

The Need for Speed: Efficient Exact Simulation of Silicon Dangling Bond Logic

Jan Drewniok*, Marcel Walter*, and Robert Wille*[‡]

*Chair for Design Automation, Technical University of Munich, Germany

[‡]Software Competence Center Hagenberg GmbH, Austria

Email: {jan.drewniok, marcel.walter, robert.wille}@tum.de

<https://www.cda.cit.tum.de/research/fcn/>

Abstract—*The Silicon Dangling Bond (SiDB) logic platform, an emerging computational beyond-CMOS nanotechnology, is a promising competitor due to its ability to achieve integration density and clock speed values that are several orders of magnitude higher compared to current CMOS fabrication nodes. However, the exact physical simulation of SiDB layouts, which is an essential component of any design validation workflow, is computationally expensive. In this paper, we propose a novel algorithm called *QuickExact*, which aims to be both, efficient and exact. To this end, we are introducing three techniques, namely 1) Physically-informed Search Space Pruning, 2) Partial Solution Caching, and 3) Effective State Enumeration. Extensive experimental evaluations confirm that, compared to the state-of-the-art algorithm, the resulting approach leads to a paramount runtime advantage of more than a factor of 5000 on randomly generated layouts and more than a factor of 2000 on an established gate library.*

I. INTRODUCTION & MOTIVATION

As the limits of Moore’s Law are approached and processing technologies reach the top of the S-curve, there is a growing interest in the *Silicon Dangling Bond* (SiDB) logic platform, an emerging computational beyond-CMOS nanotechnology [1]–[7]. One of the key features of SiDB logic are its sub-nanometer elementary devices, which can significantly improve integration density by several orders of magnitude compared to current CMOS fabrication nodes [3], [4], [6]–[11]. With contemporary *transistor-based* technology requiring vast amounts of energy and being limited to low GHz speeds, SiDB technology reduces energy consumption for similar processes by over 99% while enabling an increase in processing speeds of several orders of magnitude [12]. This makes SiDB logic promising for achieving ultra-low energy dissipation and establishes it as a much-anticipated green competitor in the beyond-CMOS domain [3], [8], [11], [13]–[15]. In addition, SiDB logic has been proposed as a candidate for the integration of quantum computers with conventional CMOS circuitry [7], [16].

The growing interest in the SiDB logic platform has led to proposals from the scientific community for gate and circuit libraries [17]–[20] and design automation solutions [17], [18], [21]–[25]. Furthermore, the technology is gaining commercial momentum, as evidenced by the recently formed research enterprise *Quantum Silicon Inc.*, which aims to be one of the first industry adopters of SiDBs and has secured multi-million dollar investments [26], [27]. These developments are also putting pressure on design and simulation tools to keep pace, especially with the rapid advances in

SiDB fabrication capabilities [3]–[6]. In particular, physical simulation is an essential component of any design validation workflow, providing the basis for verifying the behavior of circuit layouts before fabrication.

However, physical simulation of SiDBs is challenging: To accurately predict the behavior of a system of n SiDBs, it is necessary to solve a high-dimensional optimization problem that takes into account all relevant physical effects. This involves enumerating up to 3^n charge configurations and checking their physical validity. Due to this exponential runtime complexity, such an algorithm is limited to small instances.

In the literature, two approximate algorithms for solving this problem in polynomial time were proposed, namely *QuickSim* [28] and *SimAnneal* [18]. In contrast, only one exact algorithm that determines *all* physical valid charge configurations has been presented: *ExhaustiveGS* (ExGS) [29]. It performs an exhaustive search of the entire (exponential) search space. Therefore, *ExGS* is severely limited in its runtime efficiency and, thus, is not even applicable to medium-sized SiDB layouts like standard gates. This urgently calls for a more sophisticated and efficient exact simulation approach for SiDBs, which is indispensable for many specialized tasks, i. e., defect and temperature robustness analyses [30], [31].

To this end, we propose a novel exact simulation algorithm denoted *QuickExact*, which is designed with a special focus on runtime efficiency. Our approach utilizes three main techniques, namely 1) Physically-informed Search Space Pruning, 2) Partial Solution Caching, and 3) Effective State Enumeration. Extensive experimental evaluations, that are covering the simulation of established gate libraries as well as random instances, confirm that the resulting approach leads to a paramount runtime advantage of more than a factor of 5000 on randomly generated layouts and more than a factor of 2000 on an established gate library compared to the state-of-the-art algorithm *ExGS*. An implementation of *QuickExact* is publicly available as part of the *Munich Nanotech Toolkit* (MNT).

