Design Automation for Cryogenic CMOS Circuits

Victor M. van Santen[1], Marcel Walter[2], Florian Klemme[1], Shivendra Singh Parihar,[1, 3]
Girish Pahwa[4], Yogesh S. Chauhan[3], Robert Wille[2, 5], and Hussam Amrouch[2, 6]
[1]University of Stuttgart, Germany, [2]Technical University of Munich, Germany, [3]IIT Kanpur, India,
[4]University of California, Berkeley, USA, [5]Software Competence Center Hagenberg GmbH, Austria
[6]Munich Institute of Robotics and Machine Intelligence, Email: {marcel.walter, robert.wille, hussam.amrouch}@tum.de

***Abstract*—Cryogenic CMOS circuits operate at temperatures close to absolute zero and are essential in many applications such as controllers for quantum computing but also medical engineering, space technology, or physical instruments. However, operating circuits at cryogenic temperatures fundamentally changes the underlying semiconductor physics that governs the CMOS transistor—rendering existing design automation approaches infeasible. In this work, we propose and implement the first end-to-end approach that enables design automation for cryogenic CMOS circuits. To this end, we (1) perform the first-of-its-kind measurements of commercial $5\,\mathrm{nm}$ FinFET transistors from $300\,\mathrm{K}$ down to $10\,\mathrm{K}$, (2) use the results to validate and calibrate the first cryogenic-aware industrial-standard compact model for FinFET technology, (3) create cryogenic-aware standard cell libraries that are compatible with the existing EDA tool flows, and (4) propose an initial cryogenic-aware logic synthesis approach that re-uses established design automation expertise but optimizes it for cryogenic purposes. Evaluations, comparisons, and discussions of all these novel contributions confirm the applicability and validity of the resulting cryogenic-aware design automation flow.**

I. Introduction

Cryogenic CMOS circuits [1]–[3] describe electronic circuits (realized in "classical CMOS technologies") that operate at cryogenic temperatures. Although universal definitions endorse $120\,\mathrm{K}$ or colder as a threshold to distinguish "cryogenics" from conventional refrigeration, cryogenic circuits are assumed to operate at temperatures close to absolute zero, e.g., at some few Kelvin or even a fraction of a Kelvin. This is motivated by several applications in which staying close to absolute zero is key for reliable and successful execution.

Most importantly, cryogenic circuits are essential as an interface between classical computing and quantum computing [1]. Quantum computing [4] promises to solve a wide range of computational problems that are fundamentally challenging if not impossible in classical computing (including synthesizing new materials, optimizing drugs, or simulating quantum systems). The corresponding qubits operate at near absolute zero (e.g., $10\,\mathrm{mK}$) to ensure they stay in a superimposed and entangled state for as long as possible. At the same time, CMOS circuits are prerequisites to effectively connect the classical domain (where information is provided and processed) with the quantum domain (where certain computational steps are performed by qubits).

However, these control circuits currently work at room temperature. Since the corresponding qubits are very sensitive to noise, the heat radiated from such circuits may severely disturb the quantum computations. Hence, they have to be located far from the qubits. This, in turn, creates a serious bottleneck that obstructs increasing the number of qubits, which is inevitable in scaling up quantum computers. As an example, in recent experiments such as the one presented in [2], [5] the problem became evident: Here, engineers needed approximately 200 wideband coaxial cables along with 45 bulky microwave circulators and a rack of circuits to control merely 53 qubits [2], [5]. The challenge is that despite isolation, the significant temperature gradient ($300\,\mathrm{K} \leftrightarrow 0.1\,\mathrm{K}$) induced at the two ends of every wire can create a heat flux that leaks from the control circuits (outside the refrigerator) towards the qubits (inside the refrigerator)—jeopardizing the entire quantum system. In addition, the use of long cables introduces large latencies that make meeting the tight timing constraints, imposed by the short coherence time of qubits, also profoundly challenging. Having the control circuits working at cryogenic temperatures would resolve many of these problems—providing an important basis for practically relevant quantum computing. Besides that, further applications of cryogenic circuits include nuclear spin and magnetic resonance imaging in medical engineering [6], electronics for space applications where the temperature is close to absolute zero [3], or measurement instruments for particle physics [7]. However, operating CMOS circuits at cryogenic temperatures imposes tough power constraints on them because of the limited power dissipation capability that is feasible at such extreme temperatures. For example, a cryogenic quantum controller working at $10\,\mathrm{K}$ must operate within a power budget of merely $100\,\mathrm{mW}$. Otherwise, the generated heat would disturb or simply destroy the fragile state of the involved qubits. Therefore, *when operating CMOS circuits at cryogenic temperatures, power must become the primary optimization objective*.

This is in stark contrast to established design automation for CMOS. In fact, current EDA tools and flows are unaware of the fundamental changes that cryogenic temperatures cause in the underlying semiconductor physics that govern CMOS transistors. For instance, considering cryogenic temperatures, leakage current decreases, transistor sub-threshold slope decreases, and carrier mobility improves (while transistor threshold voltage increases). This makes the existing design automation flow lack the necessary information on how the delay and power of standard cells are impacted when operating at cryogenic temperatures. Hence, circuits that do not satisfy the respectively needed power and/or delay constraints result—severely impacting their applicability in the above-mentioned applications.

