincidental hitting of very low Ft amplitudes. For ease of reading, very low values have been highlighted in blue and very high ones have been in red.

Additionally, Table 3.2 gives qualitative indication on which shapes induce low amounts of crosstalk. The blue and red highlights respectively indicate the lowest and highers Fourier transform amplitudes. It is important to note that, for a single shape, different Rabi frequencies can produce very different types of results. This is due to the existence of "sweet spots" in the Fourier transform of the pulse, where its amplitude is almost zero. When one of these spots is close to the resonance frequency of a neighbouring qubit, the drive is expected to produce very little crosstalk (if any at all). The coincidental hitting of such a spot cannot be relied upon, as the exact resonance frequencies of the qubits are unpredictable. Nevertheless, as the FT values have been averaged on small ranges of frequency centered on the considered detunings, the unrepresentative low points have been eliminated. Consequently, the general trends highlighted by the table are relevant. Shapes like SFT4F, SFT5F, Kaiser (high α) and Papoulis show encouragingly low FT amplitudes at the resonance frequencies of other qubits, for most of the amplitudes considered. However, the values of shapes such as Kaiser (low α) and half-Gaussian indicate that, although promising in Table 3.1, they might not perform as well as expected. Some other shapes that seemed promising in Table 3.1 (*e.g.* Hann, Hamming or sine) show intermediate values, so they might still perform well in practice. Finally, shapes of the HFTxD family, characterised by a very wide main lobe, indeed show very high FT values at the resonance frequencies of neighbouring qubits in most cases. They are thus expected to produce high levels of crosstalk.

Recalling the two methods to define driving pulses envisioned in Section 3.2.3, fixing either the Rabi frequency or the duration as the initial parameter from which the others are computed, it is insightful to investigate how the peak amplitude required by a pulse relates to that of the non-shaped one of identical duration. Indeed, a lower peak amplitude will induce lower crosstalk on the neighbouring qubits.

Table 3.3 presents the ratio between the Rabi frequency required by each shape and that of the rectangle pulse of same duration: $\Omega_{\text{shape}}/\Omega_{\text{rectangle}}$. The pulses considered include the rise time introduced in Section 3.2.2 and its influence on the needed amplitude. It is interesting to note that the value of this ratio is independent of the desired rotation angle or of the actual duration of the pulse, and is thus an intrinsic characteristic of the shape.

Shape	$\Omega_{\text{shape}}/\Omega_{\text{rectangle}}$
Rectangle	1.000
Triangular	1.974
Sine	1.550
Hann	1.973
Hamming	1.830
Blackman	2.348
Optimised Blackman ($c = 0.10$)	3.287
Optimised Blackman ($c = 0.12$)	3.792
Optimised Blackman ($c = 0.14$)	4.482
Papoulis	2.433
Gaussian ($\sigma = 1.0$)	2.564
Half-Gaussian ($\sigma = 1.0$)	2.564
Half-Gaussian2 ($\sigma = 1.0$)	2.662
SFT3F	3.717
SFT4F	4.542
SFT5F	5.242
SFT3M	3.485
SFT4M	4.075
SFT5M	4.703
HFT90D	4.689
HFT116D	5.115
HFT169D	5.888
Kaiser 0 ($\alpha = 0.643$)	1.253
Kaiser 0 ($\alpha = 1.929$)	1.982
Kaiser 0 ($\alpha = 3.857$)	2.768
Kaiser 0 ($\alpha = 5.786$)	3.378

Table 3.3: Ratio of the peak amplitude of the shapes with the amplitude of a rectangle pulse of same duration.

The fact that all the shapes have a ratio $\Omega_{\text{shape}}/\Omega_{\text{rectangle}}$ greater than one suggests that, at fixed pulse duration, the unshaped pulse may significantly outperform the shaped ones. The rectangle modulation is however closely followed by the sine, triangular, Hann, Hamming and Kaiser ones, which have a ratio between 1 and 2. They are thus also expected to be amongst the best-performing shapes in the fixed-duration configuration. Furthermore, the lower values of the free parameters of the Kaiser and optimised Blackman shapes seem to require lower Rabi frequencies. This is coherent with the fact that, with the increase of their free parameter, these shapes tend to narrow (see Figures A.7 and A.13 in Appendix A). As a consequence, at fixed pulse duration, their maximum amplitudes need to increase for the area under the curve to remain constant. These low values of the shapes' free parameters are then likely to be preferred over larger ones in the fixed duration configuration.

Combining all these considerations, the Hann, Hamming, Blackman, Kaiser (with large α) and Papoulis shapes appear as very promising candidates. Other shapes like the sine, Gaussian, SFT4M and SFT5M could also be efficient choices. However, shapes belonging to the HFTxD family are expected to perform poorly due to their very wide main lobe. The case of the rectangle and sine modulations is slightly special. Indeed, in a fixed amplitude configuration, they are expected to perform averagely at best, but at constant duration they require the lowest amplitudes and are thus likely to induce the least amount of crosstalk among the considered shapes.

4 Implementation details

This chapter describes the software that was developed for this project. The tools used in the development, as well as the main features of the software are first listed. Then, a more detailed look at the architecture is given. A comprehensive documentation of the code is furthermore given in Appendix F.

4.1. General information

The development effort of this project was conducted using Python. Two main external packages were used: QuTip (Johansson et al., 2013), an open source software library for simulating quantum systems, and DM_solver (Philips), a C++ library developed in QuTech, wrapped to be accessible as a Python package. QuTip is used for pre- and some of the post-processing, while DM_solver handles the computation of the time evolution of the qubits. DM_solver implements a solver for both the Von-Neumann and the Lindblad master equations. As only unitary evolutions have been considered, this work exclusively used the Von-Neumann solver.

All the crosstalk phenomena described in Chapter 2 have been implemented, and modeling of the drives using and outside of the RWA are available (with a possibility to toggle from one to the other easily).

The software developed during this projects has the following features:

- n-qubits m-driving pulse configurations. The software models an n-qubits linear array that is driven by m different and independent pulses. The number of qubits and pulses is fully scalable, and modifying it does not require any significant coding effort.
- Independent driving pulse configurations. All the m driving pulses can have a different frequency, shape, peak amplitude and phase. Furthermore, each of them can implement a different arbitrary qubit rotation.
- Several different pulse shapes. Furthermore, the addition of new shapes by the user is straightforward.
- Modeling of resonance frequency shifts and the associated correction algorithm.
- Object-oriented implementation, allowing easy use of these features.
- Everything wrapped in a ready-to-use Python package.

4.2. Architecture

The architecture of the software is presented in Figure 4.1. The simulation relies on two main classes: *n_qubit_simulation* and *shaped_pulse*. The *n_qubits_simulation* class stores the configuration of the qubit array (number of qubits and their respective resonance frequencies). It also holds the solver object (which is part of DM_solver) that will compute the system's time evolution. This class also has methods to compute the fidelity (both state and unitary) of the rotations it simulates.

The *shaped_pulse* class represents the driving pulses. Because the aim of this project is to study how to implement gates on multiple qubits at the same time, it was decided that this class would represent m driving pulses together (usually with $m = n$, meaning that all the qubits are driven simultaneously). The class stores all the values defining the pulses (*i.e.* pulse envelope and the values of the drive frequencies at each time slices (see Section 3.1.2)). It also holds pulse driving parameters such as the target rotation angle, the pulse maximal amplitude and the possible free parameters defining the modulation shape. Each of the m pulses can have different characteristics. In addition, class methods allow to add rise time to the pulses envelopes and to include the resonance frequency shifts, as well as the aforementioned correction algorithm.

In terms of implementation, a pulse is defined by its targeted rotation angle and either by its duration or its amplitude. The other parameter is then computed using Equation 3.6. An additional *arbitrary_pulse* class (not represented in Figure 4.1) enables the user to specify both the amplitude and the duration, allowing to model an arbitrary driving.

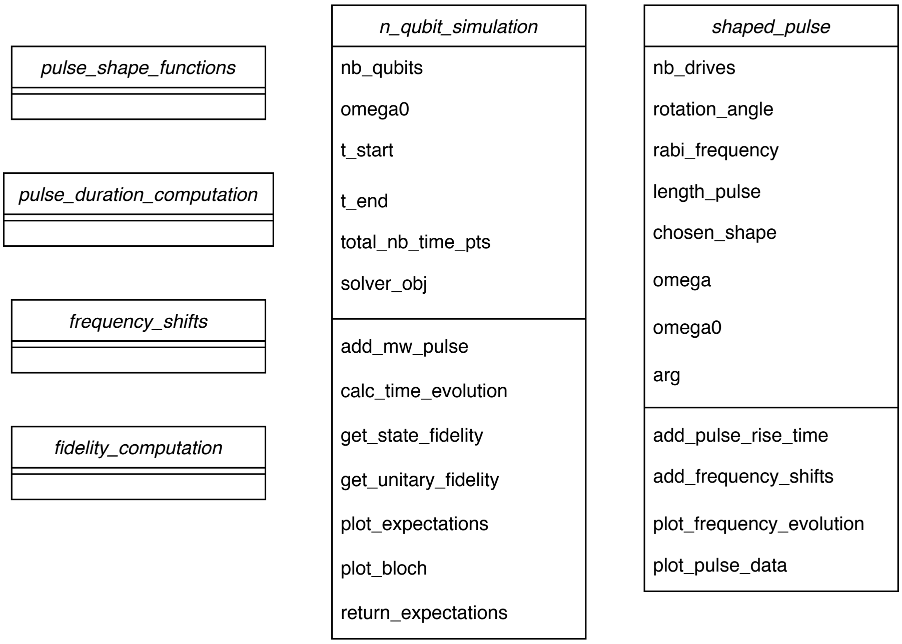

Figure 4.1: Class diagram of the software implemented in this work. The four boxes on the left of the figure represent utility functions that are not part of a class structure. They have been added here for the sake of completeness.

Each class also has plotting methods to easily visualise for instance the expectation values of specified operators or the evolution of the drive frequencies through the pulse when correcting for the resonance frequency shifts.

The software developed for this project also comprises several utility functions that are not encompassed in a class structure. These functions handle the definition of pulse shapes, the computation of the duration required to perform a rotation of angle θ using a specific shape (which may include rise time or not), the computation of fidelity, the computation of the Fourier transform of a pulse, as well as the computation of the resonance frequency shifts caused by the driving pulses. For completeness' sake, these functions have been represented by the *pulse_shape_functions*, *pulse_duration_computation*, *frequency_shifts* and *fidelity_computation* boxes in Figure 4.1. No additional detail is given here to preserve the clarity of the figure, but these functions are extensively documented in Appendix F.

5
Results

This chapter presents the results of the analyses conducted in this work. It first focuses on the performance achieved by the frequency shift correction algorithm (Section 5.1) and then describes how well pulse shaping allows to handle unwanted driving (Section 5.2).

The results presented are for $\pi/2$- and π-rotations. These rotation angles have been selected because they are widely used in quantum computation, and thus constitute a good benchmark for the performance of the considered techniques. Furthermore, the analysis was conducted using rotations around the $\hat{x}$-axis, since changing the rotation axis is just a matter of adding a constant phase to the driving pulses and has little impact on the overall results.

5.1. Frequency shifts correction algorithm

The performance of the frequency shifts correction algorithm has been evaluated by considering an array of qubits submitted to the nominal driving configuration (*i.e.* each qubit driven simultaneously) while setting the algorithm to correct for a specific set of shifts. Corrections accounting for only the AC Stark shift, for only the Bloch-Siegert shift, but also for both shifts simultaneously have been investigated. The configuration where the qubits are shifted but no correction is applied has been used as a reference for quantifying the improvement brought by the algorithm. Please note that the aim of this section is to discuss the performance of the frequency shift correction algorithm only. The relative performances of the different modulation shapes, although they are mentioned here, are investigated in detail in the following Section 5.2.

As it was discussed in Section 3.2.5, some shapes show very specific spectral “sweet spots" where they induce almost no crosstalk. However, since targeting the exact frequencies corresponding to these sweet spots is not a practical solution, steps have been taken to limit their impact on the results. Random variables centered on each qubit's resonance frequency and following a Gaussian distribution with a standard deviation of 10 MHz have been sampled. Several simulations have then been conducted using the results of this sampling as the resonance frequencies of the qubits. The obtained fidelities have then been averaged, thus greatly reducing the impact of potential spectral sweet spots.

Regardless of this manipulation, the results directly depend on the number of driving pulses involved, and hence on the number of qubits in the array. Consequently, this analysis has been performed in both a two- and a three-qubits configuration. A preliminary analysis was also conducted in a five-qubits setup using a subset of modulation shapes (see Section 5.3).

The following two subsections highlight and explain the trends common to all the shapes, using the results of the geometrical and raised-cosine shape families as illustration. The configurations that present unique features are also discussed. The detailed results of the remaining shapes are given in Appendix C. For ease of navigation, Table 5.1 summarises all the studied configurations and provides references to the Figure showing the corresponding results.

Configuration		Figure reference	
		2 qubits	**3 qubits**
$\boldsymbol{\pi/2}$-rotation,	**Geometrical** & **raised-cosine** shapes	Figure 5.2, page 31	Figure 5.5, page 34
$\boldsymbol{\pi/2}$-rotation,	**Kaiser** shapes	Figure C.1, page 77	Figure C.11, page 81
$\boldsymbol{\pi/2}$-rotation,	**Optimised Blackman** shapes	Figure C.3, page 78	Figure C.13, page 82
$\boldsymbol{\pi/2}$-rotation,	**Flat top** shapes	Figure C.5, page 79	Figure C.15, page 82
$\boldsymbol{\pi/2}$-rotation,	**Gaussian** shapes	Figure C.7, page 79	Figure C.17, page 83
$\boldsymbol{\pi/2}$-rotation,	**HFTxD** shapes	Figure C.9, page 80	Figure C.19, page 84
$\boldsymbol{\pi}$-rotation,	**Geometrical** & **raised-cosine** shapes	Figure 5.3, page 32	Figure 5.6, page 34
$\boldsymbol{\pi}$-rotation,	**Kaiser** shapes	Figure C.2, page 78	Figure C.12, page 81
$\boldsymbol{\pi}$-rotation,	**Optimised Blackman** shapes	Figure C.4, page 78	Figure C.14, page 82
$\boldsymbol{\pi}$-rotation,	**Flat top** shapes	Figure C.6, page 79	Figure C.16, page 83
$\boldsymbol{\pi}$-rotation,	**Gaussian** shapes	Figure C.8, page 80	Figure C.18, page 83
$\boldsymbol{\pi}$-rotation,	**HFTxD** shapes	Figure C.10, page 80	Figure C.20, page 84

Table 5.1: Summary of the different driving configurations investigated, and where the results are represented.

5.1.1. Case of a two-qubit array

The results obtained in a two-qubits configuration are described below. First of all, Figure 5.1 compares the fidelity of the qubit rotations when using the frequency shifts correction with those obtained by a simulation entirely without frequency shifts. The objective is to give a quick overview of the overall performance of the algorithm. The detailed improvements and their explanations are given below. This test was conducted using a subset of modulation shapes, for both $\pi/2$ and π rotations (Subfigures 5.1a and 5.1b, respectively). Each subfigure presents the difference between the case where no shift is simulated and the one where they are simulated but corrected for.

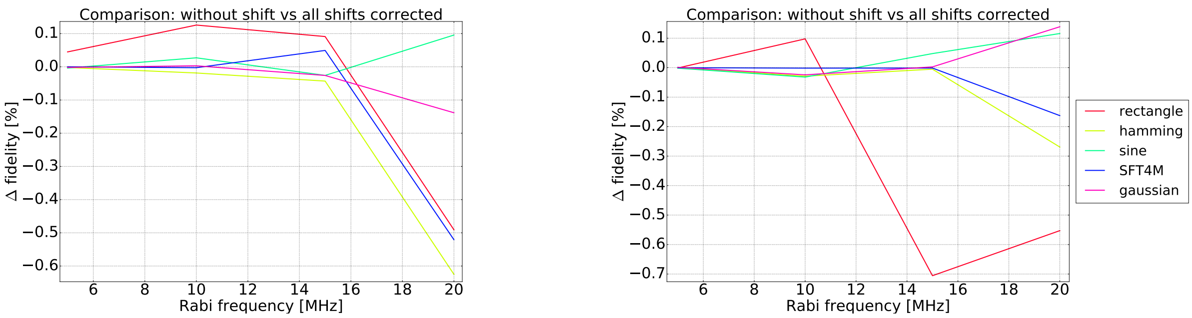

(a) $\pi/2$ rotation

(b) π rotation

Figure 5.1: Comparison between the unitary fidelities of rotations simulated either entirely without frequency shifts or with shifts and the correction algorithm. The analysis was conducted for both $\pi/2$ and π rotations of several representative Rabi frequencies. A few modulation shapes were investigated, to introduce no loss of generality by exploiting the characteristics of one in particular.

For both rotation angles, the difference in fidelity between the two test cases is very low (it never exceeds 0.14%). Furthermore, for most of the shapes and Rabi frequencies the difference is actually negative, meaning that the fidelity in the corrected case is higher than the one where the shifts are not simulated. This surprising feature can be explained by the fact that the drive frequencies selected by the algorithm are slightly further apart than the initial ones. As a consequence, the spectral separation between the qubits is slightly larger and the level of crosstalk each pulse induces on its neighbour is on average sightly lower, resulting in higher fidelities. In some specific configurations, the absolute value of the difference is much larger. This is caused by the frequency of the neighbouring qubit falling very close to either a zero (large negative difference) or a peak (large positive difference) of the pulse's Fourier transform, thus producing either a lot less or a lot more crosstalk, respectively.

Despite these outliers, the main takeaway of this experiment is that the fidelities achieved using the frequency correction algorithm are very comparable to those of the shift-free simulations, if not better. This seems to indicate that the correction algorithm is very beneficial to the fidelity of the rotations.

The performance of the correction will now be investigated and discussed in more details. In this subsection and the following one, the focus has been put on describing and understanding general trends rather than on pointing out the specific shapes that benefit a lot from the correction algorithm.

Figure 5.2 and 5.3 present the improvements in fidelity obtained by the different correction types over the uncorrected configuration, for $\pi/2$- and π-rotations, respectively. In these two figures, as well as in the other ones designated in Table 5.1 and the similar ones in Section 5.3.1 , the top three panels show the fidelity improvement brought by correcting for either only the Bloch-Siegert shift, only the AC Stark shift, or both simultaneously. The bottom panel gives the unitary infidelities achieved by the fully corrected rotations.

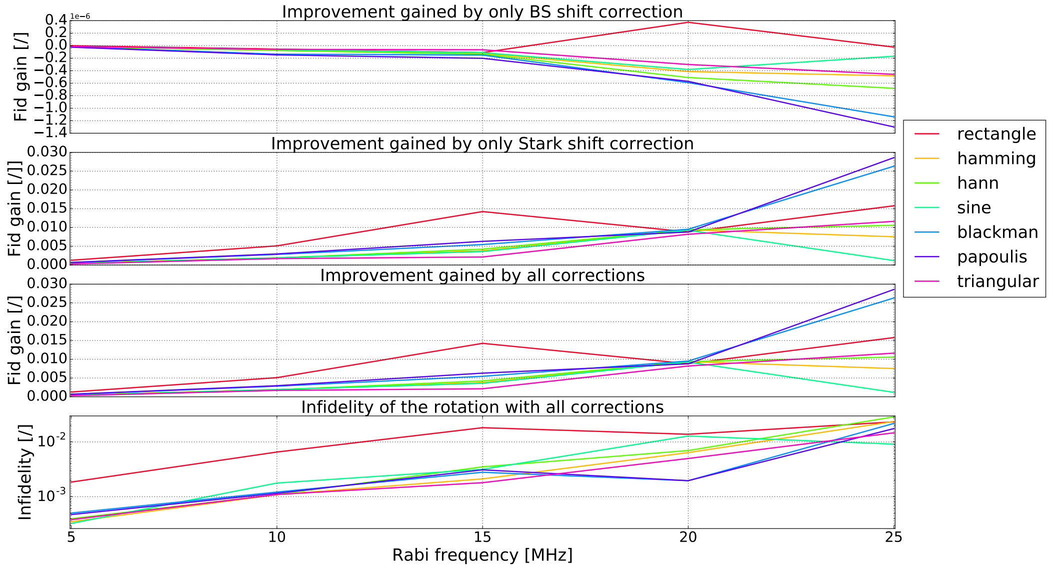

Figure 5.2: Gain of unitary fidelity provided by the frequency shift correction algorithm on geometrical and raised-cosine $\pi/2$ pulses in the 2-qubits 2-drives configuration, for different combinations of shift corrections.

Several trends are visible in Figures 5.2 and 5.3 (as well as in the configurations covered in Appendix C). First of all, it appears that the frequency shift correction algorithm improves the unitary fidelity of the qubit operations significantly (by up to 55% in the case of π-pulses modulated using flat top shapes, *cf.* Figure C.6). It also appears that the fidelity improvement in the case of π-pulses is larger than for $\pi/2$, which was expected. Indeed, a longer rotation benefits more from the correction because a detuned drive over a longer duration accumulates more phase and leaves the spin further from its intended destination, and so the correction makes a larger difference.

Furthermore, in most cases, this fidelity improvement increases with the Rabi frequency of the pulse. The increase is even approximately linear for π- and most of the $\pi/2$-pulses. This increase is again an expected result: a larger peak amplitude means larger shifts and thus a more perturbed qubit evolution. Consequently, correcting for the shifts yields a higher fidelity increase.

The behaviour of some shapes in $\pi/2$-pulses is however slightly different. While the overall trend is of a linear increase of the fidelity gain with the Rabi frequency, for some values of Ω the improvement is either unexpectedly high or low. It is for example the case of the rectangle shape at $\Omega = 15.0$ MHz in Figure 5.2. This data point shows a higher fidelity increase than what the rest of the results would suggest. This difference can be explained by a change in the crosstalk environment of the qubits between the corrected drives and the uncorrected ones. Indeed, for the other values of Rabi frequency, the level of crosstalk caused by the corrected and uncorrected drives are approximately the same. However, for $\Omega = 15$ MHz, the crosstalk intensity at the frequency of the neighbournig qubit is higher for the uncorrected drive than for the corrected one. In summary, the rectangle pulse of amplitude $\Omega = 15$ MHz is improved by both the frequency correction and the pulse spectral characteristics. On the other hand, for the other amplitudes the spectral features remain similar between corrected and uncorrected pulses. The difference is illustrated in Figure C.22 (page 85).

The fact that crosstalk causes these outliers can be confirmed by computing the performance of the correction algorithm in a crosstalk-free environment, composed of only one qubit driven by one pulse, but with

Figure 5.3: Gain of unitary fidelity provided by the frequency shift correction algorithm on geometrical and raised-cosine π pulses in the 2-qubits 2-drives configuration, for different combinations of shift corrections.

the same frequency shifts as with two drives. Because this configuration is highly unrealistic, the results of this test case are presented in a separate section of appendix C, in Figure C.23 (page 86). Despite not representing any realistic setup, this test case allows to isolate the contribution of the correction algorithm to the fidelity. In Figure C.23, the evolution of the improvement with the Rabi frequency does not show any outliers, which further proves the crosstalk-related origin of these outliers. A similar behaviour has been observed for other shapes (*e.g.* in Figures C.1 and C.16) but is difficult to predict beforehand due to the large number of parameters needed to *a priori* evaluate the crosstalk intensity that a specific setup will produce.

The results for $\pi/2$- and π-rotations in Figures 5.2 and 5.3 also clearly show that different shapes benefit differently from the correction algorithm. This time, it is not related to the crosstalk environment of the qubits because the same differences appear in the test case with only the shifts and no crosstalk. When comparing the order of performance with the shape characteristics (Table 3.1), it appears that the shapes that experience the largest improvement are those which have the widest FT main lobe and the passband flatness closest to 0 dB. This result is surprising at first, because one would expect to observe the opposite: on shapes with narrow and pointy main lobes, a specific value of shift will make the drive more off-resonant than on wider shapes. However, when reasoning in terms of energy, a pulse with a wide FT main lobe has a similarly wide main lobe in its energy spectral density (defined as $s(f) = \hat{x}(f) \times \hat{x}^{\dagger}(f)$, where $\hat{x}(f)$ is the DFT of the signal and $\hat{x}^{\dagger}(f)$ its complex conjugate). Furthermore, compared to pulses with narrow FT main lobes, those with wide and flat main lobes have higher average energy over the shifted frequency values encountered throughout the pulse. This then logically translates into the fact that they cause larger frequency shifts, hence why the correction improves their performance further.

Some remarks can also be made concerning the contribution of each of the correction individually. Only correcting for the AC Stark shift brings the most significant improvement, in agreement with the fact that it is the largest of the two shifts. Adding the correction of the Bloch-Siegert shift nevertheless increases the fidelity gain very slightly. On the other hand, correcting for the Bloch-Siegert only seems to decrease the fidelity in most cases, albeit by always less than 10^{-2}%. This is due to the fact that the Bloch-Siegert shift is always positive and that some qubits actually experience a shift that is negative. Consequently, the correction pushes the corresponding drives further away from the appropriate frequencies. The fidelity loss is so small because the Bloch-Siegert shift is very small over qubits with resonance frequencies of the order of 10 GHz. While it is interesting to observe these phenomena and to understand them, it is not entirely relevant to the main objective of this study. Indeed, driving qubits while correcting for only one of the causes of frequency shifts is not a realistic scenario. The most important outcome is that the frequency correction algorithm does improve the fidelity of the rotations, and that the largest improvement is obtained when correcting for the combination of all the shifts.

The case of the HFT90D, HFT116D and HFT169D shapes, presented in Figure C.9 and C.10 (page 80), is slightly special. While for the other shapes the performance of the correction algorithm improves with the amplitude of the pulse, for the HFTxD shapes an opposite trend appears. The improvement is indeed significant for small values of Ω but decreases as the Rabi frequency increases, until the uncorrected pulses actually outperform the corrected ones. Furthermore, the evolution of the improvement for the $\pi/2$-pulses is rather irregular. The outliers can be explained by a similar phenomenon as those of the rectangle shape described above. However, the decreasing trend requires more analysis. First of all, the HFTxD shapes all have extremely wide FT main lobes, which can easily encompass the resonance frequencies of the neighbouring qubits, resulting in bad performance. Increasing the Rabi frequency shortens the pulses, which further widens their main lobes. As a result, above a certain value of Ω (for which the FT main lobe starts to significantly overlap the frequency of the neighbouring qubits), the level of crosstalk produced by the pulse drastically increases. This very high level of crosstalk means that the fidelity of the rotations is always very low (see Figure 5.4 for the evolution of the fidelity with Ω for all the HFTxD shapes, using different types of frequency shift corrections). The fidelities of the uncorrected rotations are especially low ($< 20\%$), but vary little with Ω. In the case of corrected pulses, the fidelity of the corrected rotations decreases with increasing Ω due to the increased crosstalk (as with every other shape). This decreasing trend in the corrected case, combined with the constant fidelity in the uncorrected one leads to the downward trend observed in Figure C.9 and C.10. The relevance of HFTxD shapes as a good pulse modulation is discussed in Section 5.2.

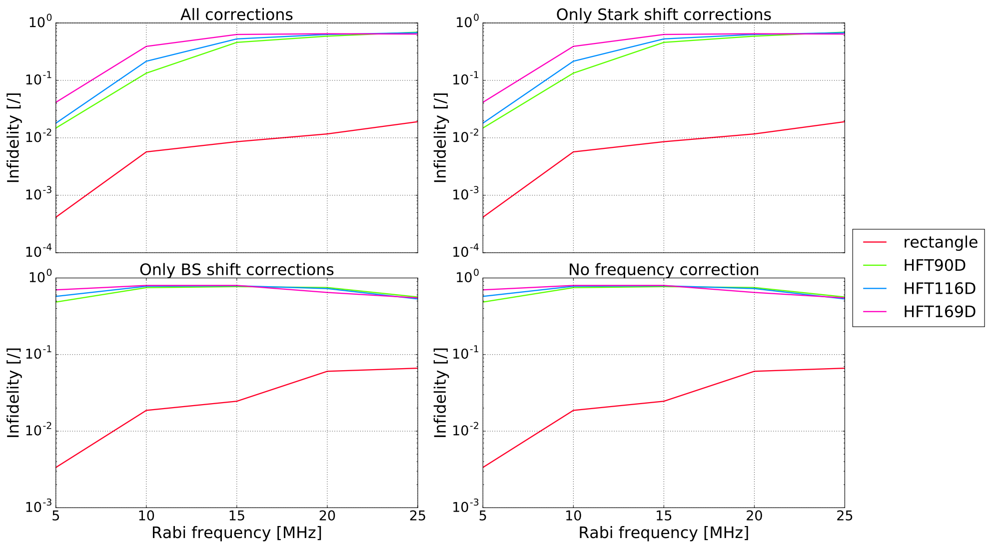

Figure 5.4: Evolution of the unitary infidelity of the π-rotation of 2 qubits driven by pulses modulated using the HFTxD shape family. Each panel uses a different type of frequency shift correction, the most notable ones being the fully corrected and uncorrected cases (top left and bottom right panels, respectively).

5.1.2. Case of a three-qubit array

The results obtained with a three-qubits array are described below. In a similar fashion to the previous subsection, the general trends are highlighted, using the results for the geometrical and raised-cosine shape families to demonstrate and explain them. Figures 5.5 and 5.6 show the fidelity improvement obtained by the frequency shift correction algorithm in the case of a $\pi/2$- and π-rotation, respectively. The performance of the other shapes is again detailed in Appendix C.

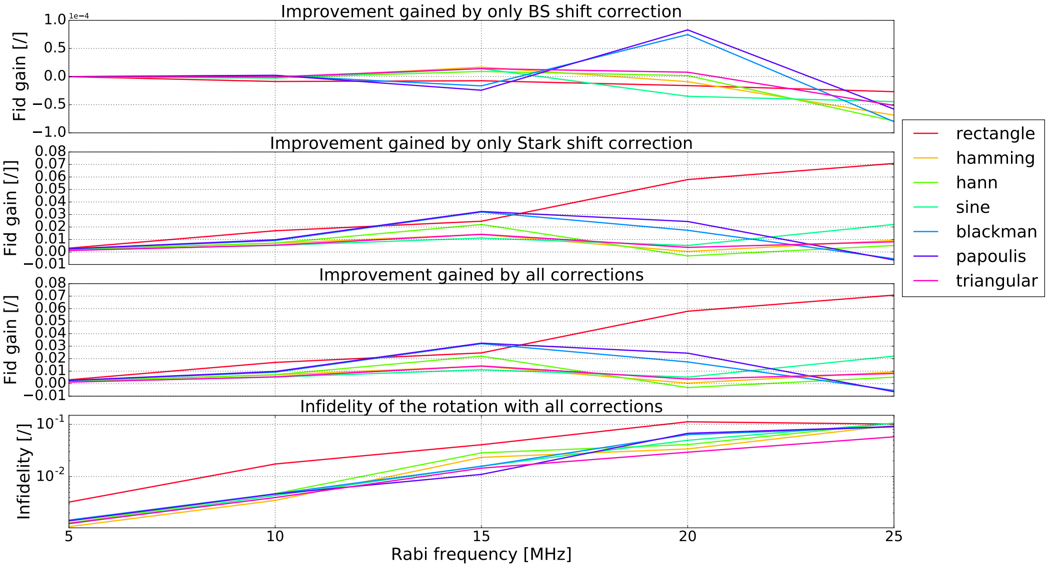

Figure 5.5: Gain of unitary fidelity provided by the frequency shift correction algorithm on geometrical and raised-cosine $\pi/2$ pulses in the 3-qubits 3-drives configuration, for different combinations of shift corrections.

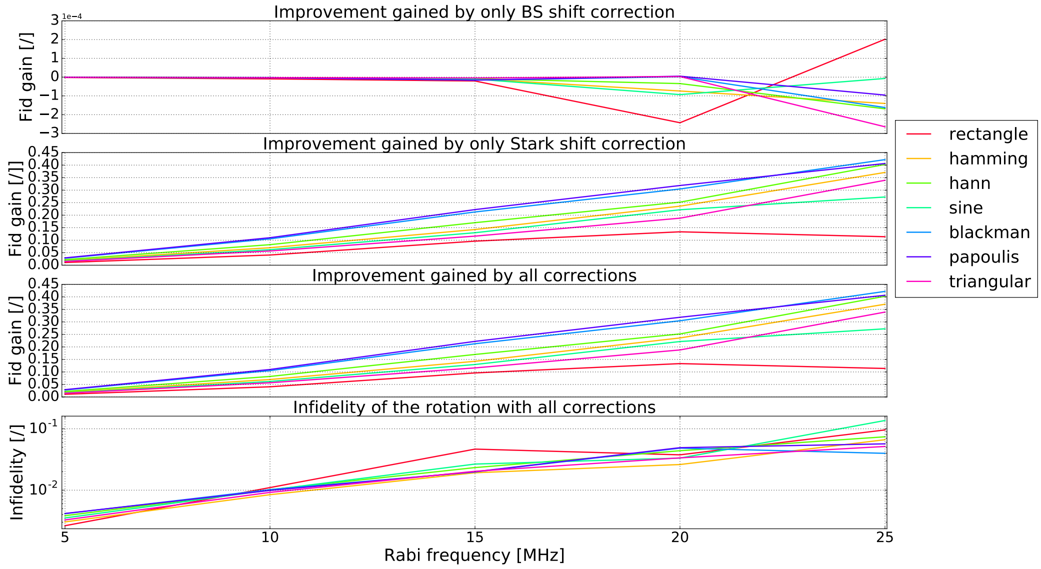

Figure 5.6: Gain of unitary fidelity provided by the frequency shift correction algorithm on geometrical and raised-cosine π pulses in the 3-qubits 3-drives configuration, for different combinations of shift corrections.

First of all, the trends identified in the case of a two-qubits array are also found here, the most important being that the frequency correction algorithm improves the fidelity of the qubits rotations. The most noticeable difference is that the values of the fidelity improvements achieved by the algorithm are larger (sometimes more than twice as high). This can be explained by the fact that having more simultaneous drives induces more frequency shifts. Thus, the rotations benefit more from the corrections. Some configurations exhibit different behaviour (such as the Blackman- and Papoulis-shaped pulses with Ω = 25 MHz, for which the correction actually decreases the fidelity), but these can entirely be explained by unfortunate crosstalk environments specific to these configurations. They consequently do not invalidate the rest of the observed performances, but the existence of these specific bad performing configurations must be kept in mind.

Similarly as to what was discussed for two qubits, the evolution of the improvement with respect to the Rabi frequency is most of the time a linear increase. However, for some shapes (for instance the optimised Blackman and flat top shapes, *cf.* Figures C.14 and C.16 page 82 and 83), the trend seems to saturate at some

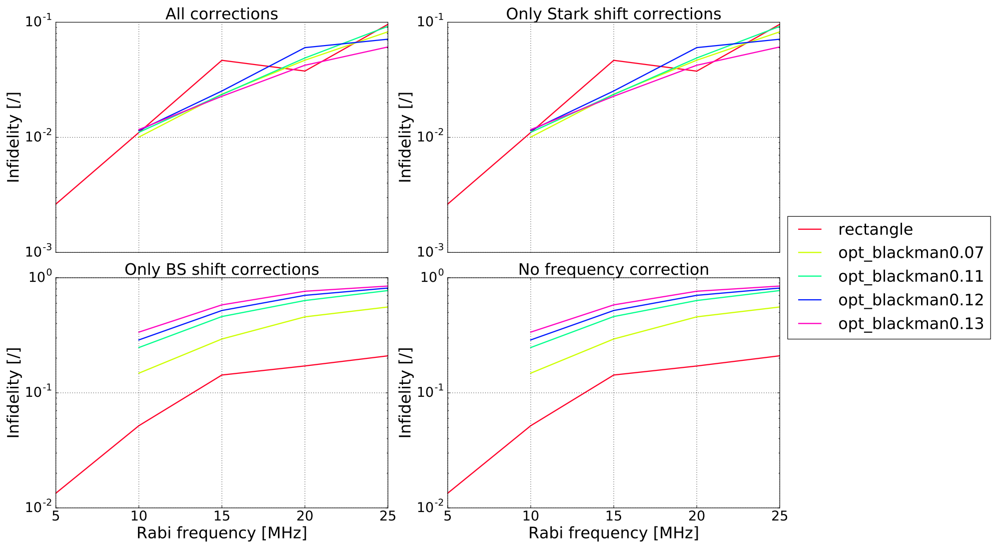

Figure 5.7: Evolution of the unitary infidelity of the π-rotation of 3 qubits driven by an optimised Blackman shaped pulse with the Rabi frequency. Each panel uses a different type of frequency shift correction.

point. It is mostly visible for the results of π-rotations, and in particular for shapes that require long pulse durations. The raw values of the fidelities gives the explanation of this saturation. Using the aforementioned shapes, the fidelity of the rotations without frequency shift correction becomes extremely low for small Rabi frequencies (less than 15% at around $\Omega = 10$ MHz) but does not degrade much further after that. On the other hand, the fidelity with the correction still decreases with increasing Ω due to the increase of crosstalk. This results in this saturation of the fidelity improvement. Figure 5.7 illustrates this for the case of the optimised Blackman modulation. As a rule of thumb, all shapes with very high crosstalk levels at the uncorrected driving frequencies exhibit this feature. This was actually already observed for the HFTxD shapes in the two-qubit array. The drastically low fidelities can be understood by examining the evolution of the qubits on the Bloch sphere, presented in Figure 5.8. In panel 5.8b, which shows the rotation of the qubits using uncorrected drives, the rotations clearly appears to be far from complete for qubits 1 and 3 (the qubits also acquire a phase due to being off-resonantly driven).