To establish this paper as a self-contained work, Section II provides an overview of the basic concepts of SiDB systems and their physical simulation. Afterwards, related work from the literature is discussed in Section III. Section IV constitutes the main contribution of this work as it provides a detailed introduction to the proposed algorithm *QuickExact*. Section V presents exhaustive experimental evaluations demonstrating the advantages of *QuickExact* over state-of-the-art techniques. Finally, Section VI concludes the paper.

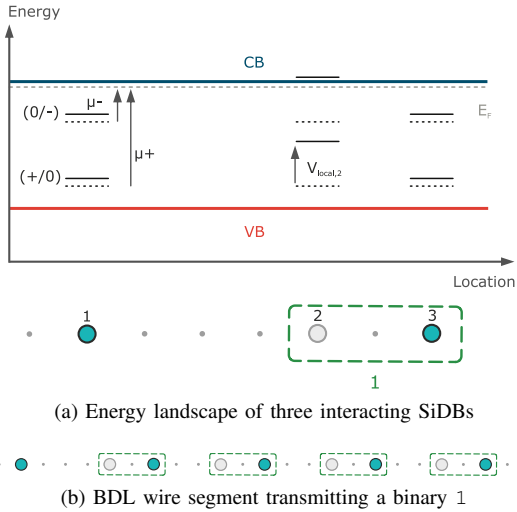


Fig. 1: Elemental SiDB logic components

II. PRELIMINARIES

In this section, preliminaries are established that are necessary for the comprehension of the remainder of this work. First, an overview of the SiDB logic platform is given in Section II-A, followed by a comprehensive presentation of the physical simulation of SiDBs in Section II-B.

A. The SiDB Logic Platform

Silicon Dangling Bonds (SiDBs) are created on a *Hydrogen-passivated Silicon* surface (H-Si(100)-2×1) by utilizing a *Scanning Tunneling Microscope* (STM) with an atomically sharp tip [4]–[6], [32]–[34]. Applying a voltage through the STM disrupts the covalent bond between a silicon and a hydrogen atom, causing the hydrogen to be removed from the sample and desorbed to the tip, leaving behind a *dangling bond*. The tip is then moved to a new silicon dimer and the process is repeated to create another SiDB. Each SiDB acts as a chemically identical quantum dot that can trap a maximum of two electrons in its sp^3 -orbital [35]. Therefore, each SiDB is confined to one of three different states: negatively, neutrally, or positively charged.

The charge state of an SiDB depends on the charge-transition energy levels, which are shown in Fig. 1a for three SiDBs. When an SiDB is undisturbed (i.e., no other SiDBs or external electrostatic potentials are present), the charge transition levels (0/−) and (0/+) (dashed lines) lie roughly 0.3 eV and 0.8 eV below the conduction band (CB) [10], [36]. The location of the energy transition levels with respect to the *Fermi energy* E_F defines the charge state of the SiDB. If (0/−) is below E_F , the SiDB is negatively charged. Vice versa, if only (+/0) is below E_F , the SiDB is neutrally charged, and in the last case, the SiDB contains no more electrons and is therefore positively charged. In n-doped systems, E_F is close to the CB and thus, SiDBs are negatively charged when not disturbed. However, if several SiDBs are placed close together as illustrated in Fig. 1a, the single charge transition levels can be shifted with respect to E_F by electrostatic potential interaction among SiDBs. As illustrated,

the two outermost SiDBs generate a local electrostatic potential at the second SiDB of $V_{local,2} > 0.4$ eV, which shifts (0/−) above the Fermi energy, which then leads to a neutrally charged second SiDB. This yields a *Binary-Dot Logic* (BDL) pair with a discrete charge state realizing the binary state 1 in this case, as illustrated by the green rectangle. This specific property is exploited to utilize SiDBs for the fabrication of logic devices [8].

A sequence of BDL pairs, depicted in Fig. 1b, can operate as a binary wire that transmits information solely through the coupling of electric fields. This principle can also be employed to create an OR-gate. In fact, both a BDL wire consisting of eight SiDBs and an OR-gate (with a total size smaller than 30 nm²) have already been realized experimentally by Huff et al. [8].

