Forecasting timelines of quantum computing

Jaime Sevilla*

Aberdeen University, Aberdeen, UK and Center for the Study of Existential Risk, Cambridge, UK

C. Jess Riedel[†]

NTT Research, Inc., Physics & Informatics Laboratories, Sunnyvale, CA, USA

We consider how to forecast progress in the domain of quantum computing. For this purpose we collect a dataset of quantum computer systems to date, scored on their physical qubits and gate error rate, and we define an index combining both metrics, the generalized logical qubit. We study the relationship between physical qubits and gate error rate, and tentatively conclude that they are positively correlated (albeit with some room for doubt), indicating a frontier of development that trades-off between them. We also apply a log-linear regression on the metrics to provide a tentative upper bound on how much progress can be expected over time. Within the (generally optimistic) assumptions of our model, including the key assumption that exponential progress in qubit count and gate fidelity will continue, we estimate that that proof-of-concept fault-tolerant computation based on superconductor technology is unlikely (<5% confidence) to be exhibited before 2026, and that quantum devices capable of factoring RSA-2048 are unlikely (<5% confidence) to exist before 2039. It is of course possible that these milestones will in fact be reached earlier, but that this would require faster progress than has yet been seen.

In this article we consider the problem of anticipating progress on Quantum Computing. The following is an outline of the contents.

Section I contextualizes this article as part of our ongoing investigation of the transformative impact of quantum computing.

- In Section IA we discuss previous related work on anticipating quantum computing timelines and technological forecasting.
- In Section IB we discuss shortcomings of our methods.

Section II describes how this article quantifies progress in quantum computing.

- In Section II A we introduce the distinction between quantum annealers, Noisy Intermediate Scale Quantum (NISQ) computers, and fault-tolerant quantum computers (FTQC).
- In Section II B we discuss what concrete milestones we intend to forecast. Drawing from [1], we establish the milestone of the first large-scale faulttolerant quantum computer capable of breaking the modern cryptographic scheme RSA 2048 as the question of interest in this report.
- In Section IIC, we discuss metrics for tracking progress so far. We establish the number of physical qubits and the average two-qubit error rate as our metrics of interest. We combine these into a single

metric, the number of "generalized logical qubits," which estimates how many logical qubits a faulttolerant computer with a given number of physical qubits and two-qubit gate errors could support, but which is defined even for current-generation devices which are yet too noisy to support a single (true) logical qubit.

• In Section II D, we relate our forecasting question to our metrics of interest. As a precise operationalization of our question of interest, we choose to forecast when we will be able to showcase proofof-concept quantum error correction — specifically when a quantum computer will be powerful enough to break the commonly used RSA-2048 cryptographic scheme. This is operationalized as a function of the number of generalized logical qubits, concretely at 4100 logical qubits.

Section III reviews progress so far in the metrics of interest. Accompanying this section we present a publicly available database of major contributions in quantum computing hardware we have curated.

Section IV presents our main model and results.

- In Section IV A, we apply a multivariate log linear regression model to study the relationship between the number of physical qubits and the error rate in a system. When conditioning on time we find a positive correlation between the metrics, with more qubits being associated with higher error rates, suggesting a frontier of development that trades-off one against another.
- In Section IV B, we *assume* exponential progress of QC and apply a log linear multivariate model to provide an upper bound of likely progress of QC

^{*} j.sevilla.20@abdn.ac.uk

[†] jessriedel@gmail.com

based on superconductors. We find that it is 95% likely that proof-of-concept FTQL will be not be developed before 2026, and that RSA-2048 will be safe from quantum computation by 2039 with the same confidence. Note however that our model is very limited in data and makes strong assumptions about the statistical nature of the underlying metrics.

- In Section IV C we explore the robustness of our extrapolation model in the previous section, discussing alternative modelling choices and how they would affect our conclusions.
- In Section IV D we provide minimal validation of the extrapolation model we propose.

Section V collects our conclusions and suggests possible improvements and open questions for future work.

The <u>data</u> [2] we collected and the <u>code for our models</u> [3] is freely available online.

To aid the reader in quickly navigating this article, each section begin with the central question it intends to answer.

I. INTRODUCTION

Why bother thinking about progress in quantum computing, and what's already known?

In this article we engage with the task of predicting the arrival of major milestones in quantum computing. We will not describe the basics of quantum computing here, but instead will refer the reader to introductions at the popular [4, 5], semi-technical [1, 6–8], and technical levels [9–13].

Instead, we mention only these brief ideas: Like traditional (classical) computers, quantum computers are information-processing devices that physically implement mathematical computations. Unlike classical computers, quantum computers harness physical phenomena associated with quantum mechanics which generally can only be observed on microscopic scales. Although very difficult to build, such devices will be able to solve some specific math problems which are intractable to solve on classical computers.

However, even when quantum computers are very advanced they are not expected to fully replace classical computers because, for most mathematical problems, quantum effects do not offer a computational speed-up; it is only useful on certain kinds of tasks. Indeed, within realistic architecture proposals, the core quantum processing device relies on pre- and post-processing by classical computers for its basic operation. At the *theoretical* level, quantum computers represent a true paradigm shift in our basic understanding of physical computation, but at the *practical* level quantum computers are likely to be analogous to existing specialized computational modules like graphics-processing units and tensor-processing units, i.e., they will be dedicated physical devices optimized to solve a particular subset of computational problems.

The most anticipated applications of quantum computing include constrained optimization, efficient simulation of chemical reactions and other large interacting quantum mechanical systems, and cryptanalysis. (Improvements to machine-learning computations is also possible, although currently more speculative.) It is as of yet unclear how important or far reaching these applications will be. Some further reading about applications can be found in [1, 14–16].

This article is part of a broader investigation we are conducting, trying to anticipate the transformative impact that quantum computing will have on society. Understanding better the major milestones of this technology will give us better insight into what quantum computing is capable of, by when and how many warning signs and room-for-maneuver will society have to adapt to unexpected developments.

A. Previous work

To our knowledge, the most comprehensive assessment of future progress in quantum computing is the recent report by the National Academies of Sciences, Engineering, and Medicine [1]. They propose several milestones and metrics to track progress in Quantum Computing [17]. They find that RSA 2048 cryptography schemes will be safe during the next decade [18]. We borrow from this work to set our milestones and metrics of interest in Section II [19].

Ref. [20] survey experts in quantum computing about their projected timelines. They find that 22.7% of the experts they surveyed think it is likely or highly likely that quantum computers will be able to crack RSA-2048 keys by 2030, and 50% think that is likely or highly likely that we will be able to crack RSA-2048 keys by 2035. We use their work as a baseline for comparison in Section IV B.

Ref. [21] conducted a quantitative study of progress in several technologies, finding evidence of exponential progress in the cost of many technologies such as transistors, genomic sequencing, DRAMs, etc. Their findings lend support to our modelling assumption of exponential progress in quantum computing, which is detailed in Section IV.

On the topic of trend robustness, [22] examines 37 technological trends, and finds a robust discontinuity in 32% of these trends, with a base yearly rate of 0.1% chances of a discontinuity per year for each trend. This result is indicative of the base chance for modelling error through Section IV.

Ref. [23] also grapples with the question of how to combine different aspects of technology to find regularities in progress, focusing on military tech as a case study. His approach is conceptually similar to our overall analysis, with two marked differences: a) Kott's figure of merit is automatically derived, while ours is motivated by expert knowledge; b) Kott's approach is closer to machine learning, in that he fits his models according to a loss function while we use classical statistical analysis.

B. Limitations

Young technologies have meager track records, and predictions about the future rely heavily on educated guesses by experts (sometimes supplemented by more reliable arguments, e.g., constraints from basic physics). As a technology matures, data accumulates and predictions become more strongly driven by the extrapolation of past trends.