In this work, we raise awareness of these challenges and provide the fundamental basis for design automation for cryogenic CMOS circuits for the first time. To this end, we provide insights and solutions starting at the transistor level (where cryogenic temperatures alter the semiconductor physics and, thus, the transistor's behavior), to the standard cell level (where drifts in the transistor electrical parameters manifest themselves as changes in the delay and power of logic gates), and all the way up to the algorithms at the logic synthesis level (where circuits are optimized to fulfill a certain objective). More precisely, we performed the *first-of-its-kind measurements* of commercial $5\,\mathrm{nm}$ FinFET transistors from $300\,\mathrm{K}$ down to $10\,\mathrm{K}$. The results (providing the basis for real cryogenic-aware design automation) are then used to validate and calibrate the *first cryogenic-aware industrial-standard compact model of FinFET technology*. Our models are then employed to create the prerequisite *cryogenic-aware standard cell libraries*. In an effort to prevent "reinventing the wheel", the resulting cell libraries are designed to be fully compatible with existing design and synthesis tools. However, by investigating how the existing logic synthesis algorithms perform when generating cryogenic CMOS circuits, we reveal that, due to the lack of awareness, the synthesized circuits are sub-optimal for cryogenic purposes. Accordingly, we create an initial *cryogenic-aware logic synthesis approach* that primarily optimizes power (rather than established synthesis objectives) using open-source software—leading to circuits that noticeably dissipate (up to 28%) lower power compared to conventional synthesis.

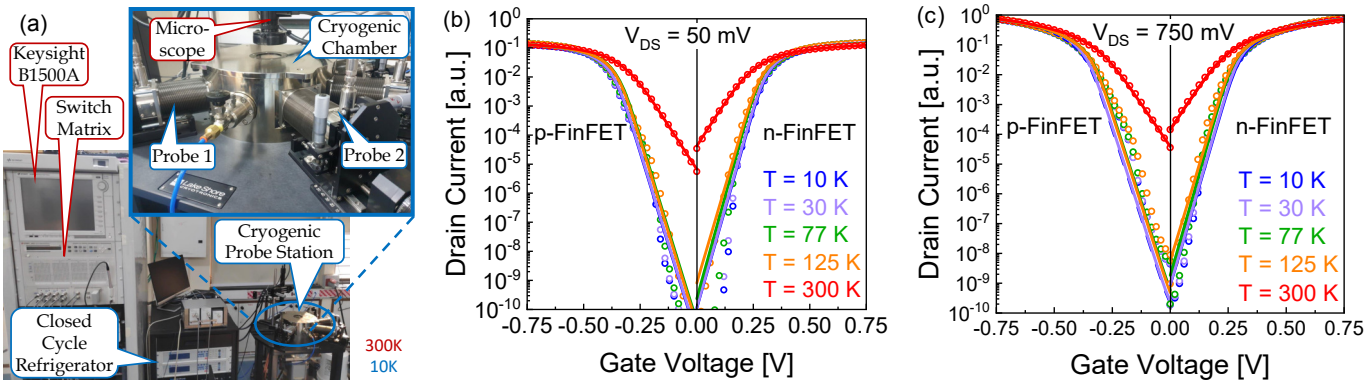

Fig. 1: (a) The used transistor measurement setup. It consists of a cryogenic probe station "Lakeshore CRX-VF", a closed-cycle refrigerator, and a semiconductor measurement station "Keysight's B1500A". (b and c) Measured drain current I_{DS} for various gate-source voltages V_{gs} for both low V_{ds} of 50 mV and high V_{ds} of 750 mV, respectively. (b and c) demonstrate, additionally, the validation of the proposed cryogenic-aware FinFET (BSIM-CMG) compact model. As observed, SPICE results (line) using the proposed model come with an excellent agreement with the obtained measurements (dots) for both n-FinFET and p-FinFET transistors.

Overall, all this leads to the following novel contributions:

- A cryogenic-aware physics-based transistor model that is capable of accurately capturing the transistor's behavior at such extreme temperatures (covered in Section II). The model is validated against measurements from commercial 5 nm FinFETs.
- A cryogenic-aware standard cell library that characterizes the combinational and sequential logic gates based on the above cryogenic-aware transistor model (covered in Section III).
- A cryogenic-aware logic synthesis in which power is the primary optimization objective and the internal cost function of algorithms is customized for that purpose (covered in Section IV).

Evaluations, comparisons, and discussions (covered in all the above-mentioned sections and later combined in Section V) confirm the applicability and validity of all these contributions. *By that, this is the first work to propose and implement an end-to-end approach enabling design automation for cryogenic CMOS circuits.*

II. Cryogenic-Aware Transistor Modeling for Commercial 5 nm FinFET Technology

Our development of a cryogenic-aware FinFET model consists of three major steps that are summarized in the following subsections.

A. Reflecting Cryogenic Temperatures the Transistor Compact Model

First, we aim at augmenting the FinFET compact model with the necessary physics-based equations to accurately reflect the effects that cryogenic temperatures induce in the electrical characteristics of p-type and n-type transistors. The need for an updated FinFET model was evident in several recent studies. For instance, [8] showed that, when 22 nm FinFETs operate at 77 K, their propagation delay become 37 % smaller, while the consumed power decreases by 50 % compared to the operation at 300 K. [9] presented cryogenic measurements from advanced 10 nm FinFETs showing 58 % mobility improvement, 100x leakage current reduction, and 39 % smaller delay. [10] reported a decrease in transistor transconductance due to the increase in diffusion resistance, leading to hysteresis effects that are detrimental. Last but not least, several works (e.g., [11]) reported a noticeable increase in V_{th} at cryogenic temperatures. A higher V_{th} requires the transistor to be biased at a higher V_{dd} to form its channel leading to higher power dissipation, which is critical to be accurately captured in any cryogenic CMOS circuit.