B. Physical Simulation

In order to validate the designs of SiDB assemblies and facilitate the advancement of this technology, the utilization of physical simulation is deemed imperative. The objective of such simulation is to accurately simulate the charge distribution of a given arrangement of SiDBs. To this end, electrostatic potential simulation is conducted. The electrostatic potential $V_{i,j}$ at position i generated by an SiDB in the state $n_j \in \{-1, 0, 1\}$ at position j is given by [8], [33]

$$V_{i,j} = -\frac{q_e}{4\pi\epsilon_0\epsilon_r} \cdot \frac{e^{-\frac{d_{i,j}}{\lambda_{tf}}}}{d_{i,j}} \cdot n_j, \quad (1)$$

where λ_{tf} defines the *Thomas-Fermi screening length* and ϵ_r the *dielectric constant*, which were experimentally extracted to be 5 nm and 5.6, respectively [8]. Moreover, ϵ_0 , q_e , and $d_{i,j}$ are the *vacuum permittivity*, the *electron charge* ($q_e = -e$; e : *elementary charge*), and the *Euclidean distance* between position i and j , respectively. The layout's total electrostatic potential energy E is thus given as

$$E = -\sum_{i<j} V_{i,j} \cdot n_i \cdot q_e. \quad (2)$$

The expected state of an SiDB layout is determined by the charge configuration with the lowest energy. Therefore, in order to determine said state, Eq. (2) has to be minimized. However, a solution (i.e., a valid charge distribution) has to meet the physical constraint of *metastability*, which can be described by the combination of two criteria, namely *population stability* and *configuration stability*. Both must be obeyed to guarantee simulation result validity.

a) *Population Stability*: As highlighted before in Fig. 1a, the charge state of each SiDB must be consistent with the energy charge-transition levels (0/−) and (+/0) relative to E_F . The energy of the SiDB's charge transition level at position i depends on the local electrostatic potential $V_{local,i}$ at that position defined by

$$V_{local,i} = \sum_{j,j \neq i} V_{i,j}. \quad (3)$$

Intuitively, this means that an SiDB is preferably neutrally charged when adjacent SiDBs are negatively charged and vice

versa. This relation is formally expressed by the following conditions: SiDB- if $\mu_- + V_{local,i} \cdot q_e < 0$, SiDB+ if $\mu_+ + V_{local,i} \cdot q_e > 0$, and SiDB0 otherwise.

b) *Configuration Stability*: If there is no feasible single-electron hop event between arbitrary SiDBs leading to a lowering of the system's total electrostatic potential energy, the charge state fulfills the configuration stability. In other words, if this constraint is not satisfied, there would exist a state of lower electrostatic potential energy that the system would spontaneously converge to.

Overall, the charge configuration with the lowest electrostatic potential system energy that is also physically valid, i. e., satisfies metastability, is called the system's *ground state*. It represents the charge configuration of a given SiDB system at low temperature and, thereby, its physical behavior.

III. RELATED WORK

In this section, the exact SiDB simulator *ExhaustiveGS* (*ExGS*) [29] from the literature is discussed. As described above, SiDBs can be present in one of three different charge states. Since their charge state depends on the location of the energy transition level with respect to E_F , which is in turn defined by the charge states of the surrounding SiDBs (cf. Fig. 1a), a high-dimensional interwoven problem has to be solved in order to determine the ground state and all other physically valid charge configurations of the system.

Mastering the complexity of simulation is a critical aspect of any design automation process for SiDB-based logic. However, *exact* physical simulation of SiDBs is a computationally complex task for which there is only *ExGS* [29].

ExGS enumerates all possible charge configurations of a given SiDB layout, computes the local electrostatic potential and the system's electrostatic potential energy with Eq. (2) and Eq. (3) from scratch for each single charge configuration. At the same time, it determines whether each charge configuration is physically valid. This eventually allows obtaining not only the ground state, but all physically valid charge configurations of a given SiDB layout. Since *ExGS* considers *all* possible charge configurations (and does not apply search space pruning or similar methods to reduce the resulting computational complexity), already small gate layouts cannot be simulated in a reasonable amount of time. This shortcoming is caused by the inherent exponential complexity of the SiDB simulation task.

Example 1. In order to simulate a layout consisting of 32 SiDBs in 3-state simulation (SiDBs are also allowed to be positively charged), a total of $3^{32} \approx 10^{15}$ charge configurations have to be considered and checked for physical validity.

Ultimately, an algorithm that is both efficient and exact (i. e., able to find all physically valid charge distributions) is desired. This algorithm is presented in this paper, with the general idea as well as the details elaborated in the next section.

IV. PROPOSED ALGORITHM: *QuickExact*

In this section, the proposed SiDB simulation algorithm *QuickExact* is introduced. First, the general idea is outlined in Section IV-A and the techniques that lead to the performance

advantage of more than three orders of magnitude compared to state of the art are elaborated, with their precise impact analyzed in Section IV-B. Finally, the implementation details of *QuickExact* are presented in Section IV-C.