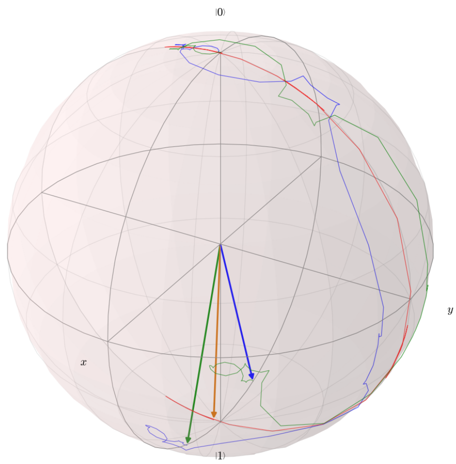

(a) All the frequency shifts are corrected for.

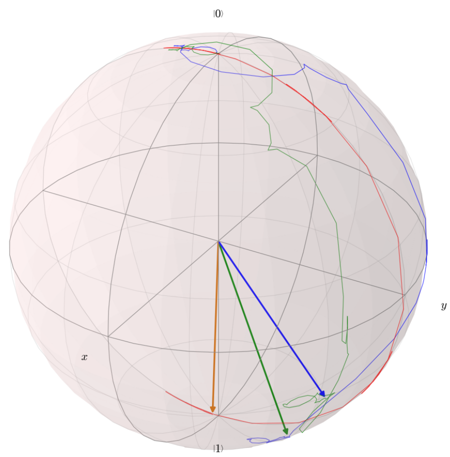

(b) No frequency shift correction is applied.

Figure 5.8: Evolution of the qubits on the Bloch sphere during a π-pulse modulated using the optimised Blackman shape. For convenience, all the qubits are represented simultaneously, but each one is observed in its own rotating frame. The blue, red and green curves correspond to the evolution of the first, second and third qubit respectively. The arrow mark the end state of each one. The backward rotation of the qubits at the beginning and end of the drive is due to the pulse shape starting and ending with negative values (*cf.* Appendix A).

Similarly to the 2-qubits simulations, the degree to which the shapes benefit from the frequency tracking is determined by the width and flatness of their FT main lobe: the wider and flatter it is, the more the shape benefits from the correction.

Regarding the contribution of the individual corrections, the AC Stark shift correction is again largely dominant. Adding the Bloch-Siegert shift correction to it does improve the fidelity a bit further, but similarly to the two qubits case, this contribution could almost be neglected. However, it is worth noting that in some 3 qubits configurations, correcting for only the Bloch-Siegert shift does bring a small improvement (usually inferior to 0.5×10^{-2}%). The most likely explanation for this phenomenon is that in these configurations, the qubits' shifted frequencies coincidentally fall in close to low points of the FT of the corrected pulses. This phenomenon will however not be investigated further because correcting for only one of the shifts is not a realistic use case, as discussed above.

5.2. Pulse shaping

The effect of pulse shaping on the performance of the rotations has been investigated through a so-called grid search, *i.e.* the regular and systematic parsing of a predefined design space. The aforementioned design space is rather simple, as it is defined by only two free parameters: the shape of the pulse and its Rabi frequency (or equivalently its duration). Rabi speed and duration are considered as the same parameter because both of them are connected through Equation 3.6. Furthermore, the initial phase of the pulse is not used as a free parameter. Modifying it only yields differences in the direction of the rotation and does not seem to have significant effects on the fidelity. This also justifies investigating X rotations only, as a Y gate can be implemented just by changing the initial phase of the driving pulse.

The simulation setup is described in Section 2.1, but it is briefly reminded here for ease of reading. The simulations are conducted on linear arrays of qubits. The qubits' resonance frequencies are of the order of 10 GHz, and each one is separated from its neighbours by 100 MHz. Additionally, a residual exchange interaction of rate $J = 100$ kHz is modeled between adjacent qubits. All of them are driven simultaneously by amplitude modulated pulses. In each of the cases discussed below, the driving frequencies are corrected throughout the pulse to compensate resonance frequency shifts (using the algorithm described in Section 3.1).

Just as for the frequency shift correction algorithm, steps have been taken to reduce the impact of spectral sweet spots on the results. Random variations around the nominal resonance frequencies of the qubits have again been used as the actual values for multiple simulations and the results have been averaged.

Furthermore, the fidelity of the rotations is directly related to the number of driving pulses (and hence of qubits), as more pulses will cumulatively produce more crosstalk. Consequently, a separate analysis has been conducted for 2- and for 3-qubits arrays (Subsections 5.2.1 and 5.2.2, respectively). Subsection 5.2.3 then draws trends on the correlation between shape characteristics (both spectral and time-domain) and rotation performance. The case of a 5-qubits array has also been briefly investigated in order to assess the scalability of the outcomes of this study. These 5-qubits results are presented in Section 5.3.

As before, the coming subsections focus on highlighting and explaining the general trends rather than on singling out the best performing shapes. The trends are illustrated and explained using the example of the geometrical and raised-cosine shape families, but features unique to other shapes are also detailed. The complete results of the remaining shapes are given in Appendix D. For ease of navigation, Table 5.2 summarises all the studied configurations and provides references to the figures showing the corresponding results.

5.2.1. Case of a two-qubit array

The performances achieved by shaped pulses on a two-qubits array are described below. Figure 5.9 and 5.11 present the results obtained for $\pi/2$ and π rotations, respectively. The shapes are first compared by modulating pulses with identical Rabi frequencies. To give additional insights, the same shapes are then compared from a fixed duration perspective in Figure 5.13.

The first noteworthy result is that some shapes exhibit strongly oscillating evolutions with respect to the Rabi frequency (*e.g.* rectangle shape in Figure 5.9 and 5.11). These oscillations are due to the previously described “sweet spots" in the Fourier transforms of the pulses. The Rabi frequency determines the duration of the pulse, and consequently also the shape of its Fourier transform. For some Rabi frequencies, the FT of

Configuration		**Figure reference**	
Fixed Rabi frequency comparison		**2 qubits**	**3 qubits**
$\boldsymbol{\pi/2}$-rotation,	**geometrical**, **raised-cosine** & **Gaussian** shapes	Figure 5.9, page 37	Figure 5.14, page 41
$\boldsymbol{\pi/2}$-rotation,	**Kaiser** shapes	Figure D.1, page 89	Figure D.10, page 93
$\boldsymbol{\pi/2}$-rotation,	**Flat top** shapes	Figure D.3, page 90	Figure D.12, page 94
$\boldsymbol{\pi}$-rotation,	**geometrical**, **raised-cosine** & **Gaussian** shapes	Figure 5.11, page 38	Figure 5.15, page 42
$\boldsymbol{\pi}$-rotation,	**Kaiser** shapes	Figure D.2, page 90	Figure D.11, page 93
$\boldsymbol{\pi}$-rotation,	**Flat top** shapes	Figure D.4, page 90	Figure D.13, page 94
Fixed duration comparison		**2 qubits**	**3 qubits**
$\boldsymbol{\pi/2}$-rotation,	**geometrical**, **raised-cosine** & **Gaussian** shapes	Figure 5.13, page 40	Figure 5.16, page 43
$\boldsymbol{\pi/2}$-rotation,	**Kaiser** shapes	Figure D.6, page 91	Figure D.15, page 95
$\boldsymbol{\pi/2}$-rotation,	**Flat top** shapes	Figure D.8, page 92	Figure D.17, page 96
$\boldsymbol{\pi}$-rotation,	**geometrical**, **raised-cosine** & **Gaussian** shapes	Figure D.5, page 91	Figure D.14, page 95
$\boldsymbol{\pi}$-rotation,	**Kaiser** shapes	Figure D.7, page 92	Figure D.16, page 96
$\boldsymbol{\pi}$-rotation,	**Flat top** shapes	Figure D.9, page 92	Figure D.18, page 96

Table 5.2: Summary of the different driving configurations investigated, and where the results are represented.

the pulses are very close to zero at the resonance frequency of the neighbouring qubit, while for some others it is close to the peak of a sidelobe. Steps have been taken to reduce the impact of accidentally hitting these sweet spots on a very specific setting, but it was expected to still observe it when the sidelobes are large, as it is the case for the rectangle shape. This behaviour exists for all the shapes, but as the block shape has among the largest sidelobes, this variation results in large differences in crosstalk levels. This then causes large differences in fidelity. The difference in FT level is illustrated in Figure 5.10, which shows the Fourier transform of two block pulses with different Rabi frequencies. The resonance frequency of the neighbouring qubit is marked by a red vertical line.

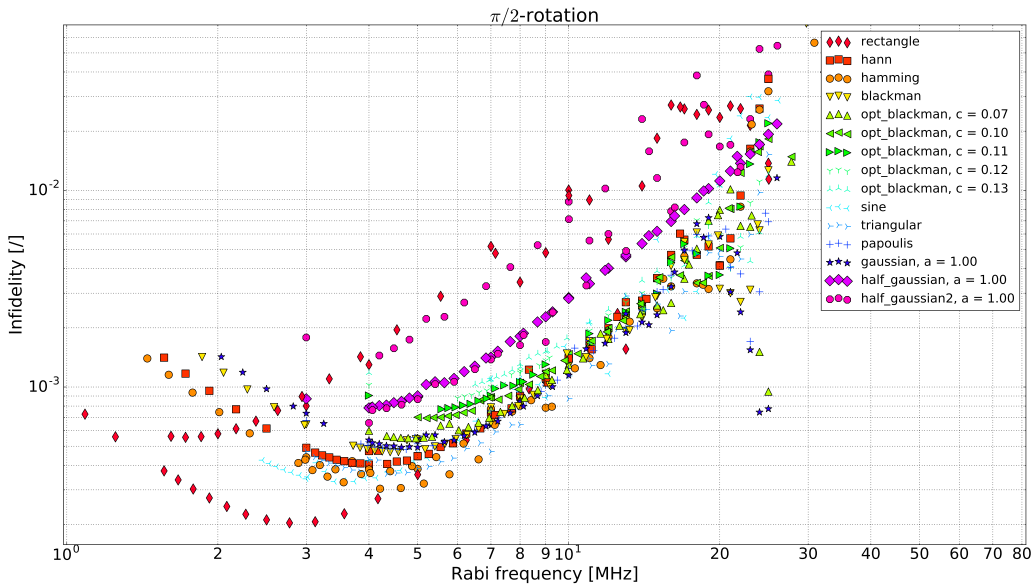

Figure 5.9: Unitary infidelity of $\pi/2$ rotations in two-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and used representative values of Rabi frequency.

The oscillations appear much less pronounced in the case of π rotations (Figure 5.11), but this is mainly due to the actual values at which the fidelity has been computed, between which the variations of FT amplitude are smaller.

The consequence of these large oscillations is that, although the rectangle shape can show very good performance, it cannot be relied upon, as needing to target such a precise resonance frequency is not advisable.

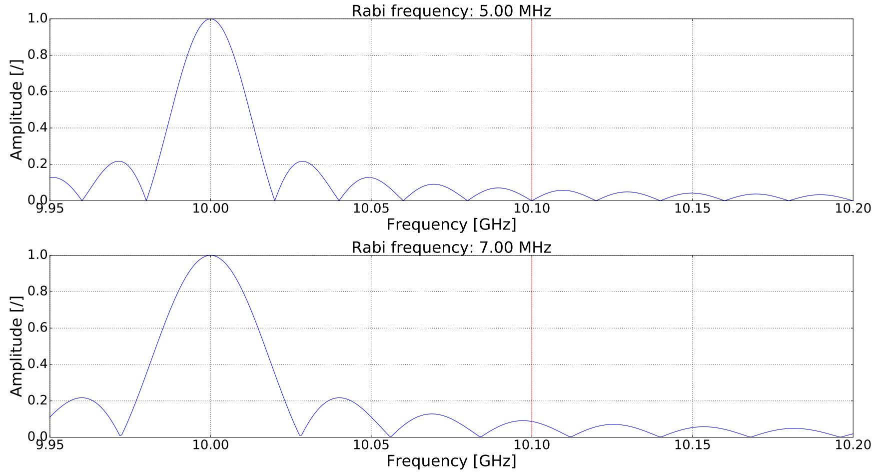

Figure 5.10: Fourier transforms of $\pi/2$-pulses modulated using the rectangle shape, for two values of Rabi frequencies corresponding to high and low performance, respectively. The resonance frequency of the neighbouring qubit is indicated by the red vertical line. It clearly appears that in the top panel the qubit's frequency is very close to a zero of the FT, and in the other it is close to the peak of a sidelobe.

As a consequence, the performance of the oscillating shapes will be considered in terms of average fidelity value in the following.

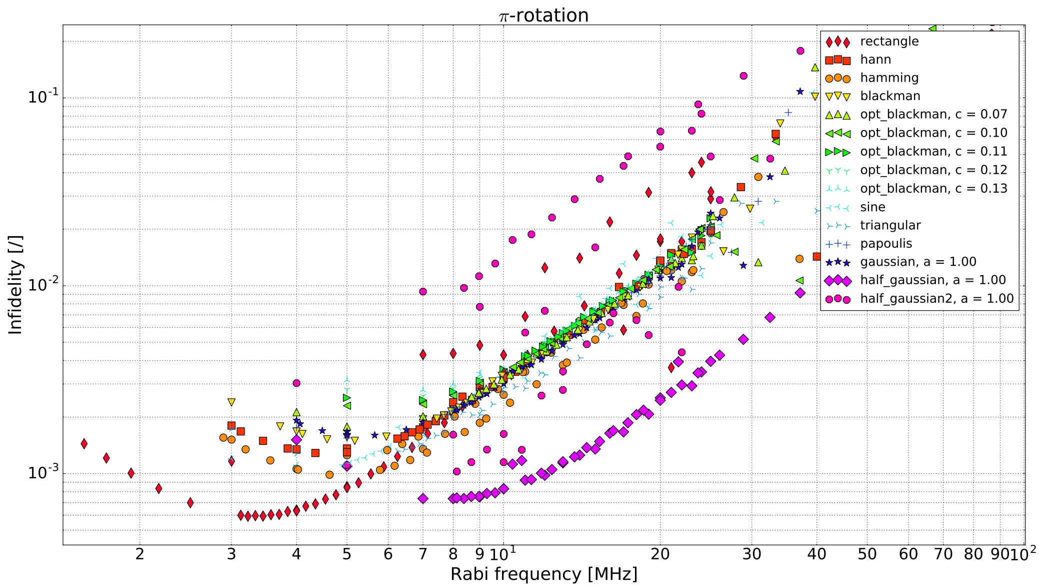

Figure 5.11: Unitary infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and used representative values of Rabi frequency.

When considering the average evolution of the fidelity instead of the oscillatory one, it appears that pulse shaping does improve the achievable fidelity of simultaneous single-qubit rotations. Furthermore, several shapes actually produce pulses with fidelities greater than 99.9%, albeit for different Rabi frequencies. The shapes with high sidelobe levels tend to require their Rabi frequencies to be lower in order to achieve such high fidelity. This makes sense, as lower pulse amplitude means overall lower levels of crosstalk.

In most cases, the fidelity seems higher for $\pi/2$ than for π rotations. This is most likely due to phase accumulated through the rotation because of the remaining off-resonant driving of the shaped pulse. At fixed Rabi frequency, π rotations require longer durations than $\pi/2$ ones, and consequently accumulate more phase. The Z error of the end state can easily be corrected in experiments by digitally changing the origin of phase, which corrects the state fidelity. However it is less straightforward to reflect this correction in the unitary fidelity. Indeed, this metric compares the whole qubit evolution with the ideal rotation, and not only the final state. As a consequence, to correct for this phase over the whole pulse, the unwanted Z rotations would need to be dynamically corrected at each time point, which was not done here. Should such a correction be implemented, the fidelity of π rotations is expected to become higher than (or at least equivalent to) that of $\pi/2$ ones at fixed Rabi frequency. Indeed, π-pulses are longer and consequently have narrower FT main lobes, which results in lower crosstalk levels at the frequency of neighbouring qubits.

The exact order of performance for the modulations varies between the different values of Rabi frequency, but a subset of shapes nonetheless provide higher fidelity. In particular, among the shapes of the raised-cosine and geometrical families, Hamming-, sine-, triangular- and Hann-shaped pulses achieve high fidelity for all investigated rotation angles. For π rotations, the half-Gaussian shape also produces rotations with very high fidelities (much higher than the other π-pulses). These high performance shapes are characterised by having the narrowest FT main lobe of those considered. This is in agreement with what was anticipated: a narrow main lobe means an overall low level of crosstalk at the frequency of the neighbouring qubits.

The large variation in performance of the half-Gaussian shape between the different rotation angles is difficult to understand at first, because Fourier transform arguments do not seem to be applicable here. Indeed, because of the large discontinuity this shape exhibits in the time domain, its Fourier transform has large sidelobes, that do not explain the shape's performances. However, a major difference seems to lie in the phase accumulated throughout the pulse. Whereas for most pulses the $\pi/2$ rotations accumulate less phase than the π ones, it appears to not always be the case for the half-Gaussian shaped pulses. This might explains why the $\pi/2$ half-Gaussian pulses does not perform very well compared to other shapes. Furthermore, the amount of Z error accumulated by the qubits during a half-Gaussian π rotation is amongst the lower ones of all the π-pulses. This might explain its surprisingly good performance compared to the other shapes.

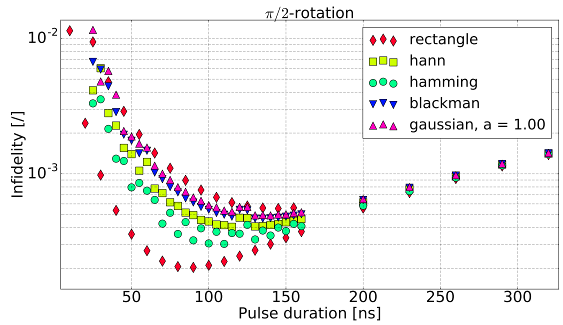

(a) $\pi/2$-rotations, with exchange

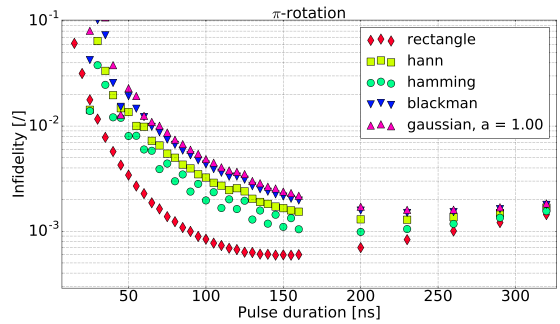

(b) π-rotations, with exchange

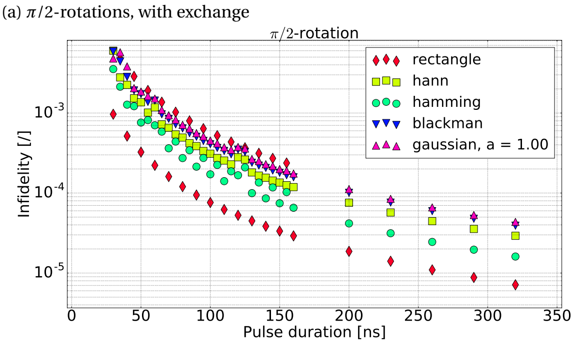

(c) $\pi/2$-rotations, without exchange

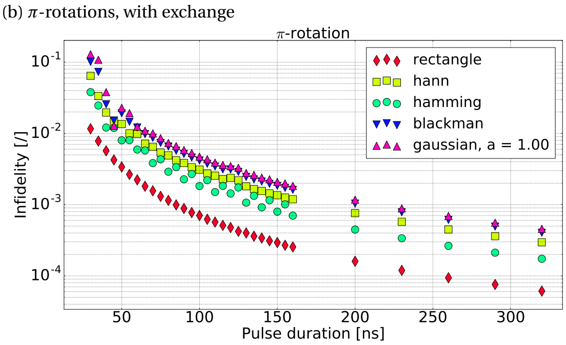

(d) π-rotations without exchange

Figure 5.12: Unitary infidelity of two-qubit rotations over longer durations, modulated using a subset of shapes. The fixed-duration approach has been chosen because the fidelity decrease phenomenon is very easily visible there.

The overall evolution of the fidelity is also interesting. For most values of Rabi frequencies, the fidelity increases with decreasing Ω. This is exactly what was predicted: a lower Rabi frequency produces less crosstalk and so the fidelity increases. However, contrarily to what was expected, a maximum fidelity is observed. Indeed, for very low Ω the trend is reversed: the fidelity decreases with decreasing Rabi frequency. Furthermore, the decreasing fidelities seem to converge towards a value common to all shapes (*cf.* Subfigures 5.12a and 5.12b). This is caused by the exchange interaction between the qubits over long drives. A pulse with a very low Rabi frequency needs to be very long in order to perform the desired rotations. However, during such a long pulse the errors due to exchange coupling between the qubits become dominant over those caused by the crosstalk. Their increase as the pulses get longer results in this fidelity degradation. This is demonstrated in Figure 5.12. It presents the evolution of the unitary infidelity with respect to the pulse duration, with and without interqubit exchange coupling (top and bottom panels, respectively). The shapes are compared in terms of pulse duration because the maximum fidelity appears more clearly in this way. In the evolution without crosstalk, the fidelity keeps increasing with longer pulses, as was initially expected. The exact Rabi frequency (or duration) of the maximum fidelity is different for each shape and corresponds to the point where the crosstalk error becomes smaller than the exchange coupling error. Consequently, the better performances a shape has, the higher the threshold Rabi frequency is.

Lastly, the performance of some shapes seems to sharply decreases for large Rabi frequencies. This is mostly visible for flat top shapes (Figures D.8 and D.9), starting around 15-20 MHz for $\pi/2$ rotations, but around 10 MHz higher for π rotations. These Rabi frequencies correspond to the point at which the FT main lobe of the pulses starts to overlap the frequency of the neighbouring qubit, causing very high levels of crosstalk and consequently greatly reducing the fidelity. This is further confirmed by the fact that this threshold Rabi frequency is higher for the shapes with narrower FT main lobes.

Investigating the results from the perspective of pulses of fixed durations is also insightful, as it corresponds to a different use case. Figure 5.13 shows the evolution of the unitary infidelity of $\pi/2$-rotations, with respect to the duration of the pulses. Similar figures detailing the results of the other shapes and rotation angles can be found in Appendix D. From this perspective, the evolution of the fidelity appears much smoother, and a clear ranking of the shapes by performance is possible. As expected, shapes with the lowest ratio $\Omega/\Omega_{\text{rectangle}}$ perform the best: requiring a lower Rabi frequency to perform a specific rotation means the pulse produces less crosstalk than the others.

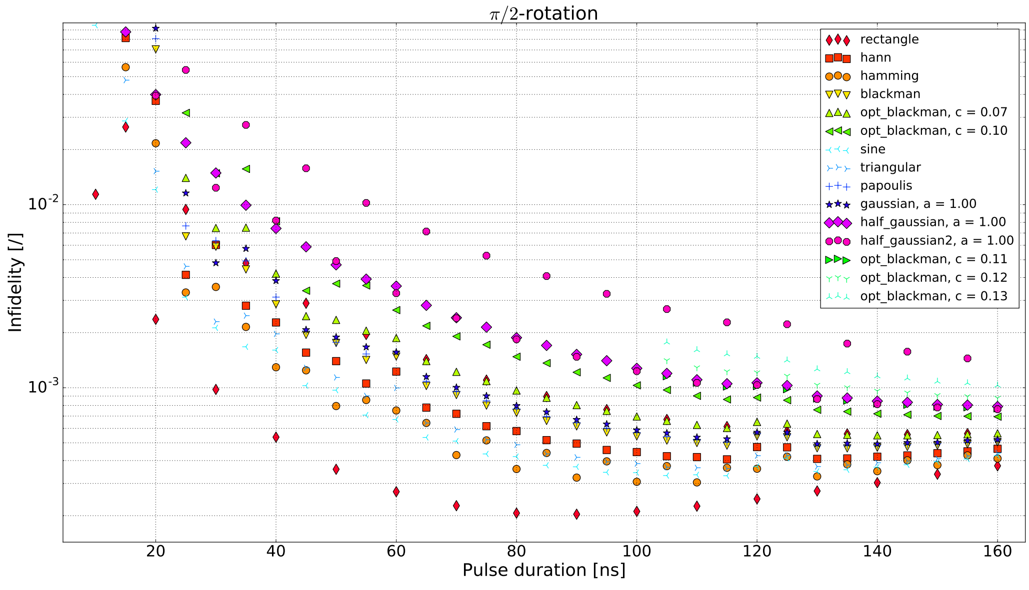

Figure 5.13: Unitary infidelity of single-qubit rotations in two-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and set to have various durations.

Overall, the same conclusions as when comparing the shapes from the point of view of the Rabi frequency apply: fidelity superior to 99.9% is achievable for some shapes, and $\pi/2$ rotations have generally higher fidelities than π ones. Furthermore, the same maximum fidelity as before is logically observed.

The rectangle shape seems to outperform every other one, but its performances show the same oscillatory behaviour as in the fixed Rabi perspective. Thus, the extremely good performance of the shape's sweet spots is unreliable and it is wiser to reason in terms of the average evolution of the fidelity. Doing that, even in this fixed duration approach, pulse shaping does improve the fidelity compared to the rectangular pulse.

5.2.2. Case of a three-qubit array

This subsection describes the performance achieved by shaped pulses on a three-qubits array. As before, the general trends are discussed and explained using the results of the geometrical and raised-cosine shapes as example, but they apply to all shapes unless specified otherwise. Figures 5.14 and 5.15 show the results those shapes obtain for $\pi/2$ and π rotations, respectively, and compare them from a fixed Rabi frequency perspective. Then, Figure 5.16 investigates the same results but for pulses of fixed duration.

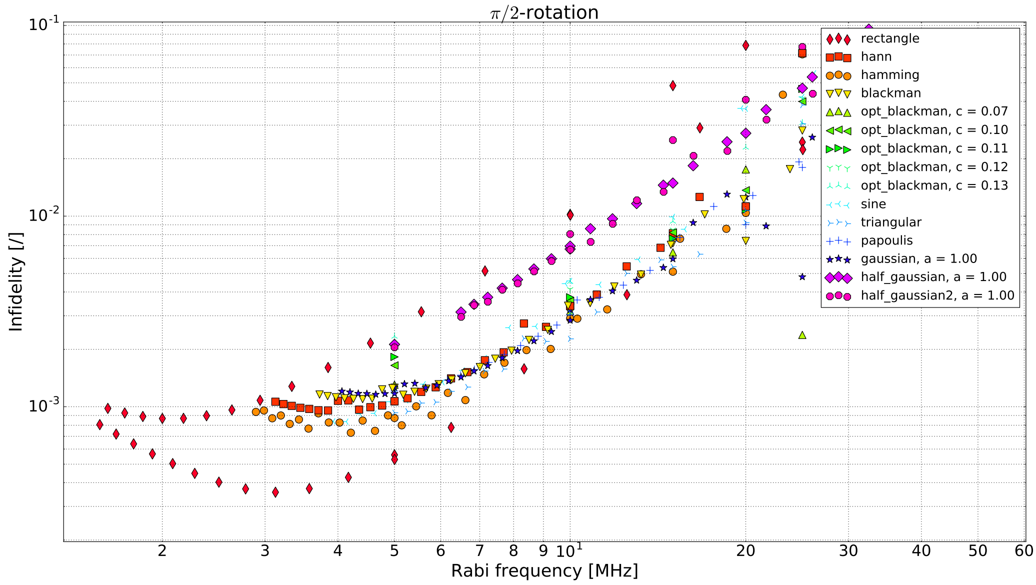

Figure 5.14: Unitary infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and used representative values of Rabi frequency.

The results in a three-qubits array are very similar to those of a two-qubits array. The most notable common trends are summarised in the following. Firstly, the fidelity increases with decreasing Rabi frequency until the exchange coupling errors become larger than the crosstalk-related ones, which decreases the fidelity. As a consequence, there is a maximum achievable fidelity. The Rabi frequency at which this maximum happens is related to the performance of the shape: it is higher for shapes with overall higher fidelity. One can observe that this feature is less visible for three qubits than it was for two. This is due to the fact that, in three-qubits configurations, the crosstalk errors are globally larger, which means that the errors due to the exchange interaction are less visible. This is also the reason why they are less visible on π rotations than on $\pi/2$ ones, regardless of the number of qubits.

Another trend that is common to two- and three-qubits configurations is the oscillatory behaviour of some shapes, caused by the frequency of the neighbouring qubits falling alternatively very close to a zero or to a sidelobe peak in the FT of the drive. This feature is logically most visible for shapes with high and sharp sidelobes like the rectangle.

The last noteworthy shared trend is that the fidelity of the π rotations is lower than that of $\pi/2$ ones, because of phase accumulated by the off-resonant driving pulses each qubit experiences. As the π rotations are longer, they accumulate more phase, which lowers their unitary fidelity.

The three-qubits configuration nonetheless presents major differences. First of all, the achieved fidelities are much lower, to the point where most shapes do not achieve fidelities higher than 99.9%. Having lower fidelities was expected, as using more simultaneous driving pulses means having more intense crosstalk on all the qubits. It is however worth noting that the crosstalk levels in the current study are possibly overestimated. Indeed, in some experimental implementation, the simplifying assumption to consider driving pulses of constant amplitude over space may be unrealistic (see Section 3.2.2 for discussion on this point). The fidelities obtained in such a setup may then be higher. This is also true in the case of a two-qubits array, but because more drives produce more crosstalk, it is expected to be increasingly visible with larger numbers of qubits.

A second major difference is comparative performance of the shapes with one another. In the case of $\pi/2$ rotations, a few shapes achieve fidelities higher than 99.9%. In terms of the overall benefit of amplitude

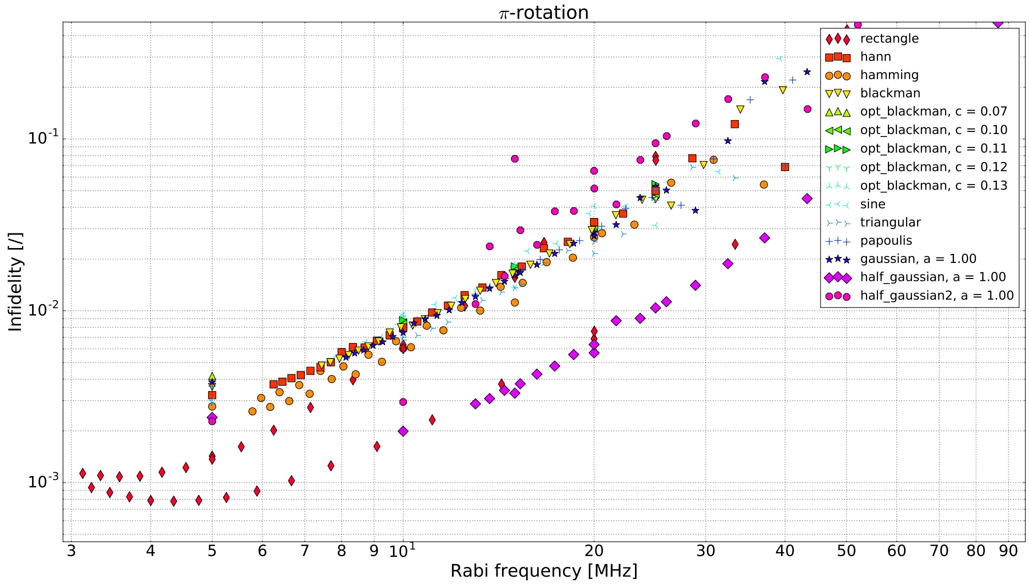

Figure 5.15: Unitary infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and used representative values of Rabi frequency.

modulation, some shaped pulses do achieve higher fidelity than the unshaped one for Rabi frequencies larger than 3.5 MHz. For lower values, the unshaped pulse actually performs the best, despite its oscillating performances. The fidelity of this pulse is indeed higher than 99.9% regardless of its oscillations (see Figure 5.14). Otherwise, the ranking of the shapes closely follows the one that was highlighted in two-qubits arrays. For π rotations, fidelity higher than 99.9% is only achieved by the good oscillations of shapes such as Kaiser 0 and rectangle (see Figures 5.15 and D.11). The performance of the other shapes is greatly degraded compared to two-qubits configurations, while still following the same ranking dependent on the FT main lobe width. Even the half-Gaussian shape, which is the second best performing shape for π rotations at constant Rabi frequency, does not seem to achieve 99.9% fidelity. Furthermore, in the case of two qubits the fidelities of the shaped pulses mostly stand between the two extremes of the rectangle pulse (see Figure 5.11). This is however not the case for three qubits. No satisfying explanation has yet been found for this phenomenon.

The fact that less shapes actually outperform the block pulse is most likely due to a bad combination of increased crosstalk and location of spectral sweet spot, but no definitive conclusion have been reached thus far.

When comparing the results in a fixed pulse duration perspective, it again appears that unshaped pulses outperforms most of the shaped ones. At low durations, the only exceptions are the Kaiser 0 shapes with $\alpha < 1$. Otherwise, some pulses with a duration longer than 100 ns achieve similar fidelities as the rectangle one. This result was expected: considering most shapes perform worse than the unshaped pulses at constant Rabi frequency, it is logical that it is also the case at fixed duration, where the rectangle shape requires the lowest Rabi frequency.

5.2.3. Correlation between shape characteristics and rotation fidelity

This subsection draws trends relating characteristics of the modulation shapes with the fidelity of rotations they produce. Two types of characteristics are considered: spectral and time domain ones. The spectral characteristics are computed from the Fourier transform of the shape, while the time domain ones are measured on its time evolution.

The spectral metrics include:

- the **FT value at the detuning of the neighbouring qubit**. The value is averaged over a small frequency range centered on the exact detuning value, and normalised to the maximum of the FT of the current pulse. The average aims at limiting the impact of coincidentally hitting a zero of the FT, while the normalisation allows to draw trends independent of the actual Fourier transform amplitudes. The conclusions are then more general, and are not undermined by the lack of straightforward correlation between FT amplitude and crosstalk levels.