In this work, we employ the existing industry standard compact model for FinFET technologies (i.e., BSIM-CMG [12]), and we extend it to incorporate the additional physics-based equations required [13] to account for the changes induced at cryogenic temperatures. The resulting cryogenic-aware BSIM-CMG takes into account the short-channel effects caused by quantum effects that become apparent at small feature sizes like 5 nm. Furthermore, it captures how extremely low temperatures change the density/surface potential for band-tail states [13]. It is noteworthy that, since our model is based on the industry-standard compact model, it can be straightforwardly employed in SPICE tools to perform circuit simulations.

B. Measuring 5 nm FinFET Transistors for the Temperature Range

In this work, we perform measurements for commercial 5 nm FinFET (both n-type and p-type) transistors covering the entire temperature range from 300 K down to 10 K. The measured FinFETs feature a minimum channel length and have multi-fins and multi-finger. The measurement setup is depicted in Fig. 1(a) which shows the "Lakeshore CRX-VF" station that is used. The cryogenic probe station itself is then shown within the zoomed-in area, in which the probes are placed on top of the measurement pads for each transistor terminal (i.e., gate, source, drain, bulk). The cryogenic probe station features a 51 mm diameter sample stage, a vacuum pump, a two-stage closed-cycle refrigerator unit, and a probe handler/positioner. The heat flux from the probes introduces thermal fluctuations in the range of 3.5 K to 8.5 K. Therefore, the temperature 10 K is chosen as the lower limit because it is the smallest stable temperature. Lastly, a semiconductor measurement station "Keysight's B1500A" is utilized to stimulate each transistor under measurement (i.e., apply the required voltage biases) and, then, measure its electrical response.

C. Calibrating and Validating the Cryogenic-Aware FinFET Model

Finally, we calibrate and validate our developed cryogenic-aware BSIM-CMG against the obtained 5 nm FinFET measurement data. This step is performed for transistor transfer characteristics (i.e., I_{ds}-V_{gs}) under various voltage biases. This is essential to ensure that the underlying physics-based models incorporated within the FinFET compact model reproduce the technology node well. Fig. 1(b and c) present the measured drain current I_{DS} for various gate-source voltages V_{GS} for both low V_{DS} of 50 mV and high V_{DS} of 750 mV, respectively. As observed, SPICE results obtained using our cryogenic-aware FinFET model, come with an excellent agreement with the obtained measurements for both n-FinFET and p-FinFET transistors. It is noteworthy that in Fig. 1(b and c): (i) SPICE results are represented by lines, while the measurements are represented

by dots. (ii) The validation of our FinFET model is done not only for cryogenic temperatures (i.e., 77 K to 10 K) but for the entire temperature range starting from room temperature (300 K).

In summary: We augmented the industry-standard FinFET model with the required physics-based equations to accurately capture the impact of cryogenic temperatures. Then, we validated the resulting model against measurements for a commercial 5 nm FinFET. The validations demonstrated an excellent match between our developed model and experimental data for the entire target temperature range.

III. Cryogenic-Aware Standard Cell Libraries

To estimate the overall impact of cryogenic temperatures on the delay and power of circuits, cryogenic-aware standard cell libraries are prerequisites. In the following, we illustrate how we have created them. Then, we present a comparative analysis w.r.t. the critical figure of merits at the cell/circuit levels. Finally, we briefly discuss how cryogenic temperatures can challenge the existing logic synthesis.

A. Creation of the Standard Cell Libraries

In this work, we employ the cryogenic-aware FinFET transistor model (presented in Section II) which is validated against 5 nm measurements. Because our models are fully compatible with the existing SPICE tools, we can deploy them directly to perform cell library characterizations. To this end, we utilize post-layout SPICE netlists for a wide range of combinational and sequential logic gates obtained from the open-source ASAP7 PDK [14]. Every standard cell is characterized using *Synopsys SPICE* under different operating conditions, which are (7×7) input signal slews and output load capacitances. This is consistent with other academic and commercial standard cell libraries. During the characterization process (which has been automated using the *Synopsys SiliconSmart* tool flow), SPICE accurately measures each of the propagation delays, the switching energy, and the leakage power of each individual standard cell. Then, the obtained results are written using the so-called "liberty"-format, which is the industry-standard format for cell libraries. Note that the layout of standard cells within the ASAP7 PDK was done for 7 nm FinFETs, which is geometrically very close to our target 5 nm FinFET. In practice, more than 10^6 SPICE simulations have been performed to characterize a whole standard cell library, which consists of 200 combinational and sequential logic gates.

B. Comparative Analysis for Power and Delay

The cell library characterization process mentioned above has been performed at both room temperature (300 K) and cryogenic temperature (10 K). This enables us to fairly compare the figure of merits (e.g., propagation delay, switching power, internal power, etc.) when cells/circuits are operated at cryogenic temperature versus room temperature. Fig.2(a) presents the distribution of the propagation delay of *all cells* throughout the library for both 300 K and 10 K. As shown, the two distributions largely overlap. This demonstrates that the delay of standard cells is marginally impacted when the temperature is extremely reduced to 10 K. This is due to the fact that the ON current of n-FinFET and p-FinFET remain almost the same for the entire temperature range from 300 K to 10 , as Fig. 1(b and c) illustrate. Fig. 2(b) demonstrates the distribution of switching energy of *all cells* across the entire library. As shown, the standard cells exhibit slightly less energy at 10 K. This can be attributed to the lesser dynamic power due to changes in gate capacitance caused by shifts in the surface potential at cryogenic temperatures [13].