A. General Idea

To reiterate, *ExGS* is the state-of-the-art *exact* physical SiDB simulator. However, it is highly inefficient and is, thus, not applicable to even simulate medium-sized SiDB assemblies like gates. Therefore, there exists a *Need for Speed*. To simulate efficiently while being exact at the same time, we propose the three techniques: 1) Physically-informed Search Space Pruning, 2) Partial Solution Caching, and 3) Effective State Enumeration. The following segments elaborate on their respective ideas.

1) *Physically-informed Search Space Pruning*: If the charge state of d SiDBs (in a physically valid charge distribution) is known prior to simulation, the search space is reduced by a factor of 2^d or 3^d for 2-state and 3-state simulation, respectively. In fact, SiDBs that only weakly interact with other ones, have a tendency to be always negatively charged under realistic μ_- values. With respect to the population stability discussed in Section II-B, this can be expressed as the following relation: the i -th SiDB has to be negatively charged for all physically valid charge distributions if the maximal possible local electrostatic potential $V_{local,max,i}$ at position i is smaller than μ_- . This is summarized by the following constraint:

$$SiDB_{i,-} \iff \mu_- < -V_{local,i}^{max} \cdot q_e \quad (4)$$

According to Eq. (3), $V_{local,i}$ is maximized if each $V_{i,j}$ is positive. Due to Eq. (1), this is the case if $\forall i : n_i = -1$. Ultimately, in order to detect all SiDBs that have to be negatively charged to fulfill the metastability, all SiDBs are set to negative. Afterwards, the SiDBs that fulfill Eq. (4) are determined. This whole process is sketched by the pseudocode in Alg. 1, where all SiDBs are set to be negatively charged in Line 1. Afterwards, all SiDBs that fulfill Eq. (4) (Line 5) are collected and returned (Line 9).

To reduce the search space further by an additional factor of 2 (or 3 for 3-state simulation), a so-called *dependent* SiDB is selected. It is an arbitrarily chosen SiDB k (but not among the d pre-determined negatively-charged SiDBs) whose charge state is always deduced by the charge states of all other SiDBs in the layout due to the population stability constraint. In other words, when enumerating all electron distributions for the remaining $n-d$ SiDBs in the layout, we can always stop at iteration 2^{n-1-d} (or 3^{n-1-d}), because the final SiDB's charge state is inherently determined by the charge states of all others. Hence, the charge state of SiDB k can be described depending on $n_i, i \neq k$ as follows:

$$SiDB_{k,-} \iff \mu_- + \underbrace{\sum_{i,i \neq k} V_{k,i}(n_i)}_{V_{local,k}} \cdot q_e < 0 \quad (5)$$

$$SiDB_{k,+} \iff \mu_+ + \sum_{i,i \neq k} V_{k,i} \cdot q_e > 0 \quad (6)$$

$$SiDB_{k,0} \iff \text{otherwise} \quad (7)$$

Algorithm 1: Negative SiDB Detection

Input: SiDB layout L comprised of n dangling bonds
Input: Physical parameters $P = \{\mu_-, \mu_+, \lambda_{ef}, \epsilon_r\}$
Output: A set S of SiDBs that must be negatively charged

```
1 Electron distribution  $D \leftarrow [D_1 = -1, \dots, D_n = -1]$ 
2  $V_{local,i} \leftarrow \sum_{j,j \neq i} V_{i,j}$  given  $D$  // Eq. (3)
3  $S \leftarrow \emptyset$ 
4 foreach  $db \in L$  do
5   if  $\mu_- + V_{local,db} \cdot q_e < 0$  then
6      $S \leftarrow S \cup \{db\}$ 
7   end if
8 end foreach
9 return  $S$ 
```

Algorithm 2: Update Potential Cache

Input: SiDB layout L comprised of n dangling bonds
Input: Physical parameters $P = \{\mu_-, \mu_+, \lambda_{ef}, \epsilon_r\}$
Input: Original electron distribution D
Input: Updated electron distribution D'

```
1  $C \leftarrow D' \setminus D$ 
2 foreach  $db' \in L \cap C$  do
3   foreach  $db \in L \setminus (L \cap C)$  do
4      $V_{local,db} \leftarrow V_{local,db} + V_{db,db'} \cdot C[db']$  given  $P$ 
5   end foreach
6 end foreach
```

Thereby, it becomes non-imperative to enumerate the charge state of the k -th SiDB, which reduces the search space by another factor of 2 in 2-state (or 3 in 3-state) simulation.