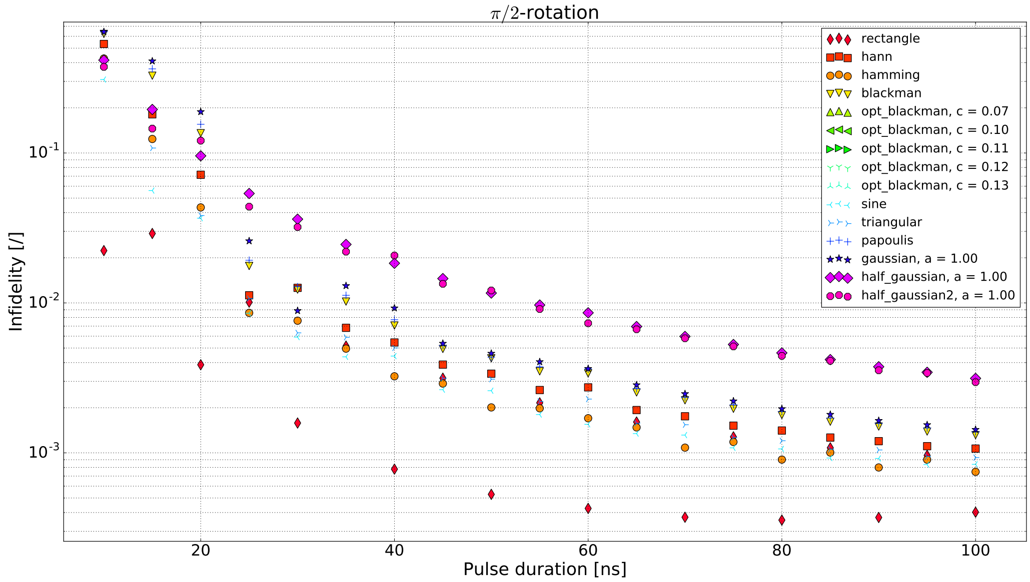

Figure 5.16: Unitary infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and set to have durations.

- the **-3 dB width of the FT main lobe** (in frequency bins, computed as the difference between the two bin numbers at which the FT value is half of its maximum),
- the **spectral flatness of the main lobe** of the energy spectral density (computed over a 400 kHz range centered on the pulse frequency).

The time domain metrics include:

- the **Rabi frequency** of the pulse,
- the **duration** of the pulse,
- the **shape's width at 50% of its maximum amplitude** (expressed in percentage of the the pulse duration so that the value depends only on the shape and not on the specific driving configuration).

The figures in this subsection present the evolution of the unitary infidelity of the qubit rotations as a function of different couples of shape metrics. The metrics span the x- and y-axes of the figures, while the infidelity is represented by the colour of the points. As a general rule, the higher the fidelity, the more blue the corresponding data point is. All figures show the infidelity of $\pi/2$ rotations in a two-qubits linear array, but the trends they highlight are also valid for other rotation angles and higher numbers of qubits being simultaneously driven. The only major difference is logically the fidelity values. As discussed previously, the fidelity decreases with higher numbers of qubits, and $\pi/2$ rotations are more accurate than π ones. For completeness sake, Appendix E regroups the same plots presented here but for π rotations and for 3 qubits. The data points in each of the following figures correspond to pulses of various shapes and of durations ranging from 10 to 160 ns.

Figure 5.17 shows the evolution of the infidelity as a function of the pulse duration and of the Rabi speed. As can be expected based on the definition of the pulse duration (Equation 3.6), the Rabi frequency decreases as the inverse of the duration. Figure 5.17 also illustrates that a combination of long duration and low Rabi frequency produces the rotations with the highest fidelities. This is one of the observations made in previous sections. To some extend, shorter pulses with Rabi frequencies still low (*e.g.* inferior to 10 MHz) also perform well. The factor with the most significant impact of those two then appears to be the Rabi frequency.

Next, Figure 5.18 relates the infidelity of the rotations to the pulse duration and to two spectral characteristics: the width of the main lobe of the Fourier transform (Subfigure 5.18a) and the FT amplitude at detuning Δ = 100 MHz (Subfigure 5.18b). Some interesting insights can be derived from this, which corroborate the

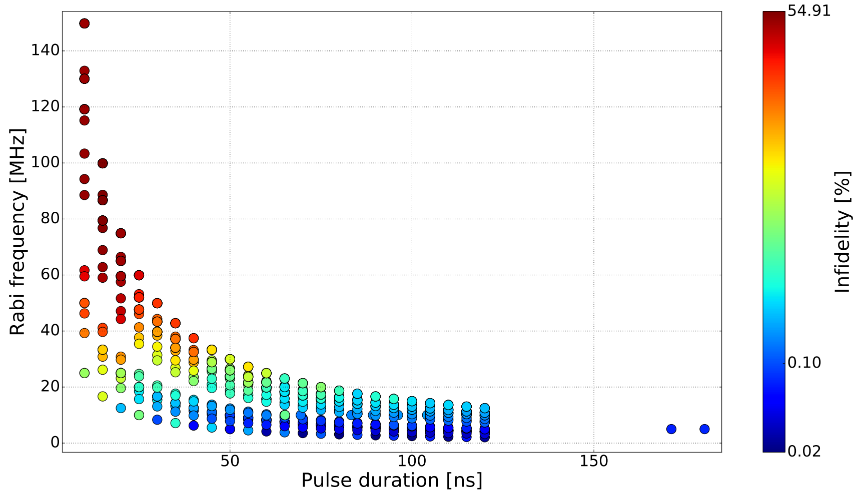

Figure 5.17: Unitary infidelity as a function of the pulse duration and its Rabi frequency.

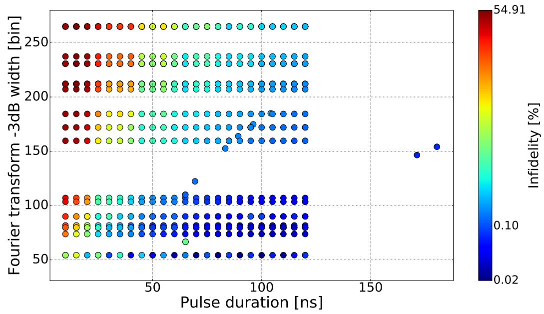

(a) Infidelity as a function of the pulse duration and the width of the FT main lobe.

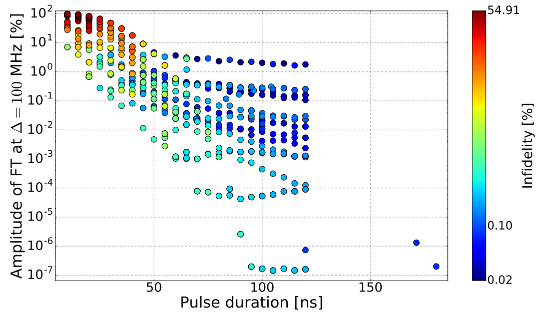

(b) Infidelity as a function of the pulse duration and the FT amplitude at the frequency of the second qubit.

Figure 5.18: Unitary infidelity as a function of the pulse duration and of either the 3dB width of the FT main lobe (panel 5.18a) or of the amplitude of the FT of the pulse at the frequency of the second qubit (panel 5.18b).

conclusions made previously. Note that here we reason in terms of pulse duration as it provides the easiest visualisation, but the correspondence with Rabi frequency is straightforward (*cf.* Figure 5.17 and Eq 3.6).

The data in Subfigure 5.18a shows clear horizontal lines. Each of them corresponds to a different modulation shape (the width of the FT main lobe is constant when expressed in frequency bins, as it is the case here, but not when expressed in units of frequency). One can clearly see that the highest fidelities are achieved by pulses with long durations and, at constant duration, by those whose Fourier transform have a narrow main lobe. This is expected since a narrow FT main lobe usually means that the crosstalk level at the frequency of the neighbouring qubits is lower. Furthermore, a narrow main lobe also usually corresponds to a longer pulse, which in turn requires a lower Rabi frequency. Such a pulse then induces lower crosstalk levels. As a side note, notice that the oscillatory behaviour of the rectangle shape also appears here (lowest horizontal line at a width of approximately 50 frequency bins).

On Subfigure 5.18b, we can see two interesting features. First of all, longer pulses generally have lower FT amplitude at the detuning of the neighbouring qubits. This observation is logical, as longer pulses have more narrow FT main lobes, as well as lower Rabi frequencies. Those two characteristics lead to lower FT amplitudes. The second feature is that the highest fidelities are obtained by shapes that have both low FT amplitudes at $\Delta = 100$ MHz and long durations. Moreover, at constant and low FT amplitude, the highest fidelities are achieved by the longest pulses. It is worth noting that the highest fidelities are not necessarily achieved by the points with lowest FT amplitude. These points in most cases actually correspond to configurations where the exchange interaction becomes the largest error source, leading to fidelities lower than what one would expect.

Further insight can be gathered from Figure 5.19, which links infidelity with the width of the time domain pulse at half of its maximum amplitude and the amplitude of the Fourier transform at $\Delta = 100$ MHz. Similarly to the -3dB width of the Fourier transform, the width at half amplitude of the pulse is constant for each shape, because it has been normalised to the pulse duration. This allows to clearly identify shapes and to draw trends. The rectangle shape is for instance easily recognisable by its half-amplitude width of 1.

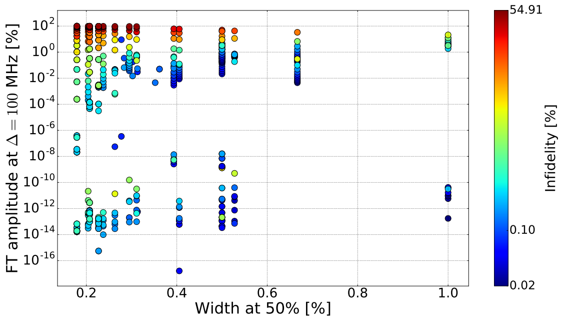

(a) Complete figure, regrouping all data points.

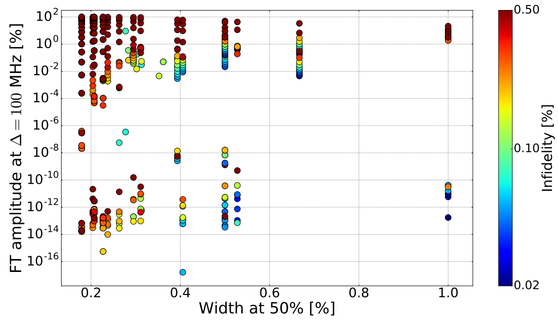

(b) Focus on high fidelity regions. The color scale is modified for better visualisation.

Figure 5.19: Unitary infidelity as a function of the pulse width at 50% and of the Fourier transform amplitude of the pulse at the frequency of the second qubit. The left panel presents the complete figure, and the right one focuses on the high fidelity regions. Consequently, the colour scale in the left panel is chosen to cover all encountered infidelity values, but it is modified in the right one to better differentiate high fidelity points.

First of all, some values of time domain width show two distinct clusters of points. These belong to the shapes that present oscillatory behaviours, corresponding to alternatively high and low FT amplitudes at the second qubit. This is for instance the case of the rectangle or Hamming shapes (at width of 1 and 0.5, respectively). These particular groups of points are more tricky to exploit, but overall it appears again that, for each shape cluster, the highest fidelities correspond most of the time to the lowest FT amplitude at the second qubit's frequency, which in turns corresponds to lower crosstalk levels. This observation is perfectly reasonable and agrees with the findings of Section 5.2. It is worth noting that for some values of time domain width, the highest fidelity is however not achieved at the lowest FT amplitude. This is due to the exchange coupling errors becoming dominant over the crosstalk-induced ones in long pulses, as highlighted in the previous section.

Additionally, the time domain width also has an influence. Indeed, as is especially visible from Subfigure 5.19b, at constant FT amplitude, pulses with wider time domain envelopes seem to produce rotations with higher fidelities. It can most likely be related to the duration or Rabi frequency of the corresponding pulses. However, a satisfying explanation has yet to be proposed.

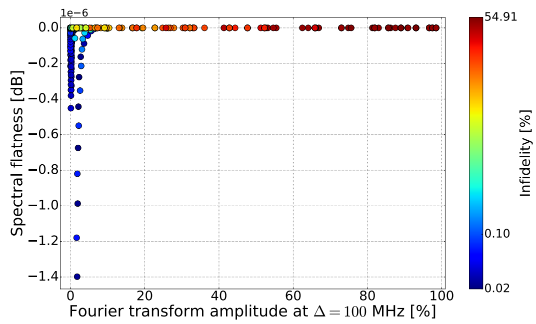

(a) Complete figure, regrouping all data points.

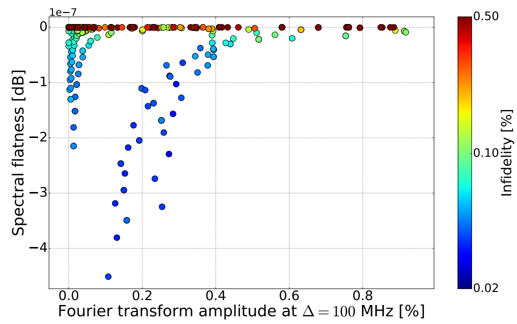

(b) Focus on high fidelity regions. The color scale is modified for better visualisation.

Figure 5.20: Unitary infidelity as a function of the Fourier transform amplitude at the frequency of the second qubit and of the spectral flatness. The left panel presents the complete figure, and the right one focuses on the high fidelity regions. Consequently, the colour scale in the left panel is chosen to cover all encountered infidelity values, but it is modified in the right one to better differentiate high fidelity points.

It can also be interesting to examine the combined influence of the spectral flatness and FT amplitude at the frequency of neighbouring qubits. This is done in Figure 5.20. As before, the left panel presents the whole figure while the right panel focuses on the high fidelity regions. Once again, as expected, shapes with low FT amplitude at the resonance frequency of the second qubit appear to drive rotations with the highest fidelity. However, Subfigure 5.20b shows that there is actually a limit to the previous statement. Indeed, the highest fidelities are actually obtained for configurations with FT amplitudes ranging between 0.1 and 0.35% of their maximum value. Below this range, the exchange coupling errors actually become larger than the crosstalk-induced ones, which lowers the fidelity.

It can furthermore be observed that the fidelities are higher in regions characterised by spectral flatness further away from 0 dB. The further the spectral flatness is from its maximal value of 0 dB, the less flat the shape is over the considered frequency range. Consequently, it means that shapes with flatter main lobes rotate the qubits with lower fidelity. This makes sense, as Fourier transform main lobes with a flat top tend to also be wide. And as we have seen previously, wider main lobes usually mean lower fidelities. While shapes with flat FT main lobes can prove more robust against phenomena such as frequency shifts or decoherence, this advantage thus needs to be balanced with their overall lower fidelities.

Lastly, let's focus on the correlation between FT amplitude at the neighbouring qubit's frequency, width of the FT main lobe and rotation infidelity. Figure 5.21 illustrates this relation. As usual, Subfigure 5.21a and 5.21b present the whole figure and a focus on the high fidelity regions, respectively.

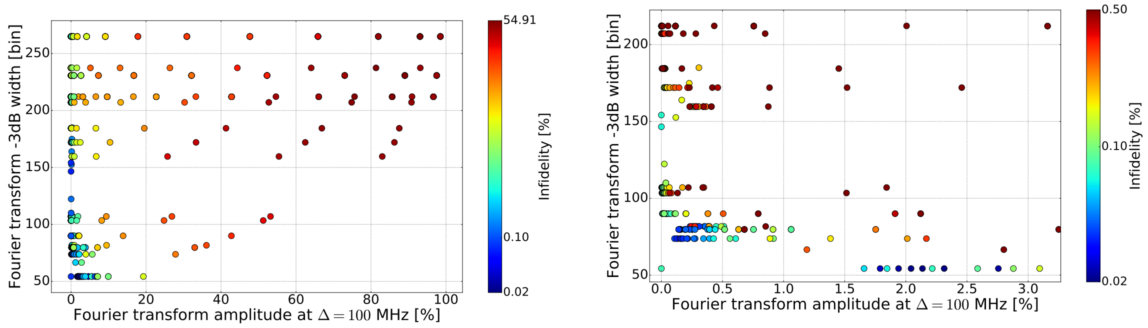

(a) Complete figure, regrouping all data points.

(b) Focus on high fidelity regions. The color scale is modified for better visualisation.

Figure 5.21: Unitary infidelity as a function of the Fourier transform amplitude at the frequency of neighbouring qubits and of the -3dB width of the Fourier transform main lobe. The left panel presents the complete figure, and the right one focuses on the high fidelity regions. Consequently, the colour scale in the left panel is chosen to cover all encountered infidelity values, but it is modified in the right one to better differentiate high fidelity points.

In that case, the highest fidelities are obtained by shapes with both low FT amplitude at detuning $\Delta = 100$ MHz and narrow FT main lobes. This is in perfect agreement with the results highlighted previously and, along with the other trends highlighted in this subsection, forms an interesting set of guidelines for pulse shape selection and/or design.

To summarise the findings of this section, the selection of a shape that performs simultaneous single-qubit rotations in multi-qubit environment should be done based on the following recommendations:

- A small -3dB width is usually better, because it generally leads to low levels of crosstalk at the frequency of neighbouring qubits.
- pulses with long durations and low Rabi frequencies produce rotations with the highest fidelities, as longer pulses naturally have narrower FT main lobes, and low Rabi frequency means overall lower FT amplitudes (and hence less intense crosstalk). The factor with the largest influence seems to be the Rabi frequency, as even pulses with rather short durations but low Rabi speed have been shown to achieve high fidelity. Furthermore, care must be taken to not select shapes with too long durations as in this case the exchange interaction may become the main error source instead of the crosstalk. This typically starts to appear for pulses longer than 120 ns, but the threshold is different for all shape.
- In case two shapes require similar durations, the one requiring the lowest Rabi frequency is to be preferred, as having a lower Rabi frequency leads to overall lower FT amplitudes, and hence to less intense

crosstalk on the other qubits (independently of the width of the FT main lobe).

- Quite logically, at constant FT main lobe width, configurations with the lowest FT amplitude at the frequency of the neighbouring qubits are to be preferred, since they induce the lowest levels of crosstalk.
- At constant pulse duration, the optimal flatness values over the 400 kHz range centered on the pulse frequency seem to lay between -1.10^{-7} and -3.10^{-7} dB. Shapes with very narrow main lobes (spectral flatness closer to 0dB) can yield high fidelities, but they may not be very robust against decoherence. On the other hand, shapes with very flat FT main lobes are best to avoid, because said main lobes are also usually wider, which contradicts the previous recommendations and produces more intense crosstalk on the neighbouring qubits.

5.3. Scalability

This section analyses the applicability of the trends highlighted previously to higher numbers of qubits. It does so by presenting the results of a brief analysis of the performance of the fidelity improvement techniques on a five-qubits array. The investigation is only conducted over the shapes which achieved satisfying results in two- and three-qubits configurations. The performance of the frequency shift correction algorithm is investigated first (Subsection 5.3.1), followed by that of pulse shaping (Subsection 5.3.2).

The methods and details of the performance evaluation were described in Sections 5.1 and 5.2, and are not modified in this Section. They will thus not be repeated here.

5.3.1. Frequency correction algorithm

The performance of the frequency shift correction algorithm over the selected shapes is presented in Figures 5.22 and 5.23, for $\pi/2$ and π rotations respectively. Their three upper panels show the fidelity gain obtained by using different types of shift correction. The bottom panel displays the values of the infidelity of each configuration when both the AC Stark and Bloch-Siegert shifts are corrected.

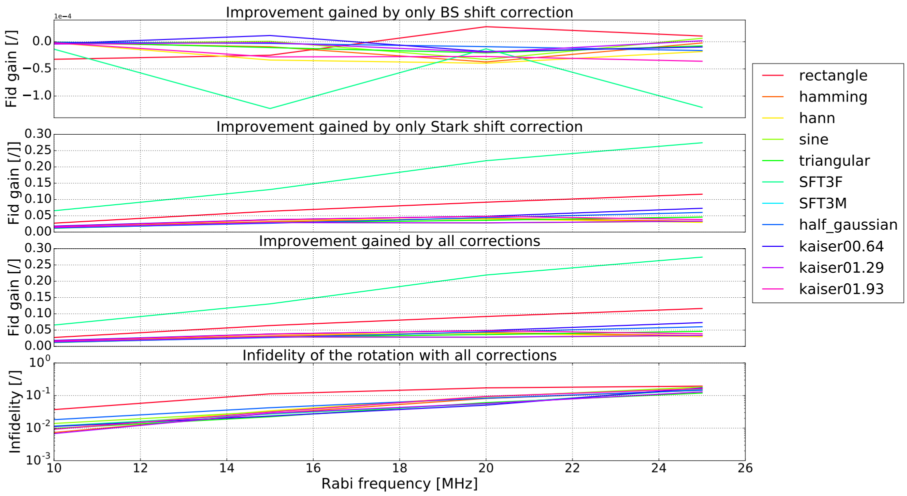

Figure 5.22: Gain of unitary fidelity provided by the frequency shift corrections on $\pi/2$ pulses in a 5-qubits array, for different combinations of shift corrections. The shapes investigated are the best performing ones identified in Section 5.2.

First of all, the fidelity gains in the case where both shifts are corrected for (3rd panel from the top in Figures 5.22 and 5.23) is positive. This shows that correcting for all the shifts concurrently continues to bring an improvement to the unitary fidelity of the rotations, be they $\pi/2$ or π. Moreover, the same trends that were highlighted in Section 5.1 are visible here. As expected, each shape's fidelity gain still increases along with the Rabi frequency. Furthermore, the fidelity increase in each driving configuration keeps getting larger with higher numbers of qubits. For instance, when correcting for all the shifts, the $\pi/2$ rectangle pulse of Rabi frequency 20 MHz gained only 0.01% of fidelity in the two-qubits configuration, but gained 0.06 and 0.09% in three- and five-qubits setups, respectively.

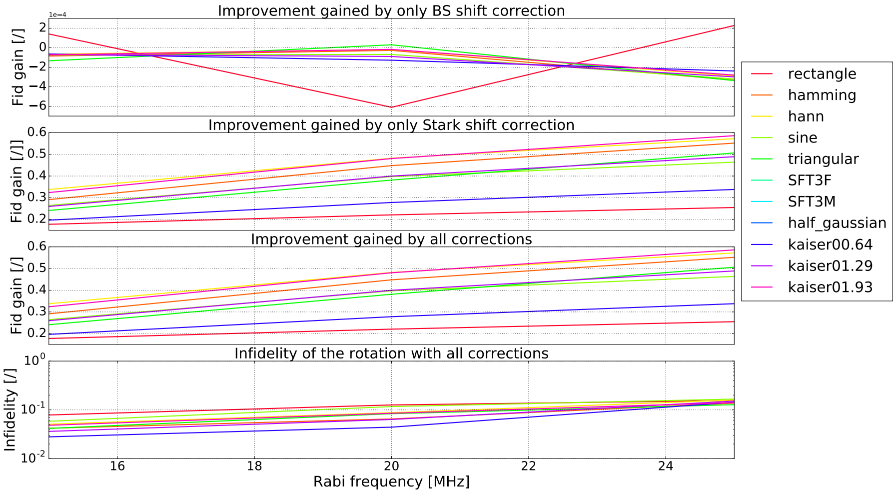

Figure 5.23: Gain of unitary fidelity provided by the frequency shift corrections on π pulses in a 5-qubits array, for different combinations of shift corrections. The shapes investigated are the best performing ones identified in Section 5.2.

Regarding the improvement provided by each type of correction, correcting for the AC Stark shift alone still achieves much higher improvements than correcting for only the Bloch-Siegert effect. Nonetheless, adding the Bloch-Siegert shift correction does bring further improvement.

Lastly, the same saturation phenomenon that was observed for three qubits is present here. Indeed, for larger Rabi frequencies, the increase in improvement slows down. It actually even starts to decrease for some shapes, like the $\pi/2$ Kaiser 0 pulse with $\alpha = 1.93$ (pink curve in Figure 5.22). This phenomenon is however more pronounced here than in the three-qubits configuration. For instance, it starts to be visible for the Hamming shape (orange curve in Figures 5.22 and 5.23), while it was clearly not the case in configuration with less qubits. This is coherent with the the more intense crosstalk. Indeed, the saturation is caused by the uncorrected fidelities being very low but remaining fairly constant with respect to the Rabi frequency, while the corrected fidelities decrease. When considering five drives, the uncorrected fidelities reach somewhat constant values faster and thus the gain brought by the correction starts reducing earlier.

However, there are also some differences in behaviour. Indeed, in contrast to what it was with less qubits, the evolution of the fidelity gain of $\pi/2$ rotations with respect to the Rabi frequency appears almost linear. This would logically be caused by the overall higher levels of crosstalk: because they are so high at the qubits frequencies, the configuration specific variations between corrected and uncorrected drives are not visible anymore. This results in the disappearance of the outliers.

Another change is noticeable in the order in which the shapes benefit from the correction: the width and flatness of the pulse's FT main lobe appears to not be the only factor anymore. This change is due to the previously discussed saturation phenomenon. As long as it is not there, the order in which the shapes appear is the same as for lower numbers of qubits. However, when the Rabi frequency is high enough, the fidelity gain of the shapes with wider FT main lobes (and thus which produce higher crosstalk intensity) starts to diminish, which disturbs the previously established order.

The evolution of the fidelity gain provided by frequency shift correction on $\pi/2$ pulses, with respect to the number of qubits, is summarised in Figure 5.24. A similar figure for π rotations can be found in Appendix C. The increase of the fidelity gain with the number of qubits is clearly visible here, just as well as the increase with the Rabi frequency.

5.3.2. Pulse shaping

The selected modulation shapes are here compared using pulses of identical Rabi frequencies. Figures 5.25 and 5.26 present the unitary infidelity of shaped $\pi/2$ and π pulses, respectively, for various values of Rabi frequency. As before, increasing the number of qubits decreases the overall fidelities of the rotations. As a consequence, no meaningful configuration exceeds 99.9% of fidelity. However, it is worth repeating that the levels of crosstalk in this study may in some experimental setups be overestimated, and that in these configurations, the fidelity is expected to be higher.

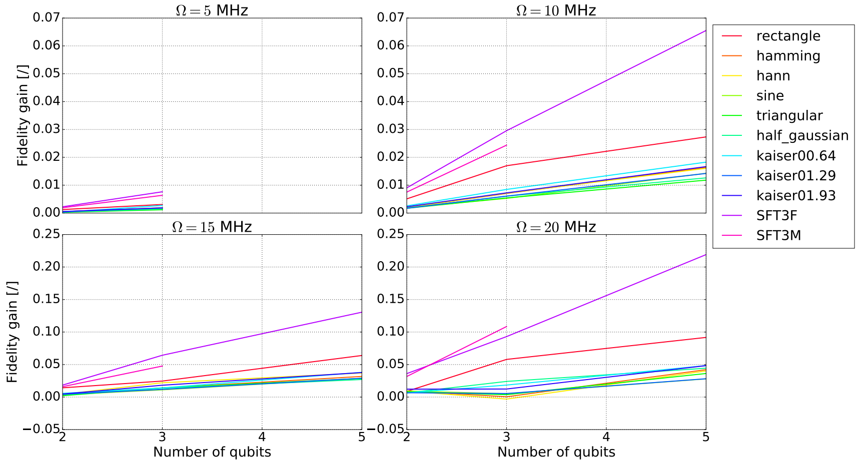

Figure 5.24: Evolution of the fidelity gain provided by frequency shift correction on $\pi/2$ pulses, with respect to the number of qubits. The evolution is presented for four representative values of Rabi frequency. Some values have been omitted, because very computationally intensive to obtain.

In addition, $\pi/2$ rotations still achieve higher fidelities than π ones, due to the off-resonantly driven qubits accumulating phase.

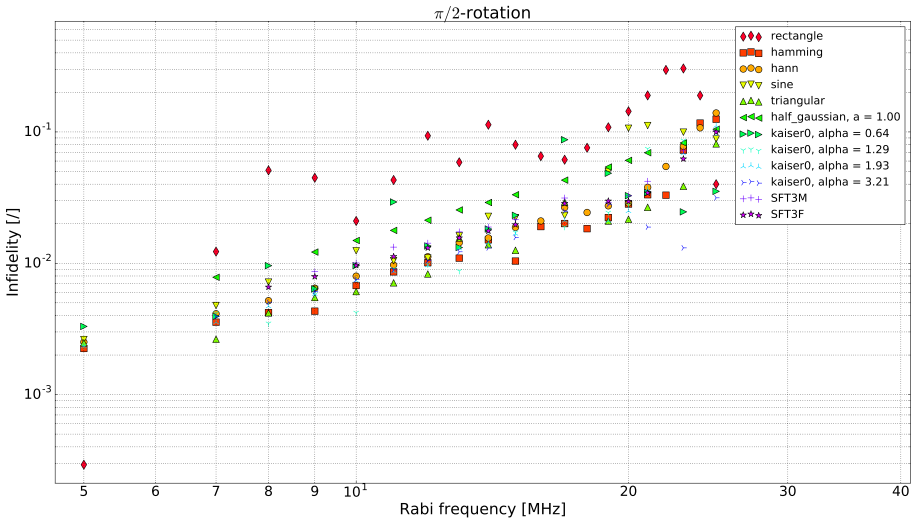

Figure 5.25: Unitary infidelity of $\pi/2$-rotations in five-qubits configurations. The driving pulses are modulated by the best performing shapes identified in the previous subsections, and used representative values of Rabi frequency.

Contrarily to the 3-qubits configuration, shaped pulses clearly outperform unshaped ones regardless of the rotation angle. This is related to the evolution of the performance of unshaped pulses. It appears that their oscillations are much slower and their average infidelity is much higher. Consequently, the evolution does not benefit from extremely high values anymore and is on average worse than that of shaped pulses. The much smaller oscillations themselves are likely due to the overall increase of the crosstalk intensity caused by all the pulses in the setup: the difference in FT amplitude between Rabi frequency is smaller than with less pulses, which results in more regular (and bad) performances. This is clearly visible in Figure 5.25.

The order of performance of the shape is determined once again by the width of the FT main lobe of the pulses: shapes with narrow FT main lobes generally achieve higher fidelities. Consequently, the conclusions highlighted for two- and three-qubits seem to scale to higher numbers of qubits.

It is worth noting that the maximal achievable fidelity is not visible here, and likely occurs at much lower Rabi frequencies. This is coherent with what was observed before. Because the fidelities are lower in this configuration, the crosstalk errors remain dominant over those caused by exchange coupling for the considered range of Rabi frequency.

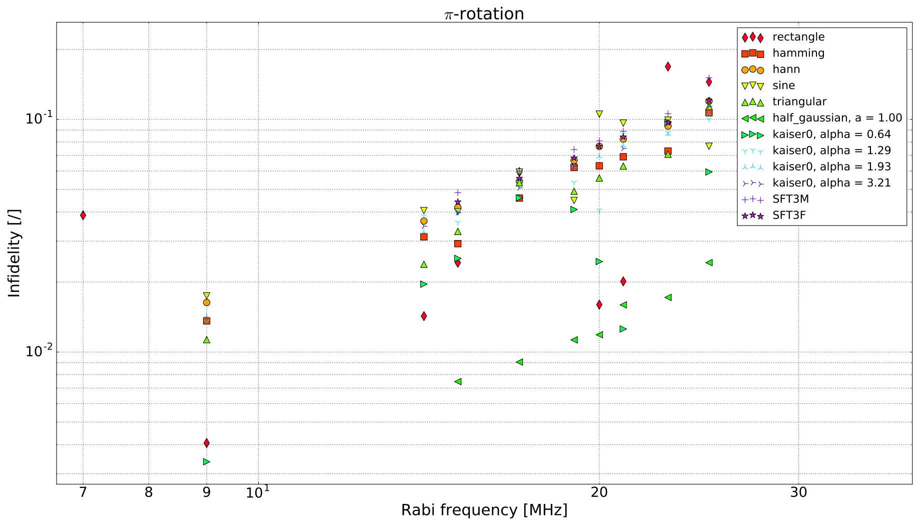

Figure 5.26: Unitary infidelity of π-rotations in five-qubits configurations. The driving pulses are modulated by the best performing shapes identified in the previous subsections, and used representative values of Rabi frequency.

For more clarity, the evolution of the fidelity of $\pi/2$ rotations with respect to the number of qubits is summarised in Figure 5.27. The same evolution but for π rotations can be found in Section D.3 of Appendix D.

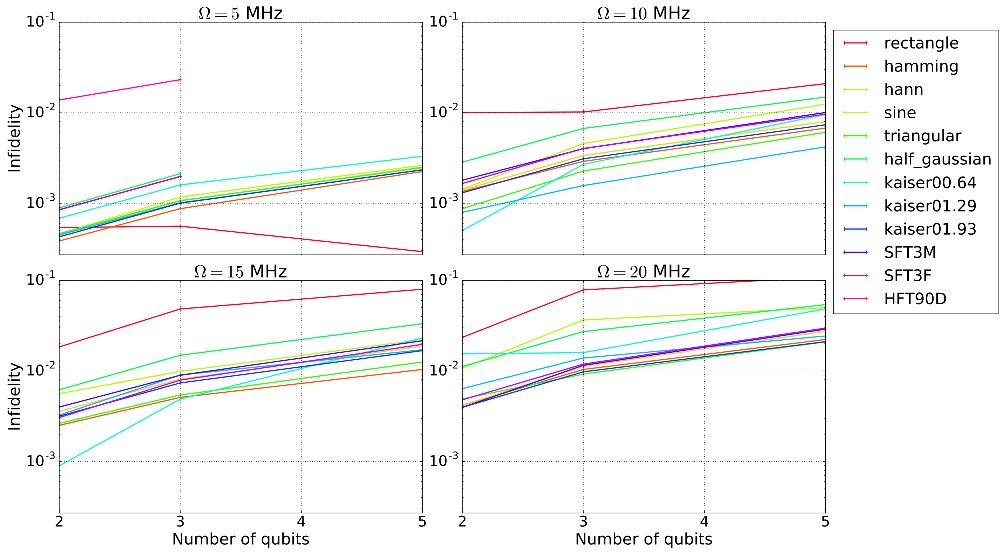

Figure 5.27: Evolution of the unitary infidelity of shaped $\pi/2$ rotations with respect to the number of qubits, for four representative values of Rabi frequency. The shapes investigated are those that achieved satisfactory performance with two and three qubits.

It appears very clearly that shaped pulses outperform unshaped ones, regardless of the number of qubits. The fact that the difference is much clearer here than in previous figures is due to the choice of the Rabi frequencies. With the exception of $\Omega = 5$ MHz, they are configurations where the neighbouring qubits' frequencies fall close to peaks of the FT sidelobes. The $\Omega = 5$ MHz configuration exemplifies the case where they fall close to zeros of the Fourier transform.

The similarities in the order of performance of the shapes is also visible: shapes with the narrowest main lobes achieve the highest performance. It is for instance the case of the Hamming, triangular and Kaiser 0 ($\alpha = 1.29$) shapes.

Lastly, Figure 5.27 effectively summarises the capability of the shapes to achieve fidelities greater than 99.9%. Keeping in mind that the crosstalk levels in this study may be overestimated compared to some experimental setups, it nonetheless shows that pulse shaping is able to achieve high fidelity where the unshaped pulses is not.

6
Conclusions and discussions

This section provides the detailed conclusions that follow from the results presented in the previous chapters. It finally reflects on certain aspects and assumptions of the work and discusses their relevance.

6.1. Conclusions

The answers to the research questions introduced in Section 1.1 are summarised here. First, each sub-question is addressed individually, which then leads to a complete answer of the principal question which has guided this work.

- **Which phenomena prevent the implementation of high fidelity single-qubit rotations in multi-qubit systems, and how do they affect gate fidelity?**
  Microwave pulses in multi-qubit environments generate unwanted driving on the qubits, as they are affected by pulses targeted at their neighbours. This phenomena gets increasingly detrimental with the number of driven qubits (due to the growing number of simultaneous pulses). While the smallest in terms of error value, this phenomenon is very significant because it prevents the fidelity from reaching 99.9% in many cases (see Section 5.2).

  In addition, the off-resonant microwave pulses couple to the spins and effectively dress them, shifting their resonance frequencies. Two effects have been identified and modeled, namely the AC Stark and Bloch-Siegert shifts. They are respectively caused by the co- and counter-rotating components of the driving fields. Due to these shifts, the driving pulses are no longer resonant with the qubits, and consequently cannot drive them correctly. This phenomenon has been shown to be responsible for the largest errors (see Section 2.2.3).

  Finally, the residual exchange coupling has been found to cause errors when the qubits are driven slowly. These are visible because they become dominant over the low crosstalk errors generated by long pulses. In other cases (*e.g.* short pulses), these errors are negligible compared to the crosstalk-induced ones. The magnitude of these errors as well as their impact depends on the intensity of the residual exchange. A lower one would induce smaller (if any) errors of our considered experiment parameters.

- **Which techniques can be applied to mitigate the impact of these phenomena?**
  The unwanted driving produced by off-resonant pulses on qubits can be mitigated by modulating the amplitude of the time envelope of the drive. This modifies the spectral characteristics of the pulse and, if the modulation shape is chosen appropriately, it can reduce the energy transmitted at the frequency of the other untargeted qubits. In a decoherence-free environment, pulse shaping has been shown to increase the fidelity of pulses of identical Rabi frequencies by up to 0.95% compared to the unshaped drives. Because the unshaped pulses already achieve fidelities close to 99% in certain driving configurations, an improvement of 0.95% is actually significant.