Quantum computing is still a very young field. Experimentation with quantum control of information has been occurring for decades, but devices advanced enough to be interpreted as non-trivial "computers" with quantifiable and comparable parameters are quite new. Because of this paucity of evidence, essentially all existing predictions about the future are based on expert wisdom (with all of its known flaws).

We consider this work to be merely a first attempt at systematically gathering and extrapolating data. At best, our results should be considered one piece of relevant evidence that can supplement expert opinion, and perhaps a reason to somewhat decrease one's credence in more extreme predictions.

Our data sources are opportunistic. We include industry blog announcements, not just academic journal articles. We are exposed to significant noise, e.g., from what numbers researchers choose to report, and what numbers we (as non-experts) are able to interpret. Some of this noise is random and will average out as the dataset accumulates, but much of it is biased and will not [24].

We consider our modeling assumptions to be optimistic in at least three ways. First, we are extrapolating from the *best* reported qubit numbers and error rates and ignore other values that have been reported. Second, we ignore both qubit connectivity and the trade-off between the number of physical qubits and the error rate. Third, we presume that constant-rate exponential growth will continue in the long term, which is disputed, especially for the average two-qubit gate error rate. These assumptions are discussed in more detail later.

For these reasons, the confidence intervals predicted by our model are quite different from the credences one might get from a wise and holistic Bayesian analysis. Still, we do expect the former to be useful input for the latter.

We are partially motivated by aggressive predictions often made informally and in the media about when quantum computers will have important real-world effects (e.g., by threatening the security of cryptography systems). We hope our work will advance the discussion about the future of this field by compelling commentators to acknowledge that aggressive predictions require making assumptions about the rate of progress increasing above the current trend, when in fact the expert wisdom seems to point, if anything, toward future progress falling below the current trend, especially for gate fidelity.

II. OPERATIONALIZING PROGRESS

How should we measure progress in quantum computing?

In this section we define the concrete milestone we are trying to predict, and relate it to metrics of current quantum computing systems.

For this purpose we define our own index, the *generalized logical qubits* (GLQs), which estimates the number of qubits that will be available for computation after accounting for error-correction overheads, and which extends to fractions less than 1 for machines (e.g. all extant ones) that fall short of the important technical threshold of fault tolerance, described further below.

We then explain why the thresholds of 1 and 4100 GLQs roughly correspond to the very important milestones of showcasing fault tolerance and compromising the widely used online security standard RSA 2048.

Readers uninterested in the motivations of our metric definitions and subsequent operationalization can just take these as given and skip to Section III.

A. Types of quantum computation

What approaches are there towards large-scale quantum computing?

Very roughly, approaches to computing (both quantum and classical) can be divided into analog and digital devices. Within classical computing these categories are now strongly associated with representing information using continuous and discrete variables, respectively, but they are also deeply connected with the degree to which the dynamics of the physical system implementing the computation are abstracted away from the mathematical form of the computation itself. In short: analog computers take advantage of natural physical dynamics that "look like" the computation one wants to perform, and are often *limited* to such computations, while digital computers are *universal* — in the sense that they are capable of performing any reasonable computation given enough resources — but require greater technological sophistication.

The analog-digital distinction is not a perfect dichotomy, but it has key implications for the use of error correction and, more broadly, fault tolerance [25]. In general, analog devices are less suitable for error correction. In the long run this leaves them vulnerable to noise as that problem becomes progressively worse for longer computations on larger devices.

Historically, for both quantum and classical computers, rudimentary analog devices without error correction became available before digital devices. In the case of classical computers, digital devices ultimately won out as technological capabilities improved. Experts disagree whether quantum analog computers will enjoy a time period where they are more practically useful than both quantum digital devices and classical devices. By far the most advanced form of analog quantum computing are so-called quantum annealers [26] (though alternative approaches are being pursued, e.g., [27]).

For these reasons, one can usefully identify three broad categories of quantum computing [1]:

- Analog quantum annealers;
- Noisy Intermediate Scale Quantum (NISQ) computers; and
- Scalable, gate-based, digital fault-tolerant quantum computers (FTQC).

Strictly speaking, the category of NISQ computers includes both analog and digital devices, but it can be useful, if crude, to simply think of them as digital devices that are not fault tolerant, and hence are limited in the scale of the computations they can perform by noise. Most experts agree that FTQC devices will dominate in the long-term [1] (though there are notable dissenters), and the timescale on which this comes about is very uncertain.

B. Milestones

What are the most important technical milestones on the road to scalable quantum computing?

For each of the three categories in the previous subsection we can identify three relevant milestones beyond the theoretical.

1. **Prototype**: A proof-of-concept implementation of the most basic machine operations.

- 2. Quantum supremacy (or "quantum advantage" [28]): The resolution of a mathematical problem by the quantum computer that would be unfeasible to solve classically [29].
- 3. **Practical demonstration**: Application to problems that are of interest *independent* from quantum computing.

Progress along these miles stones is summarized in Table I. Proof-of-concept machines showcasing analog quantum annealers [30] and NISQ computers [31] have existed for several years. Google recently claimed the quantum supremacy milestone [32] with a NISQ computer, showcasing a quantum processor capable of sampling in 0.02 seconds a probability distribution that would take the current fastest supercomputer (IBM's Summit as of this writing) 2.5 days to sample [33].

However, the highest-quality quantum computing experiments can apply only of order 1,000 gate operations before an error is likely to occur [34]. In order to perform usefully long *logical* computations, it will be necessary to use error-correction on the *physical* qubits. Fault-tolerant quantum computers will be able to drive logical errors down arbitrarily low, with overhead requirements that scale logarithmically with the desired logical error rate (which is inversely proportional to the average length of computation that can be performed without error). The theory behind error correction and FTQC is well-developed [35] but no functional experimental prototype exists as of the writing of this paper.

Our question of interest is when either of these technologies will be good enough to solve problems inaccessible to classical computers.

Question of interest: When will we have quantum computers capable of solving problems which cannot be feasibly be solved classically and which have noticeable real-world implications?

This prompts the related question, which of the three candidate technologies will first achieve this milestone?

Analog quantum annealers are not seen as a practical solution for large scale quantum computing due to the lack of a practical error correction scheme to enable computations of arbitrary length [1].

Similarly, while some small applications on certified randomness have been proposed using NISQ [36] [37], no major applications are envisioned yet [1].

Thus for the purposes of this report we will choose to study progress in FTQC.

Proxy for question of interest: When will we have large scale, fault-tolerant quantum computing?

Technology	Analog Quantum Annealers	NISQ Computer	\mathbf{FTQC}
Stage	Prototype	Quantum supremacy	Theoretical
Example	D-Wave's 2000Q [38]	Google's Scamore [32]	Fault-tolerant quantum computation [35]

TABLE I. A map of milestones in quantum computing. In this article we focus on forecasting progress in FTQC.

C. Metrics

How can we quantitatively track progress toward the relevant milestones?

The performance of current systems can be monitored via several useful quantities. Some of the most relevant are [1]:

- Number of physical qubits in a system. This is the simplest measure we have. It is the number of two-state physical subsystems in the computer in which coherent quantum information can be stored, such as the orientation (up or down) of a particle, or the direction (clockwise or counterclockwise) that a current in a superconducting wire is flowing.
- Average two-qubit-gate error rate. This is measured as the probability that when a quantum logic gate [39] (similar to a classical AND and OR gate) is applied in the system it fails to produce the correct output. An analogous metric exists for one-qubit gates, but the single-qubit error rate is typically much lower, so the overall error rate per operation is dominated by the two-qubit-gate error rate. In this work we will often simply refer to this as the error rate.
- Coherence time. The computational power of quantum computers relies on the physical qubits remaining highly isolated from their surrounding environment. When the qubits interact with their environment, information about the state of the qubits "leaks out", a process known as *decoherence*, inhibiting the necessary quantum coherent effects. Longer coherence times allow more gates to be performed before this happens, allowing for more complicated quantum computation.
- Qubit connectivity. Two physical qubits are *connected* when a gate operation can be achieved by inducing the qubits to physically interact with each other, which often is only possible if they are a short distance apart. The connectivity of a computer is the (graph) structure of these connections. To apply a gate to unconnected qubits, multiple costly interactions between intermediate qubits in a chain must be used. One quantitative measure

of a computer's connectivity is the average number of connections it takes to link two qubits in the system.