Hence, via the introducing of the two *Physically-informed Search Space Pruning* concepts (detecting negatively charged SiDBs to satisfy population stability and introducing the *dependent* SiDB) mentioned above, the search space is reduced by a factor of 2^{d+1} (or 3^{d+1}) without losing exactness.

2) *Partial Solution Caching*: As mentioned before, it is runtime-expensive to recompute the local electrostatic potential at each SiDB position every time the charge distribution changes. To reduce the runtime of this calculation step, *Partial Solution Caching* is introduced. This means that the entire calculation according to Eq. (3) is not carried out every time the charge distribution changes. Instead, the previous local electrostatic potentials are reused and updated while only considering the SiDBs at which the charge state has changed.

This can be conducted as summarized in the pseudocode shown in Alg. 2. The SiDB layout L , the physical parameters P as well as the original and updated charge distributions, D and D' , respectively, are used as the algorithm's input. As a first step, all SiDBs that have undergone a change in their charge state are identified and collected as C (Line 1). Subsequently, the local electrostatic potentials of the SiDBs that have remained unchanged in their charge state are updated (Line 4). Thus, the number of SiDBs that have changed in their charge state (i. e., the number of elements in $L \cap C$) determine the number of iteration steps required to update the local electrostatic potentials.

3) *Effective State Enumeration*: The previously presented technique, *Partial Solution Caching*, is the most effective when the number of SiDBs that are changing their charge state between each iteration is small (i. e., $|L \cap C|$ is small). This third technique focuses on the effective iteration of electron

distributions such that the least amount of SiDBs change their charge state in each step.

Electron distributions D can be interpreted as b -ary digit strings of length n , where $b \in \{2, 3\}$ is the *base*, i. e., the number of SiDB states.¹ To this end, value v at position $D[i]$ indicates that SiDB i is assigned charge state v . It follows that if the *Hamming distance* of D to D' (from one iteration step to the next) is minimal, the number of SiDBs that have changed their state is also minimized. It is further known that the average Hamming distance between increments of binary strings converges to 2, i. e., when simply iterating through electron distributions in binary order, 2 SiDBs change their polarity in each step on average. By using *Gray code* as the enumeration scheme, we can reduce this number to 1, since in Gray code order, only a single digit changes its value from one step to the next. This approach allows for maximally efficient updating of the local electrostatic potential (as described in Alg. 2), since only one charge state in the charge distribution changes at a time ($|L \cap C| = 1$), reducing computational cost and improving efficiency by roughly another factor of 2.

B. Resulting Performance Improvements

The introduced optimization techniques elaborated on in the previous section, improve the performance of *exact* physical simulation by several orders of magnitude. For 2-state simulation, the total theoretical performance gain given as the achieved search space reduction is composed as follows:

$$\underbrace{2^{d+1}}_{\text{Section IV-A1}} \cdot \underbrace{\frac{n}{|L \cap C|}}_{\text{Section IV-A2}} \cdot \underbrace{2}_{\text{Section IV-A3}} \approx 2^{d+1} \cdot n,$$

since $|L \cap C| \approx 2$.

Example 2. Consider an exact 2-state simulation of a layout consisting of 30 SiDBs ($n = 30$), where five SiDBs ($d = 5$) are detected to be negatively charged in a physically valid charge configuration. While the state-of-the-art simulator *ExGS* would require 2^{30} simulation steps, using the techniques introduced above allows to conduct the simulations approximately $2^{5+1} \cdot 30 \approx 2000$ times faster.

C. Implementation Details

The techniques proposed above can be implemented on top of any combinatorial exact physical simulator for SiDBs. In this work, we took *ExGS* as a basis and realized the proposed concepts on top of it. This led to the efficient physical simulator *QuickExact*. How *QuickExact* works as a whole is sketched in Alg. 3. The algorithm starts with an SiDB layout comprised of n SiDBs and the physical parameters P as input. First, all SiDBs that must be negatively charged in a physically valid charge distribution are detected with Alg. 1 as explained before (Line 2). All remaining SiDBs (i. e., SiDBs that can also be neutrally charged) are collected as L_S in Line 3. Then, the first SiDB of L_S is arbitrarily chosen as the *dependent* SiDB \bar{db} . Subsequently, all possible charge distributions of the SiDBs in $L_S \setminus \{\bar{db}\}$ are enumerated in *Gray code order*.

¹For the sake of simplicity, we assume $b = 2$ in the following. However, the proposed technique works for arbitrary values of b .