To analyze the power at the circuit level, we synthesized using *Synopsys Design Compiler* the available open-source RTL circuits within the EPFL benchmark suite [15]. Then, we performed power analysis using the *Synopsys PrimeTime* signoff tool to accurately estimate the leakage, internal, and switching powers in every circuit for both 300 K and 10 K. Fig. 2(c) summarizes the average contribution of each power category relative to the overall power. Importantly, it reveals that the leakage power contribution becomes *negligible* when circuits are operated at 10 K (merely 0.003%). This is due to the fact that cryogenic temperatures significantly decrease the transistor OFF current by several orders of magnitude, as evident in Fig. 1(b and c).

C. Impact of Cryogenic Temperatures on Logic Synthesis

In conventional scenarios when circuits operate at room temperature, the leakage power has a considerable contribution (15% as shown in Fig. 2(c)). Therefore, logic synthesis algorithms must consider the leakage power during netlist optimization. In contrast, the contribution of leakage power at cryogenic temperature becomes negligible. Therefore, optimization algorithms can safely ignore it and instead increase efforts to optimize other objectives such as dynamic power and delay. In the next section, we present our cryogenic-aware logic synthesis approach, which exploits the aforementioned observations to develop new cost functions that aim to mainly reduce dynamic power dissipation.

IV. Cryogenic-Aware Logic Synthesis

The results presented thus far demonstrate that a potential cryogenic-aware logic synthesis can and should focus on other objectives than the existing solutions available for conventional circuits that operate at room temperature. Nevertheless, in an effort to not entirely "reinvent the wheel", it makes sense to build any kind of cryogenic-aware solution on top of what has been developed in the past decades by the logic synthesis community. Hence, in this section, we first review selected techniques in conventional logic synthesis and discuss their adjustment to the cryogenic domain. On the basis of that, an initial cryogenic-aware synthesis flow emerges, in which power dissipation is prioritized over area and delay cost metrics.

A. Conventional Logic Synthesis

The chip design stage of logic synthesis takes an RTL description and converts it into a corresponding gate-level netlist. Logic synthesis can be broken down into (1) *RTL Synthesis*, i.e., creating an initial technology-independent logic network from a high-level description, (2) *Logic Optimization*, i.e., improving a given logic network w.r.t. to some technology-(in)dependent cost metrics, and (3) *Technology Mapping*, i.e., translating a logic network into a gate-level netlist that exclusively utilizes gates offered by the target technology. In this work, we are most interested in the steps of logic optimization and technology mapping because the changes in delay and power induced by cryogenic temperatures are reflected only in these stages.

1) Logic Optimization: The representation and manipulation of large-scale Boolean logic require the conceptualization of efficient data structures. One of the most commonly used ones today is the *And-Inverter Graph* (AIG) [16]. An AIG is a directed graph where each node represents the Boolean conjunction (AND gate) of exactly two predecessors. Via optional inverter labels at the edges, AIGs allow for a highly memory-efficient universal representation of Boolean logic. Over the years, various logic optimization algorithms based on AIG data structures have been proposed. Their goals are always similar: transforming a given AIG into a logically equivalent one with improvements with regard to a given cost function. The general consensus is that a *better* AIG will ultimately lead to a *better* design w.r.t. the target objective. At this stage of the design flow, most cost functions are not technology aware, i.e., AIGs are optimized through proxy criteria that are *assumed* beneficial for many technologies. The most common are *size* (number of nodes) and *depth* (number of levels), as they are strongly related to the resulting chip area and the critical path length. Some important techniques for this purpose are *restructuring* [17], *rewriting* [18],