  Furthermore, it has been demonstrated that the impact of coupling-induced frequency shifts on the driving scenarios considered here can be almost entirely negated using a simple optimisation method. To this end, the pulse is first divided in 1 ns slices. For each of these slices, initially off-resonant drive frequencies are selected so that the shifted qubits are resonantly driven. The process is repeated for each time slice to account for the changing shifts throughout the pulse. The corrected drives have been shown to achieve very similar unitary fidelities as rotations where no frequency shift is simulated at all. This corresponds to a fidelity improvement of up to 50% over the shift-inducing but uncorrected drives. The larger improvement compared to pulse shaping is due to the fact that the frequency shifts make the drives off-resonant, which makes the target rotations unachievable.

Mitigating steps for the exchange errors have not been investigated in this work, but a simple way to limit their impact is to favour driving pulses of reasonably low durations.

- **How can the aforementioned techniques be fine-tuned to provide the highest possible fidelity?**
  The performance of the frequency shift correction algorithm on real experimental implementations could be improved by including frequency shifts that have not been covered in this work (*e.g.* correlation between Rabi frequency and qubit frequency induced by spin-orbit coupling, effects of micromagnet), or by refining the models of those that have been. The only implementation effort required for this is the model modification, the algorithm in itself need not be altered. Taking into account other, less significant, frequency shifts will make the drive frequency match that of the qubit even more closely in a real experiment.

  Regarding pulse shaping, several guidelines have been identified for the choice of the shape. First of all, pulses whose Fourier transform have a narrow main lobe have been shown to lead to the highest fidelities, independently of the other parameters.
  A second recommendation is to favour shapes that require long pulses and low Rabi frequencies in order to perform a rotation. Care must however be taken to not select a Rabi frequency too low, as then the exchange interaction may reduce the fidelity. A minimum ratio of Rabi frequency over residual exchange Ω/J between 30 and 55 is advised, but this value can be lowered if the obtained fidelity is too low. Furthermore, in case several options requiring similar durations are available, the one with the lowest Rabi frequency is to be preferred.
  If further selection is needed, the Fourier transform constitutes a decent analysis tool. It does not accurately represent the dynamics of the spins in the configurations explored in this work, but it nonetheless provides decent qualitative insights into it. In order to decide between two otherwise similar drive configurations, the FT amplitude of the pulses at the frequency of the other qubits can be used as a discriminating factor. Indeed, configurations with lower amplitudes tend to produce rotations with higher fidelities.
  Lastly, shapes whose FT main lobes have very high spectral flatness are best avoided, as they lean towards generating more crosstalk. As an indication, the highest fidelity rotations have been achieved for shapes with a spectral flatness ranging from -3.10^{-7} to -1.10^{-7} dB (computed over a 400 kHz interval centered on the drive frequency).

How can we design driving strategies to perform high fidelity single-qubit rotations on multiple qubits simultaneously?
This work has provided promising driving design guidelines to achieve high fidelity single-qubit rotations on multiple qubits simultaneously. The recommended driving strategy is based on two techniques. The first one is the correction of the frequency shifts induced by the coupling of the pulses with the qubits. The second is to use amplitude modulation to reduce the crosstalk the pulses produce on the qubits they are not targeting.

The detrimental effect of frequency shifts on the accuracy of the rotations has been shown to be almost entirely negated by appropriately choosing and shifting the pulse frequencies so that they match the shifted ones of the qubits throughout the pulse. A time slice length of 1 ns has furthermore proved sufficient to follow the shifting frequencies closely enough to have a satisfactory correction. Furthermore, the proposed algorithm can straightforwardly be adapted to account for other frequency shifting phenomena.

The positive effect of using shaped pulses envelopes has also been clearly demonstrated. Providing general numerical guidelines is not straightforward, as the optimal characteristics are heavily dependent on the qubit configurations. However, this work highlighted general qualitative recommendations. In particular, a shape whose Fourier transform main lobe is narrow and sharp is likely to produce the rotations with the highest fidelities. In addition, the best performances have been observed for shapes requiring durations between 90 and 140 ns for both $\pi/2$ and π rotations in a two-qubits system, all the while keeping a Rabi frequency as low as possible. This aims at generating the lowest possible levels of crosstalk while not being significantly affected by exchange-induced errors. As an indication, these requirements would for instance recommend shapes similar to the Hann, Hamming, sine or Kaiser (wih low α) modulations. The conclusions have lastly been demonstrated to scale to higher numbers of qubits.

Using these techniques, fidelities larger than 99.9% have been shown to be achievable (in the simplified, decoherence-free context of this simulation). It however gets increasingly difficult when increasing the number of qubits in the array.

6.2. Discussions

This subsection goes over some important modeling decisions or simplifications that have been made during this work, as well as some future improvement points. Each is briefly described, and its relevance as well as potential impact are discussed, when applicable.

- The simulations conducted in this work have not included decoherence. The principal motivation for this decision was that the necessary features were not yet available in the DM_solver library at the time the simulation was set up. As a result, the fidelities obtained in this work are not representative of what experiments may yield. The real fidelities are indeed expected to actually be lower. Most of the trends and guidelines highlighted in this study are however expected to remain valid. The only ones that could vary are those regarding the recommended flatness of the FT main lobe of the pulses. Indeed, the higher robustness to decoherence of shapes with flatter tops may compensate the higher crosstalk amplitudes they tend to induce. Nevertheless, it is important to note that all the considered shapes have been shown to be already quite flat over the expected linewidth of a real qubit (*e.g.* 100 kHz). The advantage of shapes with very flat tops may then not be as significant as initially expected. As a consequence, all the trends and guidelines currently highlighted may actually remain valid.

- In order to reduce the implementation effort and be able to test the method on a slightly simpler case, an important type of frequency shift has been omitted from the analysis. Indeed, during the driving of the spins, an effective time-dependent magnetic field arises due to spin-orbit coupling, proportional to the alternative electric field and to the static magnetic field. (Debald and Emary, 2005, Flindt et al., 2006, Golovach et al., 2006). As a consequence, the resonance frequencies of the qubits actually vary proportionally to the Rabi frequency of the drive. However, taking this phenomenon into account in the frequency shift correction algorithm should be very straightforward: it should only be a matter of updating the computation of the shift to include this new phenomenon. More generally, this would also be the case for any other effect causing shifts in frequency.

- Additionally, a simplifying assumption was made when modeling the driving pulses. It was decided that the pulses would have a constant amplitude over space. Depending on the experimental implementation of the qubits, this may actually be realistic (*i.e.* setups where all qubits are driven using a single waveguide (Takeda et al., 2020)). Otherwise (*i.e.* qubits are driven by separate waveguides), the amplitude of each pulse at the location of the qubits it is not targeting is expected to be lower. Consequently, the crosstalk and frequency shifts experienced by all the qubits are likely to be lower than what was presented in this work. Higher achievable fidelities could then be expected. However, this should not significantly influence the general outcomes of this study, nor the guidelines for designing high fidelity driving schemes.

- Phase errors during the pulse have been shown to lower the fidelities in some cases. In particular, they are likely responsible for the fact that the performance of π rotations is significantly lower than that of $\pi/2$ ones. Because the frequency shifts of the qubits were compensated by directly modifying the frequency of the pulse, the rotation actually leaves the ideal rotation plane during the driving. If using the state fidelity, the remaining Z error at the end of the rotation can very easily be removed by modifying the origin of phase. However, the unitary fidelity is computed on the unitary of the whole rotation and not only on its final state, so making such a correction *a posteriori* is not as straightforward. However, this phase error could easily be corrected by implementing the frequency shift correction as a variation of the phase of the pulse instead of its frequency.

- No precise value has been given for the crosstalk amplitude above which the unitary fidelity of the rotations decreases below 99.9%. This is due to the fact that no practical tool has been proposed that accurately represent the spin dynamics. The analogue used in this work is the Fourier transform, which at best gives decent qualitative insights into them. A more precise tool would nevertheless be very useful. It would indeed allow us to better understand why shapes perform well or not, as well as to predict the performance of a driving pulse. It would also enable us to directly compute the effective Rabi frequency perceived by a qubit in a given configuration, accounting for all the crosstalk, and therefore to precisely tailor the driving strategy to minimise the effects of crosstalk.

- Lastly, the analysis has showed that an important factor in the performance of the driving methods proposed by this study is the rotation angle. As a benchmark, $\pi/2$ and π rotations have been investigated. However, quantum computing processes (such as for instance quantum phase estimation (Nielsen and Chuang, 2011)) may require rotations of other angles, and especially smaller ones. Consequently, it is interesting to discuss the expected applicability of our findings to such angles. First of all, in the case of smaller angles, a more precise analysis than what was performed here is possible. Indeed, for low rotation angles, the spin dynamics can be approximated to be linear. In that case, the Fourier transform is actually an accurate estimate of the spin evolution (Freeman, 1998). As a consequence, it then becomes possible to precisely predict the performance of a driving configuration, but also to tailor it on each qubit to minimise the produced levels of crosstalk. Furthermore, assuming the pulse duration remains fixed for all target rotation angles, a smaller angle requires a much slower drive. As a consequence, the performance of small rotations is expected to be higher than what is shown in this work. The guidelines for choosing the drive configurations that have been highlighted are expected to remain the same, but to be less relevant. Indeed, as the drive can be made slower, the overall crosstalk intensity at the frequency of the other qubits is expected to greatly decrease. However, as a consequence of the lower crosstalk errors, the exchange-induced ones may become relevant at shorter durations. Concerning rotations of larger angles, the reverse is expected. Still assuming the pulses have a duration similar to those of the $\pi/2$ and π rotations, the guidelines of this study should be followed more stringently in order to still achieve high fidelities.

7

Recommendations for future work

The results of this thesis work have highlighted some points of interest which still remain to be explored. They are presented here, and are divided in two categories. The first one contains short scope improvements to the methods and software, with the objective to further improve the performance of what has already been implemented in this study. The second proposes more extensive work to further validate the presented methods and to enable more thorough analyses.

Below are the short scope recommendations:

- Phase errors during the pulse have been shown to lower the fidelities in some cases. This is due to the rotation deviating from its ideal plane during the drive because of the time-dependent pulse frequency. These Z errors can easily be suppressed by implementing the frequency shift correction as a variation of the phase of the pulse rather than of its frequency. In the current software implementation, the phase of the driving pulses is a free input parameter. Consequently, modifying the algorithm to implement this type of correction should be straightforward: instead of changing the frequency at each time slice, the algorithm simply has to update the phase. This modification is expected to greatly reduce the phase-related errors and thus to make the fidelities of π and $\pi/2$ rotations at least similar. The same kind of improvement is of course expected for other rotation angles as well. This method of correction is furthermore more akin to what is currently done in experimental setups.
- The missing correlation between Rabi speed and qubit frequency could also be modeled and included both in the simulation and, more importantly, in the frequency shift correction algorithm. The former would make the simulation more realistic, while the latter would allow a more precise selection of the drive frequency for experimental settings. As a result, the so-corrected rotations could theoretically achieve higher fidelities.

The suggestions given in the following are of larger scope and may require more substantial effort to implement. However, their pay-offs are also expected to be higher and could greatly improve upon the analysis conducted in this work.

- The performance of both the frequency shift correction and the pulse amplitude modulation have been analysed using the Fourier transforms of the drives. However, it does only provide qualitative or over-estimations of the dynamics of the spins, and so cannot be used for precise analysis. Consequently, the development of a more precise estimation tool would be greatly beneficial to the design of high fidelity driving strategies. Indeed, it would then be possible to predict the performance of a pulse, but also to compute the effective driving strength received by a qubit, which would encompass all the crosstalk intensity. This would in turn allow to fine tune the drive on all the qubits to minimise the crosstalk on each of them and thus maximise the fidelity.
- In addition to the insight given by this accurate analytic tool, it would be interesting to evaluate the impact of decoherence on the techniques studied here. The more accurate simulations would better predict the fidelities that different driving pulses are susceptible to achieve. It would furthermore lead

to a more educated discrimination of the shapes based on their duration and the flatness of their FT main lobe.

- In order to broaden the scope of the insights and guidelines provided by this study, it could be interesting to evaluate the performance of both pulse shaping and frequency shift correction on two dimensional arrays of qubits. This type of configuration is a promising solution to scale the number of qubits on a chip to numbers relevant to tackle complex problems (Meunier, 2019). As it presents distinct characteristics and challenges from 1D arrays (*e.g.* different crosstalk profile for each qubit and more exchange couplings to take into account), this analysis is expected to provide valuable insights.

- Finally, the performance of the schemes and guidelines described in this work could be tested using a randomised benchmarking experiment (Knill et al., 2008, Magesan et al., 2012) on multiple qubits at the same time. This would prove that our techniques are efficient and would also allow to quantify how much the rotation fidelities vary with and without the corrections.

A

Definition of the considered pulse shapes

This appendix introduces and describes the different pulse modulation shapes used in this work. For each of them, an analytical expression is given, and some of its spectral characteristics are discussed. All the pulses are defined here to perform a similar qubit rotation. The exact rotation angle is decided by the value of the pulse length τ (*c.f.* Section 3.2.3).

The shapes presented here are described in more details in Heinzel et al. (2002), Prabhu (2013), Salvatore and Trotta (1988), Bauer et al. (1984), Warren (1984) and Geen and Freeman (1991).

A.1. Geometrical shapes

A.1.1. Rectangular

The rectangular (or block) pulse is the simplest, non-modulated pulse. It is given here as a mean of comparison and is studied in the same way as the other shapes to ensure that pulse shaping actually improves the gate fidelity as we expect it does.

The shape is defined in Equation A.1, where τ is the pulse length.

$$s(t) = \begin{cases} 1, & \text{if } |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.1}$$

A rectangular-shaped pulse of frequency $\omega = 10.0$ GHz, as well as its Fourier transform (FT) are plotted in Figure A.1. In this appendix, the Fourier transforms are normalised to the maximum amplitude of the block pulse, for ease of comparison between them. In Figure A.1b, high sidelobes and slow sidelobe decrease rate are clearly visible, which suggests a poor spectral selectivity and high crosstalk. No precise threshold above which the crosstalk is deemed too high will be given, as drawing a correlation between the amplitude of the FT and the level of crosstalk is not a trivial task and has not been the focus of this work.

Looking at the inset of panel A.1b, which presents a zoom on the low intensity components of the FT, we can in particular notice the amplitude at a detuning of $\Delta = 100$ MHz from the pulse (which, as was discussed in Section 2.1.2, would correspond to the nominal resonance frequency of a second qubit). At exactly $\Delta = 100$ MHz, the magnitude of the sidelobe is very close to zero, which would result in an extremely small amount of parasitic driving on a neighbouring qubit. However, since the slopes of the sidelobe are sharp, it is expected that a real qubit with a resonance frequency subject to variations and shifts will not be driven as well as one could predict from this figure. To take this into account, in the following we will reason on the FT value averaged over a frequency interval ranging from 15 MHz below to 15 MHz above the detuning of the neighbouring qubit. These averaged values are summarised in Table 3.2, for various representative configurations. For the sake of clarity, in the remainder of the section we may refer to this average value as simply the value at detuning 100 MHz. In this case, averaging the FT value brings it very high. Consequently, the fidelity of the driving produced by this configuration is likely to be rather low. This highlights the fact that such spectral "sweet spots" cannot be relied upon to select appropriate driving configurations, as these features are only significant for a very narrow range of frequencies and there would be no way to ensure a qubit would fall precisely within it

The same kind of reasoning can be made regarding the peaks of the FT. At these specific frequencies, the crosstalk levels are expected to be very high and thus the fidelity of the driving to be rather low, but for shapes with high sidelobes this can change drastically with a very small frequency shift. Furthermore, it is worth noting that the frequency of the peaks (as well as that of the spectral “sweet spots") is dependent on both the frequency and the duration of the pulse. For instance, a pulse with a longer duration will see the frequency of its peaks shifted down.

It is important to note that while it is likely that the rectangular shape will not be amongst the best-performing shapes when comparing pulses of identical Rabi frequency, it is the shortest possible pulse with unit amplitude. This means that when comparing drives of the same duration, this shape will require the lowest Rabi frequency of all the ones considered here, and hence is likely to produce the lowest frequency shifts and crosstalk.

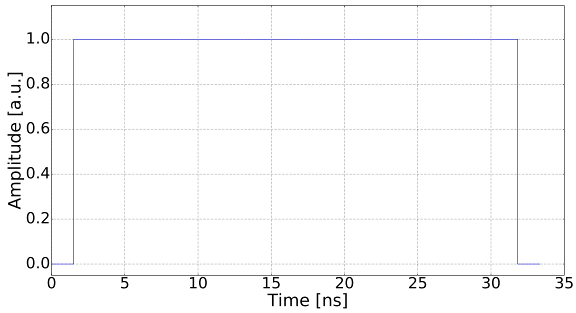

(a) Time-domain evolution of the shape, with normalised amplitude.

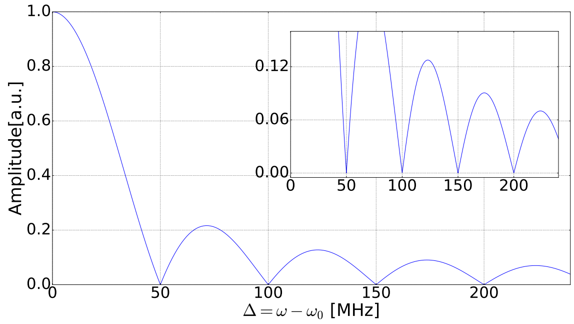

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.1: Rectangular modulation shape

A.1.2. Triangular

The triangular pulse (also known as the Bartlett modulation shape) is the first available compromise to reduce sharp edges at the beginning and end of the pulse, which will reduce the frequency leakage induced by the non-infinite pulse.
This shape is described in Equation A.2, where τ is the pulse length.

$$s(t) = \begin{cases} 1 - \frac{|t|}{\tau}, & \text{if } |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.2}$$

A triangular-shaped pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz is plotted, as well as its FT, in Figure A.2. We can notice in the right panel (FT of the pulse) that, although the sidelobe decrease rate is similar to the one of the rectangle pulse, the sidelobe levels are overall much lower. The main lobe also appears to be wider.

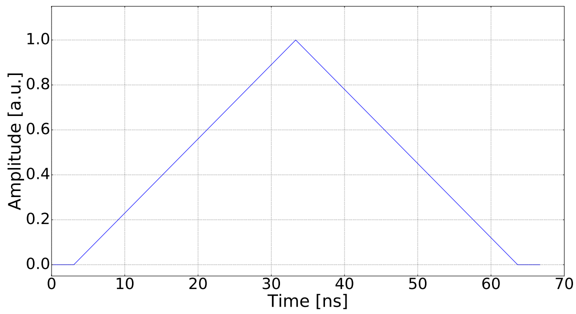

(a) Time-domain evolution of the shape, with normalised amplitude.

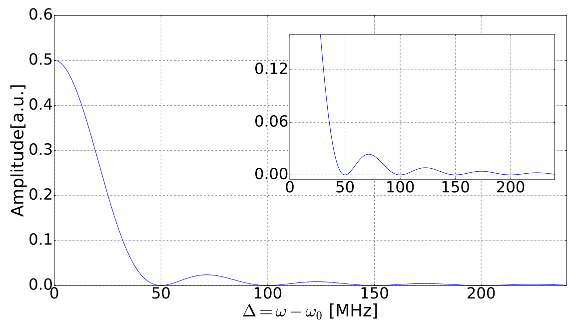

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.2: Triangular modulation shape

Similarly to the block pulse, the amplitude of the sidelobes at a detuning of $\Delta = 100$ MHz from the pulse frequency is extremely close to zero. When taking the average over an interval ranging from 75 to 115 MHz, the value is a little higher than for the rectangle shape but remains very low. It is however worth noting that the intensity of the drive does not vary as fast as for the rectangle pulse when getting away from the “sweet spot", so this feature could be very visible in a real qubit.

A.1.3. Sine

The sine modulation shape is classical in digital signal processing. Its expression is given in Equation A.3.

$$s(t) = \begin{cases} sin\left(\frac{\pi t}{\tau}\right), & if|t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.3}$$

A sine-shaped pulse of frequency $\omega = 10.0$ GHz and amplitude $\Omega = 25$ MHz is plotted, as well as its Fourier transform, in Figure A.3.

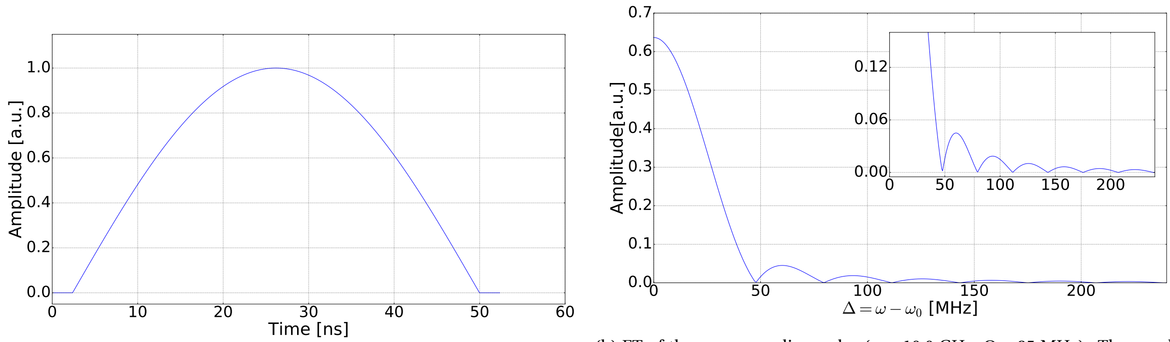

(a) Time-domain evolution of the shape, with normalised amplitude.

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.3: Sine modulation shape

This shape presents low sidelobes and shows a correct sidelobe decrease rate (see panel A.3b). Since for this pulse configuration the intensity at a detuning of 100 MHz is close to the peak of a sidelobe, this pulse will produce a higher amount of parasitic driving. However, other configurations with slightly longer pulse durations may bring the zeros of the FT closer to $\Delta = 100$ MHz. In these cases, the fidelity of the rotations may still be satisfying.

A.2. Generalised-cosine shapes

A.2.1. Hann

The Hann shape significantly reduces the sidelobes compared to a block pulse, but the counterpart is a much wider central peak. Its expression is as follows (Equation A.4):

$$s(t) = \begin{cases} 0.5 + 0.5\cos\left(\frac{\pi t}{\tau}\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.4}$$

A Hann-shaped pulse of frequency $\omega = 10.0$ GHz and amplitude $\Omega = 25$ MHz is plotted, as well as its FT in Figure A.4.

The main lobe of the FT indeed appear much wider than that of the block pulse, but it is also overall of lower amplitude. Additionally, the low sidelobes are clearly visible in panel A.4b. The resonance frequency of the first neighbouring qubit (10.1 GHz) corresponds once again to a very small driving intensity. This is still the case when considering the average over an interval of 30 MHz around the nominal frequency, so this pulse configuration is likely to drive the qubits with high fidelity, even in the case of noise- or coupling-induced shifts of the resonance frequencies.

A.2.2. Hamming

The Hamming shape can be thought of as an optimized version of the Hann shape, which minimises the first sidelobe level (*i.e.* the power density of the first sidelobe). It is given by the following equation (A.5)

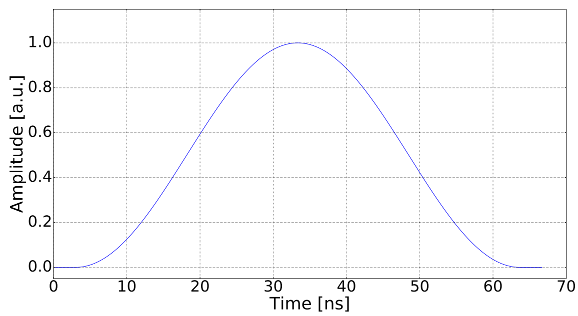

(a) Time-domain evolution of the shape, with normalised amplitude.

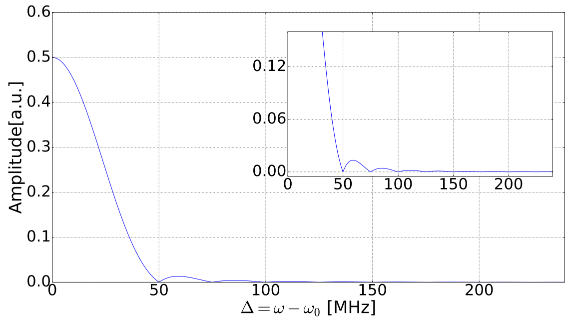

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.4: Hann modulation shape

$$s(t) = \begin{cases} 0.54 + 0.46\cos\left(\frac{\pi t}{\tau}\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.5}$$

A Hamming-shaped pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz is plotted, as well as its FT, in Figure A.5. We can see that indeed the first sidelobe level is lower than that of the Hann shape, but as a counterpart the sidelobe decrease rate is lower and the overal amplitude is higher. This can easily be explained by the discontinuities at the beginning and the end of the pulse (see panel A.5a) that introduce spectral leakage, similarly to what happens with the block pulse.

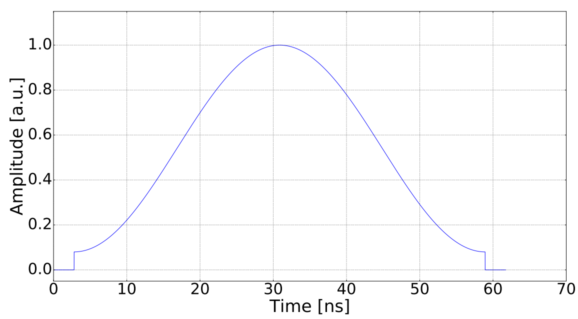

(a) Time-domain evolution of the shape, with normalised amplitude.

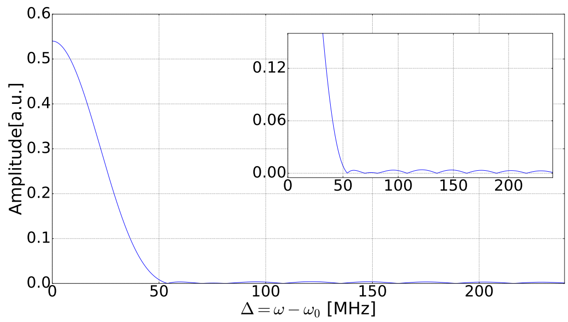

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.5: Hamming modulation shape

The minimised first sidelobe is not very useful in the configuration given as example, because the resonance frequency of the others qubits are much higher than the frequency corresponding to this minimum. However, a pulse of shorter duration would have a more spread FT, and hence may have this very low first sidelobe corresponding to the resonance frequency of another qubit, leading to very low crosstalk. However, a shorter pulse requires a higher peak amplitude, which will increase the amplitude of the Fourier components over the whole spectrum, leading to higher crosstalk. Because the illustration in Figure A.5b already has a high amplitude of Ω = 25.0 MHz, the optimal driving configuration for this shape is probably to not leverage the minimum first sidelobe level. Nonetheless, even in that case, the shape is still likely to perform very well.

A.2.3. Blackman

The expression of the Blackman shape is given in Equation A.6.

$$s(t) = \begin{cases} 0.42 + 0.5\cos\left(\frac{\pi t}{\tau}\right) + 0.08\cos\left(\frac{2\pi t}{\tau}\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.6}$$

A Blackman-shaped pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz is plotted, as well as its FT, in Figure A.6. We can see that the main lobe is slightly wider than that of the shapes presented so far (especially at its base). This will however not be problematic because it is not wide enough to induce high off-resonant drive for the other qubits: the high-intensity components that compose the main lobe do not correspond to the expected resonance frequencies of the other qubits. A wider main lobe can on the contrary be beneficial as it will allow the drive to still be efficient in the case of a shift of a few Megahertz in the resonance frequency of the targeted qubit.

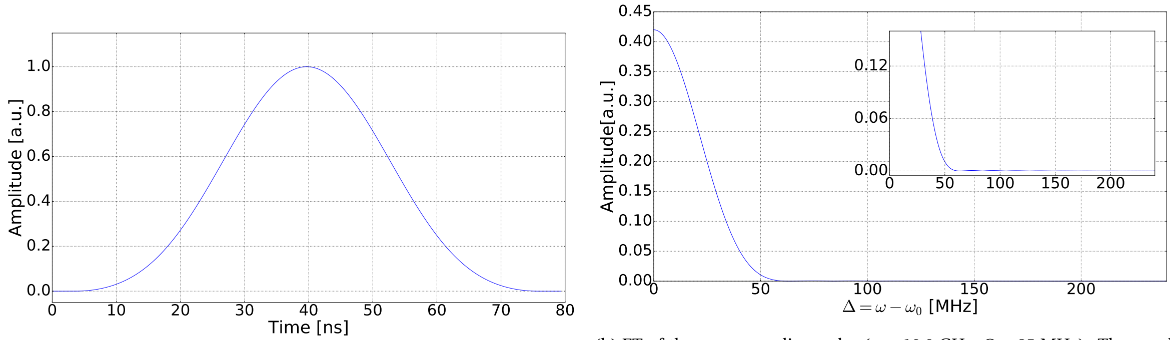

(a) Time-domain evolution of the shape, with normalised amplitude.

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.6: Blackman modulation shape

The inset of Figure A.6b shows a zoom on the low-intensity Fourier components of the pulse. We can see that the main lobe has a very sharp falling edge, followed by low sidelobes. This is useful for quick sidelobe rejection. Furthermore, we can see that the intensity at Δ = 100 MHz is very low. This makes the Blackman shape a sensible candidate for good performance.

A.2.4. Optimised Blackman

The optimised Blackman shape has lower sidelobes than the non-optimised one, while in theory keeping the same main lobe, which makes this family of shapes an interesting choice for applications that require immediate sidelobe rejection, as it is the case here. Its expression is given in equation A.7, where C is a free parameter. Notice that when $C = 0.04$, we get back the non-optimised shape (Equation A.6) and when $C = 0$, we get the expression of a Hann shape (Equation A.4).

$$s(t) = \begin{cases} (0.5 - 2C) + 0.5\cos\left(\frac{\pi t}{\tau}\right) + 2C\cos\left(\frac{2\pi t}{\tau}\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.7}$$

The shape and Fourier transform of several 10.0 GHz optimised Blackman-shaped pulses are presented in Figure A.7. A representative range of values for C as been explored, while keeping the pulse duration arbitrarily fixed (for easier visualisation). In theory however, varying the free parameter while keeping the pulse frequency and amplitude constant modifies the pulse duration.

One notable feature of panel A.7a is that increasing the free parameter increases the fraction of the pulse envelope that has negative values. As a consequence, the higher the C values, the longer the pulses have to be in order to perform a specific rotation.

The inset of Figure A.7b presents a zoom on the low-intensity components of the pulse. We can see that (with the exception of the first one) the sidelobe levels are indeed lower than those visible in Figure A.6. Furthermore, the main lobe actually appears narrower than that of the Blackman-modulated pulse. The higher first sidelobe could in theory produce high levels of crosstalk, but it is unlikely to happen, as it is located much lower than Δ = 100 MHz. Shorter pulses will have this peak at higher detuning, but such a configuration is not advisable as it would require very fast driving.

Furthermore, increasing the C parameter yields a narrower main lobe with a similar sidelobe decrease rate. The narrower main lobe shifts lowers the frequeny of the zeros of the FT, which overall results in lower intensity around Δ = 100 MHz. However, the narrow main lobe can also negatively affect the fidelity because a too selective pulse might miss a qubit whose resonance frequency is slightly shifted. Consequently, there is likely to be an optimal C value which balances low sidelobe levels and appropriate selectivity.

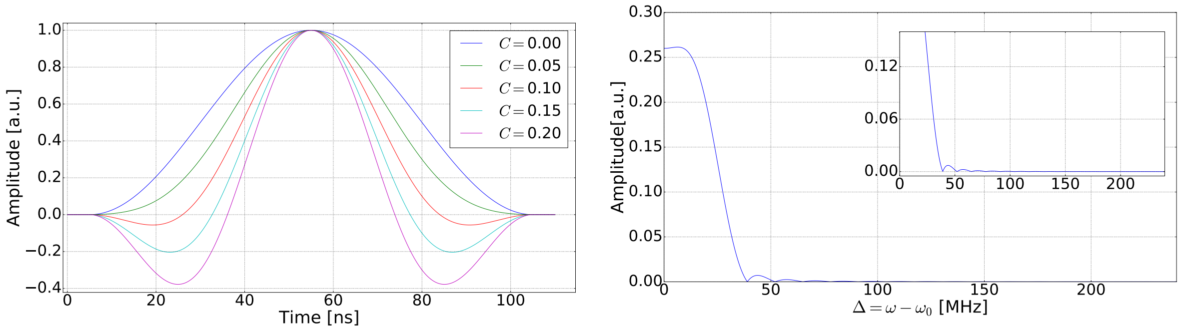

(a) Time-domain evolution of the shape, with normalised amplitude, for different values of the shape's free parameter.

(b) FT of the corresponding pulse ($\omega = 10.0$ GHz, $\Omega = 25$ MHz, $C = 0.12$). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.7: Optimised-Blackman modulation shape

A.2.5. Papoulis

The Papoulis shape is characterised by a main lobe similar to that of the Blackman shape, with a slightly higher first sidelobe but a much faster sidelobe decrease rate. Its expression is given in Equation A.8.

$$s(t) = \begin{cases} \frac{1}{\pi}\left|\sin\left(\frac{\pi t}{\tau}\right)\right| + \left(1 - \frac{|t|}{\tau}\right)\cos\left(\frac{\pi t}{\tau}\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.8}$$

A Papoulis-shaped pulse of frequency $\omega = 10.0$ GHz and amplitude $\Omega = 25$ MHz is plotted, as well as its FT, in Figure A.8. Its high sidelobe decrease rate ensures that the amplitude at the expected resonance frequency of the other qubits will be strongly attenuated.

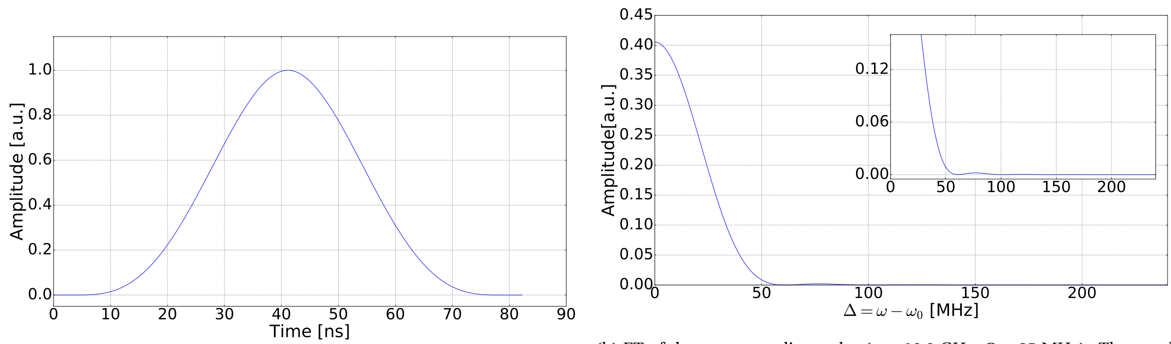

(a) Time-domain evolution of the shape, with normalised amplitude.

(b) FT of the corresponding pulse ($\omega = 10.0$ GHz, $\Omega = 25$ MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.8: Papoulis modulation shape

We can see in the zoom presented in panel A.8b that, although the first sidelobe is comparatively high, the next one is already extremely low. The amplitude at the resonance frequency of a second qubit ($\Delta = 100$ MHz) is among the lowest so far. The high sidelobe decrease rate also ensures that at the resonance frequency of the next qubits, the intensity of the pulse will be extremely low, which would in theory yield an almost negligible amount of unwanted driving.

A.3. Gaussian-based shapes

Gaussian and Gaussian-based shape functions are widely used in NMR experiments, making them very suitable candidates in our context.

A.3.1. Gaussian shape

Although the Fourier transform is not a perfect approximation for the spin dynamics near resonance, it still gives an accurate indication when further away. Consequently, FT-based arguments are to some extend applicable to our case.