• Number of logical qubits in a system. A quantum computer with error correction and a given number of noisy *physical* qubits can emulate a noiseless computation on a smaller number of *logical* qubits. As the noise in a device gets greater, it can generally support fewer logical qubits for the same number of physical qubits. If the noise is greater than some threshold, *zero* logical qubits can be supported. Very few extant devices have achieved noise levels below this threshold, and none of them have created a single logical qubit (in part because of insufficient, and insufficiently connected, physical qubits).

In this report we choose to focus on (1) the physical number of qubits in a system and (2) the average twoqubit error rate, because they are of high importance and are widely reported in the literature [40].

Beyond these immediate metrics, we can try to combine them to produce a metric that captures progress more accurately. IBM opts for this approach and tracks the **quantum volume** [41].

In a similar spirit, we construct our own figure of merit, the number of **generalized logical qubits**. This figure approximates the number of logical qubits that a device theoretically would be able to simulate using quantum error correction at a given average error rate with a given number of physical qubits using the surface code [42–44], the leading choice error correction scheme for systems based on superconducting qubits. Unlike the actual number of logical qubits, this metric extends smoothly and sensibly to fractional qubits. It is well-defined for devices that have achieved two-qubit gate error rates below the fault-tolerance threshold mentioned above.

Formally, we define the number of generalized logical qubits to be $N_L = N_P f_{\text{QEC}}$, where N_P is the number of physical qubits, p_P is the two-qubit gate error rate,

$$f_{\rm QEC} = \left[4 \frac{\log \left(\sqrt{10} \, p_{\rm P} / p_{\rm L} \right)}{\log \left(p_{\rm th} / p_{\rm P} \right)} + 1 \right]^{-2} \tag{1}$$

is an overhead factor $(0 < f_{\rm QEC} < 1)$ accounting for error correction, $p_{\rm th} \approx 10^{-2}$ is the approximate threshold error under which fault-tolerance becomes possible for the surface code [43–45], and $p_{\rm L} = 10^{-18}$ is the acceptable logical error rate. A contour plot of this function is



FIG. 1. Contour plot of the generalized logical qubit map. Our primary thresholds of interest, 1 and 4100, are delineated with dotted curves. The contours asymptote at the fault tolerance threshold $p_{\rm th} \approx 10^{-2}$. Note that our definition depends on the underlying error correction scheme, in this case we are assuming the surface code.

Milestone	Operationalization
Quantum fault tolerance	1 generalized logical qubits
RSA 2048 quantum attack	4100 generalized logical qubits

TABLE II. Milestones of interest operationalized in terms of generalized logical qubits.

D. Operationalization

How can we convert our primary forecasting questions into statements about generalized logical qubits?

In this subsubsection, we develop our preferred operationalization of progress milestones, as summarized in Table II.

Question of interest: When will we have fault-tolerant quantum computing?

As discussed in Section II A, the first step to that achievement is proof-of-concept quantum error correction. One reasonable operationalization of this intermediate milestone based on quantum computing is straightforward. **Operationalization:** When will the number of generalized logical qubits exceed 1?

The next step would be to scale up fault-tolerant quantum computers to a size where they can perform useful computation. As a landmark of disruptive application we chose the development of a quantum computer capable of threatening a commonly used cryptographic algorithm, RSA 2048.

Question of interest: When will a large faulttolerant quantum computer that can run Shor's algorithm to break RSA 2048 be developed?

This operationalization is based on one of the most anticipated uses of a quantum computer. However, this operationalization does not allow a ready estimation based on short term metrics. We address that now:



FIG. 2. Reported number of physical qubits of systems between 2003 and 2020. n = 52 data points.



FIG. 3. Reported average two-qubit gate error rate of systems between 2003 and 2020. n = 40 data points. (This is smaller than Fig. 2 because fewer papers reported qubit counts than error rates.)

First, note that running Shor's algorithm to break RSA

2048 will almost certainly require error correction [46];

the computation is of sufficient length that many physical errors will reliably occur in any hardware available in the foreseeable future. Next we observe that the number of logical qubits needed to run Shor's algorithm to break RSA 2048 has been estimated at ~4098 logical qubits [47–49] [50]. We can thus use our generalized logical qubits metric to operationalize this question.

Operationalization: When will a quantum computer with more than ~ 4098 generalized logical qubits be built?

III. CURRENT STATE OF THE ART

How do current systems perform on the metrics of interest?

Using Refs. [1, 51] as a helpful starting point, we compiled a large historical dataset available <u>here</u> [52]. Altogether, it is based on data points reported by Refs. [31, 32, 42, 53–81].

The metrics we are plotting in Figs. 2, 3, 4 are the physical qubits, average two-qubit gate error rate and our own index, the generalized logical qubits. Explanation of these metrics can be found in Section II C. They are plotted against the date when the system specifications were made public. [82] Our data runs from the year 2003, the date of the earliest source we could find satisfying our requirements, to the first half of 2020, when we froze our dataset for the purpose of our analysis.

In total we have n = 52 data points for which we know the number of physical qubits and the date, n = 40 where we know the physical qubits, the error rate and the date, n = 12 where we know the physical qubits, the error rate, the date and the GLQs are well defined.

The datapoints collect systems using different quantum computing technologies, including superconductorts, trapped ion qubits, spin qubits and silicon qubits.

IV. MODELLING FUTURE PROGRESS

When will various quantum computing abilities be achieved, and how certain are we?

In previous sections we have identified the key metrics that indicate progress (physical qubits and gate error rate), constructed an index that combines the two (the generalized logical qubit, abbreviated as GLQ) and casted the milestones we are interested in terms of our index (1 GLQ for the proof-of-concept error correction and 4100 GLQs for a quantum Shor attack on RSA 2048).

Here we statistically analyze the historic trends on these metrics.

First in Section IV A we show that quantum computer designs so far exhibit a trade off between a large number of physical qubits and low gate error rate.

Then, specializing to superconducting qubit devices, in Section IV B we extrapolate the current trends in the best performance on these metrics to predict the trend in GLQs and make predictions about our milestones of interest. Section IV C and Section IV D explore our extrapolation model in depth, respectively looking at alternate modelling choices and a rolling validation of the model.

Through this section we are assuming a continued trend of exponential progress for physical qubits and average gate error.

Ref. [22] examines 37 technological trends, and finds a robust discontinuity in 32% of these trends, with a base yearly rate of 0.1% chances of a discontinuity per year for each trend. This work suggests a substantial chance of modelling error that we are not explicitly accounting for.

A. A technological development frontier

How do the physical qubits and average two-qubit gate error rate relate to each other?

To answer this question, we will model how papers score in both metrics as a multivariate log linear model [83] that takes as input a date and outputs a distribution for the combination of metrics that papers around that date are likely to produce. In particular, we are not just looking at the total correlation between these two metrics over our full dataset, since steady but independent progress in each would induce a positive correlation; rather we are looking at the correlation between them *conditional* on the year of publication to see whether there is a trade-off between these two metrics at any given level of technological maturity.