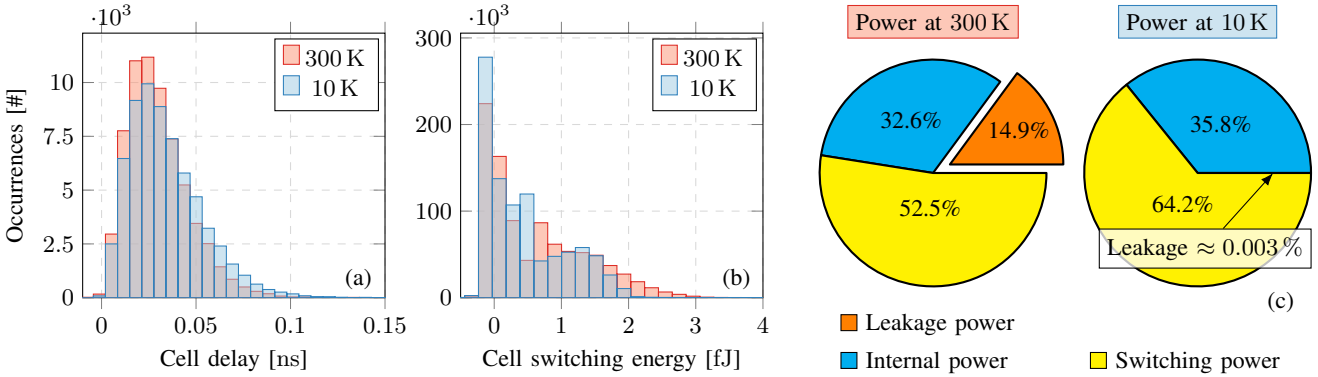

Fig. 2: (a and b) The distribution of propagation delay and switching energy of all cells (200 in total) across the entire standard cell library at both 300 K and 10 K. As shown, the cryogenic temperature has a negligible impact on the cells' delay. (c) The average contribution of leakage, internal, and switching powers to the overall power of various circuits from the EPFL benchmark suite. As shown, the contribution of leakage power at cryogenic temperature becomes negligible, contrary to the conventional scenario in which circuits operate at room temperature.

refactoring [19], and *resubstitution* [20], some of which are driven by powerful reasoning engines such as SAT solvers [21] and guided by a priority list of the cost functions aforementioned.

2) Technology Mapping: The input to technology mapping is a logic network and a cell library. Its output is a functionally equivalent mapped netlist that exclusively utilizes the available gates in the library. Optimizations in standardized technology-independent networks, such as AIGs, often translate directly into improvements for technology-dependent gate netlists as well. To this end, separating logic optimization from technology mapping was a means of handling complexity because target technologies could differ greatly. In its essence, technology mapping enumerates overlapping chunks of nodes in a given logic network—so-called *cuts*—that are manageable in size so that their truth tables can be computed; a technique that is highly efficient for small cuts [22]. Each cut is then matched against standard cells from the library. When faced with the decision on which cut to replace with which standard cell, a technology mapper relies on a pre-defined cost function that conventionally evaluates area or delay estimations [23].

B. Proposed Cryogenic-Aware Logic Synthesis

For cryogenic CMOS circuits, power optimization is of utmost importance due to the limited power dissipation capability in the cryogenic refrigerator. Consequently, a conventional logic synthesis that optimizes mainly for area and delay might yield sub-optimal or even unusable circuits that cannot fulfill the tough power constraints. Although some research was conducted on power-aware logic synthesis at both technology-independent and technology-dependent levels [24]–[26], conventional power consumption estimates might fail at cryogenic temperatures due to the exclusion of cell leakage power. Thus, we propose the adaptation of existing mature and proven logic synthesis techniques to incorporate cryogenic-aware cost functions.

For our studies, we use *ABC* [27], an industrial-strength, open-source logic synthesis and verification tool, as a running example and experimentation framework. *ABC* incorporates optional power-aware optimization strategies into a selection of its built-in algorithms, e.g., SAT-based resubstitution with "don't-care" optimization (`mfs`), technology-independent k-LUT mapping (`if`), structural choice computation (`dch`), and standard-cell technology mapping (`map`) [24]. Algorithms `if` and `map` use priority lists of cost functions to optimize the size of the given network first. If the size of two choices is equal within a threshold, the delay is utilized as a tie-breaker. Further down the hierarchy, costs like the fan-out limit and fan-in number are considered. However, since network size has been established as such a strong proxy criterion for area, delay, and power characteristics, even when it is particularly requested to optimize for power consumption, *ABC* refuses to give up on network size as its main optimization target. Alternatively, *ABC* simulates the switching activity of each node in the given network assuming a certain activation rate for each primary input. Consequently, nodes with high switching costs can be optimized by the `mfs` algorithm that back-annotates cost values between AIG and mapped representations of the same network to exploit structural properties. *The point that can be made by exploring mature EDA tools such as ABC is clear: conventional logic synthesis flows have not been built with cryogenic awareness in mind and hence they might yield sub-optimal netlists.*

A simple toggle that can be flipped or an input file that can be specified to tune logic synthesis more toward power consideration and away from aggressive size optimization does not exist. One has to change the source code of available tools to achieve such a goal. We propose to make power estimations the number one priority in all considerations of cryogenic-aware logic synthesis. Therefore, we modified the cost function priority lists in *ABC* in two ways (in descending priority): (i) $power \rightarrow delay \rightarrow area$, and (ii) $power \rightarrow area \rightarrow delay$. This requires some substantial changes in the corresponding code, which, due to page limitations, are not covered here.

V. Evaluation, Comparisons, and Discussion

All contributions described above provide the fundamental basis for design automation for cryogenic CMOS circuits. Aspects such as the model and the resulting cell libraries have already been validated and evaluated in Section II and Section III, respectively. In this section, we now confirm the applicability and validity of *all* these combined contributions. To this end, we utilize the proposed cryogenic-aware logic synthesis (as described in Section IV-B and based on the other contributions) and compare the results obtained by this with those generated by established (i.e., conventional) logic synthesis. By this, we evaluate, compare, and discuss (1) how much power savings does cryogenic-aware design automation provide compared to state-of-the-art conventional power-aware optimization and (2) how much are the circuits' delays impacted by the proposed optimization (relevant to ensure a fair comparison).

