A 1.2-V 90-MHz Bitmap Index Creation Accelerator with 0.27-nW Standby Power on 65-nm Silicon-On-Thin-Box (SOTB) CMOS

Xuan-Thuan Nguyen*, Trong-Thuc Hoang†, Katsumi Inoue†, Ngoc-Tu Bui†, Van-Phuc Hoang‡, and Cong-Kha Pham§
*University of Toronto, 40 St. George St., Toronto ON M5S 3H7, Canada
†§University of Electro-Communications (UEC), 1-5-1 Chofugaoka, Chofu-shi, Tokyo 182-8585, Japan
‡Le Quy Don Technical University, 236 Hoang Quoc Viet St., Bac Tu Liem Dist., Hanoi, Vietnam
Email: *xuanthuan.nguyen@utoronto.ca; §phamck@uec.ac.jp

***Abstract*—Although bitmap index (BI) can surmount complex and multi-dimensional queries, the creation of BI itself is a time-consuming task. Many studies exploit the highly parallel processing capabilities of multi-core CPUs, graphics processing units (GPUs), or field-programmable gate arrays (FPGAs) to overcome this obstacle. This study, on the other hand, proposes a 65-nm silicon-on-thin-buried-oxide (SOTB) hardware accelerator dedicated to BI creation. The fabricated chip could operate at different supply voltages, from 0.45-V to 1.2-V. Concretely, in the active mode with the supply voltage of 1.2-V, this chip was fully operational at 90-MHz and consumed approximately 88.1-pJ/cycle. In the standby mode with the supply voltage of 0.45-V and clock gated, the power consumption was only 476.1-nW. Moreover, when the reverse back-gate bias voltage of -2.5-V is supplied, the standby power sharply dropped to 0.27-nW or approximately 1,763 times. This achievement is vitally essential for the energy-efficient applications, where the performance should be maximized during peak workload hours and the power should be minimized during off-peak time.**

I. Introduction

Recent years have witnessed a massive growth of global data generated from web services, social media networks, and science experiments, as well as the explosion of Internet-of-Things devices. According to Cisco's forecast, total data center traffic would be projected to hit 20.6 zettabytes (ZB) by the end of 2021 [1]. Gaining insight into a vast amount of data is highly important because valuable data are the driving force for business decisions and processes, as well as scientists' exploration and discovery. However, analyzing such the massive data is a time-consuming task.

In order to tackle this challenge, data are usually indexed in advance. Depending on the applications, specific indexing frameworks are employed. For example, B-tree and hash indexing are widely used in online transaction processing systems, where the number of querying and updating processes is nearly equal. Bitmap index (BI), besides, is commonly deployed in the online analytical processing (OLAP) systems, where updating processes infrequently happened. Furthermore, most queries in the OLAP system are highly complex, involve aggregations, and requires limited execution time. These problems, fortunately, can also be addressed by BI.

BI was first defined by Wong *et al.* [2] and was later popularized by O'Neil *et al.* [3]. Generally speaking, BI is a bit-level array, where the number of bits is equal to the dataset length. With BI generated, answering multi-dimensional queries is simply done by using only the bitwise operations, e.g., AND, OR, XOR, and NOT. Because each operation can be executed independently, the BI-based system is suitable for the highly parallel processing platforms, such as multi-core CPUs [4], [5], graphics processing units (GPUs) [6], and field-programmable gate arrays (FPGAs) [7].

This work would like to further exploit the implementation of BI in an application-specific integrated circuit (ASIC). Specifically, a hardware accelerator dedicated to BI creation (BIC) is our primary focus because generating BI is considered as the most expensive task in query processes, as verified by J. Chou *et al.* [4]. The authors pointed out that indexing 50-TB dataset cost as high as two hours while answering a multi-dimensional query using generated BIs took only 12 seconds. The accelerator then was implemented in the 65-nm silicon-on-thin-buried-oxide (SOTB) CMOS process. The SOTB technology is targeted because it allows the accelerator to operate at either high performance or low power by simply applying some specific back-gate biasing voltages [10]. The measurement results showed that the BIC chip was fully operational at 90 MHz and consumed 7.75-mW when the supply voltage is 1.2-V. The maximum energy point, therefore, was 88.1-pJ/cycle. In the standby mode, when the clock was utterly cut off, and a supply voltage of 0.4-V and a reverse back-gate bias voltage of -2.5-V were applied, the leakage power reduced as low as 0.27-nW.