For a quick and informal introduction to multivariate models, see [84]. For a more formal treatment, see Chapter 6 of Ref. [85].

Our model has the form:

$$Y = XB + \Xi; \qquad \Xi \sim N(0, \Sigma) \tag{2}$$

where each column in $Y \in \mathbb{R}^{n \times 2}$ corresponds to the logarithm of our metrics of interest (the physical qubits and the average two-qubit gate error rate) for each of the papers in our dataset, and $X \in \mathbb{R}^{n \times 2}$ is the date as a fractional year stacked with a constant intercept 1.



FIG. 4. Calculated GLQs for systems between 2003 and 2020. No data points are shown prior to 2014 because the metric is not defined for devices with gate error rates that do not satisfy the fault-tolerance threshold. n = 12 data points.



FIG. 5. Physical qubits plotted against gate error rate for systems between 2003 and 2020. n = 40 data points.

The model is characterized by a matrix $B \in \mathbb{R}^{2 \times 2}$ of drift parameters and a covariant noise matrix $\Sigma \in \mathbb{R}^{2 \times 2}$.

We can estimate their values with the max likelihood estimators \hat{B} and $\hat{\Sigma}$, calculated as:

$$\hat{B} = \left(X^{\top}X\right)^{-1}X^{\top}Y \tag{3}$$

$$\Sigma = \frac{1}{n-2} \left(Y - XB \right)^{\top} \left(Y - XB \right)$$
(4)

Using all data from 2003 onwards the maximum likelihood estimation of the parameters are displayed in Table III.

The off-diagonal parameter of the symmetric matrix $\hat{\Sigma}$

estimates the covariance between the two metrics.

We can provide a confidence interval for the covariance using the naive bootstrap procedure [86] (see Chapter 11 of Ref. [87]).

Using all data from 2003 onward (n = 40 data points) we estimate that the covariance is positive with 98.8% confidence, with a 90% CI of (0.13, 0.76). We note that the physical qubits metric is positively oriented while the error rate metric is negatively oriented, so the positivity of the covariance suggests the existence of a robust trade-off between both metrics.

	\hat{B}		$\hat{\Sigma}$	
	Phys. qubits	Gate error	Phys. qubits	Gate error
Intercept	-181	256	0.76	0.44
Slope	0.090	-0.13	0.44	2.02

TABLE III. Point estimates of the parameters of the multivariate log linear model for data between 2003 and 2020, n = 40 data points. The first row of \hat{B} are the intercepts, while the second row corresponds to the yearly log slope. Similarly, the first column of \hat{B} corresponds to the parameters associated with the physical qubits, while the second corresponds to the parameters associated with the average two-qubit gate error rate. To cast these values in a more amenable way, these parameters indicate that the median value by the year 2020 are 9.22 physical qubits and 0.02 error rate. The first diagonal entry of $\hat{\Sigma}$ is the estimated variance for the log physical qubits, while the second diagonal entry is the variance of the log error rate, and the off-diagonal symmetrical entries are the estimated covariance between the two metrics.



FIG. 6. Visualization of the Gaussian pdf of the model with the parameterization in Table III conditioned on the year 2023. That is, the model predicts that a certain percentile of papers published in that year will report metrics inside the respective ellipse. (Note that this is a prediction about the the distribution of *all* sources satisfying our selection criteria, *not* a prediction about the best values obtained or the technological frontier.) The size and orientation of the ellipse represents the covariance of the metrics. Each curve delineates an x% probability region, centered around the median.



FIG. 7. Extrapolated progress of generalized logical qubits at 5%, 50%, and 95% bootstrapping quantiles (red, orange, and blue, respectively). The trend lines are based on n = 39 data points corresponding to superconducting quantum computers developed between 2007 to 2020 and constitute the main result of our paper. Also shown are the 12 data points from our dataset for which the number of GLQs are defined. Directly extrapolating these data points would be very unreliable because of their low number and position near the divergence at the fault-tolerance threshold.

In plain English, this suggests that quantum computer designers face a trade-off between trying to optimize for quantum computers with many physical qubits and quantum computers with very low gate error rate.

If we focus on data on the most promising substrate so far, superconducting qubits, from 2003 onward (n = 31 data points) we find that the correlation weakens, with a 90% CI of (-0.11, 0.48) and a confidence of 84.5% that the covariance is positive. The significantly lower confidence indicates that our finding is not robust.

B. Extrapolating current trends

When will we reach the relevant quantum computing milestones?

Now, to extrapolate current trends of growth, we assume continued exponential progress on the best physical qubit and error rate metrics achieved so far. For reasons described below, we consider this to provide a soft *upper bound* on the likely rate of quantum computing progress.

We have chosen to focus on predicting the development of superconductor-based quantum computers for two reasons: First it is generally regarded as the single most promising substrate for large-scale FTQC. Second, and more practically, there are substantially more papers in the literature quoting physical qubit numbers and gate error rates in the superconducting category.

We now describe our main modeling method. We begin by considering our two primary metrics — the number of physical qubits and the error rate — separately. For each metric, we select only the subset of data points corresponding to papers reporting a new best achieved value of that metric (as of the year of publication). We then take the two distinct dataset and, to each, we fit a loglinear minimum square error model with log-Gaussian noise.

In other words, for each of our two metrics we consider just the set of new "world records" and perform traditional linear regression on that data in log space.[88] Although this maximization procedure loses some information before feeding it to our model, it is necessary in order to correctly track our topic of interest: the bleeding-edge capabilities of the field. Notably, this means our model resists being influenced by a glut of papers describing devices that are not intended to be competitive on our chosen metrics.

The median trajectory of each model can be extrapolated forward in time, and we combine these to produce a prediction for the progress in generalized logical qubits. Here is a summary of our model:

$$Y_1 = XB_1 + \Xi_1; \qquad \Xi_1 \sim N(0, \sigma_1^2)$$
 (5)

$$Y_2 = XB_2 + \Xi_2; \qquad \Xi_2 \sim N(0, \sigma_2^2)$$
 (6)

$$Z = f\left(\exp\left(Y_1\right), \exp\left(Y_2\right)\right) \tag{7}$$

X = (1 t) stacks an intercept 1 and the date t. Y_1 is the maximum log physical qubits at date t. Y_2 is the minimum log average gate error rate at date t. Z is the number of generalized logical qubits given the log physical qubit metric Y_1 and log gate error rate Y_2 .

To produce our confidence intervals we are going to use a bootstrapping procedure (see Chapter 11 of Ref. [87]).

Our dataset of n = 39 papers is resampled with repetition B = 1000 times and log scaled.

Then the data for each of the metrics is aggregated via a sparse maximum, where only the years where an actual improvement is made are registered.

This produces two derivative datasets (X_1, Y_1) and (X_2, Y_2) for the trend of max log physical qubits and the trend of min log error rate respectively.

We use these datasets to compute the maximum likelihood estimates for the parameters (intercept, slope, and noise variance in log space) of the log-linear model.

$$B_i = \left(X_i^{\top} X_i\right)^{-1} X_i^{\top} Y_i \tag{8}$$

$$\sigma_i^2 = \frac{1}{n-1} \left(Y_i - X_i B_i \right)^\top \left(Y_i - X_i B_i \right)$$
(9)

Note that unlike in Section IV A here we are fitting two separate log linear models instead of a multivariate one. Furthemore, we emphasize that we are fitting only the subset of data corresponding to new best metric values instead of all reported values like in Section IV A.