A. Experimental Setup

We consider all circuits available in the state-of-the-art EPFL benchmark suite [15] that was proposed to enable a contemporary

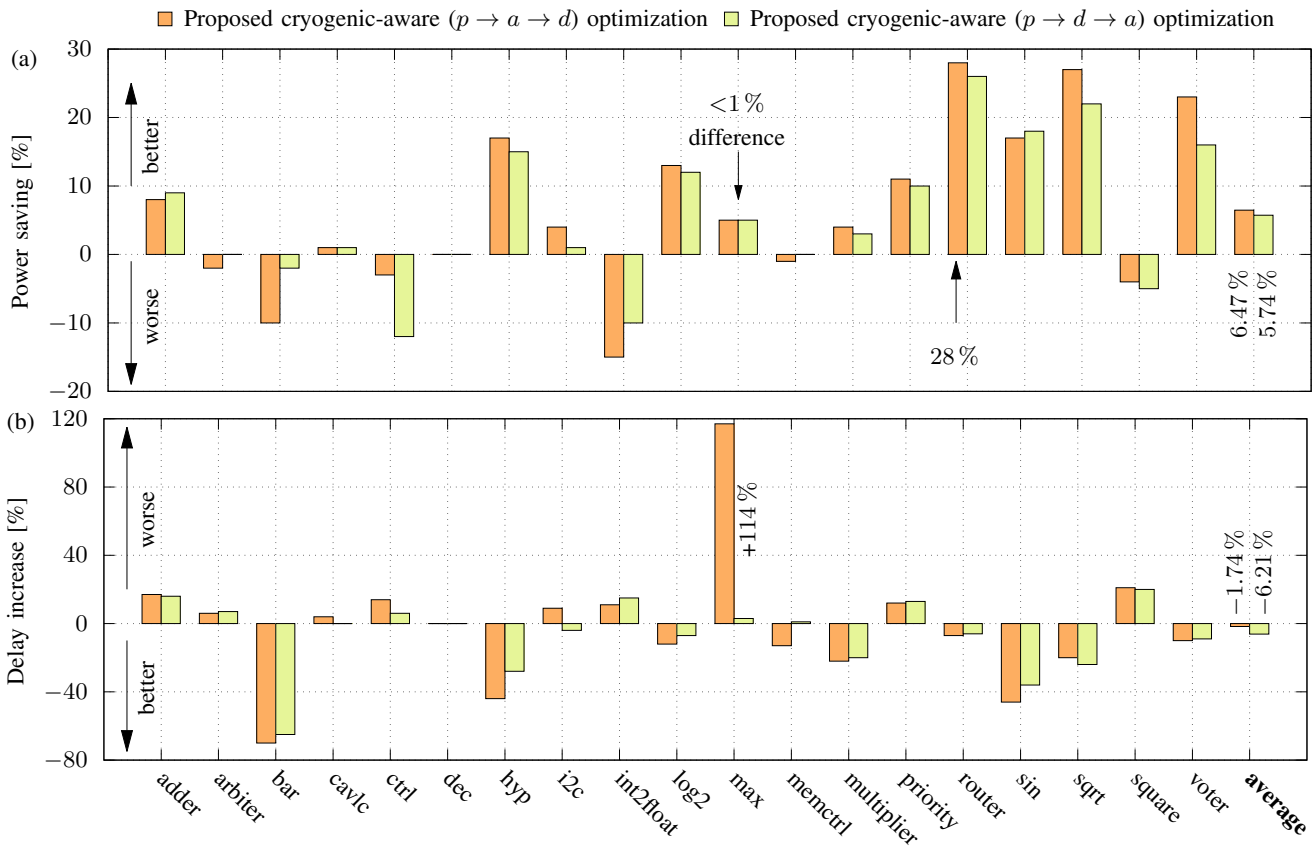

Fig. 3: Power and delay analysis for the EPFL benchmark circuits. We compare the effectiveness of the proposed cryogenic-aware logic synthesis against state-of-the-art power-aware logic synthesis where the best power optimizations that *ABC* offers are enabled. The proposed cost function ($p \rightarrow a \rightarrow d$) prioritizes $power \rightarrow area \rightarrow delay$, while ($p \rightarrow d \rightarrow a$) prioritizes $power \rightarrow delay \rightarrow area$.

comparative standard for the logic optimization and synthesis community. As discussed in Section IV, we utilize *ABC* [27], which is an open-source logic synthesis tool. Note that evaluating our work against commercial logic synthesis tools is not possible because such tools are not open source and, additionally, the end-user license agreement prohibits that. The recent version of *ABC* implements the state-of-the-art power-aware logic synthesis against which we compare the proposed cryogenic-aware optimization. For technology mapping, we employ the created cryogenic-aware standard cell libraries (described in Section III). To accurately quantify the power and delay of the gate-level netlists obtained by *ABC*, we use *Synopsys PrimeTime*, which is a signoff tool for power and delay analysis. Because our created libraries are fully compatible with existing commercial EDA tool flows, we can directly deploy them within *Synopsys PrimeTime* to perform the required power and delay analysis[1].

B. Logic Synthesis Analysis

We create a three-stage pipeline, with the input network being an AIG and the output being a technology-mapped circuit based on the provided cryogenic-aware standard cell library.
(1) *Technology-independent AIG optimization:* To initially compress the size of the input AIG, we apply a synthesis script consisting of a series of Boolean resubstitution, rewriting, and refactoring. It is predefined in *ABC* as a shortcut named `c2rs`.

[1]For fair comparisons, the power consumption of every circuit is estimated with a clock period set to the propagation delay of the slowest resulting circuit variant (e.g., the slowest *adder* determines the timing for all *adder*s but not other circuits). Otherwise, the faster circuit variants would be reported to consume considerably more power due to their higher clock speed.