The remainder of this paper is organized as follows. Section II briefly mentions the background of BI and related works. Section III describes the proposed hardware architecture. Section IV shows the measurement results of BIC chip on the 65-nm SOTB CMOS. Finally, Section V gives the conclusion.

II. Background and Related Works

Fig. 1 depicts an attribute named $Sale$, which is extracted from a typical relational database. This attribute stored eight sale records numbered sequentially from zero to seven. For example, a BIC command "*(Sale $<$ \$50k) OR (Sale $>$ \$250k) OR (Sale = [\$100k, \$200k])*" is needed, then three corresponding sub-BIs must be calculated beforehand. Taking the first statement of *(Sale $<$ \$50k)* as an example, the 2^{nd} and 6^{th} bits turn into ones because both $Sale$[2] and $Sale$[6] are smaller than \$50k. The master BI is then obtained by performing bitwise OR in three sub-BIs above.

978-1-7281-0397-6/19/$31.00 ©2019 IEEE

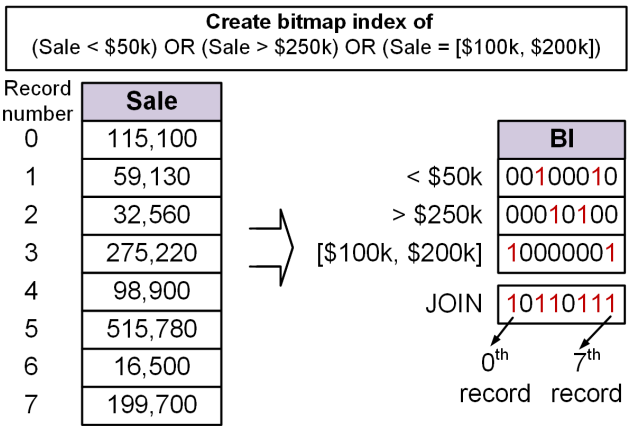

Fig. 1: An example of BIC.

As generating BI requires huge computation cost, multi-core CPUs, GPUs, and FPGAs are exploited to reduce the indexing time. J. Chou *et al.* [4] proposed an efficient BI framework that indexed massive datasets using modern supercomputing platforms. The indexing throughput reached 6.1 GB/s and 15.7 GB/s when 2,304 and 11,520 cores were utilized, respectively. C. H.-Te *et al.* [5] further improved that work by taking advantage of the memory technology. Specifically, the authors proposed an in-memory system that allowed the indexing throughput to reach as high as 28 GB/s and 510 GB/s when using 1,250 cores and 20,000 cores, respectively.

F. Fusco *et al.* [6] exploited the highly parallel capability of NVIDIA Geforce GTX-670 GPU for indexing data in real-time. According to the authors, the GPU-based design could index 20 million random 16-bit numbers within 122.6-μs or reached a throughput of 163 million words per second. X.-T. Nguyen *et al.* [7] proposed an energy-efficient BIC design that could index up to 32,768 16-bit words at every 10-ns clock cycle. The experimental results on an Intel Arria V FPGA confirmed this design achieved approximate indexing throughput of 1.46 GB/s and only consumed 3.28% of energy compared with the CPU-based and GPU-based approaches.

Oracle recently designed the data analytics accelerator engines [8] that were integrated into their processors to speed up analytic queries, as well as to performed real-time data decompression. With such in-memory query acceleration, Oracle's processors could deliver performance that was up to 10 times faster compared to other processors. Following the trend, our work proposed a 65-nm BIC accelerator especially aiming to the OLAP systems.

III. Proposed Designs

Fig. 2(a) gives an example of BIC command consisting of three statements. Each statement returns a separate BI depending on the given *threshold*, such as "$<$ \$50k". The BI length is identical to the length of $Sale$ attribute. The final result is obtained by joining three BIs with two OR *operations*.