We use these estimates to estimate T, the date when the critical threshold of 4100 qubits is crossed by the median trajectory.

$$T = \min\{t : Z > 4100\}$$
(10)

$$\approx \min\left\{t: f\left(\exp\left(\hat{B}_1\begin{pmatrix}1\\t\end{pmatrix}\right)\right),$$
 (11)

$$\exp\left(\hat{B}_2\begin{pmatrix}1\\t\end{pmatrix}\right) > 4100$$
 (12)

The result is three representative trajectories corresponding to the quantiles 0.05, 0.5 and 0.95 of T. For more details on our calculations, please refer to <u>our code</u> [89].

We consider this model optimistic in three ways. First, it is being extrapolated over the maximum data instead of the whole dataset. Second, it ignores the trade-off between both metrics we uncovered in Section IV A. Third, it presumes that the regime of exponential growth will continue in the long term, which is a dubious assumption — especially for the average two-qubit gate error rate [90].

This model predicts that proof-of-concept FTQL will be developed between early 2026 and early 2033 with 90% confidence with the median in early 2030, and that RSA-2048 Shor attacks will become feasible between mid 2039 and mid 2058 with 90% confidence with the median in early 2050.

Because we are not completely confident in the model, these do not exactly coincide with our Bayesian credences for when these events will take place, but we do find them a strong starting point for further deliberation.

These results are more pessimistic but broadly comparable to those produced through the survey of experts in [20]. We emphasize that 22.7% of the experts they surveyed think it is likely or highly likely that quantum computers will be able to crack RSA-2048 keys by 2030, and 50% think that is likely or highly likely that we will be able to crack RSA-2048 keys by 2035.

Note that our models have two levels of modelled uncertainty: the uncertainty introduced by bootstrapping and the gaussian noise of each estimated model.

However, the gaussian noise is negligible compared to the bootstrapping uncertainty. This is evident when we plot the quantiles of gaussian noise for each bootstrapped mode, see figure 9. Thus we ignore the gaussian uncertainty in our analysis.

C. Model robustness

How robust are the conclusions of our extrapolation model to various alternative choices we might have made?

As we analyzed the data, we also considered the following variations on our model. Although our main model is substantially better justified than these variations, we will examine how they would affect things.

Choices made:

- Using a max procedure vs. using a best fit line for all data points
- All data available vs using data of quantum computers with more than two qubits



FIG. 8. Extrapolated progress of physical qubits (upper plot) and average two-qubit gate error rate (lower plot), from n = 39 data points corresponding to superconducting quantum computers developed between 2007 to 2020. There are compelling expert assessments that qubit error rates will not and need not fall as dramatically as depicted here (though note that these error rates are still far larger than the rates experienced by modern classical computers).

If we do not aggregate the data through a maximum at all we unsurprisingly get far more pessimistic predictions, see Fig 10. Worth noting that in this case we are not estimating the *best* results each year, but rather the *typical* results each year.

Since many papers will report numbers of physical qubits and gate error rates without attempting to push

the technological frontier on those particular metrics, we believe our (sparse) max procedure allows our model to better track the hypothetical technological level of the field as a whole. (In particular, the max is not affected by whether lots of non-frontier papers are published.) A constraint imposed by using a max procedure is that it greatly reduces the effective size of our dataset.



FIG. 9. An optimistic extrapolation of current progress on GLQs, based on n = 39 data points corresponding to superconducting quantum computers developed between 2007 to 2020, where we have plotted the 5%, 50% and 95% quantiles of the gaussian noise credence intervals for the 5%, 50% and 95% quantile bootstrapped trajectories.



FIG. 10. Extrapolated progress of generalized logical qubits based on n = 39 data points corresponding to superconducting quantum computers developed between 2007 to 2020 and fitted to the raw data, no maximum.

Our main model utilizes data from 2007 to 2020. If we include data from 2014 onwards the resulting predictions are substantially more optimistic, see fig 11. We have

selected this time threshold as 2014 marks the appearance of superconducting quantum computers with more than two qubits, which are arguably more representative of smooth progress.

Our model is quite simple. A more sophisticated approach would likely rely on modelling progress with a geometric drift model as in [21]. Unfortunately, this approach was not workable because (1) we have insufficient GLQ data points to apply the geometric drift model directly to univariate GLQ data and (2) the geometric drift model has not yet be generalized to the multivariate setting, which would require significant modifications [91]. Therefore, we have decided to use our simpler approach instead.

On the other hand, we could have applied an even simpler model by just fitting the raw GLQs. However, the GLQ metric is not well defined for most of the dataset, as the gate error rate of most quantum computers is below the 1e-2 threshold. Thus our dataset is quite limited and the model very noisy. See figure 13 below.

Overall, we observe that the modelling choice makes a significant difference to the predicted time when quantum computing will reach the milestone of being able to crack RSA 2048, which reduces our confidence on the main model.

In particular, the much faster timelines for the 5% quantile of the model fitted to data from 2014 to 2020 suggests that the trend of development may have sped up in the last decade. However, given the paucity of data and the dramatic differences with the 50% and 95% quantile trajectories in this model this might be just the result of noise.

D. Model validation

Does our extrapolation model perform well on historic

data?

In this subsection, we attempt to validate our model by applying it to a subset of our data and using it to predict the rest.

In figure 14 we plot the predictions when using data from 2007 to 2018, and check the predictions for 2019 against the actual values. In Table IV we expand this procedure to a rolling validation that encompasses different years used as data. The results are consistent with a well-calibrated model, with 1 value out of 10 falling outside of the predicted 90% confidence intervals.

V. CONCLUSIONS AND OPEN QUESTIONS

A. Summary

We began this work by briefly summarizing some of the past literature on technological progress, and argued that exponential growth in hardware technology is a reasonable (though by no means assured) expectation.

We then compiled a quantitative database of achievements in quantum computing research in the last two decades, focusing especially on two metrics: the highest physical qubit count and lowest average two-qubit gate error rates.

When conditioning on year, we found an (anti-)correlation between the highest qubit counts and the lowest error rates across all technologies, in line with expectations that these metrics are in tension at a given level of technological development. However, when looking at superconducting data specifically our confidence weakens significantly, indicating that this finding might be spurious.



FIG. 11. Extrapolated progress of generalized logical qubits based on n = 33 data points corresponding to superconducting quantum computers developed between 2014 to 2020.



FIG. 12. The fit and extrapolation of the GLQ data (m = 5 years of data, n = 10 data points corresponding to superconducting quantum computers developed between 2015 to 2020) to a geometric drift model.



FIG. 13. Simple log-linear extrapolation of the GLQs of superconducting quantum computers between 2007 and 2020, n = 12.



FIG. 14. Extrapolated progress of generalized logical qubits based on n = 20 data points corresponding to superconducting quantum computers developed between 2007 to 2018.

	2019	2018	2017	2016
Actual maximum	6.31e-03	4.29e-3	4.29e-3	1.44e-4
2007-2018 data $n = 20$	1.29e-3; 1.94e-3; 7.99e-2	N/A	N/A	N/A
2007-2017 data $n = 14$	0.; 6.21e-3; 2.69e-2	0.; 1.10e-2; 5.57e-2	N/A	N/A
2007-2016 data $n = 7$	1.15e-4; 7.53e-3; 6.43e-2	5.82e-5; 4.44e-3; 3.02e-2	2.25e-5; 2.54e-3; 1-38e-2	N/A
2007-2015 data $n = 5$	3.97e-7; 5.84e-3; 1.57e-1	0.; 3.33e-3; 6.58e-2	0.; 1.18e-3; 2.55e-2	0.; 8.23e-4; 8.75e-3

TABLE IV. Actual and predicted records for GLQs in recent years, as predicted by models trained on different time spans, following the modelling choices explained in Section IVB and Section IVC. Each entry contains the estimated 5%, 50%, and 95% quantiles separated by semicolons. The confidence intervals where the actual value falls within the predicted 90% confidence interval are shaded green, while the ones where it does not are shaded red. Note that when the predicted gate error rate is above the threshold where error correction is possible we consider the generalized logical qubits to be 0., hence the null entries.