(2) *Power-aware optimization:* Afterward, we collapse the AIG into k-LUTs with structural choices and optimize it by applying SAT-based resubstitution. Structural hashing converts the k-LUTs back into AIG nodes. This stage exploits structural redundancies and restructures the network for optimized power characteristics. All algorithms utilized in this stage were adjusted to use the proposed cryogenic-aware cost hierarchy (cf. Section IV-B). The respective *ABC* commands are `dch -p; if -p; mfs -pegd; strash`.
(3) *Technology mapping:* Finally, we apply the technology mapping that we adjusted for cryogenic-aware optimization to map the optimized AIG to a gate-level circuit using the proposed cryogenic-aware cells. The *ABC* command is `map -p`.
Using this setup, we evaluated the following scenarios:

- State-of-the-art power-aware logic synthesis: We apply the outlined three-stage pipeline without any modification to *ABC*'s original code, and we enable the best power optimizations that *ABC* offers out-of-the-box. This scenario serves as a baseline against which we compare our two proposed cost functions.
- The proposed cryogenic-aware ($p \rightarrow a \rightarrow d$) logic synthesis: Here, the proposed cost function that prioritizes $power \rightarrow area \rightarrow delay$ is implemented within *ABC*.
- The proposed cryogenic-aware ($p \rightarrow d \rightarrow a$) logic synthesis: Here, the proposed cost function that prioritizes $power \rightarrow delay \rightarrow area$ is implemented within *ABC*.

C. Experimental Results and Comparisons

Fig. 3(a) visualizes the power savings achieved with the two proposed cryogenic-aware synthesis settings compared to *ABC*'s state-of-the-art power-aware logic synthesis. As depicted, the proposed

cryogenic-aware synthesis leads to less power dissipation (up to $28\,\%$) in the majority of circuits. The average power saving is $6.47\,\%$ and $5.74\,\%$ for the case of ($p \rightarrow a \rightarrow d$) and ($p \rightarrow d \rightarrow a$), respectively (note that a negative saving, i.e., an overhead, is observed for some instances; this, however, is natural in an approach which relies on heuristics). Besides that, Fig. 3(b) demonstrates the potential delay increase caused by the two cryogenic-aware syntheses compared to the state-of-the-art power-aware logic synthesis. For the case ($p \rightarrow d \rightarrow a$), the delay is a second priority and as such is slightly impacted. However, for the other case ($p \rightarrow a \rightarrow d$), the delay becomes the third priority. Hence, a large increase might occur. As observed in the circuit *max*, the power saving marginally improves by $<1\,\%$, but this comes with a significant delay penalty reaching $114\,\%$. On average, both cryogenic-aware synthesis achieve a delay *reduction*, that is, a negative overhead of $-1.74\,\%$ and $-6.21\,\%$ for ($p \rightarrow d \rightarrow a$) and ($p \rightarrow a \rightarrow d$), respectively. Importantly, a negative sign here means *performance improvements* because smaller delays allow faster clocks.

D. Discussion and Perspective

Compared to state-of-the-art power-aware logic synthesis in which the best power optimizations that *ABC* can offer are present, the proposed cryogenic-aware logic synthesis (e.g., $p \rightarrow d \rightarrow a$) makes a difference. On average, $6.47\,\%$ and, in the best case, up to $28\,\%$ power savings can be reported while the delay is additionally improved by $1.74\,\%$. As a result, circuits optimized by the proposed cryogenic-aware logic synthesis will dissipate less power as heat, which is the key prerequisite for cryogenic circuits. While, at a first glance, these savings might seem to be moderate, it still clearly demonstrates that cryogenic-aware design automation makes the difference in a well-researched domain where every single improvement does count. Moreover, we would like to point out that the cryogenic-aware logic synthesis proposed in this work constituents an *initial* solution for this domain (which is compared against *ABC*'s mature algorithms that have been highly optimized over the years and reflect the state of the art of conventional logic synthesis). We truly believe that substantial further optimizations are possible, but we left this for future work. The transistor models and cell libraries, as well as the initial synthesis approach presented in this work, provide an ideal basis for this purpose and pave the way for cryogenic-aware EDA.

VI. Conclusion

In this work, we presented the first end-to-end design automation approach for cryogenic CMOS circuits. The made contributions span across multiple abstraction levels starting from (1) presenting a cryogenic-aware compact model which was validated against the first-of-its-kind measurements of commercial $5\,\mathrm{nm}$ FinFET transistors, to (2) creating cryogenic-aware standard cell libraries, all the way up to (3) implementing the first cryogenic-aware logic synthesis approach. We demonstrated that cryogenic CMOS circuits bring new challenges to the existing EDA tool flows and provided the basis for developing novel optimization algorithms customized for this purpose. Such algorithms will be essential in the near future to enable applications such as large-scale quantum computers in which ultralow-power cryogenic circuits must be harmlessly implemented without disturbing, e.g., qubits, with their deleterious heat radiation.

Acknowledgement

This work is partially supported by Advantest as part of the Graduate School "Intelligent Methods for Test and Reliability", received funding from the European Research Council (ERC) under the Horizon 2020 research and innovation program of the European Union (grant agreement No. 101001318), and was part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

References

[1] X. Xue, B. Patra, J. P. van Dijk, N. Samkharadze, S. Subramanian, A. Corna, B. Paquelet Wuetz, C. Jeon, F. Sheikh, E. Juarez-Hernandez *et al.*, "CMOS-based cryogenic control of silicon quantum circuits," *Nature*, 2021.