Fig. 2(b) illustrates the high-level architecture of the BIC system including a direct access memory controller (DMAC) and a BIC accelerator. The DMAC is responsible for transferring input data and BI results back and forth between the external memory and accelerator with the lowest latency. The detailed architecture of the BIC system is depicted in Fig. 2(c). It is comprised of three main modules, namely Parallel Comparator (PACMP), Parallel Joiner (PAJOIN), and Post-Processor (PSPROC).

- PACMP module contains N 16-bit memory words (WORDs) and comparators (CMPs) operating in parallel. The WORD length can be easily modified, such as from 16 bits to 32 bits, according to the application requirements. In order to prepare for indexing, each *attribute* is sequentially pushed into WORDs. When all WORDs are filled up, each *operation* and *threshold* are dispatched to all CMPs in turn. Each BI is subsequently outputted pipelining after every clock cycle.
- PAJOIN module contains a set of computation units, each of which consists of two logic gates, i.e., NOT and OR, one multiplexer, and one-bit register (BIT). All BITs are cleared at the beginning of every BI request command. The first BI of *(Sale $<$ \$50k)* is stored in BITs due to OR operation. The second and third BIs are also calculated in a similar vein. After three clock cycles, the final result comes out and is immediately transferred to the next module.
- PSPROC module contains two sub-modules, namely multi-match priority encoder (MMPE) and BUFFER. The former aims to return the positions of all activated bits inside the results. Its architecture is mainly based on the so-called 1D-to-2D-based method presented in [9]. For instance, if the PAJOIN output is 01000110, the PSPROC outputs would become $\{1, 5, 6\}$. The latter, on the other hand, returns the raw BI with specific headers and footers attached.

Assume that the number of attributes, or the number of WORDs, is N and the number of statements in a BIC command is K, then the indexing time T_{INDEX} is roughly the sum of N and K. We exclude PSPROC from the estimation because this module can operate at the same time with PACMP. Nonetheless, the actual indexing time would be slightly longer due to the external memory access latency.

IV. Performance Analysis

This section briefly introduces the character of a 65-nm SOTB technology and evaluates in detail the performance of BIC chip on that technology.

A. SOTB CMOS Process

Fig. 3 describes the SOTB CMOS, which is an improvement of fully depleted silicon on insulator (FD-SOI) CMOS, with ultra-thin buried oxide (BOX) layer. Because the low impurity concentration in the channel region leads to the small parameter variations, SOTB is a good candidate for low-voltage operations. Furthermore, the ultra-thin BOX layer allows efficient transistor control by applying an appropriate back-gate bias voltage V_{BB}, where $V_{BB} = V_{BN} = V_{DD} - V_{BP}$. In fact, the reverse back-gate biasing (RBB) can deeply reduce the leakage current, whereas the forward back-gate biasing (FBB) can significantly increase the operating frequency, which correspondingly targets either low-power or high-performance requirements, respectively [10], [11].

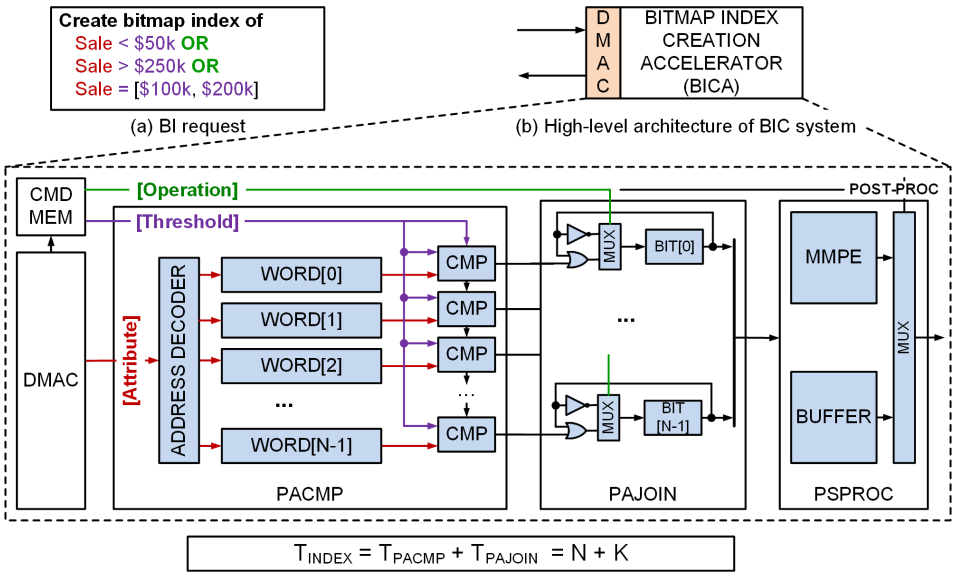

Fig. 2: The block diagram of the proposed BIC system.