We proposed a single scalar figure of merit to measure progress towards large-scale fault-tolerant quantum computing, the generalized logical qubit (GLQ), which combines our two metrics. Using GLQs, we operationalized two milestone of fault-tolerant quantum computing:

- 1. Realization of a single GLQ, roughly corresponding to the beginning of *scalable* quantum computation.
- 2. Realization of 4100 GLQs, roughly corresponding to the arrival of computing power that has practical implications for real-world cryptographic systems.

Ultimately, we extrapolated the dataset we had available to predict a less than 5% chance that proof-ofconcept fault-tolerant computation will be achieved before 2024, and less than 5% chance that RSA-2048 Shor attacks will be feasible before 2039. These predictions are dependent on disputable modeling choices we have made, although we generally feel our assumptions have been conservative in the sense that these are reasonable upper bounds on the rate of progress.

B. Future work

There is much room for future work to improve on our methods, especially as more data points appear in the coming years:

- Other datasets, such as patents as in [92].
- Modeling of other technical metrics like coherence time, qubit connectivity, and quantum volume, and

incorporating other considerations (like the Steane code) into the primary figures of merit.

- Substrates beyond superconducting qubits, especially ion traps.
- More sophisticated (but data-hungry) models like that of [21] or [93].
- More extensive validation.
- Fitting the models to variables sensitive to global developments that may alter the effort put into quantum computing, such as cumulative inflation-adjustment capital investment in quantum computing research.
- Estimate when reversible computing, a necessary but not sufficient property of quantum computers, will be developed in order to maintain progress in *classical* computing. (We briefly expand on this in Appendix 7.2.)

Our dataset is freely available <u>here</u> [94], and the code for our models is here [95].

C. Acknowledgements

We thank J. Doyne Farmer, Chris Ferrie, Steve Howard, Adrian Hutter, Francois Lafond, John-Clark Levin, Pablo Moreno, Florian Pein, Ehud Reiter, Daniel Sank, and Graeme Smith for thoughtful feedback and discussion.

This work is partly supported by a grant from the Center on Long Term Risk. The Center on Long Term Risk and the employers of the authors did not have any role in the design of the study; in the collection, analysis and interpretation of data; in the writing of the article; or in the decision to submit the article for publication..

- E. Grumbling and M. Horowitz, eds., Quantum Computing: Progress and Prospects (National Academies Press, Washington, D.C., 2019).
- Full URL: https://docs.google.com/spreadsheets/ d/1pwb4gf0Fx1xgfVhtXTaqEGS9b7FwsstsJ 0v7Zb1naQ0/.
- [3] Full URL: https://colab.research.google.com/ drive/1XkWcsUy-CiNDDJffPC3dNJbgvR6iDAqh.
- [4] S. Aaronson, Scientific American **298**, 62 (2008).
- [5] S. Aaronson, The New York Times (2011).
- [6] S. Aaronson, *Quantum Computing Since Democritus* (Cambridge University Press, 2013).
- [7] A. Perry, R. Sun, C. Hughes, J. Isaacson, and J. Turner, arXiv:1905.00282 (2019).
- [8] T. Rudolph, Q is for quantum (Terence Rudolph, 2017).
- [9] A. Y. Kitaev, A. Shen, M. N. Vyalyi, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, 2002).
- [10] A. Matuschak and M. Nielsen, "Quantum country," (2019).
- [11] N. D. Mermin, Quantum Computer Science: An Introduction (Cambridge University Press, 2007).
- [12] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information (Cambridge University Press, Cambridge; New York, 2000).
- [13] M. Ozols and M. Walter, "The quantum quest," (2018).
- [14] R. de Wolf, Ethics and Information Technology , 1 (2017).
- [15] A. Majot and R. Yampolskiy, Futures Confronting Future Catastrophic Threats To Humanity, 72, 17 (2015).
- [16] M. Möller and C. Vuik, arXiv:1705.07413 (2017).
- [17] These milestones include (1) experimental quantum annealers, (2) small (tens of qubits) computers, (3) gate-based quantum computers that demonstrates quantum supremacy, (4) quantum annealers that demonstrate quantum supremacy, (5) implementation of QEC for improved qubit quality, (6) commercially useful quantum computers and (7) large (>1,000 qubits), fault-tolerant, modular quantum computers.
- [18] "Given the current state of quantum computing and recent rates of progress, it is highly unexpected that a quantum computer that can compromise RSA 2048 or comparable discrete logarithm-based public key cryptosystems will be built within the next decade." [1].
- [19] As one interesting point, the Academies report highlights Rock's law and contextualizes it as a warning that quantum computing may fail to achieve exponential progress unless it secures increasing funding: "Moore's law is the result of a virtuous cycle, where improvements in integrated circuit manufacturing allow the manufacturer to reduce the price of their product, which in turn causes them to sell more products and increase their sales and profits. This increased revenue then enables them to improve the manufacturing process again, which is harder this time, since the easier changes have already been made.".
- [20] M. Piani and M. Mosca, Quantum Threat Timeline Report, Tech. Rep. (Global Risk Institute, 2019).
- [21] J. D. Farmer and F. Lafond, Research Policy 45, 647–665 (2016).
- [22] K. Grace, "Discontinuous progress in history: an up-

date," (2020).

- [23] A. Kott, The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology, 154851291987552 (2019).
- [24] For example, one may expect that papers may only report the metrics they were explicitly trying to improve.
- [25] There are important technical distinctions between error correction and fault tolerance, but we will not address them in this work.
- [26] A. Das and B. K. Chakrabarti, Reviews of Modern Physics 80, 1061 (2008).
- [27] Y. Yamamoto, K. Aihara, T. Leleu, K.-i. Kawarabayashi, S. Kako, M. Fejer, K. Inoue, and H. Takesue, npj Quantum Information 3, 1–15 (2017).
- [28] Unfortunately, terminology is not completely standardized. Some authors take "quantum supremacy" and "quantum advantage" to be synonyms, while others distinguish them based on whether the problem could be solved by *existing* classical (super)computers vs. *any* feasible classical computer, or based on the economic value of the problem.
- [29] It is worth emphasizing that the mathematical problem selected for demonstrating quantum supremacy is invariably contrived to be as easy as possible for the quantum devices and as hard as possible for the classical devices, i.e., the "quantum home-field advantage" has been maximized. In particular, quantum supremacy does not mean the quantum computer is better at all computational problems, or even a large class of them.
- [30] K. Karimi, N. G. Dickson, F. Hamze, M. H. S. Amin, M. Drew-Brook, F. A. Chudak, P. I. Bunyk, W. G. Macready, and G. Rose, arXiv:1006.4147 (2011).
- [31] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, Nature **519**, 66 (2015).
- [32] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. Mc-Clean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Nature 574, 505 (2019).
- [33] E. Pednault, J. A. Gunnels, and J. M. Gambetta, "On "Quantum supremacy"," (2019).
- [34] This is quantified by the average gate error rate, that we will explain in Section IIC. Google's average error

rate is of the order of 1e3 [32].