Algorithm 3: QuickExact

Input: SiDB layout L comprised of n dangling bonds
Input: Physical parameters $P = \{\mu_-, \mu_+, \lambda_{ij}, \epsilon_r\}$
Output: All valid electron distributions with their respective energies

```
1  $DE \leftarrow \emptyset$ 
2  $S \leftarrow \text{NEGATIVE SiDB DETECTION}(L, P)$  // Alg. 1
3  $L_S \leftarrow L \setminus S$ 
4 dependent SiDB  $\tilde{db} \leftarrow L_S[1]$ 
5 Electron distribution  $D \leftarrow [D_1 = -1, \dots, D_n = -1]$ 
6  $V_{local,i} \leftarrow \sum_{j,j \neq i} V_{i,j}$  given  $D$  // Eq. (3)
7 foreach possible  $D'$  of  $L_S \setminus \{\tilde{db}\}$  in Gray code order do
8   UPDATE POTENTIAL CACHE( $L, P, D, D'$ ) // Alg. 2
9   if  $D'$  is phys. valid given  $P$  then
10      $E \leftarrow$  energy of  $L$  given  $D'$  and  $P$  // Eq. (2)
11      $DE \leftarrow DE \cup \{(D', E)\}$ 
12   end if
13    $D \leftarrow D'$ 
14 end foreach
15 return  $DE$ 
```

To reiterate, this is an efficient way to cover all possible charge distributions by changing the charge state of only one SiDB in each iteration. The local electrostatic potential of the current charge distribution cache is updated with Alg. 2 (Line 8) to check if the physical validity is fulfilled (Line 9). If the charge distribution under consideration is physically valid, it is stored together with its corresponding electrostatic potential energy (Line 11).

Ultimately, all physically valid charge distributions with their corresponding electrostatic potential energy are returned (Line 15).

V. EXPERIMENTAL EVALUATIONS

To demonstrate the applicability and improved runtime of the resulting tool *QuickExact* in comparison to the state of the art, an exhaustive experimental evaluation is conducted, whose results are presented in this section. To this end, first the experimental setup is described in Section V-A. Afterwards, the respectively obtained experimental results are presented in Section V-B considering two different classes of SiDB layouts, namely 1) randomly-generated, and 2) established ones from the literature. Finally, the obtained results are discussed in Section V-C.

A. Experimental Setup

QuickExact has been implemented as summarized in Section IV-C. We compare the performance of *QuickExact* with *ExGS*. All code has been implemented in C++17 on top of the *fiction* framework [37], which is part of the *Munich Nanotech Toolkit* (MNT).² All code was compiled with AppleClang 14.0.0, and the experiments were carried out on a macOS 13.0 machine with an Apple Silicon M1 Pro SoC with 32 GB of integrated main memory.

B. Considered Layouts and Obtained Results

To validate the efficiency and accuracy of physical simulation, it is important to use a wide variety of SiDB layouts. To this end, in the first series of evaluations, 280

²The implementation of *QuickExact* as presented in this work is publicly available at <https://github.com/cda-tum/fiction>

TABLE I: Exact simulation of randomly generated layouts

BENCHMARK		PHYSICAL SIMULATION ^a	
#SiDBs	#inst.	ExGS [29]	QuickExact
33	20	>3600	75.67
32	20	>3600	50.69
31	20	>3600	11.42
30	20	>3600	3.24
29	20	>3600	2.83
28	20	>3600	2.95
27	20	3408.25	0.98
26	20	1700.68	0.28
25	20	827.67	0.11
24	20	414.38	0.05
23	20	207.25	0.04
22	20	103.64	0.00
21	20	48.60	0.00
20	20	25.34	0.00
<i>Sum</i> ^b		6735.81	1.49
%		100	0.02

^aAll runtime values are in seconds.

^bNot counting the >3600 rows.

TABLE II: Exact simulation of established gate layouts

BENCHMARK [17]			PHYSICAL SIMULATION ^a	
Name	#SiDBs	#inst.	ExGS [29]	QuickExact
DOUBLE WIRE	28	4	1211.85	0.48
CX	27	4	599.61	0.28
HA	24	4	71.35	0.04
AND	21	4	58.53	0.06
XOR	21	4	8.40	0.00
OR	21	4	6.84	0.00
XNOR	21	4	8.41	0.00
FO2	20	2	2.16	0.00
NOR	19	4	2.07	0.00
NAND	19	4	2.07	0.00
INV	18	4	1.04	0.00
WIRE	15	4	0.26	0.00
<i>Sum</i>			1972.58	0.91
%			100	0.05

^aAll runtime values are in seconds.

randomly-generated layouts with different numbers of SiDBs per layout (ranging from 20 to 33 SiDBs per layout) are used. In a second series of evaluations, all 12 gate layouts from the *Bestagon* gate library [17] are simulated. The *Bestagon* gate library is a universal collection of hexagonal SiDB standard tiles for circuit layout design that have been created using Reinforcement Learning with a focus on realistic design constraints imposed by fabrication capabilities.