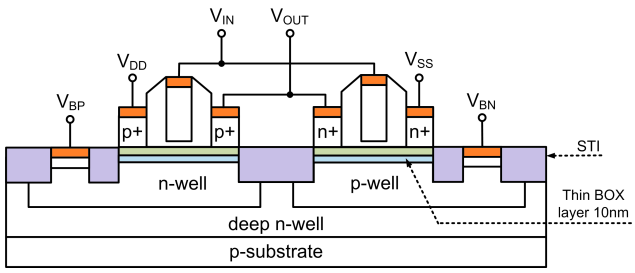

Fig. 3: The block diagram of a SOTB CMOS.

B. Measurement Results

The BIC system with N = 256 and K = 64 was fabricated in a 65-nm SOTB CMOS process. Fig. 4 shows the micrograph and chip summary. A source meter was used to supply the required voltages, such as core V_{DD}–from 0.45-V to 1.2-V–and I/O V_{DD}, as well as measure the active and leakage currents. The clock frequency was provided by an Intel Cyclone IV FPGA board. The input data were also generated and sent out at each clock cycle by that FPGA board. After receiving the input data, the chip started processing and returned the result to the FPGA for verification. The BIC chip reached the highest operating frequency of 90-MHz at 1.2-V and consumed approximately 7.75-mW. In the standby clock-gating (CG) mode, the power consumption was around 476.1-nW. Moreover, if reverse back-gate biasing voltage of -2.5-V was applied, the power significantly dropped off 1,763 times, to 0.27-nW.

Fig. 5 shows the variation of frequency when the supply voltage V_{DD} ranges from 0.45-V to 1.2-V. Without back-gate biasing, the minimum frequency was 1-MHz at 0.45-V, while the maximum frequency reached 90-MHz at 1.2-V. Accordingly, the minimum and maximum power consumptions were 0.095-mW and 7.75-mW, respectively. If a forward V_{BB} = 2.5-V was supplied, the frequency at 0.45-V was significantly picked up to 14-MHz, which was 14 times higher than the normal case. On the other hand, if a reverse V_{BB} = -2.5-V is applied, the chip could only operate at $V_{DD} \geq$ 0.7-V.

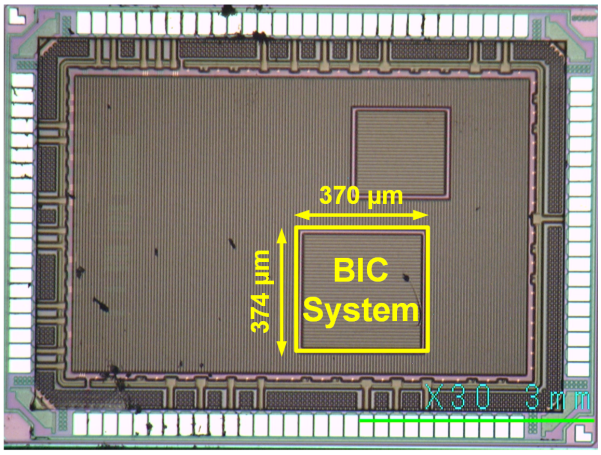

Technology	65-nm SOTB CMOS
Core size	0.14 mm^2
# of Cells	29,777
# of Transistors	367,500
I/O V_{DD}	0.7 V, 1.8 V, 3.3 V
Core V_{DD}	0.45 V – 1.2 V
Active Power	0.095 mW @ 1 MHz @ 0.45 V 7.75 mW @ 90 MHz @ 1.2 V
Standby Power	476.1 nW @ 0.45 V (CG) 0.27 nW @ 0.45 V (CG + RBB)

Measured at room temperature

Fig. 4: The die chip and its features.