- [35] J. Preskill, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 454, 469 (1998).
- [36] For a popular-level description, see [96].
- [37] A. Acín and L. Masanes, Nature 540, 213 (2016).
- [38] A. D. King, J. Carrasquilla, I. Ozfidan, J. Raymond, E. Andriyash, A. Berkley, M. Reis, T. M. Lanting, R. Harris, G. Poulin-Lamarre, A. Y. Smirnov, C. Rich, F. Altomare, P. Bunyk, J. Whittaker, L. Swenson, E. Hoskinson, Y. Sato, M. Volkmann, E. Ladizinsky, M. Johnson, J. Hilton, and M. H. Amin, Nature 560, 456 (2018).
- [39] Here we use "logic gate" simply to distinguish this machine operation from other sorts of gates, like the ones that keep in horses. We are still referring to a gate operation on physical qubits. (In this paper, we will not need to discuss logical gates, which are a counterpart to logical qubits.).
- [40] We emphasize that qubit connectivity is essentially for fault tolerance and useful computations. (Trivially, Nseparate and non-interacting quantum devices, each consisting of 2 physical qubits interacting through a gate with error rate r, could be construed as a quantum "computer" with arbitrarily high physical qubits (2N) and arbitrarily low error rate (r).) The reasonableness of our model is based on the premise that, in practice, connectivity will to some extent track qubit counts and error rates within the field overall.
- [41] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, Physical Review A 100, 032328 (2019).
- [42] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, Nature **508**, 500 (2014).
- [43] A. G. Fowler, D. S. Wang, and L. C. L. Hollenberg, Quantum Information & Computation 11, 8–18 (2011).
- [44] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Physical Review A 86 (2012).
- [45] A. Javadi-Abhari, Towards A Scalable Software Stack For Resource Estimation And Optimization In General-Purpose Quantum Computers, Ph.D. thesis, Princeton University (2017).
- [46] "Without QEC, it is unlikely that a complex quantum program, such as one that implements Shor's algorithm, would ever run correctly on a quantum computer." [1].
- [47] C. Gidney and M. Ekerå, arXiv:1905.09749 (2019).
- [48] T. Häner, M. Roetteler, and K. M. Svore, arXiv:1611.07995 (2016).
- [49] M. Roetteler, M. Naehrig, K. M. Svore, and K. Lauter, arXiv:1706.06752 (2017).
- [50] Häner et al. describe a procedure for using 2(N + 2) qubits to factor an N-bit number [48]. Recently, Gidney & Ekerå have described a procedure that uses $\sim 50\%$ more qubits, $3N + 0.002N \log_2 N$, but runs two orders of magnitude faster for these size numbers [47].
- [51] D. Finke, "Scorecards," (2020).
- [52] Full URL: https://docs.google.com/spreadsheets/ d/1pwb4gf0Fx1xgfVhtXTaqEGS9b7FwsstsJ0v7Zb1naQ0/ edit#gid=0.
- [53] Intel, "2018 CES: intel advances quantum and neuro-

morphic computing research," (2018).

- [54] J. Kelly, "A preview of bristlecone, google's new quantum processor," (2018).
- [55] J. M. Gambetta and S. Sheldon, "Cramming more power into a quantum device," (2019).
- [56] C. Nay, "IBM opens quantum computation center in new york; brings world's largest fleet of quantum computing systems online, unveils new 53-Qubit quantum system for broad use," (2019).
- [57] IBM, "IBM quantum experience," (2020).
- [58] Intel, "Intel delivers 17-Qubit superconducting chip with advanced packaging to QuTech," (2017).
- [59] Amazon, "Amazon braket hardware providers," (2019).
- [60] Rigetti Computing, "The quantum processing unit (QPU) — pyQuil 1.9.1.dev0 documentation," (2018).
- [61] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, Physical Review Letters **117**, 060504 (2016).
- [62] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. L. Heras, R. Babbush, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, E. Solano, H. Neven, and J. M. Martinis, Nature 534, 222 (2016).
- [63] J. M. Chow, J. M. Gambetta, A. D. Corcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen, Physical Review Letters **109**, 060501 (2012).
- [64] Rigetti Computing, "The rigetti 128-qubit chip and what it means for quantum," (2018).
- [65] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, Nature 536, 63 (2016).
- [66] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Nature 460, 240 (2009).
- [67] D. Finke, "Qubit quality," (2020).
- [68] N. Friis, O. Marty, C. Maier, C. Hempel, M. Holzäpfel, P. Jurcevic, M. B. Plenio, M. Huber, C. Roos, R. Blatt, and B. Lanyon, Physical Review X 8, 021012 (2018).
- [69] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, Physical Review Letters 117, 060505 (2016).
- [70] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, Nature 438, 643 (2005).
- [71] S. S. Hong, A. T. Papageorge, P. Sivarajah, G. Crossman, N. Didier, A. M. Polloreno, E. A. Sete, S. W. Turkowski, M. P. da Silva, and B. R. Johnson, Physical Review A **101**, 012302 (2020).
- [72] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland, Nature **422**, 412 (2003).
- [73] A. Nersisyan, S. Poletto, N. Alidoust, R. Manenti, R. Renzas, C. Bui, K. Vu, T. Whyland, Y. Mohan, E. A. Sete, S. Stanwyck, A. Bestwick, and M. Reagor, arXiv:1901.08042 (2019).
- [74] M. Reagor, C. B. Osborn, N. Tezak, A. Staley,

G. Prawiroatmodjo, M. Scheer, N. Alidoust, E. A. Sete,
N. Didier, M. P. da Silva, E. Acala, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, S. Caldwell, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. El Bouayadi, D. Girshovich, S. Hong, A. Hudson, P. Karalekas, K. Kuang,
M. Lenihan, R. Manenti, T. Manning, J. Marshall,
Y. Mohan, W. O'Brien, J. Otterbach, A. Papageorge,
J. Paquette, M. Pelstring, A. Polloreno, V. Rawat,
C. A. Ryan, R. Renzas, N. Rubin, D. Russel, M. Rust,
D. Scarabelli, M. Selvanayagam, R. Sinclair, R. Smith,
M. Suska, T. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. T. Rigetti, Science Advances 4, eaao3603 (2018).

- [75] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, Nature 422, 408 (2003).
- [76] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Physical Review A 93, 060302 (2016).
- [77] C. Song, K. Xu, W. Liu, C. Yang, S. Zheng, H. Deng, Q. Xie, K. Huang, Q. Guo, L. Zhang, P. Zhang, D. Xu, D. Zheng, X. Zhu, H. Wang, Y. Chen, C. Lu, S. Han, and J. Pan, Physical Review Letters **119**, 180511 (2017).
- [78] M. Steffen, M. Ansmann, R. C. Bialczak, N. Katz, E. Lucero, R. McDermott, M. Neeley, E. M. Weig, A. N. Cleland, and J. M. Martinis, Science **313**, 1423 (2006).
- [79] T. Watson, "IBM doubles compute power for quantum systems, developers execute 300K+ experiments on IBM quantum cloud," (2017).
- [80] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak, Nature **526**, 410 (2015).
- [81] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J. Chen, N. C. Pisenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe, and J. Kim, Nature Communications **10**, 5464 (2019).
- [82] When we only know the year of publication we imputed the date as 1 June that year.
- [83] We could instead have used multiple linear regression where, for example, we predict the log error rate based on the year and log physical qubits. However this would require us to assume a linear dependency between the log error rate and log physical qubits, whereas the multivariate model is agnostic about the relation of the two metrics.
- [84] R. D. V. Vega and A. G. Rai, "Multivariate regression," (2020).
- [85] A. J. Izenman, Modern Multivariate Statistical Techniques: Regression, Classification, and Manifold Learning, Springer Texts in Statistics (Springer-Verlag, New York, 2008).
- [86] For those unfamiliar with bootstrapping, we repeatedly sample n papers with replacements from our dataset and compute the estimator of the covariance $\sigma_{12}^2 = \hat{\Sigma}_{1,0}$ over the resample. The quantiles 0.05 and 0.95 of this sampling procedure converge to a 90% confidence interval for the covariance.
- [87] B. Efron and T. Hastie, Computer Age Statistical Inference: Algorithms, Evidence, and Data Science (Cambridge University Press, 2016).