All obtained results are summarized in Table I and Table II for the randomly-generated layouts and the established *Bestagon* library, respectively. In Table I, the first two columns provide the number of SiDBs and the number of simulated layouts (instance count), respectively. Subsequently, the simulation runtimes of *ExGS* and *QuickExact* are stated. Table II has the same structure; however, the name of the investigated

gate is added as first column. Since each 2-input gate can be simulated in up to four configurations, with input perturbers specifying one of the possible 2^2 input combinations, the instance count is generally 4. However, the fan-out tile (FO2) is a 1-input/2-output gate and thus only provides 2 instances. WIRE and INV, despite being 1-input tiles, have instance count 4 because two different types—straight and diagonal—exist for each.

C. Discussion

The simulation of the randomly-generated layouts and the established *Bestagon* library reveals a massive performance advantage of *QuickExact* compared to *ExGS*. For randomly-generated layouts, *ExGS* yields a timeout (one hour) for 120 layouts (40 of them need even more than 24 hours), while *QuickExact* determines valid solutions to all instances within seconds, resulting in a total runtime advantage of a factor of 5000 compared to the state of the art, *ExGS*.

These results are again confirmed by the simulation of the *Bestagon* gate library. While *ExGS* takes about 2000 s to simulate all gates in all configurations, *QuickExact* takes less than 1 s, again a reduction in runtime of more than a factor of 2000.

This runtime advantage in all cases demonstrates the clear benefit that the techniques proposed in this work have on exact physical simulation of SiDBs. We have hereby enabled the obtainment of exact simulation results for large layouts that were previously impossible to process.

VI. CONCLUSIONS

In this work, a novel approach, *QuickExact*, for exact physical simulation of SiDB layouts was presented with a special focus on runtime improvement. To this end, three techniques were proposed: 1) Physically-informed Search Space Pruning, 2) Partial Solution Caching, and 3) Effective State Enumeration. In an extensive experimental evaluation, these techniques lead to runtime improvements of up to a factor of 5000 compared to the state of the art, *ExGS*, as shown on two different sets of SiDB layouts: 1) randomly-generated ones, and 2) established standard gate tiles from the literature. In case of random layouts, *ExGS* did not yield any results within one hour for 120 of the considered 280 layouts (40 of them need even more than 24 hours), while *QuickExact* required 75.67 s at maximum. In case of the standard gate tiles, *ExGS* needed a total of around 2000 s while *QuickExact* stayed below 1 s, emphasizing its immense performance advantage. To support open research and open data, an implementation of *QuickExact* is publicly available as part of the *Munich Nanotech Toolkit* (MNT).