There is a Fourier transform theorem that states that if a function can be differentiated k times before its derivative exhibits a discontinuity, then its Fourier transform falls off as the inverse of the k^{th} power of the frequency in the tails (Freeman, 1998). In this respect, the Gaussian shape should be ideal, as it can infinitely be derived without ever introducing a discontinuity. Real implementations introduce a small complication in that it is necessary to truncate the pulse shape, as an infinite pulse is not a realistic proposition. This does reduce the sidelobe decrease rate, but it cannot be avoided if the shape is to be used in experiments. Truncating at 2 or 3 percent of the maximal amplitude already appears to not introduce too much spectral leakage (Freeman, 1998). To limit the spectral leakage even more, it has then been decided to make the truncation at 0.5% of the maximal amplitude.

The Gaussian shape is defined in Equation A.9.

$$s(t) = \begin{cases} \exp\left[-a\left(\frac{t}{\sigma}\right)^2\right], & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.9}$$

where τ is the pulse duration and σ defines the width of the curve. a is a free parameter.

A Gaussian-shaped pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz is presented in Figure A.9, as well as its Fourier transform.

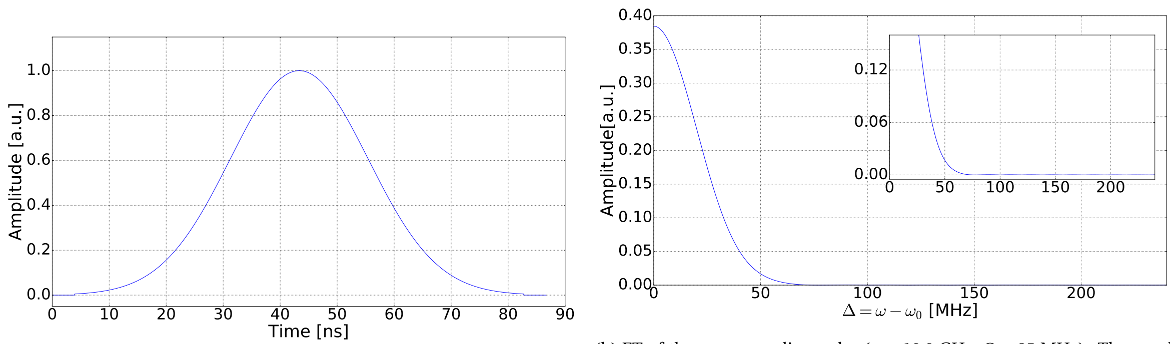

(a) Time-domain evolution of the shape, with normalised amplitude.

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.9: Gaussian modulation shape

Two important features are shown in panel A.9b. First, the main lobe appears to be among the widest so far (especially at its base), which can be beneficial in case of noise-induced variations in the qubit's resonance frequency. The second feature is that the sidelobes have extremely low amplitudes, which is consistent with what was expected. This also makes the low sidelobe decrease rate not as significant a problem, as the highest value is already very low. Consequently, the Gaussian shape should theoretically yield very little crosstalk.

A.3.2. Half-Gaussian shape

It has been noticed (Kessler et al., 1991) that most of the phase accumulated by a Gaussian pulse occurred during the second half of the pulse. It then seems natural to consider using only the first half of the Gaussian pulse.

The shape is defined exactly as the Gaussian pulse (Equation A.9), but includes a truncation at $t = \tau/2$, where τ is the duration of the similar Gaussian pulse. In our implementation, the pulse duration was actually kept identical between the two shapes, and the width of the half-Gaussian pulse was doubled, but those two definitions are equivalent.

A half-Gaussian pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz and its Fourier transform are presented in Figure A.10. Looking at the Fourier transform of such a pulse, the very wide main lobe and high sidelobes shown in Figure A.10b are expected because of the large truncation. They would tend to indicate a poorly performing shape, however literature has shown that, in NMR experiments, the frequency selectivity of this pulse is relatively good.

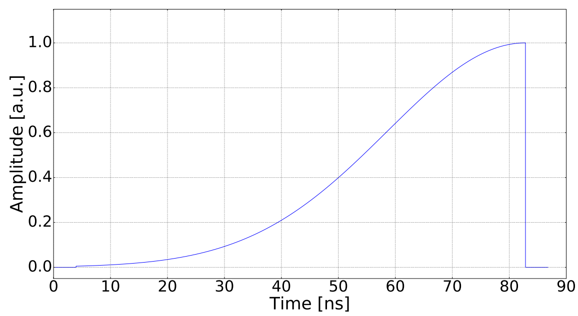

(a) Time-domain evolution of the shape, with normalised amplitude.

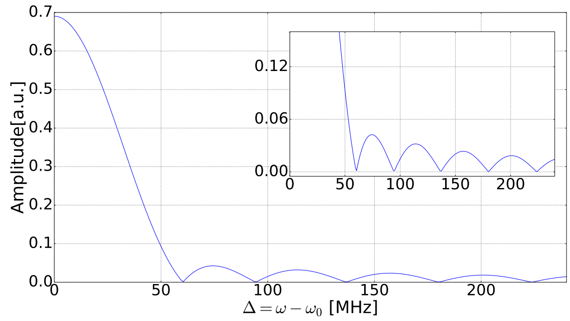

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.10: Half-Gaussian modulation shape

A.3.3. Hermite shape

One drawback of the Gaussian shape is its rather large transition region, *i.e.* the spectral range where the spin excitation is not large enough to produce good rotations but still not negligible. It has been shows (Warren, 1984) that multiplying a Gaussian by a even-order polynomial should reduce the transition region. In fact, it appears that second-order polynomials already provide a very good improvement. This defines the Hermite shape. Its expression is given in Equation A.10.

$$s(t) = \begin{cases} \exp\left[-a\left(\frac{t}{\sigma}\right)^2\right] \times \left[a + b\left(\frac{t}{\sigma}\right)^2\right], & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.10}$$

The coefficients of the polynomial are free parameters of the shape and can be varied arbitrarily. Two sets of values have been chosen based on well-performing designs in literature (Warren, 1984): ($a = 1.0, b = -0.956$) and ($a = 1.0, b = -0.667$). The former is optimised for a π-rotation while the latter is designed for a $\pi/2$-rotation.

A $\pi/2$-pulse using the Hermite shape is presented in Figure A.11, along with its Fourier transform.

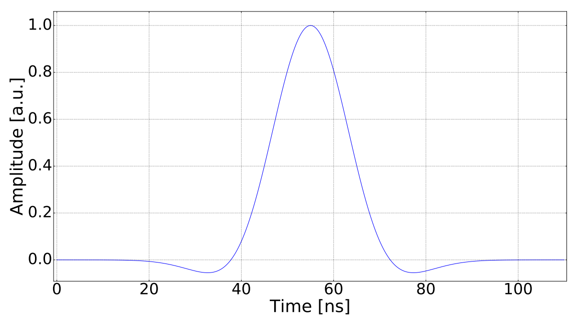

(a) Time-domain evolution of the shape, with normalised amplitude.

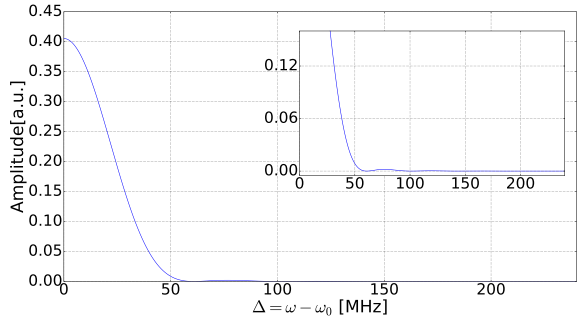

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.11: Hermite modulation shape

While the main lobe of the FT does not appear to be significantly different from the Gaussian shape in terms of width, its transition region indeed appears to be narrower. This leads to slightly higher sidelobes, but they remain of very small amplitude. This shape should then show a better frequency selectivity than the Gaussian, with similar amounts of crosstalk.

A.4. Flat top

This is once again a family of shapes. Flat top shapes are characterised by a much broader main lobe with a larger region of maximal amplitude (resulting in higher spectral accuracy, and in the appellation “flat top").

They also usually present very high sidelobe decrease rates (among the highest of the shapes described in this appendix). These characteristics can be very advantageous for our requirements, but one need to keep in mind that the very broad main lobe can induce strong off-resonance drive if the qubits do not have sufficiently separated resonance frequencies.

The functions of this family are described by the general expression:

$$s(t) = \begin{cases} c_0 + \sum_{k=1}^{m} c_k cos\left(\frac{2\pi}{\tau} kt\right), & |t| \leq \tau \\ 0, & \text{elsewhere} \end{cases} \quad \text{(A.11)}$$

	STF3F	STF4F	STF5F	STF3M	STF4M	STF5M	HFT90D	HFT116D	HFT169D
c_0	0.26526	0.21706	0.188100	0.28235	0.241906	0.209671	1	1	1
c_1	-0.50000	-0.42103	-0.36923	-0.52105	-0.460841	-0.407331	-1.942604	-1.9575375	-1.97441842
c_2	0.23474	0.28294	0.28702	0.19659	0.255381	0.281225	1.340318	1.4780705	1.65409888
c_3	0	-0.07897	-0.13077	0	-0.041872	-0.092669	-0.440811	-0.6367431	-0.95788186
c_4	0	0	0.02488	0	0	0.0091036	0.043097	0.1228389	0.33673420
c_5	0	0	0	0	0	0	0	-0.0066288	-0.06364621
c_6	0	0	0	0	0	0	0	0	0.00521942
c_7	0	0	0	0	0	0	0	0	-0.00010599

Table A.1: c_k coefficients for flat top pulse shapes included in our analysis

The values of the parameters c_k for the flat-top shapes currently included in the analysis are given in Table A.1. The naming convention for the functions comes from Heinzel et al. (2002). They are regrouped in three sub-families:

- Fast-decaying functions (with names formatted as SFTxF) are characterised by very high sidelobe decrease rates.
- Minimum-sidelobe functions (with names formatted as SFTxM) are characterised by the lowest peak sidelobe levels achievable with between 3 and 5 cosine terms.
- The last family (with names formatted as HFTxD) have been designed to have both the lowest peak sidelobe levels achievable with a defined number of cosine terms and a sidelobe decrease rate in f^{-3}.

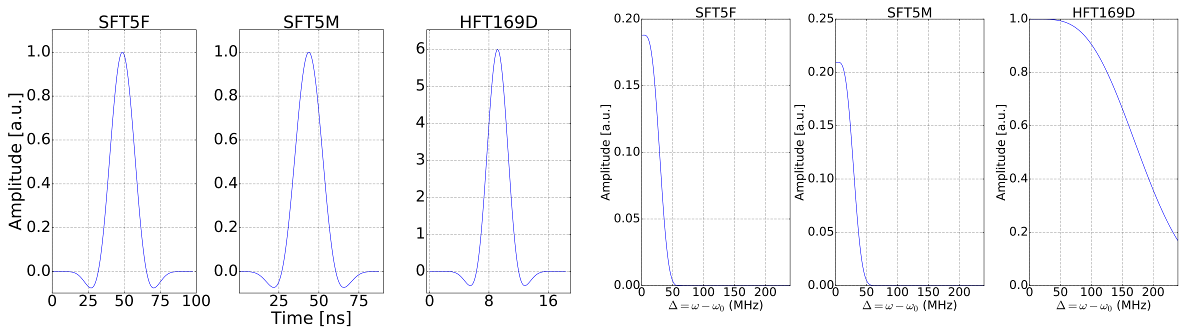

(a) Time-domain evolution of three of the the shape, with normalised amplitude.

(b) FT of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.12: Flat top modulation shapes, belonging to the fast-decaying, minimum-sidelobe and third families respectively

Plots of a 10.0 GHz pulses modulated with some of those shapes, as well as their FT, are given in Figure A.12 as an example. Note that the scale of both the X- and Y-axes of figure A.12a are different per shape. In particular, the amplitude of the HFT169D shape is about six times larger than that of the others, for a pulse duration much shorter.

The very broad main lobe that is characteristic of flat top pulses is especially visible for the HFT169D shape (rightmost panel of each figure). We can also see that the main lobe of this pulse covers detunings larger than 100 MHz. This indicates that the amount of parasitic driving will be significant. One can then expect that this shape will not allow for high fidelity driving.

For the other two sub-families of shape, the Fourier main lobe is in particular wider at the peak (the “flat top") but not significantly more at the base. This results in a narrow transition region and so in both a good tolerance to decoherence-induced movements of the qubits and in a high frequency selectivity.

We can also see the fast decay rate of the first sub-family of shapes. The fast-decaying sub-family is especially promising. It is the group of shapes that require the longest pulse durations to complete the rotations, but this may not be as detrimental as it would be for other shapes. Indeed, the flat top of the FT main lobe indicates that the pulse may still be able to efficiently drive qubits whose resonance frequency is shifted by decoherence.

The low sidelobe levels that characterise the minimum-sidelobe sub-family are also visible in Figure A.12b. This family is also interesting but, similarly to the Hamming shape, the minimum peak sidelobe is located closer to resonance than the neighbouring qubits. A shorter pulse could bring the low sidelobe to the appropriate frequency, but it would in turn increase the overall amount of crosstalk due to the very fast driving it would require. As a consequence, the optimal use of the shapes in this sub-family is to find a balance between using the lowest first sidelobe and the speed of the drive.

A.5. Other shapes

A.5.1. Kaiser’s modified zeroth-order Bessel functions

This family of functions is a so-called near-optimum shape family (*i.e.* it is a very good approximation of the Slepian shape which maximises the energy inside the main lobe of the shape with respect to the total energy). There are several such families, but here we only describe the zeroth-order family, given by:

$$s(t) = \begin{cases} \dfrac{I_0\left[\alpha\sqrt{1-\left(\frac{t}{\tau}\right)^2}\right]}{I_0(\alpha)}, & |t| \le \tau \\ 0, & \text{elsewhere} \end{cases} \tag{A.12}$$

where I_0 is the modified zeroth-order Bessel function of the first kind (defined in equation A.13), and α is a free parameter, usually in the range $[0;9]$. Some literature use $\beta = \alpha\pi$.

$$I_0(x) = 1 + \sum_{k=1}^{\infty}\left[\frac{1}{k!}\left(\frac{x}{2}\right)^k\right]^2 \tag{A.13}$$

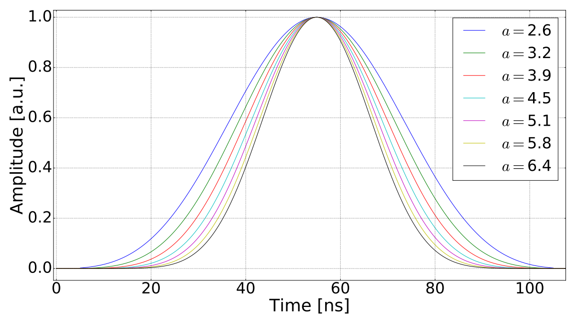

(a) Time-domain evolution of the shape, with normalised amplitude, for different values of the shape’s free parameter.

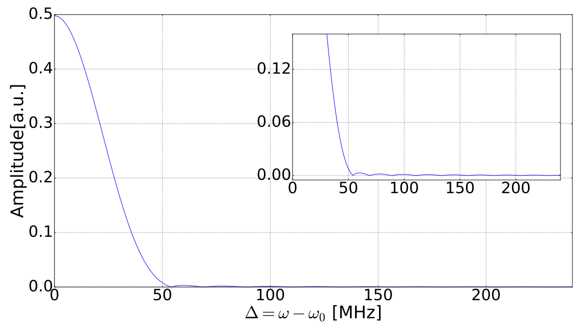

(b) T of the corresponding pulse (ω = 10.0 GHz, Ω = 25 MHz, α = 1.929). The amplitude is normalised to the maximum amplitude of the block pulse. The inset focuses on the low intensity components of the FT.

Figure A.13: Kaiser modulation shape

A pulse of frequency ω = 10.0 GHz and amplitude Ω = 25 MHz shaped using this function is plotted, as well as its FT, in Figure A.13. Several typical values of the free parameter α have been explored, all the while keeping a constant pulse duration for ease of representation. Similarly to the optimised Blackman shape, modifying α without changing ω and Ω actually yields different pulse durations.

Increasing α appears to narrow the pulse envelope, resulting in theory in higher required amplitude in order to perform a specific rotation. Logically, when comparing the pulse using fixed amplitude and variable duration, it shows that higher α values lead to longer pulses. Just as was the case for the optimised Blackman shape, there then exist an optimal value of α that balances crosstalk level and pulse duration.

Decreasing the value of α brings the shape closer to that of a block pulse (both shapes become similar for $\alpha = 0$). On the contrary, increasing the α parameter yields a shape with a much wider main lobe and extremely low sidelobe levels. The trend seems however to saturate beyond some values of α (the FT is very similar for $\alpha = 4.5$ and $\alpha = 9.0$ for instance). The sidelobe decrease rate also does not appear to change with the value of α. Some particular values are of interest: $\alpha \approx 1.59$ makes the function similar to the Hamming shape, $\alpha \approx 1.91$ to the Hann shape and $\alpha \approx 2.55$ to the Blackman shape.

Keeping the previously discussed characteristics in mind, this shape is expected to produce pulses that allow for high fidelity single-qubit rotations.

B
Validation

In this appendix, the methods used to validate various features implemented in this project are described. The results and conclusions of these validations are then given.

B.1. Pulse duration computation

Theory gives that for a driving pulse of specific Rabi frequency to perform a rotation of angle θ, it must be turned on for a duration that obeys Equation 3.6, *i.e.* the integral of the pulse envelope over its duration must equal the target rotation angle. It is then of tremendous importance to ensure that the computation of this duration is implemented correctly.

Before evaluating the correctness of the pulse duration computation, it is necessary to characterise the maximal rotation error allowed while still having a gate fidelity superior to 99.9%. This (as well as the rest of this validation) was done by considering a setup containing only one qubit driven by a resonant pulse. The rise time effect due to the imperfect AWG has been accounted for, since it has an impact on the required duration. However, frequency shifts effects have been disregarded because they are not related to the pulse duration computation and would only introduce extra errors. In this configuration, the only expected error source is the duration of the pulse.

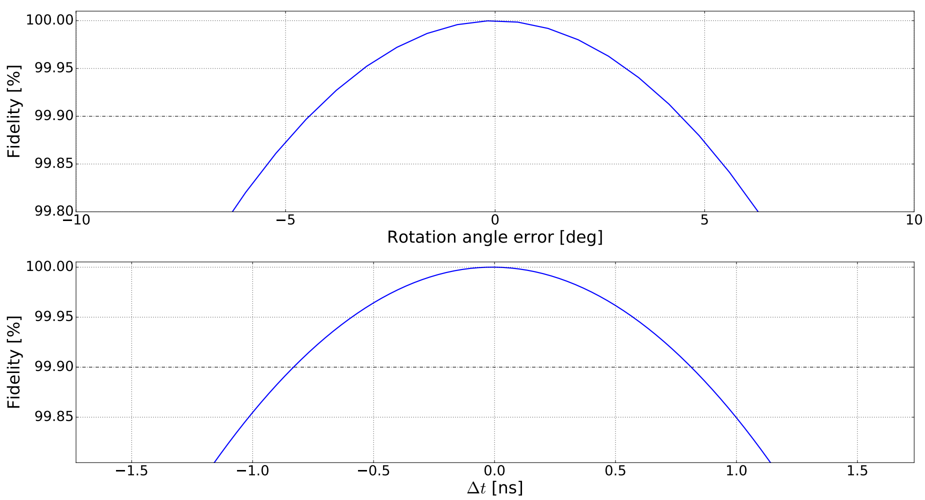

Figure B.1: Average gate fidelity of a drive as a function of the achieved rotation angle. The fidelity is expressed as a function of either the rotation angle error (top panel) or the pulse duration error (bottom panel).

The threshold rotation angle error over which the fidelity is below 99.9% was evaluated as follows. The average gate fidelity was computed between the spin rotation matrix corresponding to the target rotation angle and the one corresponding to the achieved rotation. This last angle was taken as a free variable in the range $[0, 2\pi]$.

The resulting fidelities are plotted in Figure B.1. The top panels presents the fidelity with respects to the rotation angle error, and the bottom panel does the same but with respects to the actual duration error. These results are independent of the considered target rotation angle. It appears that over- or under-rotation lower than 4.4 degrees yields an average gate fidelity superior to 99.9%. This is equivalent to the pulse duration being over- or underestimated by less than 0.81 ns.

The validation of the pulse duration computation was performed by evaluating the rotation angle errors produced by each of the pulse shapes considered in this work. The rotation error itself was evaluated as follows. The pulse duration required to perform a rotation was computed for various driving configurations (*e.g.* target rotation, Rabi speed). Then, the rotation that the corresponding pulse of this duration would theoretically achieve was computed (using Equation 3.6) and compared with the target of the operation. The process is illustrated in Figure B.2.

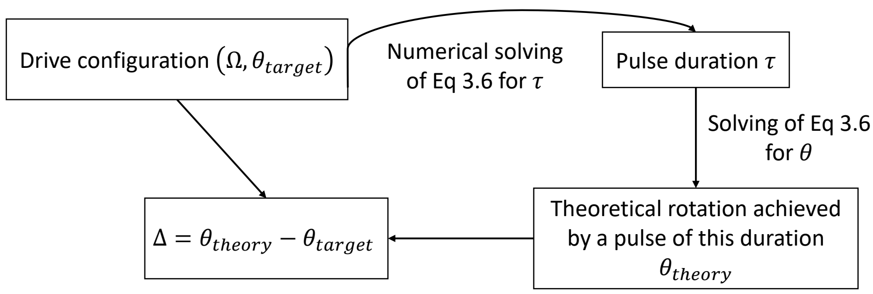

Figure B.2: Flowchart of the computation of the error in rotation angle at the end of a driving pulse

This process was conducted for both $\pi/2$- and π-rotations, for representative Rabi frequencies in the range [5.0, 25.0] MHz, and for all the shapes described in Appendix A. The obtained rotation angle errors are summarised in Table B.1 for $\pi/2$ rotations, expressed in degrees. The errors in the case of π rotations are extremely similar. For the sake of brevity, they are thus not provided.

These errors are very low (well under one degree for most of the shapes and driving speeds). The slightly large ones (*e.g.* for the Rectangle shape) can be safely attributed to numerical integration errors. According to Figure B.1, the maximal remaining angle error (-0.59600 degrees, corresponding to a rectangle pulse of amplitude 25.0 MHz) yields a fidelity error close to 0.002%, which is negligible compared to other error sources. As discussed in Appendix A, the Kaiser shape becomes very similar to the block shape when the free parameter α is close to zero. This explains why the second highest rotation errors are for the Kaiser shape with $\alpha = 0.643$.

Based on these results, the computation of the pulse duration required to perform a specific rotation appears to be implemented correctly.

As discussed in Section 3.2, it is also possible to define a driving pulse by its duration. In that case, the Rabi frequency of the pulse is the free parameter that determines the achieved rotation. A similar validation as for the Rabi-defined pulses has been performed in this case: the rotation angle theoretically achieved by a pulse with the computed Rabi frequency has been determined using Equation 3.6 and compared with the target angle. The obtained rotation angle errors in this case are even smaller than for the fixed-amplitude configurations. The actual values have been omitted for the sake of brevity, but all of them stay well below 10^{-2} degrees. Considering how little this range of errors impacts the fidelity of the rotation, this is sufficient to guarantee that the computation has also been implemented correctly for fixed-duration pulses.

B.2. Frequency shifts

The frequency shifts induced by the coupling of the driving pulses with the qubits are an important error source and must be modeled correctly. This section first presents how the implementation of the shifts in the simulation software has been validated. It then describes how the behaviour of the frequency shift correction algorithm described in Section 3.1.2 was verified.

	$\pi/2$				
Ω [MHz]	5	10	15	20	25
Rectangle	-0.12514	-0.23725	-0.35586	-0.47495	-0.59600
Triangular	-0.00030	-0.00119	-0.00267	-0.00473	-0.00743
Sine	-0.00059	-0.00237	-0.00533	-0.00948	-0.01481
Hann	0.00000	0.00000	0.00001	0.00002	0.00004
Hamming	-0.00949	-0.01898	-0.02846	-0.03794	-0.04741
Blackman	0.00000	0.00000	0.00000	0.00000	0.00001
Opt Blackman ($c = 0.10$)	0.00000	0.00000	0.00000	0.00000	0.00001
Opt Blackman ($c = 0.12$)	0.00000	0.00000	0.00000	0.00000	0.00001
Opt Blackman ($c = 0.14$)	0.00000	0.00000	0.00000	0.00000	0.00001
Papoulis	-0.00000	-0.00000	-0.00000	-0.00000	-0.00000
Gaussian	-0.00120	-0.00242	-0.00368	-0.00498	-0.00630
Half-Gaussian	-0.06486	0.11018	0.11025	0.11031	-0.06461
SFT3F	0.00000	0.00000	0.00000	0.00000	0.00001
SFT4F	0.00000	0.00000	0.00000	0.00000	0.00000
SFT5F	0.00000	0.00000	0.00000	0.00000	0.00000
SFT3M	0.00499	0.00999	0.01498	0.01998	0.02497
SFT4M	0.00064	0.00129	0.00193	0.00257	0.00321
SFT5M	0.00000	0.00000	0.00000	0.00000	-0.00000
HFT90D	0.00000	0.00000	0.00002	0.00009	0.00031
HFT116D	0.00000	0.00000	0.00001	0.00004	0.00012
HFT169D	0.00000	0.00000	0.00000	0.00000	0.00000
Kaiser 0 ($\alpha = 0.643$)	-0.05456	-0.10995	-0.16615	-0.22323	-0.28110
Kaiser 0 ($\alpha = 1.929$)	-0.00184	-0.00383	-0.00597	-0.00827	-0.01071
Kaiser 0 ($\alpha = 3.857$)	-0.00001	-0.00001	-0.00002	-0.00004	-0.00005
Kaiser 0 ($\alpha = 5.786$)	0.00000	0.00000	0.00000	0.00000	0.00000

Table B.1: Rotation angle error due to pulse duration computation, for different driving configuration attempting to implement a $\pi/2$-rotation. The angle error is expressed in degrees.

B.2.1. Bloch-Siegert shift

The implementation of the Bloch-Siegert shift uses the expression derived by Shirley (1965). More specifically, since the qubit operations occur in the weak driving regime, it uses an approximation to 6^{th} order in the pulse amplitude, derived by perturbation analysis (Equation 2.16).
The correctness of the implementation has been validated using the results described in Yan et al. (2015).

Pulse amplitude	[GHz]	1.0	3.5	6.0	8.5	11.0	13.5	16.0
$\Delta\omega_{\text{BS, Shirley}}$	[GHz]	0.063228	0.712320	1.650482	2.639255	3.641373	4.650384	5.664602
$\Delta\omega_{\text{BS, Yan}}$	[GHz]	0.063228	0.712320	1.650482	2.639255	3.641373	4.650384	5.664602

Table B.2: Validation of the analytical Bloch-Siegert shift expression against data from Yan et al. (2015). The pulse amplitudes as well as the shift values are given in GHz.The values of the last row are taken from Table I of Yan et al. (2015).

First of all, the implementation of the analytical expression from Shirley (1965) has been validated by computing values of the Bloch-Siegert shift for a wide range of pulse amplitudes. The obtained values are presented in Table B.2, along with those given in Table I of Yan et al. (2015). The values are clearly identical, validating this implementation.

However, these values correspond to a driving regime much stronger than the one considered in this work, and are thus not suitable to verify the correctness of the approximate expression given in Equation 2.16. The previously validated analytical expression has thus been used as a reference for this. Table B.3 presents the difference between the Bloch-Siegert shift values computed using Equation 2.16 and those using the expression from Shirley (1965), for driving field amplitudes corresponding to the weak driving regime.

The differences are again very low, which indicates a proper implementation of the 6^{th} order approximation used in the simulations.

Pulse amplitude	[MHz]	2.5	5.0	10.0	15.0	20.0	25.0
$\Delta\omega_{\text{BS, Shirley}} - \Delta\omega_{\text{BS}}^{(6th)}$	[Hz]	3.5848e-08	1.0079e-07	5.0637e-08	1.1301e-08	7.6028e-08	-1.3629e-07

Table B.3: Validation of the 6^{th} order Bloch-Siegert shift expression against the analytical expression from Shirley (1965). The pulse amplitudes are given in MHz and the difference in shift values in Hz.

B.3. Frequency shift correction algorithm

Two elements of the correction algorithm (presented in Section 3.1.2) require validation. The first one is the selection of the drive frequencies so that they are equal to the shifted resonance frequencies. The second one is that the sets of drive frequencies computed are indeed resonant with the shifted qubits frequencies for each of the discrete time slices that compose the pulse.

Each of those steps has been validated by considering a two qubits and two drives system. The qubits frequencies are as usual chosen equal to 10.0 and 10.1 GHz respectively. To introduce no loss of generality, the drives have been chosen to use different modulation shapes and have different peak amplitudes. Peak amplitude values of $(\Omega_0, \Omega_1) = (25.0, 35.0)$ MHz have been used, as a large pulse amplitude yields larger shifts, which makes the results more easily exploitable. One pulse has been modulated using the so-called SFT3F flat-top shape and the other with a Gaussian.

It is important to note that the choice of the number of qubits introduces no loss of generality either to the validation. Indeed, having more qubits (and thus more driving pulses) simply modifies the number of equations and variables in the system that the algorithm must solve, and not its capability to reach the desired result.

B.3.1. Selection of the drive frequencies to match the dressed qubits

The selection of the drive frequencies amounts to the minimisation of the residual between the drive frequencies and the dressed qubits'. The convergence of this minimisation has been validated by performing a Monte-Carlo analysis. A large number (approx. 1.5×10^6) of sets of driving frequencies $(\omega_{\text{drive 1}}, \omega_{\text{drive 2}})$ has been randomly generated and their fitness was computed. The fitness function used was the average of the absolute values of the detuning of each pulse compared to its target dressed qubit. The set of free parameters that lead to the lowest residual is selected as the optimum.

To ensure the robustness of the analysis, the analysis was conducted using multiple seeds to initialise the random number generator. The values of $\omega_{\text{drive 1}}$ and $\omega_{\text{drive 2}}$ selected by the analysis for each random seed are given in Table B.4, as well as their relative difference with the optimal values obtained by the algorithm: $\omega_{\text{drive 1}}^{\text{opti}} = 9.99706$ GHz and $\omega_{\text{drive 2}}^{\text{opti}} = 10.10149$ GHz.

Random seed	$\omega_{\text{drive 1}}$ [GHz]	$\omega_{\text{drive 2}}$ [GHz]	$\omega_{\text{drive 1}}^{\text{MC}} - \omega_{\text{drive 1}}^{\text{opti}}$ [%]	$\omega_{\text{drive 2}}^{\text{MC}} - \omega_{\text{drive 2}}^{\text{opti}}$ [%]
1	9.99709	10.10152	3.00090×10^{-4}	2.96986×10^{-4}
2	9.99706	10.10146	0.00000	-2.96986×10^{-4}
3	9.99704	10.10150	-2.00059×10^{-4}	9.8995×10^{-5}
4	9.99706	10.10153	0.00000	3.95981×10^{-4}
5	9.99706	10.10147	0.00000	-1.97991×10^{-4}

Table B.4: Results of the Monte-Carlo analysis of the choice of drive frequencies to minimise the difference between the dressed qubit frequencies and the selected driving ones.

The values obtained by the Monte-Carlo analysis are very close to those obtained by the algorithm (the relative difference is well below 0.01% for each random seed). Furthermore, the results appear to be independent of the random seed used. This indicates that the frequency shift correction algorithm converges, as desired, towards frequencies that match the shifted resonance frequencies of the qubits.

To ensure that the process has not converged to a local minimum only, values for the couple of driving frequencies have been sampled in a large interval around the qubits' resonance frequencies, and the fitness of each of them has been computed. The results are plotted in Figure B.3. The optimum identified by the algorithm is indicated by a red cross. It appears clearly that the minimum found by the correction algorithm is a global one.

The frequency selection part of the correction algorithm therefore appears to be working as intended.

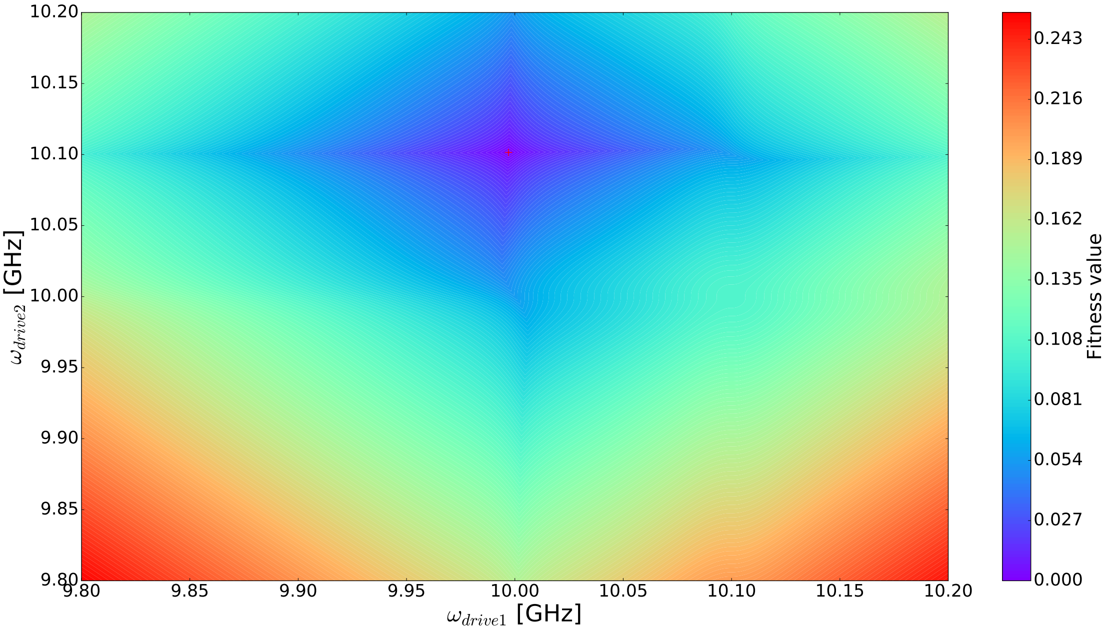

Figure B.3: The color represents the difference between the drive and the dressed qubits' resonance frequencies as a function of the selected frequency for each drive. The fitness function used in this figure is the average of the absolute value of the detuning of each pulse compared to its target qubits. The result of the optimisation algorithm is indicated by the red cross. It seems to correspond to a global minimum. All the frequencies are expressed in GHz.

B.3.2. Correctness of the frequency selection throughout a driving pulse

This section focuses on validating whether the drive frequencies selected by the algorithm remain resonant with the shifted qubits frequencies throughout the pulse. It has been done by simulating the simultaneous driving of two qubits while saving both the frequency of the drive and of the qubits for each of the time slices of a 50 ns pulse. As described before, each pulse has different shape and amplitude parameters to introduce no loss of generality.

Figure B.4 shows the evolution of both the drive and qubit frequencies throughout the pulse (blue and green curves respectively). Figure B.5 gives the difference between these two frequencies. In both cases, the top and bottom panels present the evolution of those frequencies for the first and second qubit respectively. In addition, the left and right panels present the evolution of the frequencies with and without the correction algorithm respectively. Two features can be observed. First of all, the magnitude of the shift from the nominal qubit frequencies is smaller when using the correction algorithm (left panels) than without it (right panels). This is expected, because in this case each qubit always has one drive that is resonant with it, and consequently that only induces a very small shift (Bloch-Siegert shift only). The total shift is thus lower than when the two pulses are off-resonant, as is the case in the right panels of the figure. Furthermore, the left panels of Figure B.4 show that the drive frequencies appear to match the dressed qubits' ones during the entirety of the pulse. This is confirmed by the computation of the detuning, (Figure B.5): the detuning is zero when the correction algorithm is used.

The frequency shift correction algorithm thus appears to be working as expected.