- [88] During bootstrapping, the median of "record setting" physical qubits data points was 11, and the median of "record setting" average two-qubit gate error rate data points was 6.
- [89] Full URL: https://colab.research.google.com/ drive/1XkWcsUy-CiNDDJffPC3dNJbgvR6iDAqh# scrollTo=jV29mJw0iz9W&uniqifier=3.
- [90] There are informal arguments for why error rates may fail to sustain exponential progress in the future even if the number of physical qubits grows exponentially. Once the fault-tolerant threshold is achieved, there do not seem to be fundamental physical barriers to adding more physical qubits given sufficient effort and resources, but error rates refer to a single gate, and these will plausibly hit strongly diminishing returns with respect to how perfectly they can be engineered. Even if lower error rates are experimentally possible, there are reasons to think devoting engineering resources to more physical qubits may have a higher payoff in terms of the number of logical qubits. Unlike for traditional CPUs, superconducting qubits currently use parallel control channels. Adding arbitrary numbers of physical qubits may require partially serializing this control, through shared control lines, which hasn't yet been attempted. We thank Daniel Sank for discussion on these points.
- [91] Private communication with Farmer and Lafond.
- [92] C. L. Benson and C. L. Magee, PLOS ONE 10, e0121635 (2015).
- [93] A. C. Harvey, The Journal of the Operational Research Society 35, 641 (1984).
- [94] Full URL: https://docs.google.com/spreadsheets/ d/1pwb4gf0Fx1xgfVhtXTaqEGS9b7FwsstsJ0v7Zb1naQ0/ edit#gid=0.
- [95] Full URL: https://colab.research.google. com/drive/1XkWcsUy-CiNDDJffPC3dNJbgvR6iDAqh# scrollTo=jV29mJw0iz9W&uniqifier=4.
- [96] S. Aaronson, American Scientist 102, 266 (2017).
- [97] In particular, note that once the two-gate error rate gets extremely low, below 10^{-6} , it can become advantageous to use alternative protocols like the Steane concatenated code. This will somewhat change how fast we expect the available number of generalized logical qubits to increase with time (compared to the naive assumption that the surface code will always be used. However, such low error rates are very challenging to achieve, perhaps for a very long time, and for the foreseeable future something like the surface code is likely to dominate [45].
- [98] Earlier authors take a similar form but sometimes differ by a numerical constant due to different estimates for technical complications, e.g., a factor of $\sqrt{10}$ in [43, 44].
- [99] Indeed, we believe it would be misleading for any composite metric like GLQ to include the discrete structure associated with integer jumps in code distance since we are highly uncertain about the detailed features of future error correction schemes. (Any discrete structure would be "fake detail".) Rather, GLQ is intended only to capture the roughly scaling of the surface code.
- [100] R. Landauer, Ferroelectrics 2, 47 (1971).
- [101] C. H. Bennett, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics Quantum Information and Computation, 34, 501 (2003).
- [102] C. H. Bennett, IBM Journal of Research and Develop-

ment 17, 525 (1973).

- [103] E. P. DeBenedictis, M. P. Frank, N. Ganesh, and N. G. Anderson, in 2016 IEEE International Conference on Rebooting Computing (ICRC) (IEEE, San Diego, CA, 2016) pp. 1–8.
- [104] E. P. DeBenedictis, J. K. Mee, and M. P. Frank, Computer 50, 76 (2017).
- [105] E. P. DeBenedictis and M. P. Frank, Computer 51, 69 (2018).
- [106] H. Rahaman, S. Chattopadhyay, and S. Chattopadhyay, Progress in VLSI Design and Test: Proceedings of the 16th International Symposium on VSLI Design and Test (Springer, Shipur, India, 2012).

VI. APPENDIX

This appendix is intended for readers with some technical familiarity with quantum computing.

A. Generalized logical qubits

Our definition of the number of generalized logical qubits makes use of the scaling associated with the (planar) surface error correction code [42–45]. Thus, when we project the arrival of a fault-tolerant QC for a particular year with this metric, we are implicitly assuming the device will use the surface code or something that makes similar trade-offs [97].

For the surface code, the number of physical qubits required to encode a single logical qubit is $f_{\text{QEC}} = (2d - 1)^2$ where $d = 1, 3, 5, 7, \ldots$ is the code distance, a measure of the size of the errors to which the code is robust. To a good approximation, this code satisfies [98]

$$p_{\rm L} \approx \sqrt{10} \, p_{\rm P} \left(\frac{p_{\rm P}}{p_{\rm th}}\right)^{(d-1)/2}$$
(13)

where $p_{\rm P}$ is the average two-qubit error rate on the raw physical qubits in the computer, $p_{\rm th} \approx 10^{-2}$ is the approximate threshold error at which fault-tolerance becomes possible, and $p_{\rm L}$ is the two-qubit error for the logical qubits, i.e., the *effective* error rate of the mathematical computation [45]. As long as $p_{\rm P}$ is below $p_{\rm th}$, the logical error rate can be driven arbitrarily low by choosing a sufficiently large code distance d (at the expense of requiring more physical qubits to host the computation).

The inverse $1/p_{\rm L}$ is the number of two-qubit gates that can typically be applied before an error occurs at the logical level, and that sets the maximum size of the computation. The particular value we choose is not too important because our results depend only weakly (logarithmically) on this choice, which in principle is determined by the length of the particular computation one wants to do. We follow [1] by taking $p_{\rm L} = 10^{-18}$, corresponding to computations with of order a trillion steps (not necessarily in serial). We then ignore the discreteness [99] of the code distance d, eliminate it from the equations, and derive ratio of physical qubits per logical qubits:

$$f_{\rm QEC} = \left[4\frac{\log\left(\sqrt{10}\,p_{\rm P}/p_{\rm L}\right)}{\log\left(p_{\rm th}/p_{\rm P}\right)} + 1\right]^{-2} \tag{14}$$

We emphasize that the above formula is only an approximation, and in particular ignores the discreteness of the code distance d.

B. Reversible computing

Landaurer's limit [100–102] is a fundamental thermodynamic barrier, imposing an energy cost on all computations that are irreversible (at the logical level). It is known that it can be circumvented by adopting reversible computing, where all computations are modified to eliminate (almost all) irreversible steps [101, 102].

Landaurer's limit is not yet a serious consideration for modern classical computers because current technology already wastes several orders of magnitude more energy, due to conventional engineering imperfections, than is required by Landaurer's limit. But, as part of the progress associated with Moore's law, the amount of energy wasted has been decreasing exponentially each year, and eventually Landaurer's limit *will* become important for continuing technological progress in classical computing [103, 104]. Once that happens, classical computers will likely need to become reversible (or partially reversible) to continue improving.

Reversible computing is a prerequisite for quantum computers in the sense that quantum computations are necessarily reversible and a subset of quantum operations (the classical operations with respect to some fixed computational basis) are sufficient to implement any classical computation reversibly [105]. Indeed, our impression from discussion with experts is that achieving reversibility is a large component of the engineering difficulty of quantum computing. It is unclear how much more difficult it would be to build a useful quantum computer once a useful reversible one is created.

Therefore, if we extrapolate current classical computation trends, the date at which classical computers become reversible is a good hint at when quantum computers may be available. We have not performed a thorough search of the literature, but we did find a plot, Fig. 2 on page 386 of [106], suggesting that achieving (partial) reversibility will become an important part of classical computing technological progress in the mid 2030's. A quick skim of seminars and articles by DeBenedictis & Frank seem broadly consistent with this [103–105].

Constructing a more accurate projection and drawing implications for quantum computing is left for future work. Besides being valuable by virtue of being independent (and complementary to) approaches directly measuring research progress in quantum computing, the estimation approach based on reversible computing is no-

table because it naturally incorporates the economics of classical computing investments.