Fig. 6 depicts the energy consumption of the BIC chip defined as the quotient of power and frequency. Without back-gate biasing, the maximum energy point was 88.1-pJ/cycle at 1.2-V. The minimum energy point varied according to the biasing voltages. Depending on the specific requirements, e.g., high performance or low power, either forward V_{BB} or reverse V_{BB} is provided, respectively.

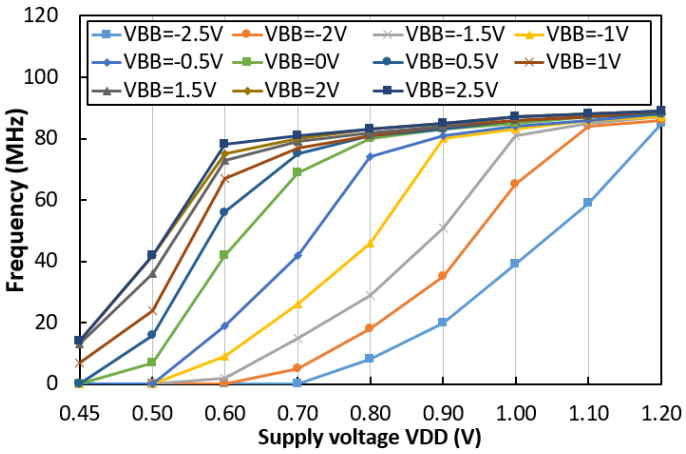

Fig. 5: Frequency versus supply voltage.

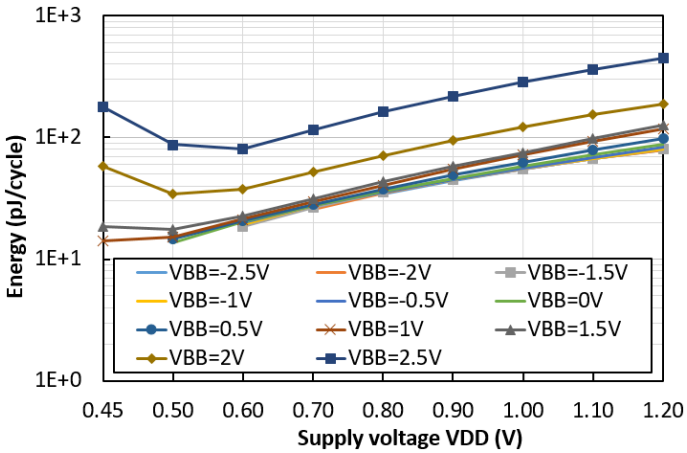

Fig. 6: Energy versus supply voltage.

Fig. 7 shows the leakage current in standby mode (I_{STB}) when reverse V_{BB} varied from -2.5-V to 0-V and V_{DD} varied from 0.45-V to 1.2-V. Without biasing, i.e., V_{DD} = 0.45-V and V_{BB} = 0-V, the standby leakage current I_{STB} was approximately 476.1-nA. However, whenever V_{BB} decreases by 0.5-V, I_{STB} was proportionally reduced up to one order of magnitude. The minimum I_{STB} = 0.6-nA achieved at V_{BB} = -2.5-V. In other words, I_{STB} drops 1,763 times (1.06-μA/0.6-nA) if V_{BB} changes from 0-V to -2.5-V. It is noted that I_{STB} could not reduce linearly mainly because of the gate induced drain leakage current (I_{GIDL}). In fact, although I_{GIDL} was theoretically suppressed by SOTB structure, it can significantly increase and completely dominate I_{STB} when V_{BB} gets smaller, and V_{DD} gets bigger [11].

V. Conclusion

This paper has proposed a 65-nm SOTB chip dedicated for BIC. This chip could operate at a wide-range supply voltage between 0.45-V and 1.2-V. Specifically, at 1.2-V, it was fully operational at 90-MHz and required 7.75-mW. By employing CG and RBB techniques, the ultra-low standby power of 0.27-nW has been attained. Due to that achievement, the proposed system can either maximize the performance during peak workload hours or minimize the power consumption during the off-peak time, which supposes to play a significant role in the energy-efficient systems.