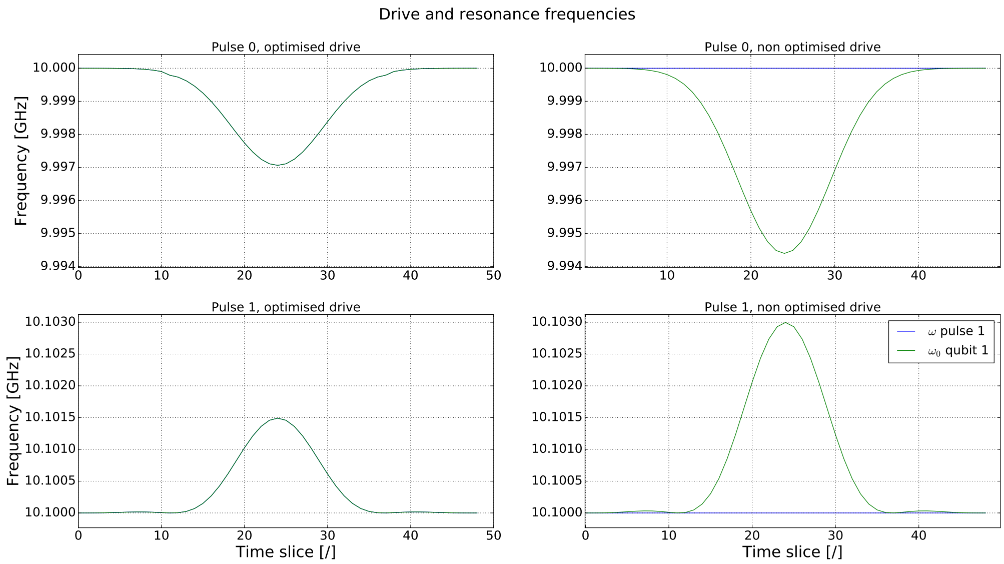

Figure B.4: Evolution of the qubits and drives frequencies of the first (top panels) and second qubit (bottom panels), when using the correction algorithm (left panels) and without it (right panels).

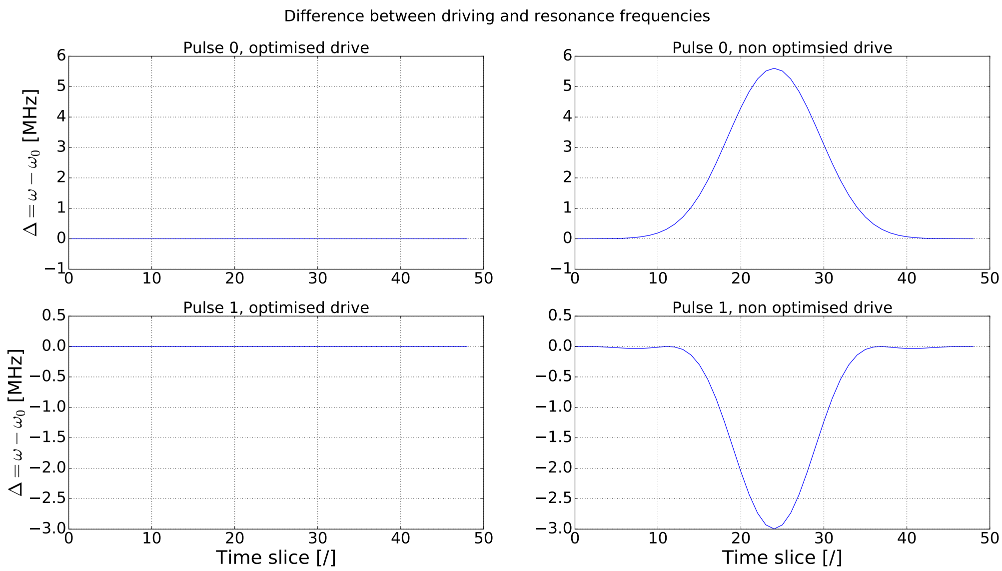

Figure B.5: Evolution of the detuning of the drives from their target qubits, when using the correction algorithm (left panels) and without it (right panels).

C

Additional figures of frequency shift correction results

This appendix contains the detailed results of the frequency shift correction algorithm. These results are summarised, commented and interpreted in Chapter 5. They are presented here for the sake of completeness. The shapes have been grouped by family and presented separately for better readability. The block pulse is represented in all figures to serve as a reference.

Results are first given for a two-qubits two-drives configuration, and then for a three-qubits three-drives one. Subsection C.3 then provides additional figures of the preliminary analysis for a five-qubits configuration. Lastly, Subsection C.4 regroups the outcomes of various test cases which provide useful insight into the results described in Chapter 5.

Each figure in this appendix shows the fidelity improvement brought by the use of the algorithm to correct either all frequency shifts, only the AC Stark shift or only the Bloch-Siegert shift. The bottom panel of each figure gives the infidelity values obtained by the fully corrected rotations.

C.1. Two-qubits two-drive configuration

The results in this section are obtained when considering an array of two qubits, each of them being driven by a resonant pulse.

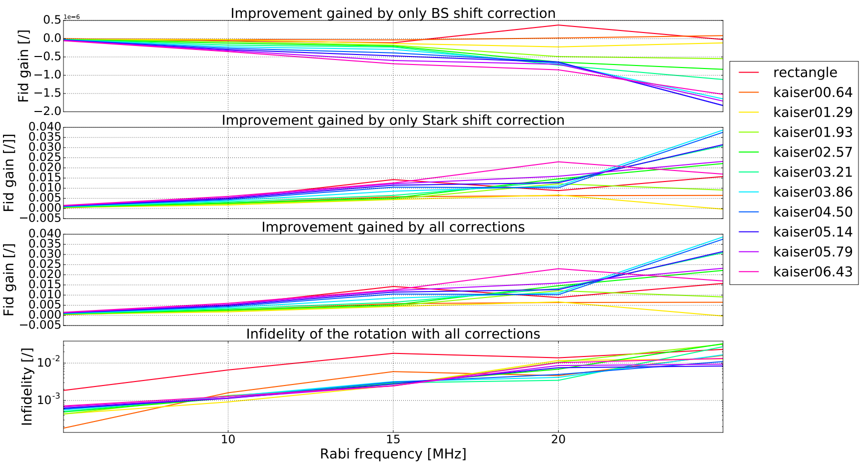

Figure C.1: Fidelity gain provided by the frequency shift correction algorithm on Kaiser-shaped $\pi/2$ rotation in the two-qubits two-drives configuration, for different combinations of shift corrections.

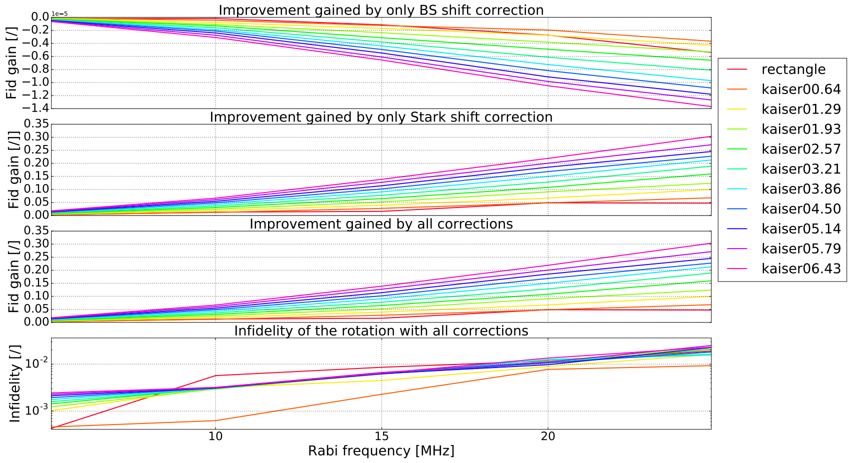

Figure C.2: Fidelity gain provided by the frequency shift correction algorithm on Kaiser-shaped π rotation in the two-qubits two-drives configuration, for different combinations of shift corrections.

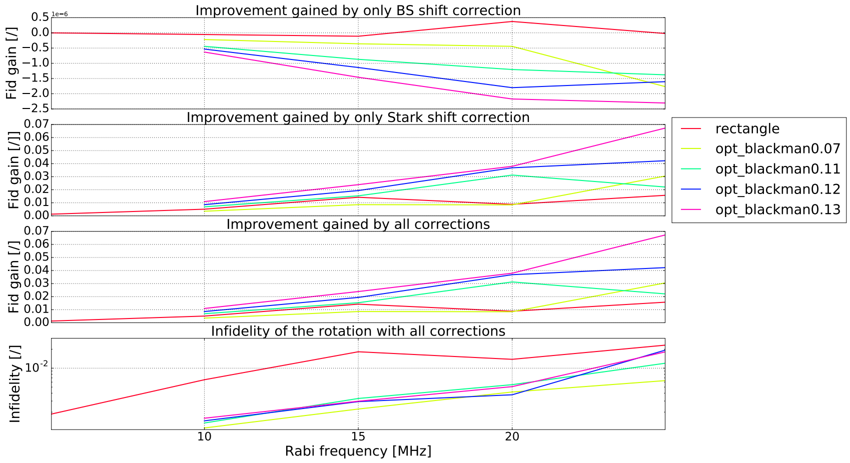

Figure C.3: Fidelity gain provided by the frequency shift correction algorithm on optimised Blackman-shaped $\pi/2$ rotation in the two-qubits two-drives configuration, for different combinations of shift corrections.

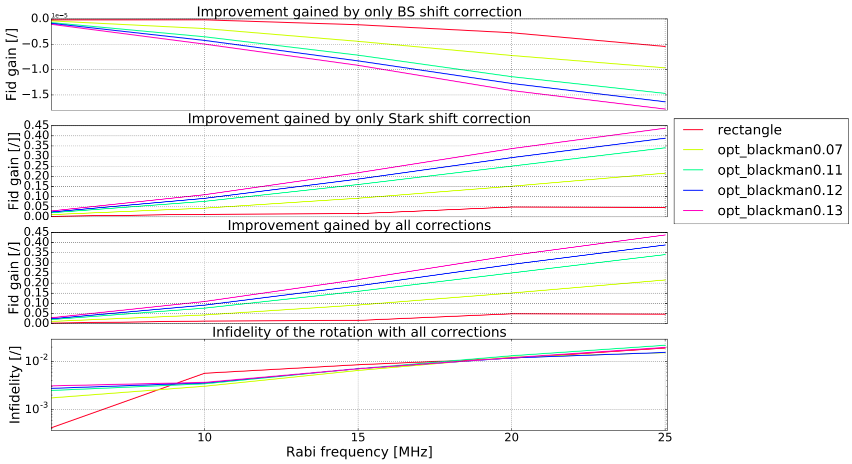

Figure C.4: Fidelity gain provided by the frequency shift correction algorithm on optimised Blackman-shaped π rotation in the two-qubits two-drives configuration, for different combinations of shift corrections.

Figure C.5: Fidelity gain provided by the frequency shift correction algorithm on flat top shaped $\pi/2$ rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

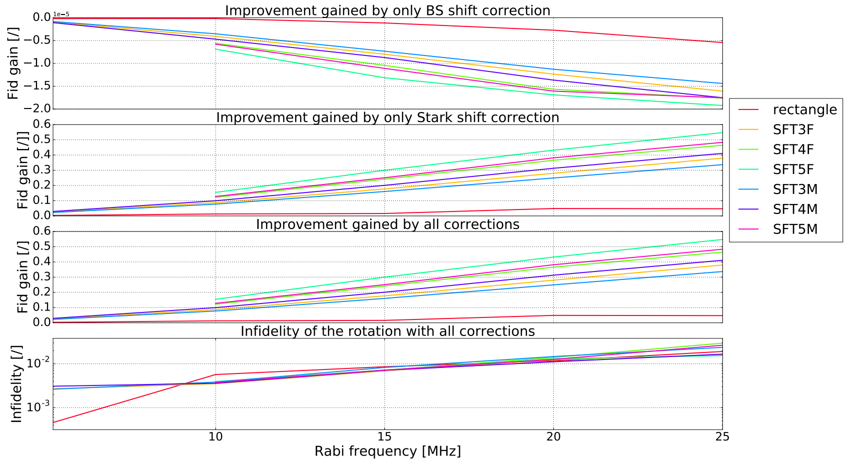

Figure C.6: Fidelity gain provided by the frequency shift correction algorithm on flat top shaped π rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

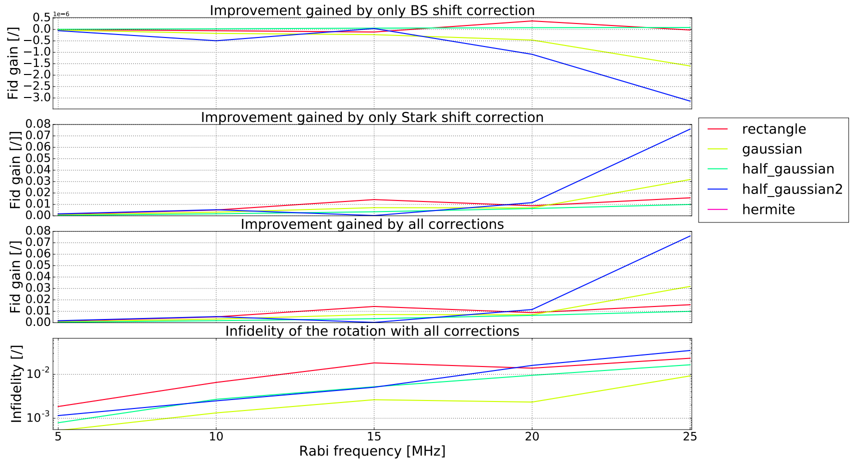

Figure C.7: Fidelity gain provided by the frequency shift correction algorithm on Gaussian-shaped $\pi/2$ rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

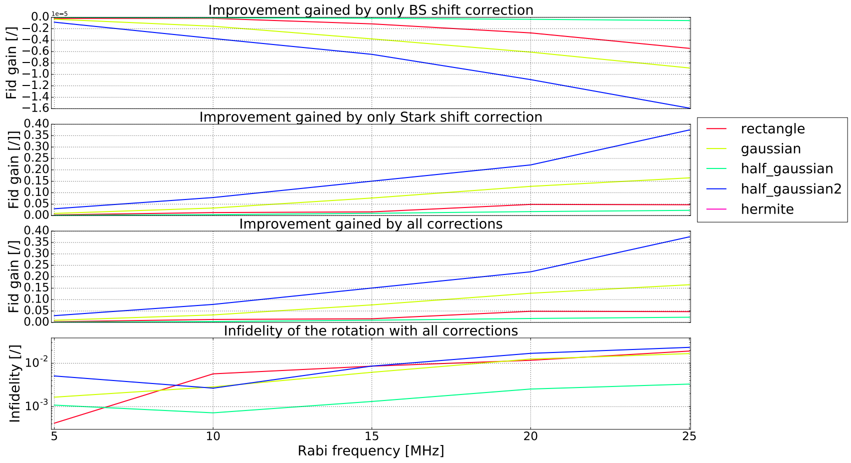

Figure C.8: Fidelity gain provided by the frequency shift correction algorithm on Gaussian-shaped π rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

Figure C.9: Fidelity gain provided by the frequency shift correction algorithm on HFTxD-shaped $\pi/2$ rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

Figure C.10: Fidelity gain provided by the frequency shift correction algorithm on HFTxD-shaped π rotations in the two-qubits two-drives configuration, for different combinations of shift corrections.

C.2. Three-qubits three-drives configuration

The results in this section are obtained when considering a similar configuration as the previous section, but with three qubits and three resonant drives.

Figure C.11: Fidelity gain provided by the frequency shift correction algorithm on Kaiser-shaped $\pi/2$ rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

Figure C.12: Fidelity gain provided by the frequency shift correction algorithm on Kaiser-shaped π rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

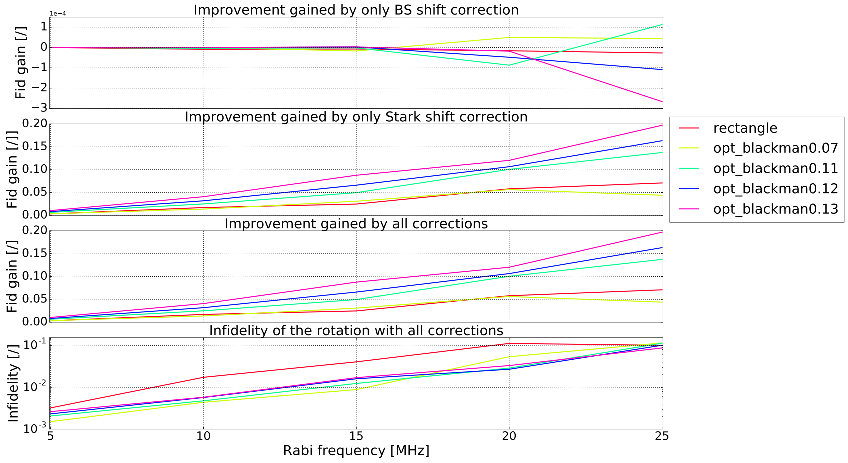

Figure C.13: Fidelity gain provided by the frequency shift correction algorithm on optimised Blackman-shaped $\pi/2$ rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

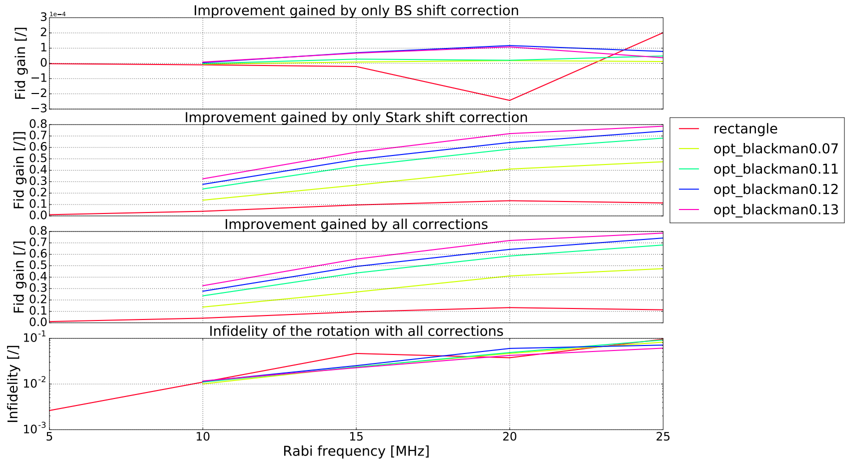

Figure C.14: Fidelity gain provided by the frequency shift correction algorithm on optimised Blackman-shaped π rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

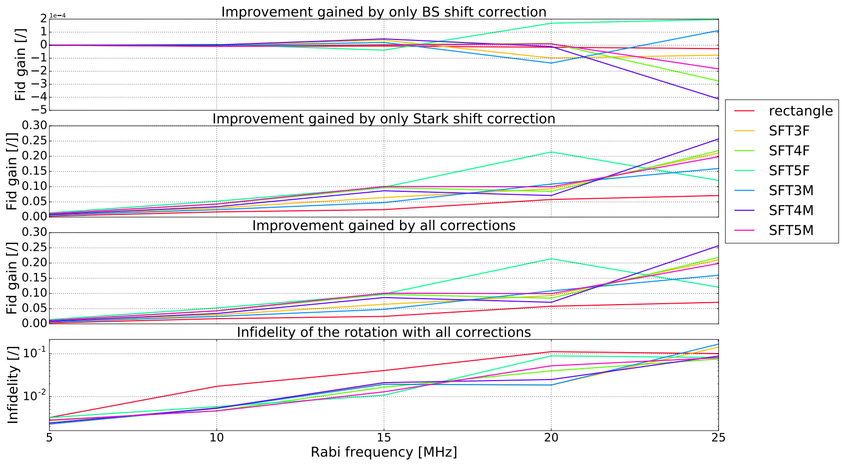

Figure C.15: Fidelity gain provided by the frequency shift correction algorithm on flat top shaped $\pi/2$ rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

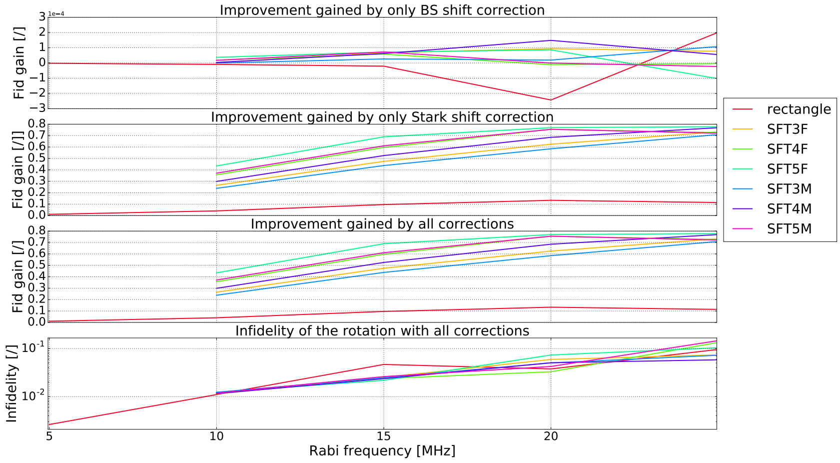

Figure C.16: Fidelity gain provided by the frequency shift correction algorithm on flat top shaped π rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

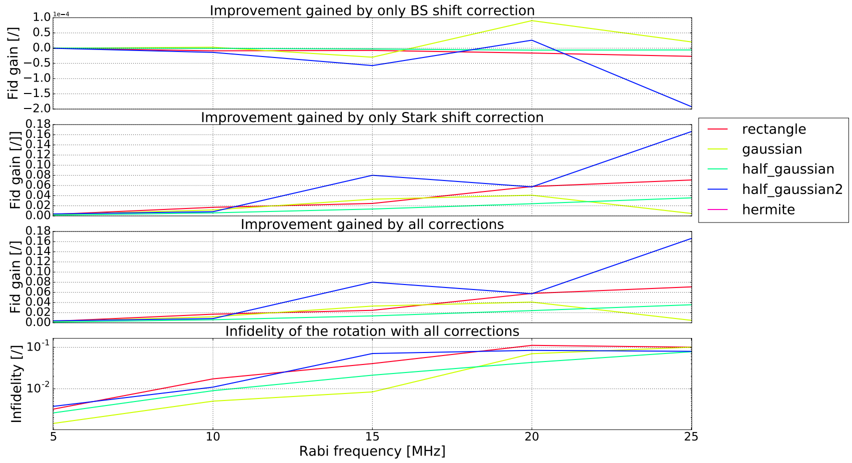

Figure C.17: Fidelity gain provided by the frequency shift correction algorithm on Gaussian-shaped $\pi/2$ rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

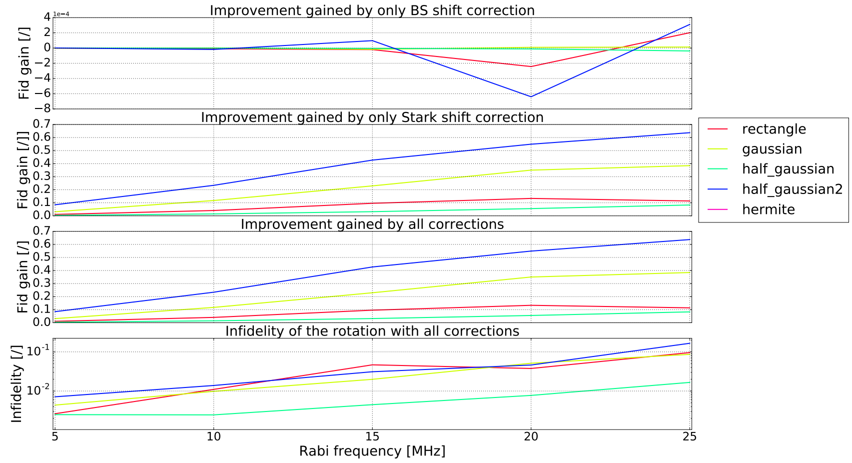

Figure C.18: Fidelity gain provided by the frequency shift correction algorithm on Gaussian-shaped π rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

Figure C.19: Fidelity gain provided by the frequency shift correction algorithm on HFTxD-shaped $\pi/2$ rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

Figure C.20: Fidelity gain provided by the frequency shift correction algorithm on HFTxD-shaped π rotations in the three-qubits three-drives configuration, for different combinations of shift corrections.

C.3. Scalability to higher numbers of qubits

The figure in this section summarises the evolution of the fidelity gain provided to π pulses by frequency shift correction, with respect to the number of qubits. Four representative values of Rabi frequency are covered.

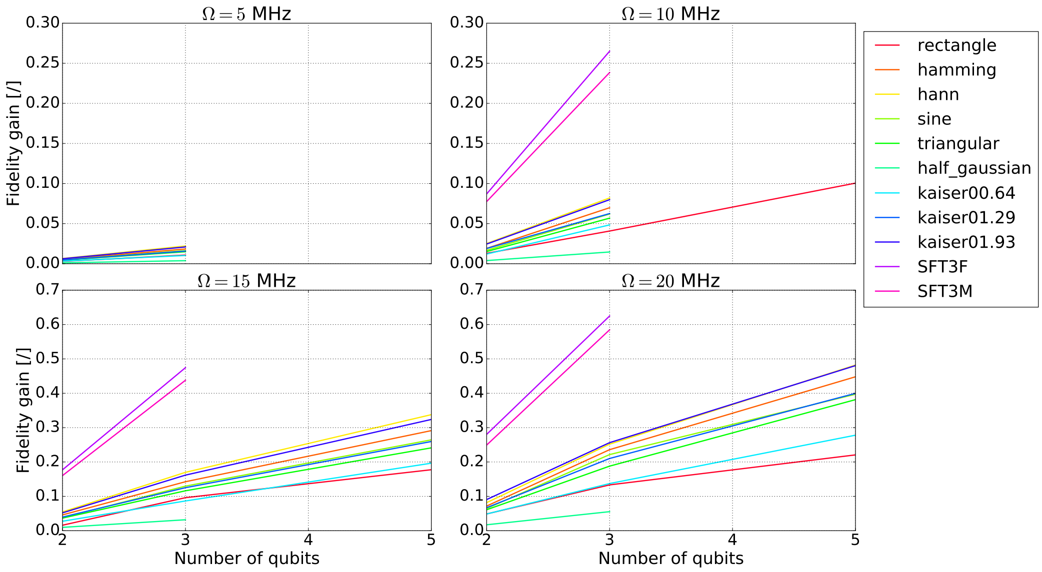

Figure C.21: Evolution of the fidelity gain provided by frequency shift correction on π pulses, with respect to the number of qubits. The evolution is presented for four representative values of Rabi frequency. Some very computationally intensive values to obtain have been omitted.

C.4. Miscellaneous and test cases

The figures in this section illustrate various test configurations, some of which too unrealistic to be labeled results. They nonetheless provide valuable insight into the difference of FT amplitude between corrected and uncorrected drives, as well as into the performance of the frequency correction algorithm in a crosstalk-free environment.

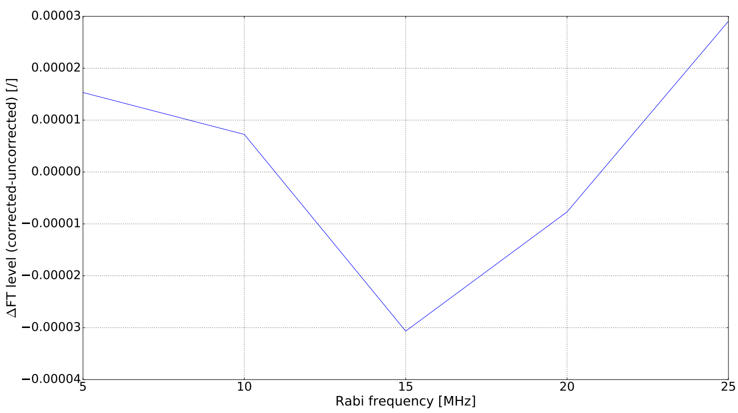

Figure C.22: Difference of the Fourier transform amplitude of a frequency corrected pulse with that of an uncorrected one, at the resonance frequency of the neighbouring qubit.
At $\Omega = 15$ MHz, the FT amplitude of the uncorrected pules is larger than that of the corrected one. At $\Omega = 25$ MHz, the opposite is true. This explains the larger fidelity improvement observed for $\Omega = 15$ MHz and the lower than expected at $\Omega = 25$ MHz.

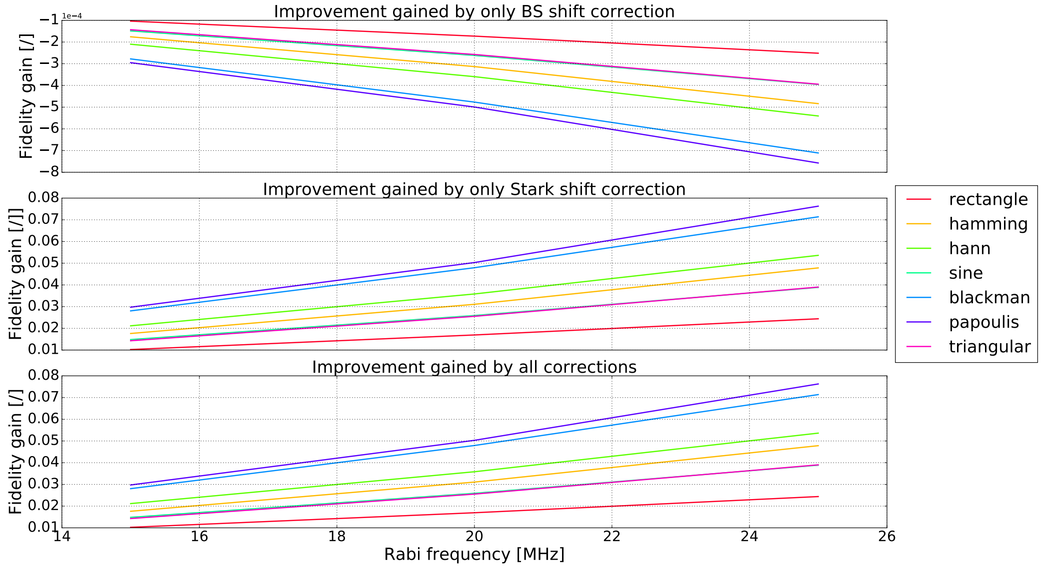

Figure C.23: Gain of average fidelity provided by the frequency shift correction algorithm on geometrical and raised-cosine π pulses in the testing crosstalk-free configuration, for different combinations of shift corrections. The single qubit is driven by a resonant pulse, but his resonance frequency experiences the shift of a second "phantom" pulse. This configuration is highly realistic and is only used here as a test case.

The two figures below present the performance of the frequency shift correction algorithm on the Blackman and Papoulis shapes used to drive a single qubit in a crosstalk-free environment. In Figure C.24, the shift-inducing pulses have frequencies higher than the qubit's. In Figure C.25, one of them is above the qubit's and the other is below. Since the configurations where the shift-inducing frequencies are both above or below the qubit's are completely symmetrical, these two figures illustrate the performance of the driving on each of the qubits of a three qubits array.

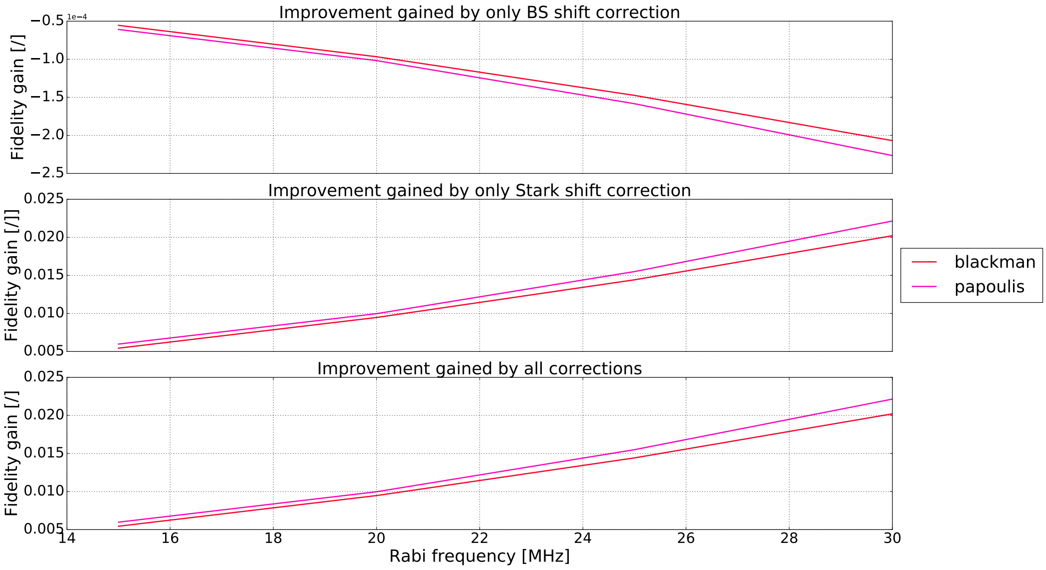

Figure C.24: Gain of average fidelity provided by the frequency shift correction algorithm on Blackman- and Papoulis-shaped $\pi/2$ pulses in the testing crosstalk-free configuration, for different combinations of shift corrections. The single qubit is driven by a resonant pulse, but his resonance frequency experiences the shift of a second and third "phantom" pulses with frequencies higher than the qubit's. This configuration is highly realistic and is only used here as a test case.

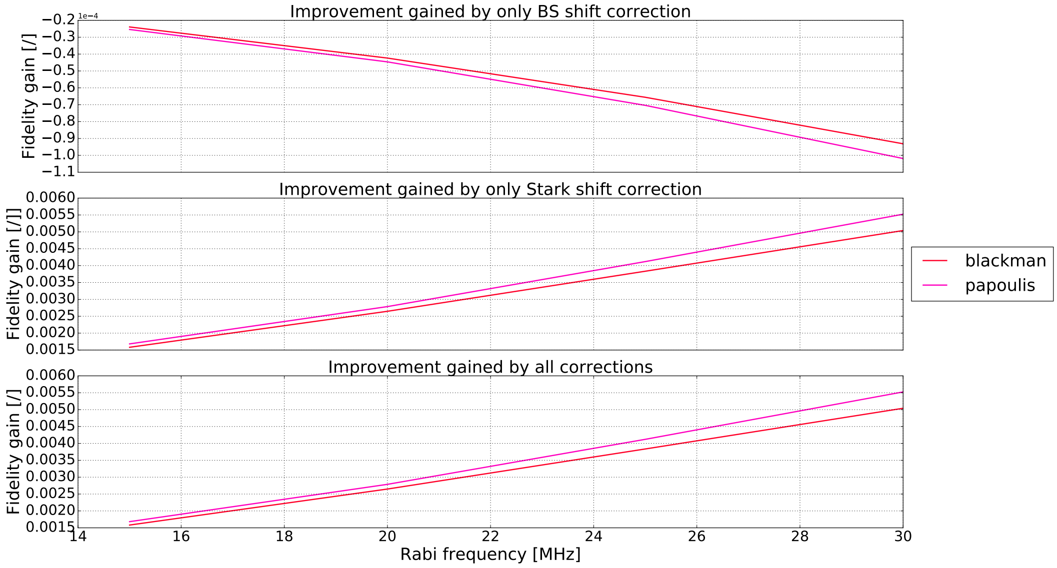

Figure C.25: Gain of average fidelity provided by the frequency shift correction algorithm on Blackman- and Papoulis-shaped $\pi/2$ pulses in the testing crosstalk-free configuration, for different combinations of shift corrections. The single qubit is driven by a resonant pulse, but his resonance frequency experiences the shift of a second and third “phantom" pulses. One of their frequencies is above the qubit’s and the other is below. This configuration is highly realistic and is only used here as a test case.

D

Additional figures of pulse shaping results

This appendix contains the detailed results of the grid search evaluating the performance of pulse shaping. These results are summarised, commented and interpreted in Chapter 5. They are presented here for the sake of completeness. The shapes have been grouped by family and presented separately for better readability. The block pulse is represented in all figures to serve as a reference.

Results are given first for a two-qubits two-drives configuration, and then for a three-qubits three-drives one. Each of these sections presents the results in the two case of pulses definition: first by their Rabi frequency and then by their pulse duration.

Each figure in this appendix shows the infidelity of either $\pi/2$- or π-pulses modulated by various shapes, for representative values of the pulse free parameter.

D.1. Two-qubits two-drive configuration

The results in this section are obtained when considering an array of two qubits, each of them being driven by a resonant pulse.