[2] S. Pauka, K. Das, R. Kalra, A. Moini, Y. Yang, M. Trainer, A. Bousquet, C. Cantaloube, N. Dick, G. Gardner *et al.*, "A Cryogenic CMOS chip for generating control signals for multiple qubits," *Nature Electronics*, 2021.

[3] T. Wada, H. Nagata, H. Ikeda, Y. Arai, M. Ohno, and K. Nagase, "Development of low power cryogenic readout integrated circuits using fully-depleted-silicon-on-insulator CMOS technology for far-infrared image sensors," *Journal of Low Temperature Physics*, 2012.

[4] M. A. Nielsen and I. Chuang, "Quantum computation and quantum information," 2002.

[5] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, 2019.

[6] D. H. Johansen, J. D. Sanchez-Heredia, J. R. Petersen, T. K. Johansen, V. Zhurbenko, and J. H. Ardenkjær-Larsen, "Cryogenic preamplifiers for magnetic resonance imaging," *TBCAS*, 2017.

[7] A. Pullia, F. Zocca, S. Riboldi, D. Budjas, A. D'Andragora, and C. Cattadori, "Cryogenic performance of a low-noise JFET-CMOS preamplifier for HPGe detectors," *TNS*, 2010.

[8] C. Auth, C. Allen, A. Blattner, D. Bergstrom, M. Brazier, M. Bost, M. Buehler, V. Chikarmane, T. Ghani, T. Glassman *et al.*, "A 22nm high performance and low-power cmos technology featuring fully-depleted tri-gate transistors, self-aligned contacts and high density mim capacitors," in *VLSIT*, 2012.

[9] H. Chiang, T. Chen, J. Wang, S. Mukhopadhyay, W. Lee, C. Chen, W. Khwa, B. Pulicherla, P. Liao, K. Su *et al.*, "Cold CMOS as a power-performance-reliability booster for advanced FinFETs," in *VLSIT*, 2020.

[10] B. Dierickx, L. Warmerdam, E. Simoen, J. Vermeiren, and C. Claeys, "Model for hysteresis and kink behavior of MOS transistors operating at 4.2 K," *TED*, 1988.

[11] R. Joshi, J. Timmerwilke, K. Tien, M. Yeck, and S. Chakraborty, "A 0.31 V Vmin Cryogenic SRAM Based Memory in 14nm FinFET Technology For Quantum Application," in *VLSI Symposium*, 2022.

[12] J. P. Duarte, S. Khandelwal, A. Medury, C. Hu, P. Kushwaha, H. Agarwal, A. Dasgupta, and Y. S. Chauhan, "BSIM-CMG: Standard FinFET compact model for advanced circuit design," in *ESSCIRC*, 2015.

[13] G. Pahwa, P. Kushwaha, A. Dasgupta, S. Salahuddin, and C. Hu, "Compact modeling of temperature effects in FDSOI and FinFET devices down to cryogenic temperatures," *TED*, 2021.

[14] L. T. Clark, V. Vashishtha, L. Shifren, A. Gujja, S. Sinha, B. Cline, C. Ramamurthy, and G. Yeric, "ASAP7: A 7-nm finFET predictive process design kit," *Microelectronics Journal*, 2016.

[15] L. Amarú, P.-E. Gaillardon, and G. De Micheli, "The EPFL combinational benchmark suite," in *IWLS*, 2015.

[16] A. Kuehlmann and F. Krohm, "Equivalence checking using cuts and heaps," in *DAC*, 1997.

[17] J. Cortadella, M. Kishinevsky, and A. Mishchenko, "Restructuring multilevel networks by using function approximations," 2003.

[18] A. Mishchenko, S. Chatterjee, and R. Brayton, "DAG-aware AIG rewriting: A fresh look at combinational logic synthesis," in *DAC*, 2006.

[19] R. Brayton and A. Mishchenko, "Scalable logic synthesis using a simple circuit structure," in *Proc. IWLS*, 2006.

[20] A. Mishchenko, R. Brayton, J.-H. R. Jiang, and S. Jang, "Scalable don't-care-based logic optimization and resynthesis," *TRETS*, 2011.

[21] J. S. Zhang, S. Sinha, A. Mishchenko, R. K. Brayton, and M. Chrzanowska-Jeske, "Simulation and satisfiability in logic synthesis," *Computing*, 2005.

[22] J. Cong, C. Wu, and Y. Ding, "Cut ranking and pruning: Enabling a general and efficient FPGA mapping solution," in *FPGA*, 1999.

[23] S. Chatterjee, A. Mishchenko, R. K. Brayton, X. Wang, and T. Kam, "Reducing structural bias in technology mapping," *TCAD*, 2006.

[24] S. Jang, K. Chung, A. Mishchenko, R. Brayton *et al.*, "A power optimization toolbox for logic synthesis and mapping," 2009.

[25] L. Benini and G. D. Micheli, "Logic synthesis for low power," *Logic Synthesis and Verification*, 2002.

[26] V. Tiwari, P. Ashar, and S. Malik, "Technology mapping for low power in logic synthesis," *Integration*, 1996.

[27] R. Brayton and A. Mishchenko, "ABC: An Academic Industrial-Strength Verification Tool," in *CAV*. Springer, 2010, pp. 24–40.