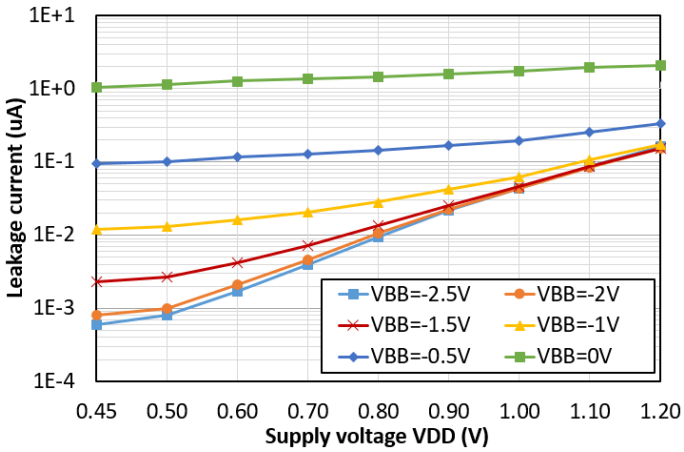

Fig. 7: The leakage current.

Acknowledgment

This work is supported by VLSI Design and Education Center (VDEC), the University of Tokyo in collaboration with Synopsys, Inc. and Cadence Design Systems, Inc.

References

[1] Cisco Global Cloud Index: Forecast and Methodology, 2016–2021. https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html. Accessed day: 2018/10.

[2] H. Wong, H.-F. Liu, F. Olken, D. Rotem, and L. Wong, "Bit transposed files," in *Proceedings of the 11th International Conference on Very Large Data Bases*, pp. 448–457, 1985.

[3] P. O'Neil and D. Quass, "Improved query performance with variant indexes," *ACM SIGMOD Record*, vol. 26, no. 2, pp. 38–49, 1997.

[4] J. Chou, R. D. Ryne, M. Howison, B. Austin, K. Wu, J. Qiang, E. W. Bethel, A. Shoshani, and O. Rbel, "Parallel index and query for large scale data analysis," in *Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 1–11, 2011.

[5] C. Hsuan-Te, J. Chou, V. Vishwanath, and W. Kesheng, "In-memory Query System for Scientific Datasets," in *Proceedings of IEEE 21st International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 362–371, 2015.

[6] F. Fusco, M. Vlachos, X. Dimitropoulos, and L. Deri, "Indexing million of packets per second using GPUs," in *Proceedings of the Conference on Internet Measurement Conference (IMC)*, pp. 327–332, 2013.

[7] X.-T. Nguyen, T.-T. Hoang, H.-T. Nguyen, K. Inoue, and C.-K. Pham, "An FPGA-Based Hardware Accelerator for Energy-Efficient Bitmap Index Creation," *IEEE Access*, Vol. 6, No. 1, pp. 16046 - 16059, 2018.

[8] Oracle SPARC M7 Processor, http://www.oracle.com/us/products/servers-storage/sparc-m7-processor-ds-2687041.pdf. Accessed day: 2018/10.

[9] X.-T. Nguyen, H.-T. Nguyen, and C.-K. Pham, "A Scalable High-Performance Priority Encoder Using 1D-array to 2D-array Conversion," *IEEE Trans. Circuits Syst. II: Express Briefs*, Vol. 64, No. 9, pp. 1102–1106, 2017.

[10] R. Tsuchiya, M. Horiuchi, S. Kimura, M. Yamaoka, et al., "Silicon on thin BOX: a new paradigm of the CMOSFET for low-power and high-performance application featuring wide-range back-bias control," in *Proc. IEEE Int. Electron Devices Meeting (IEDM) Technical Digest*, pp. 631–634, 2004.

[11] K. Ishibashi, N. Sugii, S. Kamohara, K. Usami, H. Amano, K. Kobayashi, and C.-K. Pham, "A Perpetuum Mobile 32bit CPU on 65nm SOTB CMOS Technology with Reverse-Body-Bias Assisted Sleep Mode," *IEICE Trans. Electron.*, Vol. E98C, No. 7, pp. 536–543, 2015.