D.1.1. Results of fixed Rabi frequency pulse definition

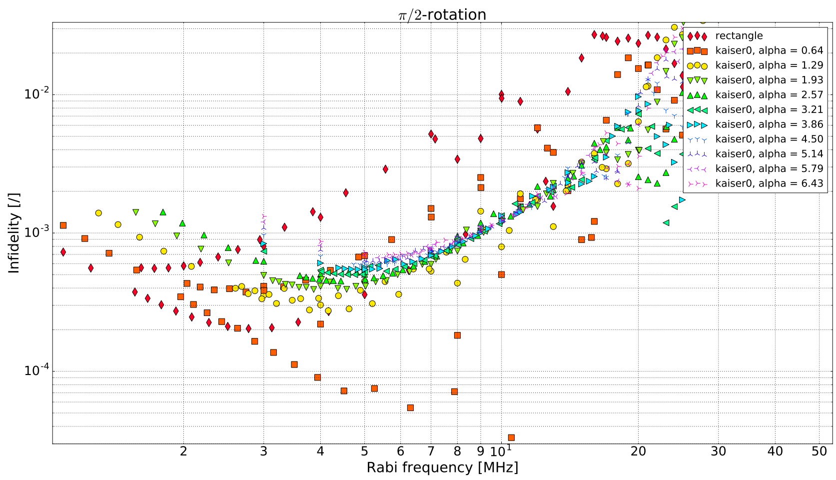

Figure D.1: Infidelity of $\pi/2$ rotations in two-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

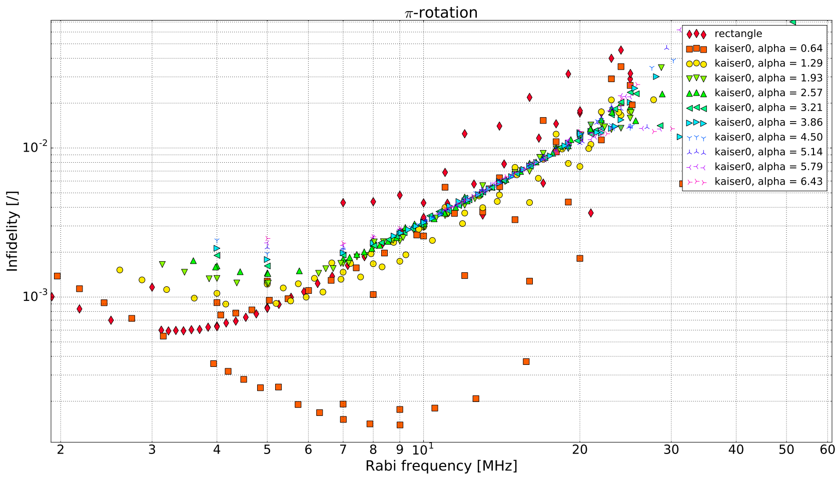

Figure D.2: Infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

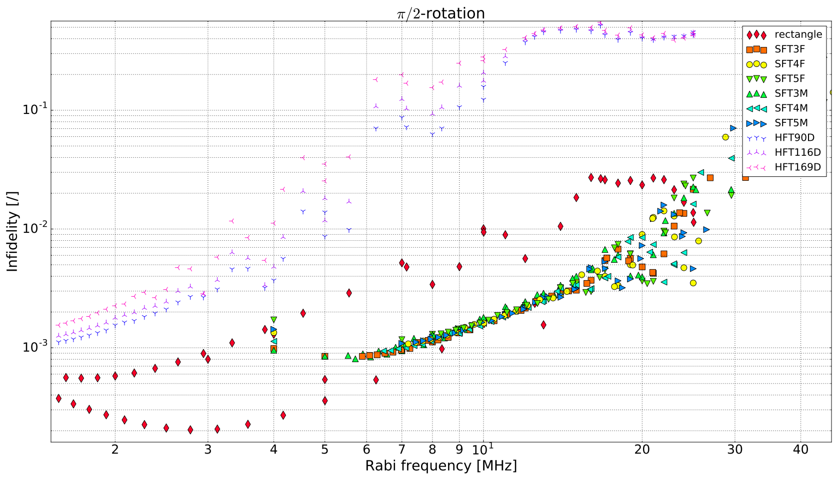

Figure D.3: Infidelity of $\pi/2$ rotations in two-qubits configurations. The driving pulses are modulated by shapes of the flat top 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

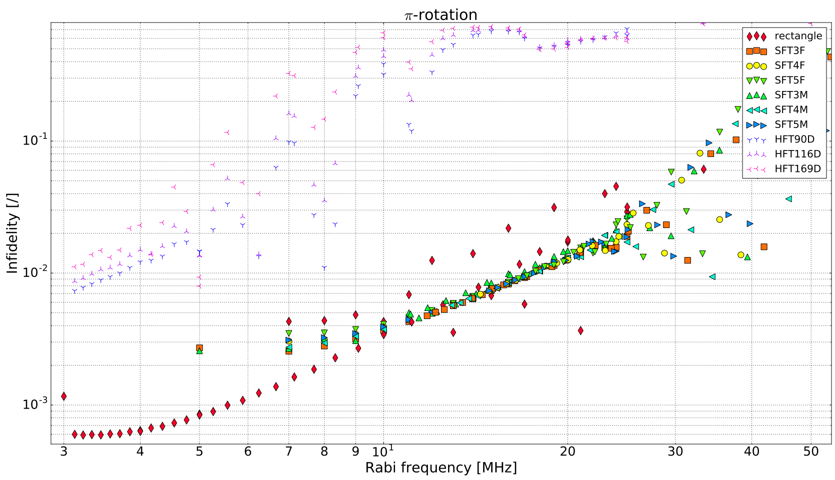

Figure D.4: Infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the flat top family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

D.1.2. Results of fixed duration pulse definition

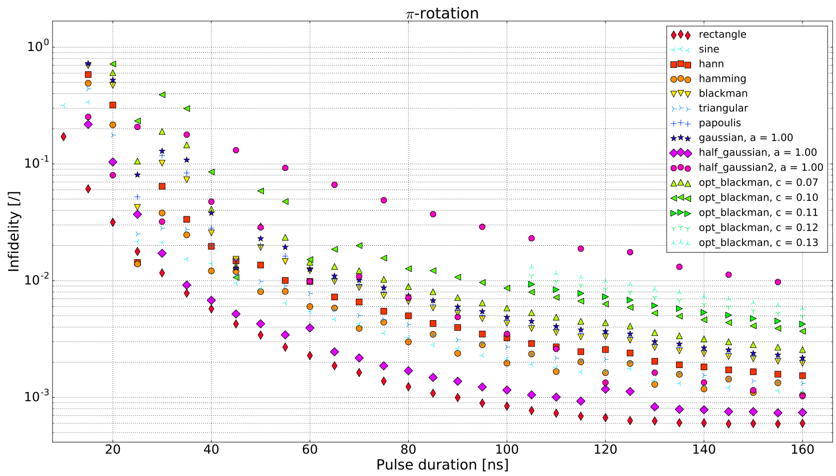

Figure D.5: Infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and set to have various durations.

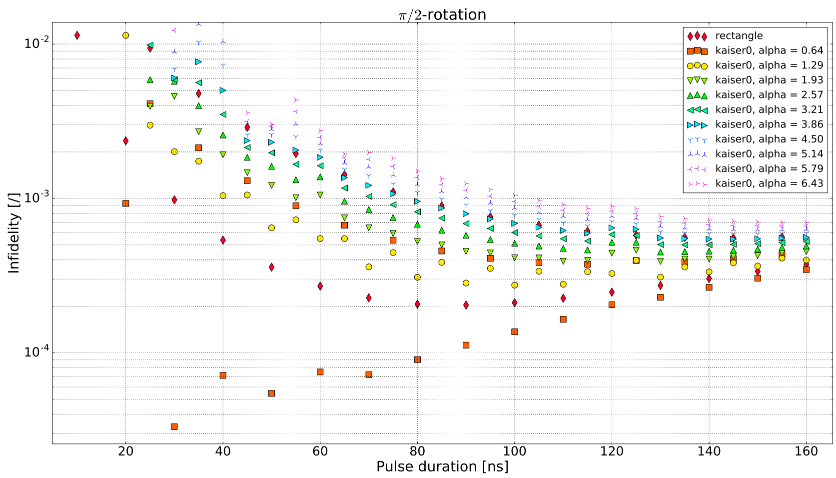

Figure D.6: Infidelity of $\pi/2$ rotations in two-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

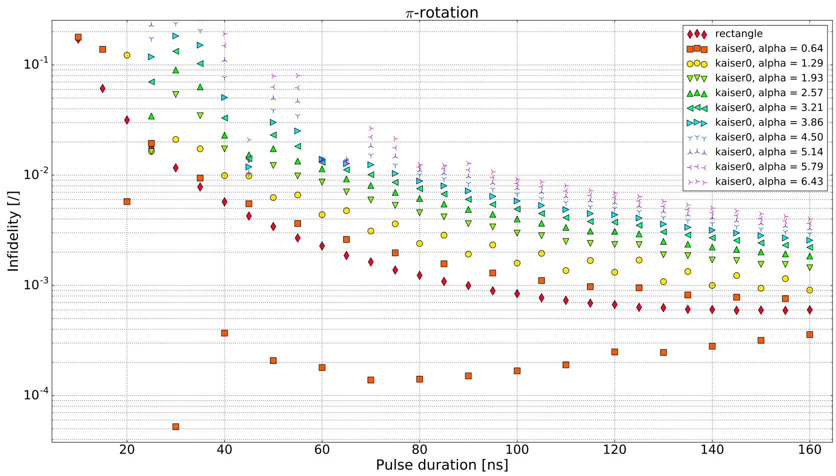

Figure D.7: Infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

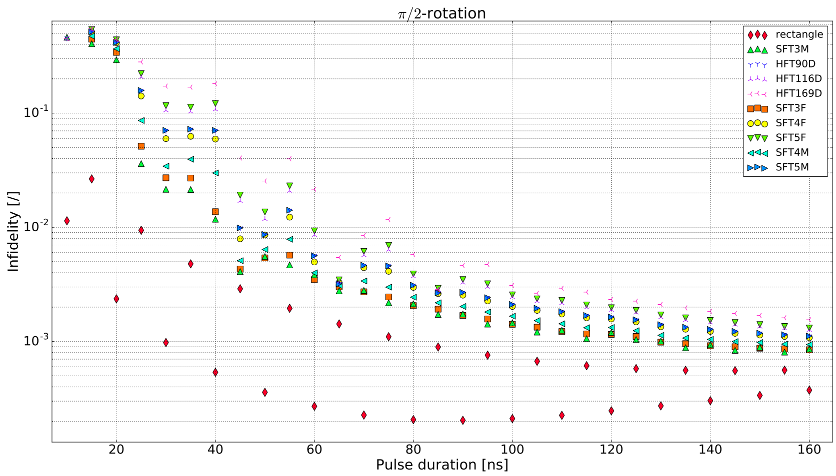

Figure D.8: Infidelity of $\pi/2$ rotations in two-qubits configurations. The driving pulses are modulated by shapes of the flat top 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

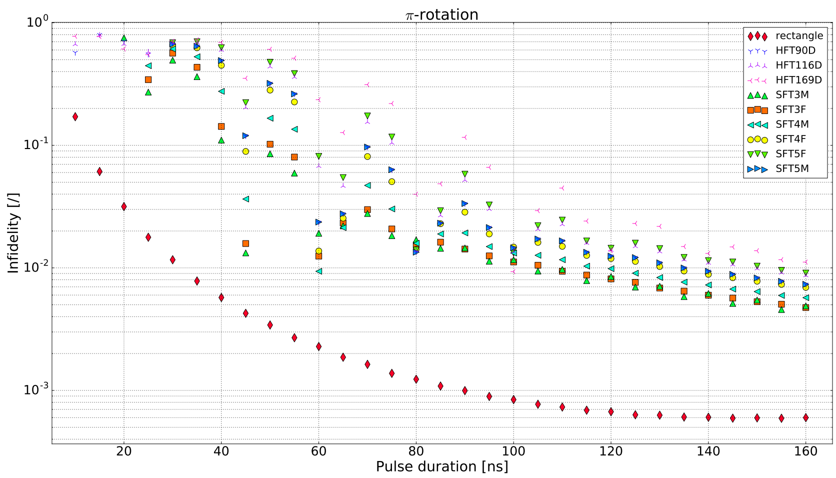

Figure D.9: Infidelity of π rotations in two-qubits configurations. The driving pulses are modulated by shapes of the flat top family, and set to have various durations. The rectangle shape is also represented as a reference.

D.2. Three-qubits three-drives configuration

The results in this section are obtained when considering a similar configuration as the previous section, but with three qubits and three resonant drives.

D.2.1. Results of fixed Rabi frequency pulse definition

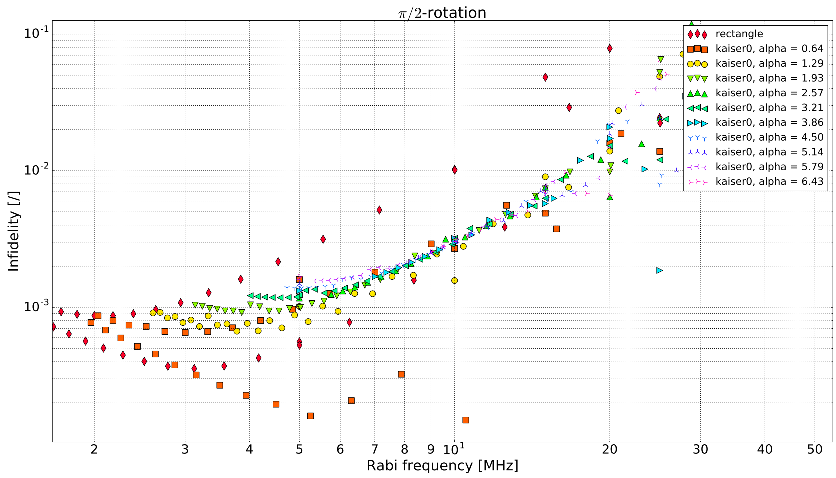

Figure D.10: Infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

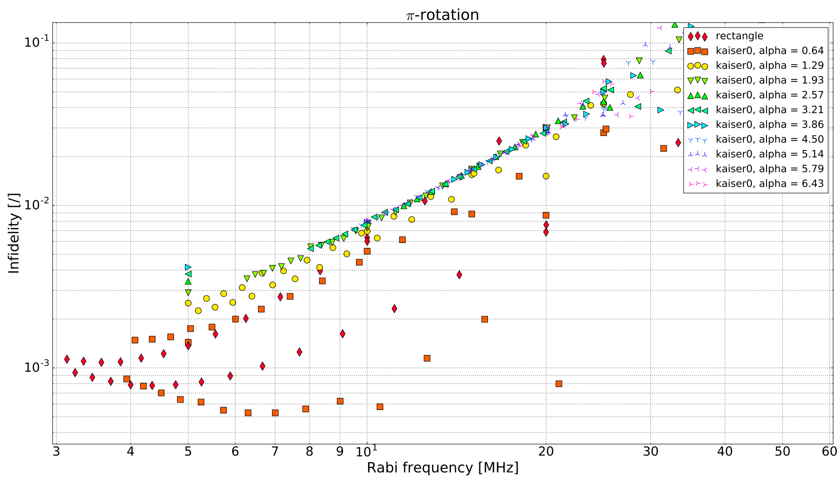

Figure D.11: Infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

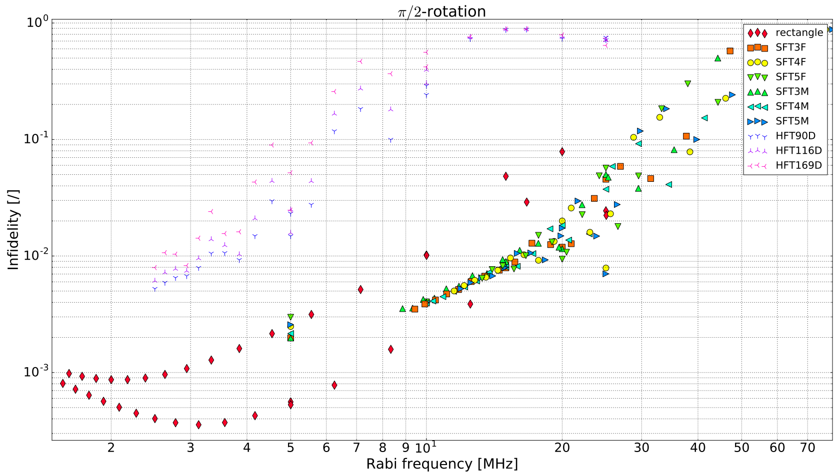

Figure D.12: Infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the flat top 0 family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

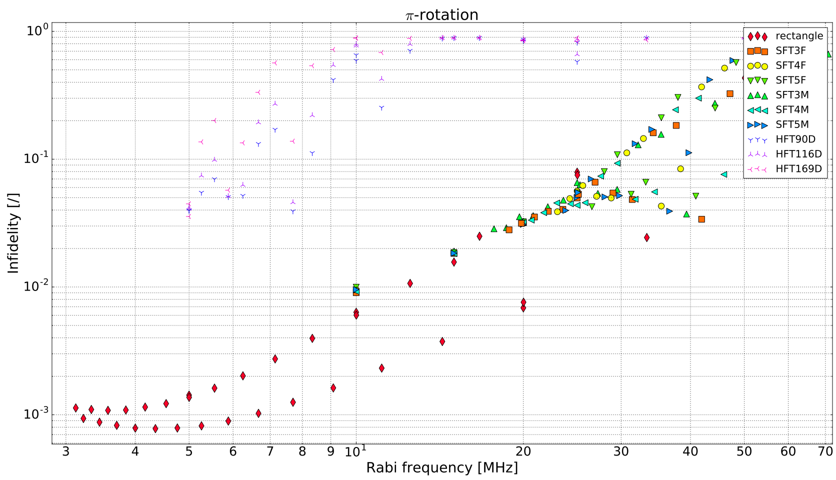

Figure D.13: Infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the flat top family, and used representative values of Rabi frequency. The rectangle shape is also represented as a reference.

D.2.2. Results of fixed duration pulse definition

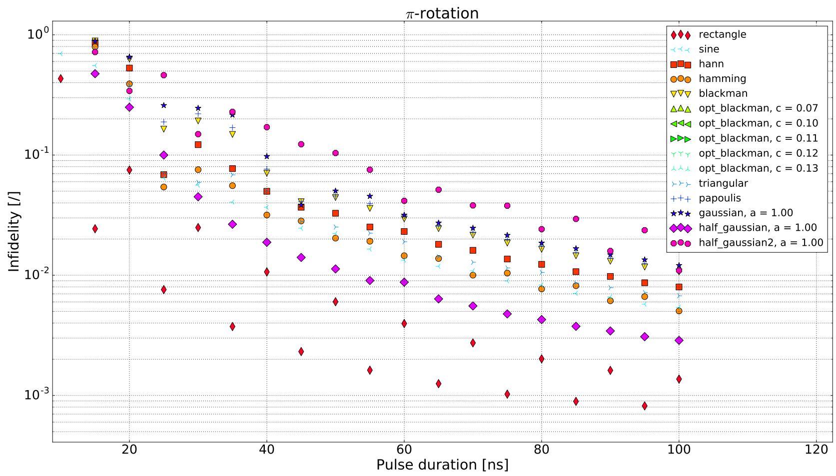

Figure D.14: Infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the geometrical, raised-cosine or gaussian families, and set to have durations.

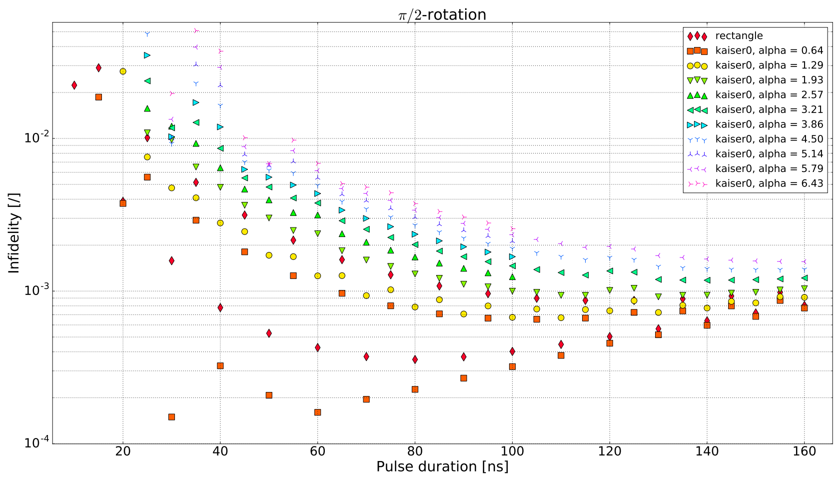

Figure D.15: Infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

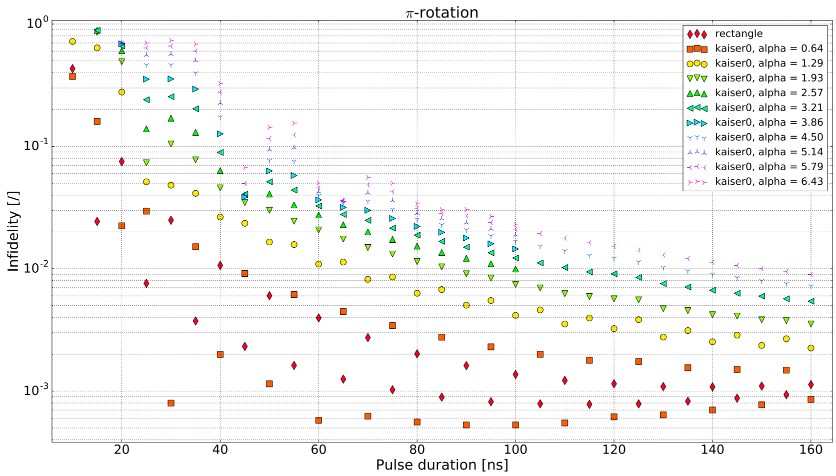

Figure D.16: Infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the Kaiser 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

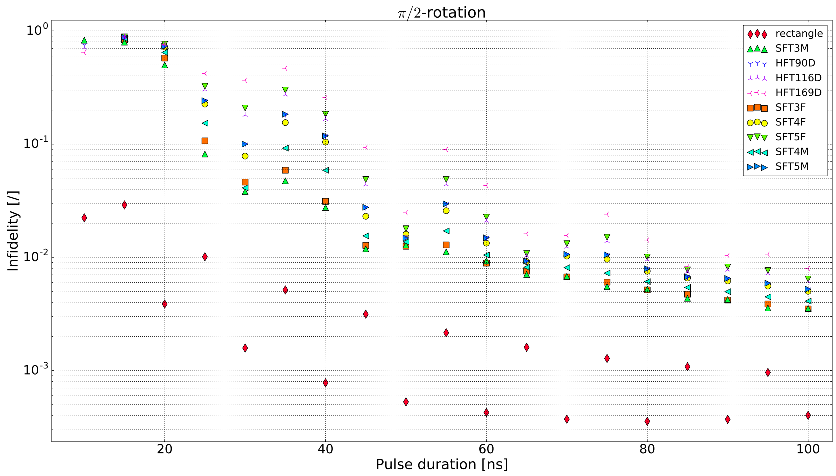

Figure D.17: Infidelity of $\pi/2$ rotations in three-qubits configurations. The driving pulses are modulated by shapes of the flat top 0 family, and set to have various durations. The rectangle shape is also represented as a reference.

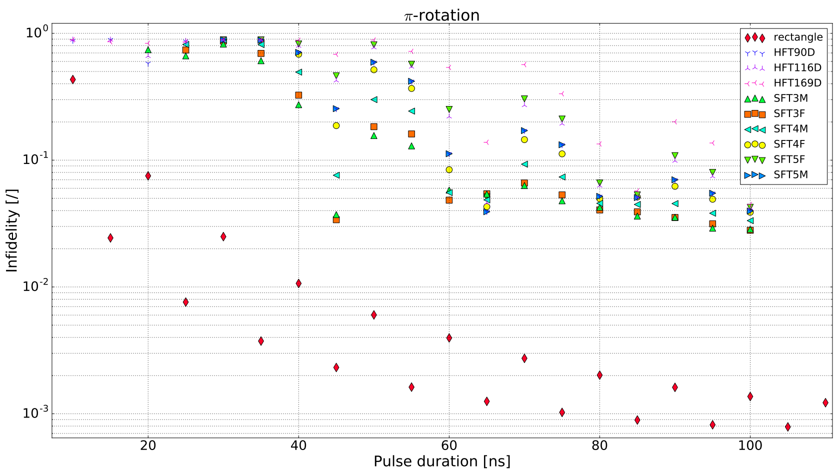

Figure D.18: Infidelity of π rotations in three-qubits configurations. The driving pulses are modulated by shapes of the flat top family, and set to have various durations. The rectangle shape is also represented as a reference.

D.3. Scalability to higher numbers of qubits

The figure in this section summarises the evolution of the unitary fidelity of π rotations with respect to the number of qubits, for various modulation shapes and representative values of Rabi frequency.

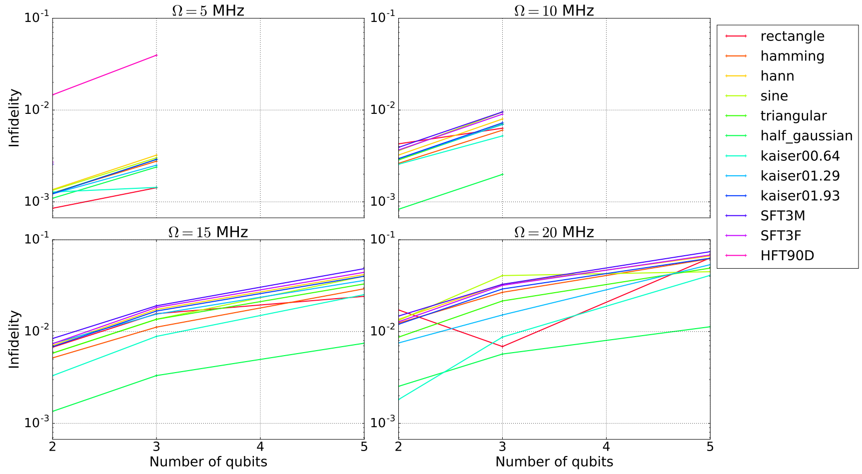

Figure D.19: Evolution of the unitary infidelity of shaped π rotations with respect to the number of qubits, for four representative values of Rabi frequency. The shapes investigated are those that achieved satisfactory performance with two and three qubits.

E

Additional figures of correlation between between shape characteristics and rotation fidelity

This Appendix completes Section 5.2.3 and regroup the same metric plots but for rotation angles and number of qubits not covered in the main body of the report. The interpretation of these figures is included in Section 5.2.3, as the trends are the same regardless of rotation angle and number of qubits.

E.1. Rotations in two-qubits array

This section presents the results for π rotations in a two-qubits array.

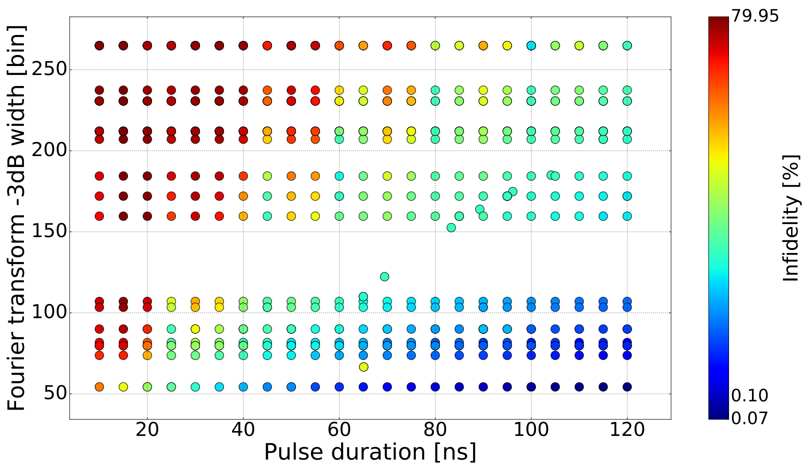

Figure E.1: Unitary infidelity as a function of the pulse duration and the width of the FT main lobe.

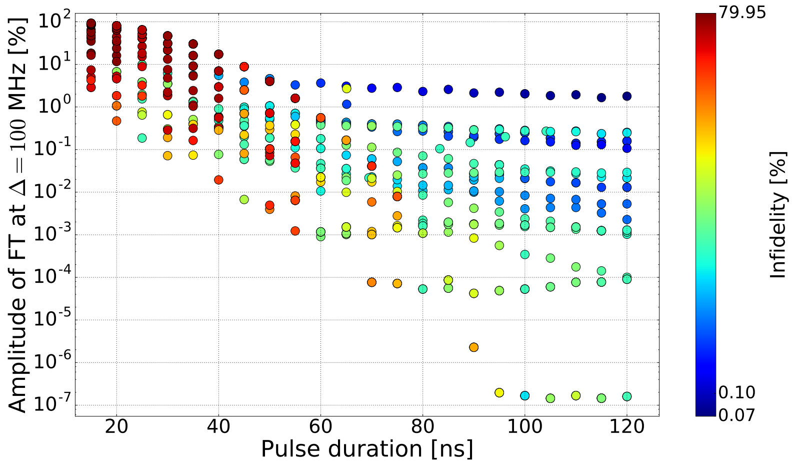

Figure E.2: Unitary infidelity as a function of the pulse duration and the FT amplitude at the frequency of the second qubit.

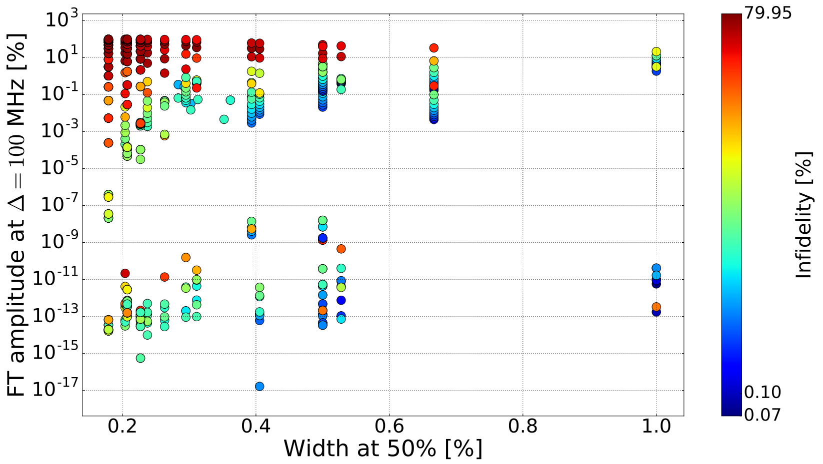

Figure E.3: Unitary infidelity as a function of the pulse width at 50% and of the Fourier transform amplitude of the pulse at the frequency of the second qubit.

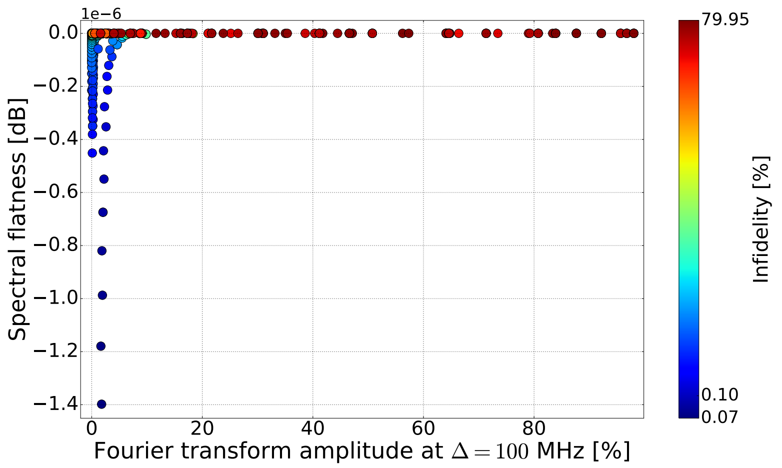

Figure E.4: Unitary infidelity as a function of the Fourier transform amplitude at the frequency of the second qubit and of the spectral flatness.

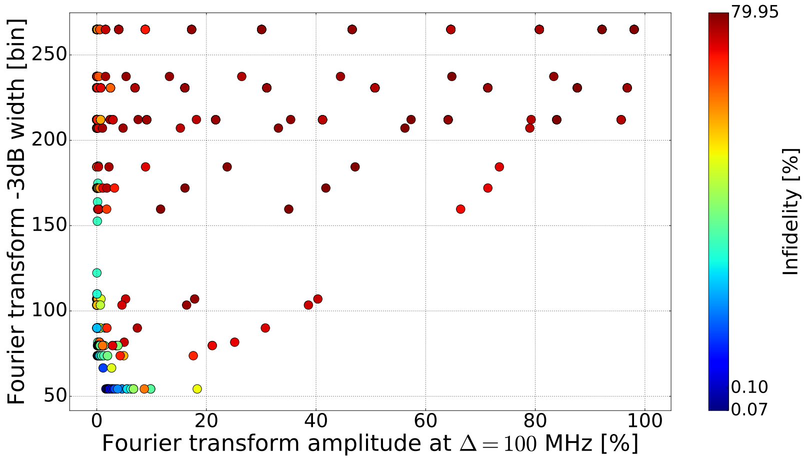

Figure E.5: Unitary infidelity as a function of the Fourier transform amplitude at the frequency of the second qubit and of the -3dB width of the Fourier transform main lobe.

E.2. Rotations in three-qubits array

In this subsection, the results for rotations in a three-qubits array are presented. In the cases where a plot has two panels, the left and right panels give the evolution for $\pi/2$ and π rotations, respectively.

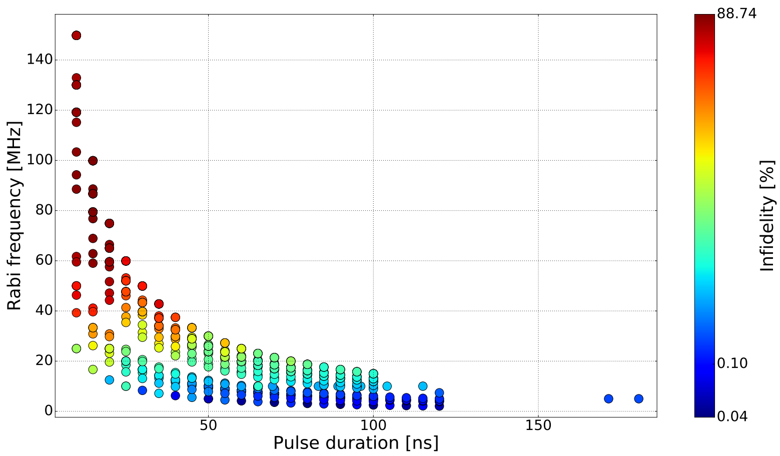

Figure E.6: Unitary infidelity as a function of the pulse duration and its Rabi frequency.

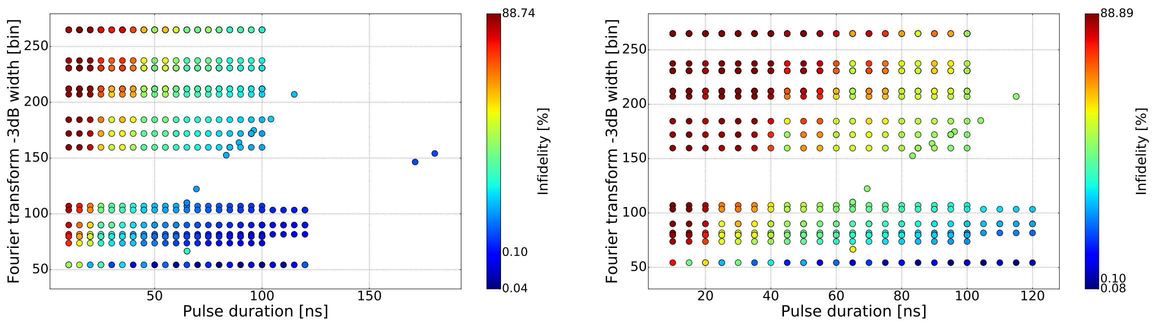

(a) $\pi/2$ rotation

(b) π rotation

Figure E.7: Unitary fidelity as a function of the pulse duration and the -3 dB width of the FT main lobe of the pulse

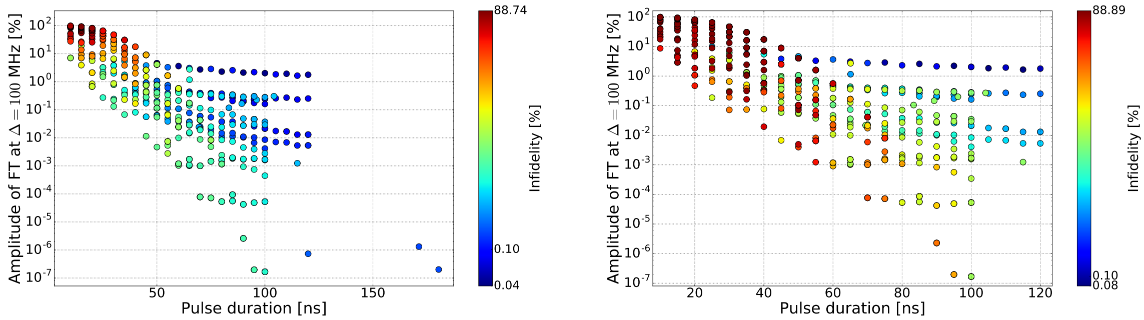

(a) $\pi/2$ rotation

(b) π rotation

Figure E.8: Unitary fidelity as a function of the pulse duration and the FT amplitude at $\Delta = 100$ MHz.