REFERENCES

- [1] J. Croshaw *et al.*, “Ionic Charge Distributions in Silicon Atomic Surface Wires,” *Nanoscale*, 2021.
- [2] J. Wyrick *et al.*, “Atom-by-Atom Fabrication of Single and Few Dopant Quantum Devices,” *Advanced Functional Materials*, 2019.
- [3] M. B. Haider *et al.*, “Controlled Coupling and Occupation of Silicon Atomic Quantum Dots at Room Temperature,” *Physical Review Letters*, 2009.
- [4] T. R. Huff *et al.*, “Atomic White-Out: Enabling Atomic Circuitry through Mechanically Induced Bonding of Single Hydrogen Atoms to a Silicon Surface,” *ACS Nano*, 2017.
- [5] N. Pavliček *et al.*, “Tip-Induced Passivation of Dangling Bonds on Hydrogenated Si(100)- 2×1 ,” *Applied Physics Letters*, 2017.
- [6] R. Achal *et al.*, “Lithography for Robust and Editable Atomic-scale Silicon Devices and Memories,” *Nature Communications*, 2018.
- [7] R. A. Wolkow *et al.*, “Silicon Atomic Quantum Dots Enable Beyond-CMOS Electronics,” 2014.
- [8] T. Huff *et al.*, “Binary Atomic Silicon Logic,” *Nature Electronics*, 2018.
- [9] X. Wang *et al.*, “Atomic-Scale Control of Tunneling in Donor-Based Devices,” *Communications Physics*, 2020.
- [10] J. L. Pitters *et al.*, “Charge Control of Surface Dangling Bonds using Nanoscale Schottky Contacts,” *ACS Nano*, 2011.
- [11] M. Rashidi *et al.*, “Initiating and Monitoring the Evolution of Single Electrons within Atom-defined Structures,” *Physical Review Letters*, 2018.
- [12] L. Livadaru *et al.*, “Dangling-bond Charge Qubit on a Silicon Surface,” *New Journal of Physics*, 2010.
- [13] R. Landauer, “Irreversibility and Heat Generation in the Computing Process,” *IBM Journal of Research and Development*, 1961.
- [14] R. W. Keyes *et al.*, “Minimal Energy Dissipation in Logic,” *IBM Journal of Research and Development*, 1970.
- [15] G. Toth *et al.*, “Quasiadiabatic Switching for Metal-Island Quantum-Dot Cellular Automata,” *Journal of Applied Physics*, 1999.
- [16] A. S. Dzurak *et al.*, “Construction of a Silicon-Based Solid State Quantum Computer,” *Quantum Information & Computation*, 2001.
- [17] M. Walter *et al.*, “Hexagons are the Bestagons: Design Automation for Silicon Dangling Bond Logic,” in *DAC*, 2022.
- [18] S. S. H. Ng *et al.*, “SiQAD: A Design and Simulation Tool for Atomic Silicon Quantum Dot Circuits,” *TNANO*, 2020.
- [19] M. D. Vieira *et al.*, “Three-Input NPN Class Gate Library for Atomic Silicon Quantum Dots,” *IEEE Design & Test*, 2022.
- [20] A. N. Bahar *et al.*, “Atomic Silicon Quantum Dot: A new Designing Paradigm of an Atomic Logic Circuit,” *IEEE Transactions on Nanotechnology*, 2020.
- [21] R. Lupoiu *et al.*, “Automated Atomic Silicon Quantum Dot Circuit Design via Deep Reinforcement Learning,” *arXiv*, 2022.
- [22] M. Walter *et al.*, “An Exact Method for Design Exploration of Quantum-dot Cellular Automata,” in *DATE*, 2018.
- [23] —, “One-pass Synthesis for Field-coupled Nanocomputing Technologies,” in *ASP-DAC*. ACM New York, NY, USA, 2021.
- [24] —, “Verification for Field-coupled Nanocomputing Circuits,” in *DAC*, 2020.
- [25] S. Hofmann *et al.*, “Scalable Physical Design for Silicon Dangling Bond Logic: How a 45° Turn Prevents the Reinvention of the Wheel,” in *2023 IEEE 23rd International Conference on Nanotechnology (NANO)*, 2023.
- [26] R. A. Wolkow *et al.*, “Multiple Silicon Atom Quantum Dot and Devices inclusive thereof,” US Patent 10 937 959, 2021.
- [27] R. Wolkow *et al.*, “Initiating and Monitoring the Evolution of Single Electrons within Atom-defined Structures,” US Patent 11 047 877, 2021.
- [28] J. Drewniok *et al.*, “*QuickSim*: Efficient and Accurate Physical Simulation of Silicon Dangling Bond Logic,” in *2023 IEEE 23rd International Conference on Nanotechnology (NANO)*, 2023.
- [29] S. S. H. Ng, “Computer-aided Design of Atomic Silicon Quantum Dots and Computational Applications,” Master’s thesis, University of British Columbia, 2020.
- [30] J. Drewniok *et al.*, “Temperature Behavior of Silicon Dangling Bond Logic,” in *2023 IEEE 23rd International Conference on Nanotechnology (NANO)*, 2023.
- [31] S. S. H. Ng *et al.*, “Charged Defect Simulation in SiDB Systems,” 2022. [Online]. Available: <https://arxiv.org/abs/2211.08698>
- [32] M. Rashidi *et al.*, “Automated Atomic Scale Fabrication,” US Patent, 2022, 17/429,443.
- [33] T. R. Huff *et al.*, “Electrostatic Landscape of a Hydrogen-Terminated Silicon Surface Probed by a Moveable Quantum Dot,” *ACS Nano*, 2019.
- [34] R. Achal *et al.*, “Detecting and Directing Single Molecule Binding Events on H-Si (100) with Application to Ultradense Data Storage,” *ACS Nano*, 2019.
- [35] M. Taucer, “Silicon Dangling Bonds Non-equilibrium Dynamics and Applications,” 2015.
- [36] M. Rashidi *et al.*, “Time-resolved Single Dopant Charge Dynamics in Silicon,” *Nature Communications*, 2016.
- [37] M. Walter *et al.*, “*fiction*: An Open Source Framework for the Design of Field-coupled Nanocomputing Circuits,” 2019.