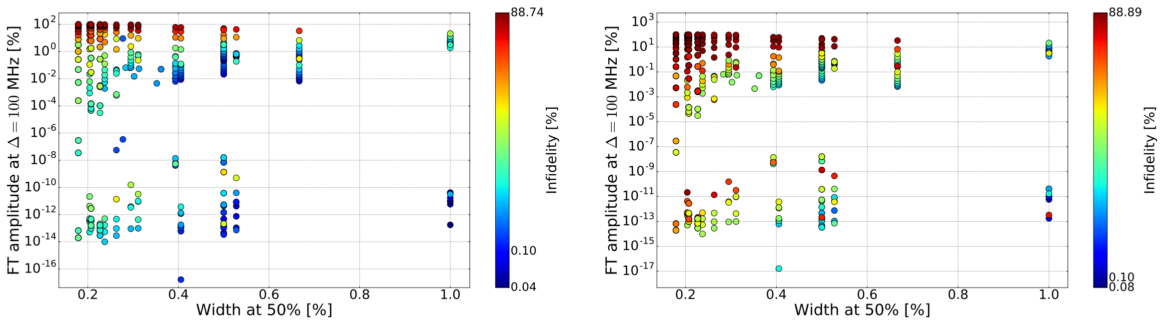

(a) $\pi/2$ rotation

(b) π rotation

Figure E.9: Unitary infidelity as a function of the pulse width at 50% and of the Fourier transform amplitude of the pulse at the frequency of the neighbouring qubits.

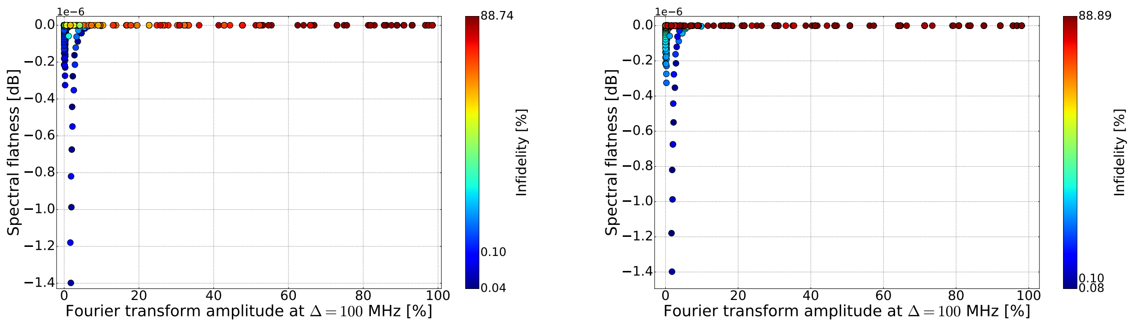

(a) $\pi/2$ rotation

(b) π rotation

Figure E.10: Unitary infidelity as a function of the Fourier transform amplitude at $\Delta = 100$ MHz and of the spectral flatness.

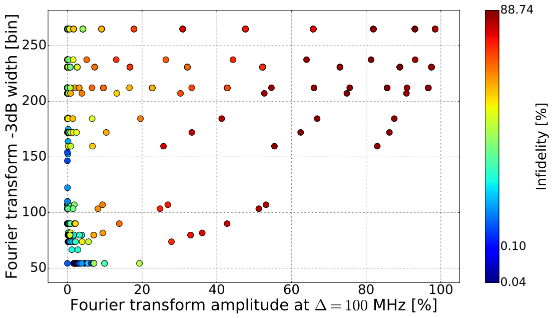

(a) $\pi/2$ rotation

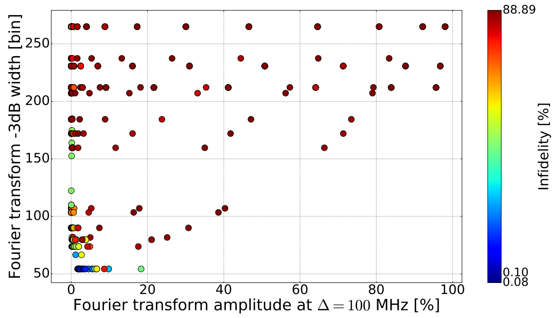

(b) π rotation

Figure E.11: Unitary infidelity as a function of the Fourier transform amplitude at the frequency of the second qubit and of the -3dB width of the Fourier transform main lobe.

F

Code documentation

The following appendix documents the main classes and methods of the software written the current work. An example of a simulation demonstrating the use of most of these is provided in *obj_driving.py*

F.1. Classes

n_qubit_simulation

General class representing a time-evolution simulation of a linear array of qubits driven my multiple microwave pulses at the same time.

- **init(self, nb_qubits, omega0, rotating_frame=True, rate_coupling=100e3)**
  Constructor. An exchange coupling of 100 kHz is simulated by default.
  <u>*Parameters:*</u>
  - *nb_qubits*: Number of qubits considered in the simulation
  - *omega0* : Array containing the resonance frequencies of the qubits
  - *rotating_frame*: Boolean indicating whether the simulation is to be done in the rotating frame of not
  - *rate_coupling*: rate of the exchange coupling between adjacent qubits, expressed in Hz.
- **add_mw_pulse(self, mw_pulse, rotating_frame=True, RWA=False)**
  Adds a microwave pulse to the simulation, while making the rotating-wave approximation (RWA) or not.
  <u>*Parameters:*</u>
  - *mw_pulse* : shaped_pulse class instance representing the drive applied to the linear array of qubits
  - *rotating_frame*: Boolean indicating whether the pulses will be added to the solver in the rotating frame or in the laboratory frame
  - *RWA* : Boolean indicating whether to make the RWA or not (True = RWA applied, False = RWA not applied)
- **calc_time_evolution(self, psi0, t_start, t_end)**
  Compute the time evolution of the system.
  <u>*Parameters:*</u>
  - *psi0* : Array containing the initial density matrix of the system
  - *t_start*: Time at which the simulation starts
  - *t_end* : Time at which the simulation ends
- **get_state_fidelity(self, target)**
  Returns the state fidelity of the operation compared to a target state.
  <u>*Parameters:*</u>

  - *target*: target state to compare with. target is an instance of QuTip Quobj.

  Returns:

  - *fid*: fidelity value (ranging between 0 and 1).

- **get_unitary_fidelity(self, target_unitary)**
  Returns the average gate fidelity of the process (unitary fidelity). The formula used is taken from Pedersen et al. (2007) (https://www.sciencedirect.com/science/article/pii/S0375960107003271?via%3Dihub).
  Parameters:

  - *target_unitary*: target unitary to compare with. target_unitary is an instance of QuTip Quobj.

  Returns:

  - *fid*: value of unitary fidelity (ranging between 0 and 1).

- **plot_expectations(self, fig_nb)**
  Plots the X, Y and Z operator expectation values for each qubit in the simulation.
  Parameters:

  - *fig_nb*: number of the matplotlib.figure object in which the evolution is to be plotted.

- **plot_bloch(self, fig_nb)**
  Plots the evolution of the simulation on the Bloch sphere. The end state of each qubit is indicated by an arrow.
  Parameters:

  - *fig_nb*: number of the matplotlib.figure object in which the evolution is to be plotted.

- **return_expectations(self)**
  Returns the X, Y and Z operator expectation values for each qubit in the simulation.
  The expectations are computed using the dedicated method built-in DM_solver. Returns:

  - *expect*: 1D array of size $3 \times n$ (with n the number of qubits in the simulation), containing the X, Y and Z expectations of each qubit.
    The values are ordered in the following way: (X_qubit-0, X_qubit-1, ..., X_qubit-n, Y_qubit-0, Y_qubit-1, ..., Y_qubit-n, Z_qubit-0, Z_qubit-1, ..., Z_qubit-n).

shaped_pulse

Class representing the driving pulses to be applied to the system. It contains the caracteristics of each of the driving pulses as well as a separate pulse data per qubit in the simulation.

- **init(self, nb_qubits, nb_drives, rotation_angle, fixed_parameter, chosen_shape, omega0, omega, phase, arg=(0.11, 5, 8, 1.0), rise_time=True, fixed_pulse_dur=False)**
  Constructor.
  Parameters:

  - *nb_qubits* : Number of qubits considered in the simulation
  - *nb_drives* : Number of driving pulses. If nb_qubits > nb_drives, the drives will be applied to the n_drives first qubits
  - *rotation_angle* : Array containing the target rotation angles of each of the driving pulses
  - *fixed_parameter*: Parameter fixed for the generation of the pulse.
    If the element is an array, it corresponds to the Rabi frequencies of each driving pulse.
    If the element is a number, it corresponds to the pulse duration of all the driving pulses.
  - *chosen_shape* : Array containing the amplitude modulation shape of each of the driving pulses
  - *omega0* : Array containing the resonance frequencies of the qubits
  - *omega* : Array containing the frequencies of each of the driving pulses

- *phase* : Array containing the initial phase of each of the driving pulses
- *arg* : Tuple containing the shape function arguments: (c, alpha, gamma, a) for optimised blackman, kaiser 0, kaiser1 and gaussian modulation shapes
- *rise_time* : Boolean determining if the effects of pulse rise time are to be applied directly by the constructor
- *fixed_pulse_dur*: Boolean determining if the pulse must be created using a pre-specified pulse duration (and hence adapt its Rabi frequency) or using a fixed Rabi frequency (and hence adapt its duration)

• **add_pulse_rise_time(self, fixed_pulse_duration=False)**
Adds finite rise time to the amplitude modulated pulse, using third order lowpass filtering.
It modifies the pulse amplitude data, as well as the pulse duration (and derived variables, e.g. number of time points, of time slices, etc) to conserve the target rotation angle.
Parameters:

- *fixed_pulse_duration*: Boolean value that indicates whether the addition of rise time must be done with or without modifying the total pulse duration (and hence the number of time points and other relevant values).

• **add_frequency_shifts(self, omega0, correction_type)**
Adds the effect of coupling between the qubits and the driving pulses on the qubits' resonance frequencies (AC Stark shift, Bloch-Siegert shift), and stores the time evolution of the resonance frequencies.
Implements a correction scheme to choose the drive frequency of each of the pulse such that the shifts are minimised, and tracks the remaining shift over time.
Parameters:

- *omega0* : Array containing the resonance frequencies of the qubits
- *correction_type*: String specifying which shift(s) must be accounted for in the correction scheme. Both AC Stark and Bloch-Siegert shifts are however applied to the qubit's resonance frequencies in each case.

'all'	: all shifts (AC Stark and Bloch-Siegert)
'stark'	: only AC Stark shift
'bs'	: only Bloch-Siegert
'none'	: no shift (no correction is applied)

• **plot_frequency_evolution(self, fig_nb1, fig_nb2)**
Plots the evolution of both the resonance frequencies of the simulated qubits and the frequencies of each of the driving pulses., as well as the phase of each pulse.
Parameters:

- *fig_nb1*: number of the matplotlib.figure object in which the evolution of the frequencies is to be plotted.
- *fig_nb1*: number of the matplotlib.figure object in which the evolution of the phases is to be plotted.

• **plot_pulse_data(self, fig_nb)**
Plots the enveloppe of the driving pulse in the time domain.
Parameters:

- *fig_nb*: number of the matplotlib.figure object in which the evolution is to be plotted.

arbitrary_pulse

Class representing a shaped pulse of arbitrary duration and Rabi frequency.

- **init(self, nb_qubits, nb_drives, rabi_frequency, length_pulse, chose_shape, omega0, omega, phase, arg=(0.11, 5, 8, 1.0), rise_time=True)**
  Constructor of the aribtrary_pulse class. A rotation angle is not specified. In order to create a pulse that implements a specific rotation angle, use the shaped_pulse class.
  Parameters:

  - *nb_qubits* : Number of qubits considered in the simulation
  - *nb_drives* : Number of driving pulses. If nb_qubits > nb_drives, the drives will be applied to the n_drives first qubits
  - *rabi_frequency* : Rabi frequencies of each driving pulse.
  - *length_pulse* : Pulse duration of all the driving pulses.
  - *chosen_shape* : Array containing the amplitude modulation shape of each of the driving pulses
  - *omega0* : Array containing the resonance frequencies of the qubits
  - *omega* : Array containing the frequencies of each of the driving pulses
  - *phase* : Array containing the initial phase of each of the driving pulses
  - *arg* : Tuple containing the shape function arguments: (c, alpha, gamma, a) for optimised blackman, kaiser 0, kaiser1 and gaussian modulation shapes
  - *rise_time* : Boolean determining if the effects of pulse rise time are to be applied directly by the constructor

F.2. Utility functions outside of class structures

Pulse shape definition functions

The functions defined below are indexed in the dictionary *shapes* for easy access. Because most of the shape functions share identical structures and signatures, only the special cases will be presented here. Table F.1 summarises all the shapes that obey to the templates given below, and give their corresponding name in the implementation.

- **rectangle_window(t, args)**
  Computes the value of a rectangle modulation shape at a given time.
  Parameters:

  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    't_end' : end time of the desired shape (in s)
    'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)

  Returns:

  - *res*: the value of the modulation shape defined by *args* at time *t*.

Modulation shape	Function name
Rectangle	rectangle_window
Step	step_window
Hamming	hamming_window
Hann	hann_window
“Concave" Hamming	hamming_window_alt
Triangular	triangular_window
Sine	sine_window
Blackman	blackman_window
Papoulis	papoulis_window
SFT3F	SFT3F_window
SFT4F	SFT4F_window
SFT5F	SFT5F_window
SFT3M	SFT3M_window
SFT4M	SFT4M_window
SFT5M	SFT5M_window
HFT90D	HFT90D_window
HFT116D	HFT116D_window
HFT169D	HFT169D_window

Table F.1: Modulation shape names and the names of the corresponding Python functions.

- **optimized_blackman_window(t, args)**
  Computes the value of an optimised Blackman modulation shape at a given time.
  Parameters:
  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    ’t_end’ : end time of the desired shape (in s)
    ’rabi_freq’ : Rabi frequency of the pulse (in Hz.rad)
    ’c’ : value of the shape’s free parameter

  Returns:
  - *res*: the value of the modulation shape defined by *args* at time *t*.
- **kaiser0_window(t, args)**
  Computes the value of a 0^{th} order Kaiser modulation shape at a given time.
  Parameters:
  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    ’t_end’ : end time of the desired shape (in s)
    ’rabi_freq’ : Rabi frequency of the pulse (in Hz.rad)
    ’alpha’ : value of the shape’s free parameter

  Returns:
  - *res*: the value of the modulation shape defined by *args* at time *t*.
- **kaiser1_window(t, args)**
  Computes the value of a 1^{st} order Kaiser modulation shape at a given time.
  Parameters:
  - *t*: time point at which to compute the value of the modulation shape

  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    't_end' : end time of the desired shape (in s)
    'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)
    'alpha' : value of the shape's free parameter

  Returns:

  - *res*: the value of the modulation shape defined by *args* at time *t*.

- **gaussian_window(t, args)**
  Computes the value of a Gaussian modulation shape at a given time.
  Parameters:

  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    't_end' : end time of the desired shape (in s)
    'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)
    'a' : value of the shape's free parameter

  Returns:

  - *res*: the value of the modulation shape defined by *args* at time *t*.

- **half_gaussian_window(t, args)**
  Computes the value of a half-Gaussian modulation shape at a given time.
  Parameters:

  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    't_end' : end time of the desired shape (in s)
    'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)
    'a' : value of the shape's free parameter

  Returns:

  - *res*: the value of the modulation shape defined by *args* at time *t*.

- **hermite_window(t, args)**
  Computes the value of a Gaussian modulation shape at a given time.
  Parameters:

  - *t*: time point at which to compute the value of the modulation shape
  - *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
    't_end' : end time of the desired shape (in s)
    'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)
    'a' : value of the gaussian's free parameter
    'hermite_a' : value of the polynomial's 0^{th} order coefficient
    'hermite_b' : value of the polynomial's 2^{nd} order coefficient

  Returns:

  - *res*: the value of the modulation shape defined by *args* at time *t*.

- **slepian_window(t, args)**
  Computes the value of a Slepian modulation shape at a given time.
  Parameters:

– *t*: time point at which to compute the value of the modulation shape
– *args*: dictionnary containing the arguments used by the shape computations. args must contain the keys:
't_end' : end time of the desired shape (in s)
'rabi_freq' : Rabi frequency of the pulse (in Hz.rad)
'alpha' : value of the shape's free parameter

Returns:

– *res*: the value of the modulation shape defined by *args* at time *t*.

Pulse duration computation functions

Similarly to the pulse shape functions, several functions here also share the same structure and signature. As a consequence, only one of them is detailed here.

- **compute_pulse_duration(rotation_angle, shape, rabi_frequency, arg)**
  Compute the duration required to perform a desired rotation with a shaped pulse.
  Parameters:

  – *rotation_angle*: desired rotation angle, in radians
  – *shape* : string specifying the name of the shape of the pulse
  – *rabi_frequency*: amplitude of the pulse, in hertz
  – *arg* : tuple containing the values of the parameters c, alpha and gamma for optimised Blackman, Kaiser 0 and Kaiser 1 shapes respectively

  Returns:

  – *pulse_duration*: Pulse duration required to perform the rotation with the given pulse, in seconds.

- **integr_rectangle(t, args)**
  Computes the residual between the integral of the modulation shape and the target rotation angle. Function used in the computation of the pulse duration.
  Parameters:

  – *x* : integration variable (duration of the pulse)
  – *arg*: tuple containing the parameters of the shapes.
  For more generality, values for the possible free parameters of the functions are also expected. In case the desired shape does not have free parameters, any dummy value can be provided.
  The tuple is organised as follows: c (optimsed blackman shape parameter), alpha (0th order Kaiser shape parameter), gamma (1st order Kaiser shape parameter), a (gaussian and gaussian-like shape parameter), rabi_frequency, rotation_angle.

  Returns:

  – *residual*: residual between rotation_angle/rabi_frequency and integral of the shape over the duration x

- **filtered_shape(end_time, shape, args)**
  Returns an array containing the value of the filtered shape function over its whole duration.
  Parameters:

  – end_time: total duration of the pulse (in ns)
  – shape : string specifying the name of th shape
  – args : tuple containing the parameters of the shapes. For more generality, values for the possible free parameters of the functions are also expected. In case the desired shape does not have free parameters, any dummy value can be provided.
  The tuple is organised as follows: c (optimsed blackman shape parameter), alpha (0th order Kaiser shape parameter), gamma (1st order Kaiser shape parameter), a (gaussian and gaussian-like shape parameter), rabi_frequency, rotation_angle.

<u>*Results:*</u>

- *filtered_shape* Array containing the time domain evolution of the filtered shaped pulse.

• **integr_shape_rt(x, *arg)**
General purpose residual computation function for the pulses simulating rise time. This function has the same objective as **integr_rectangle** and the other similar residual computation functions. It is however designed to work for all shape functions.

- *x* : integration variable (duration of the pulse)
- *arg*: tuple containing the parameters of the shapes.
  For more generality, values for the possible free parameters of the functions are also expected. In case the desired shape does not have free parameters, any dummy value can be provided.
  The tuple is organised as follows: c (optimsed blackman shape parameter), alpha (0th order Kaiser shape parameter), gamma (1st order Kaiser shape parameter), a (gaussian and gaussian-like shape parameter), rabi_frequency, rotation_angle, omega0, omega0_max, nb_time_pts.

<u>*Returns:*</u>

- *residual*: residual between rotation_angle/rabi_frequency and integral of the shape over the duration x.

• **compute_pulse_duration_rt(rotation_angle, shape, rabi_frequency, nb_time_pts, omega0, omega0_max, arg)**
Compute the duration required to perform a desired rotation with a shaped pulse, which includes the modeling of rise time effect on the AWG.
<u>*Parameters:*</u>

- *rotation_angle*: desired rotation angle, in radians
- *shape* : string specifying the name of the shape of the pulse
- *rabi_frequency*: amplitude of the pulse, in hertz
- nb_time_pts : number of time samples in the pulse
- omega0 : frequency of the pulse (in Hz.rad)
- omega0_max : largest frequency in the considered qubit array (in Hz.rad)
- *arg* : tuple containing the values of the parameters c, alpha and gamma for optimised Blackman, Kaiser 0 and Kaiser 1 shapes respectively

<u>*Returns:*</u>

- *pulse_duration*: Pulse duration required to perform the rotation with the given pulse, in seconds.

Frequency shifts and other utility functions

• **compute_FT(drive_settings, duration=0)**
Compute and returns the Fourier transform of the pulse defined by the parameters contained in the dictionary drive_settings.
<u>*Parameters:*</u>

- *drive_settings* : dictionnary containing the parameters used to define the driving pulse.
  They keys of the dictionary are:
  'rotation' : desired rotation angle (in radians)
  'rabi' : Rabi frequency of the pulse (in Hz.rad)
  'shape' : string containing the name of the shape to use
  'omega' : frequency of the driving pulse whose Fourier transform is computed (in Hz.rad)
  'arg' : Tuple containing the shape function arguments: (c, alpha, gamma, a) for optimised blackman, kaiser 0, kaiser1 and gaussian modulation shapes

<u>*Returns:*</u>

  - *fft_signal*: Real frequency components of the Fourier transform of the pulse
  - *freqs* : Frequencies corresponding to the FT elements (in Hz)

- **compute_power_density(drive_settings, duration=0)**
  Compute and returns the energy spectral density of the pulse defined by the parameters contained in the dictionary drive_settings. *Parameters:*

  - *drive_settings* : dictionnary containing the parameters used to define the driving pulse.
    They keys of the dictionary are:
    'rotation' : desired rotation angle (in radians)
    'rabi' : Rabi frequency of the pulse (in Hz.rad)
    'shape' : string containing the name of the shape to use
    'omega' : frequency of the driving pulse whose Fourier transform is computed (in Hz.rad)
    'arg' : Tuple containing the shape function arguments: (c, alpha, gamma, a) for optimised blackman, kaiser 0, kaiser1 and gaussian modulation shapes

  Returns:

  - *s*: energy spectral density of the pulse
  - *freqs* : Frequencies corresponding to the spectral density elements (in Hz)

- **compute_multiqubit_hamil(template, nb_qubits, index_template, dim_I=2)**
  Compute a Hamiltonian of the form IxIx...xHxIx...xI, where there in total nb_qubits elements in the tensor product, and the index of the one that is not I is given by index_template.
  Parameters:

  - *template*: element of the tensor product not equal to identity (a QuTip Quobj is expected).
  - *nb_qubits*: number of elements in the tensor product.
  - *index_template*: index of the element of the tensor product equal to *template*
  - *dim_I*: dimension of the identity matrix in the tensor product

  Returns:

  - *h*: computed Hamiltonian (a QuTip Quobj)

- **compute_free_evolution_hamiltonian(nb_qubits, nb_drives, omega, omega0, hbar)**
  Compute the free evolution Hamiltonian of one or more two-level systems coupled to the same number of driving microwave fields (in the rotating frame of the drive).
  Parameters:

  - *nb_qubits*: Number of qubits in the system
  - *nb_drives*: Number of drives addressing the qubits
  - *omega*: Array containing the drive frequencies (in Hz.rad)
  - *omega0*: Array containing the resonance frequencies of the qubits (in Hz.rad)
  - *hbar*: Reduced Planck constant (left here as an input so that the desired degree of precision can be applied)

  Returns:

  - *h_free_evolution*: free evolution Hamiltonian of the considered system.

- **compute_BS_shift(omega0, omega, rabi_freq, order=0)**
  Compute the shift in resonance frequency of the qubit due to Bloch-Siegert shift (formulae taken from Steck 2007 and Shirley 1965).
  Parameters:

  - *omega0* : resonance frequency of the qubit before shift (in Hz.rad)
  - *omega* : frequency of the driving pulse (in Hz.rad)

  - *rabi_freq*: Rabi frequency of the pulse causing the shift (in Hz.rad)
  - *order* : order of the correction to be computed (so far, only 0, 1, 4 and 6 are implemented)

  Returns:

  - *delta_omega_bs* : Bloch-Siegert shift to the resonance frequency

- **compute_Stark_shift(omega0, omega, rabi_freq)**
  General expression for the AC Stark shift. Projects the value from the dressed rotation axis to the uncoupled one.
  Parameters:

  - *omega0* : resonance frequency of the qubit before shift (in Hz.rad)
  - *omega* : frequency of the driving pulse causing the shift (in Hz.rad)
  - *rabi_freq*: Rabi frequency of the pulse (in Hz.rad)

  Returns:

  - *delta_omega_s* : AC Stark shift to the resonance frequency

- **compute_total_freq_shift(omega0, omega, rabi, nb_qubits, nb_drives, i_slice=1, nb_time_slices=2, order=6, correction_type='all')**
  Compute the total frequency shift caused by *nb_drives* driving pulses on each of the *nb_qubits* qubits considered.
  Parameters:

  - *omega0* : array containing the resonance frequencies of the considered qubits before shift (in Hz.rad)
  - *omega* : array containing the frequencies of the drives causing the shifts (in Hz.rad)
  - *rabi*: array containing the Rabi frequency of the pulses causing the shifts (in Hz.rad)
  - *nb_qubits*: number of qubits in the system
  - *nb_drives*: number of driving pulses causing the shifts
  - *i_slice*: index of the current time slice
  - *nb_time_slice*: total number of time slices
  - *order*: order of approximation for the Bloch-Siegert shift computation
  - *correction_type*: string indicating which type of shift should be taken into account in the computations. The values are:
    'none' : no shift is computed (returns an array of 0)
    'stark' : only the AC Stark shift is computed
    'bs' : only the Bloch-Siegert shift is computed
    'all' : both AC Stark and Bloch-Siegert shifts are computed

  Returns:

  - *total_shift* : 1D array containing the total shift per resonance frequency (ordered the same way as *omega0*)

- **objective_fn_freq_shift(omega, omega0_initial, rabi, correction_type)**
  Objective function used in the optimisation part of the frequency shift correction algorithm (numerical solving of Equation 3.1).
  Parameters:

  - *omega* : array containing the frequencies of the drives causing the shifts (in Hz.rad)
  - *omega0_initial* : array containing the resonance frequencies of the considered qubits before shift (in Hz.rad)
  - *rabi*: array containing the Rabi frequency of the pulses causing the shifts (in Hz.rad)

- *correction_type*: string indicating which type of shift should be taken into account in the computations. The values are:

'none'	: no shift is computed (returns an array of 0)
'stark'	: only the AC Stark shift is computed
'bs'	: only the Bloch-Siegert shift is computed
'all'	: both AC Stark and Bloch-Siegert shifts are computed

Returns:

- *residual* : Residual between the dressed resonance frequencies and the drive frequencies.

Bibliography

József S Bakos. Ac stark effect and multiphoton processes in atoms. *Physics Reports*, 31(3):209–235, 1977.

Christopher Bauer, Ray Freeman, Tom Frenkiel, James Keeler, and AJ Shaka. Gaussian pulses. *Journal of Magnetic Resonance (1969)*, 58(3):442–457, 1984.

MJ Biercuk, AC Doherty, and H Uys. Dynamical decoupling sequence construction as a filter-design problem. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(15):154002, 2011.

C Cohen-Tannoudji, J Dupont-Roc, and C Fabre. A quantum calculation of the higher order terms in the bloch-siegert shift. *Journal of Physics B: Atomic and Molecular Physics*, 6(8):L214, 1973.

Łukasz Cywiński, Roman M Lutchyn, Cody P Nave, and S Das Sarma. How to enhance dephasing time in superconducting qubits. *Physical Review B*, 77(17):174509, 2008.

Stefan Debald and Clive Emary. Spin-orbit-driven coherent oscillations in a few-electron quantum dot. *Physical review letters*, 94(22):226803, 2005.

S. Dubnov. Generalization of spectral flatness measure for non-gaussian linear processes. *IEEE Signal Processing Letters*, 11(8):698–701, 2004.

Christian Flindt, Anders S Sørensen, and Karsten Flensberg. Spin-orbit mediated control of spin qubits. *Physical review letters*, 97(24):240501, 2006.

Ray Freeman. Shaped radiofrequency pulses in high resolution nmr. *Progress in Nuclear Magnetic Resonance Spectroscopy*, 32(1):59–106, 1998.

Helen Geen and Ray Freeman. Band-selective radiofrequency pulses. *Journal of Magnetic Resonance (1969)*, 93(1):93–141, 1991.

Vitaly N Golovach, Massoud Borhani, and Daniel Loss. Electric-dipole-induced spin resonance in quantum dots. *Physical Review B*, 74(16):165319, 2006.

Gerhard Heinzel, Albrecht Rüdiger, and Roland Schilling. Spectrum and spectral density estimation by the discrete fourier transform (dft), including a comprehensive list of window functions and some new flat-top windows. 2002.

Christopher L Henley. Spin hamiltonians and exchange interactions, 2008. URL `http://www.lassp.cornell.edu/clh/p654/MM-Lec0.pdf`.

W Huang, CH Yang, KW Chan, T Tanttu, B Hensen, RCC Leon, MA Fogarty, JCC Hwang, FE Hudson, Kohei M Itoh, et al. Fidelity benchmarks for two-qubit gates in silicon. *Nature*, 569(7757):532–536, 2019.

J Robert Johansson, Paul D Nation, and Franco Nori. Qutip 2: A python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 184(4):1234–1240, 2013.

J. D. Johnston. Transform coding of audio signals using perceptual noise criteria. *IEEE Journal on Selected Areas in Communications*, 6(2):314–323, 1988.

Erika Kawakami, P Scarlino, Daniel R Ward, FR Braakman, DE Savage, MG Lagally, Mark Friesen, Susan N Coppersmith, Mark A Eriksson, and LMK Vandersypen. Electrical control of a long-lived spin qubit in a si/sige quantum dot. *Nature nanotechnology*, 9(9):666, 2014.

Horst Kessler, Siggi Mronga, and Gerd Gemmecker. Multi-dimensional nmr experiments using selective pulses. *Magnetic resonance in chemistry*, 29(6):527–557, 1991.

Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R Brad Blakestad, John D Jost, Chris Langer, Roee Ozeri, Signe Seidelin, and David J Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008.

Easwar Magesan, Jay M Gambetta, and Joseph Emerson. Characterizing quantum gates via randomized benchmarking. *Physical Review A*, 85(4):042311, 2012.

Tristan Meunier. Coherent control of individual electron spins in a two-dimensional array of tunnel coupled quantum dots. *APS*, 2019:P29–003, 2019.

Tristan Meunier, VE Calado, and LMK Vandersypen. Efficient controlled-phase gate for single-spin qubits in quantum dots. *Physical Review B*, 83(12):121403, 2011.

Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011. ISBN 1107002176, 9781107002173.

Line Hjortshøj Pedersen, Niels Martin Møller, and Klaus Mølmer. Fidelity of quantum operations. *Physics Letters A*, 367(1):47 – 51, 2007. ISSN 0375-9601. doi: https://doi.org/10.1016/j.physleta.2007.02.069.

S Philips. Dm solver. URL `https://github.com/stephanlphilips/DM_solver`.

M Pioro-Ladriere, T Obata, Y Tokura, Y-S Shin, Toshihiro Kubo, K Yoshida, T Taniyama, and S Tarucha. Electrically driven single-electron spin resonance in a slanting zeeman field. *Nature Physics*, 4(10):776, 2008.

KMM Prabhu. *Window functions and their applications in signal processing*. CRC press, 2013.

Luigi Salvatore and Amerigo Trotta. Flat-top windows for pwm waveform processing via dft. In *IEE Proceedings B (Electric Power Applications)*, volume 135, pages 346–361. IET, 1988.

MS Shahriar, Prabhakar Pradhan, and Jacob Morzinski. Driver-phase-correlated fluctuations in the rotation of a strongly driven quantum bit. *Physical Review A*, 69(3):032308, 2004.

Jon H Shirley. Solution of the schrödinger equation with a hamiltonian periodic in time. *Physical Review*, 138 (4B):B979, 1965.

Daniel A Steck. Quantum and atom optics. *Oregon Center for Optics and Department of Physics, University of Oregon*, 47, 2007.

Matthias Steffen, Lieven MK Vandersypen, and Isaac L Chuang. Simultaneous soft pulses applied at nearby frequencies, 2000.

Petre Stoica, Randolph L Moses, et al. Spectral analysis of signals. 2005.

Kenta Takeda, Akito Noiri, Takashi Nakajima, Jun Yoneda, Takashi Kobayashi, and Seigo Tarucha. Quantum tomography of an entangled three-spin state in silicon. *arXiv preprint arXiv:2010.10316*, 2020.

Yasuhiro Tokura, Wilfred G van der Wiel, Toshiaki Obata, and Seigo Tarucha. Coherent single electron spin control in a slanting zeeman field. *Physical review letters*, 96(4):047202, 2006.

Lieven MK Vandersypen and Isaac L Chuang. Nmr techniques for quantum control and computation. *Reviews of modern physics*, 76(4):1037, 2005.

M Veldhorst, JCC Hwang, CH Yang, AW Leenstra, Bob de Ronde, JP Dehollain, JT Muhonen, FE Hudson, Kohei M Itoh, A Morello, et al. An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature nanotechnology*, 9(12):981, 2014.

M Veldhorst, CH Yang, JCC Hwang, W Huang, JP Dehollain, JT Muhonen, S Simmons, A Laucht, FE Hudson, Kohei M Itoh, et al. A two-qubit logic gate in silicon. *Nature*, 526(7573):410, 2015.

Warren S Warren. Effects of arbitrary laser or nmr pulse shapes on population inversion and coherence. *The Journal of chemical physics*, 81(12):5437–5448, 1984.

TF Watson, SGJ Philips, Erika Kawakami, DR Ward, Pasquale Scarlino, Menno Veldhorst, DE Savage, MG Lagally, Mark Friesen, SN Coppersmith, et al. A programmable two-qubit quantum processor in silicon. *Nature*, 555(7698):633, 2018.

Yiying Yan, Zhiguo Lü, and Hang Zheng. Bloch-siegert shift of the rabi model. *Physical Review A*, 91(5):053834, 2015.

CH Yang, KW Chan, R Harper, W Huang, T Evans, JCC Hwang, B Hensen, A Laucht, T Tanttu, FE Hudson, et al. Silicon qubit fidelities approaching incoherent noise limits via pulse engineering. *Nature Electronics*, 2(4): 151–158, 2019.

Jun Yoneda, Kenta Takeda, Tomohiro Otsuka, Takashi Nakajima, Matthieu R Delbecq, Giles Allison, Takumu Honda, Tetsuo Kodera, Shunri Oda, Yusuke Hoshi, et al. A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%. *Nature nanotechnology*, 13(2):102, 2018.

David M Zajac, Anthony J Sigillito, Maximilian Russ, Felix Borjans, Jacob M Taylor, Guido Burkard, and Jason R Petta. Resonantly driven cnot gate for electron spins. *Science*, 359(6374):439–442, 2018.

Daniel Zeuch, Fabian Hassler, Jesse Slim, and David P DiVincenzo. Exact rotating wave approximation. *arXiv preprint arXiv:1807.02858*, 2018.