Hitachi and the University of Tokyo develop dynamic pruning technology that speeds up Big Data retrieval time by up to 135 times

Supporting traceability in the manufacturing industry and data utilization in the medical care and finance fields will help solve issues confronting society

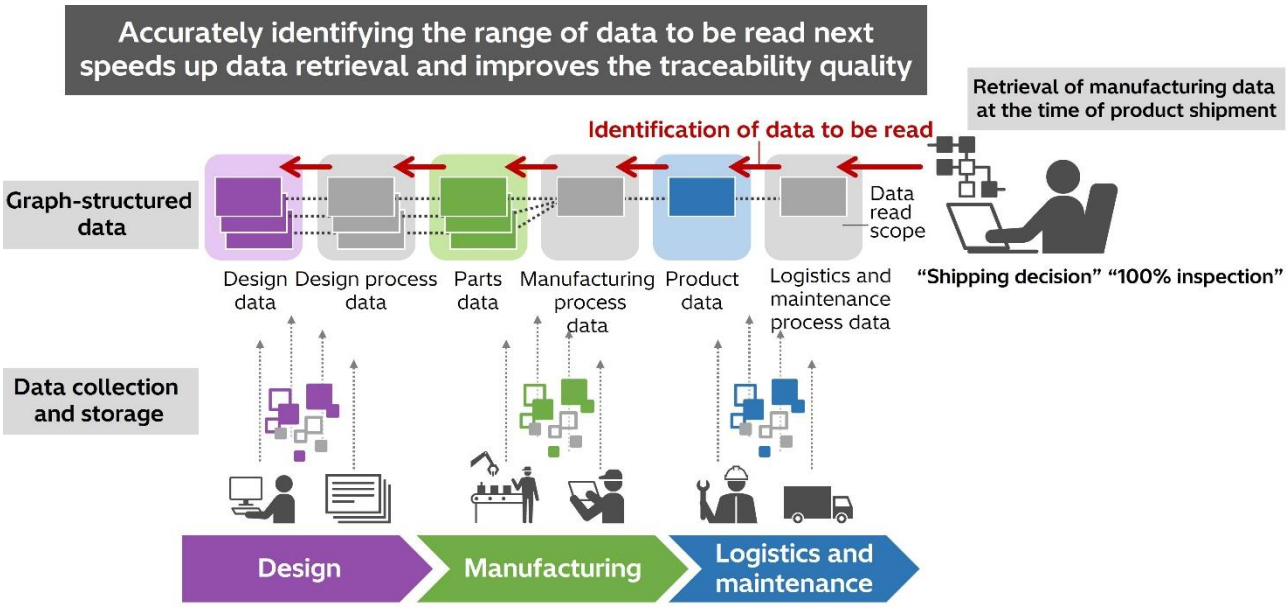

Figure 1. Typical application of the developed technology (traceability query example: exhaustive production inspection)

Tokyo, June 19, 2025 Hitachi, Ltd. (TSE:6501, "Hitachi"), and National University Corporation, The University of Tokyo ("the University of Tokyo"), toward speeding up Big Data analysis, have developed dynamic pruning[*1] technology that significantly improves the retrieval time of data having complex interconnections (hereinafter, "graph-structured data"[*2]). The conventional method of performing data analysis in a database, known as recursive query processing,[*3] involves successively traversing graph-structured data. Since this approach necessitates repeatedly reading data not relevant to the query, the retrieval speed is slowed down. The newly developed technology, based on information obtained through recursive query processing, accurately identifies in real-time the scope of the next data to be read and skips the reading of unnecessary data, significantly improving retrieval speed. In the application of verification testing to product shipment decisions in the manufacturing industry, data retrieval speed was confirmed to be up to 135 times faster[*4] than with the conventional methodology. This speed gain is expected to contribute toward improving the quality of traceability[*5] by enabling faster analysis of graph-structured data, such as in tracing processes or parts from product design to manufacturing, distribution, and maintenance (Figure 1).

Going forward, Hitachi and the University of Tokyo will aim to apply this technology beyond the manufacturing industry, such as in predicting disease risk by analyzing medical care patterns in the social insurance sector, or in detecting unauthorized access in the financial sector, thereby promoting technical innovation to solve societal problems.

The results of this development project are to be presented in part at the 2025 ACM SIGMOD/PODS International Conference on Management of Data, an international database conference, being held from June 22 to 27, 2025, in Berlin, Germany.[*6]

*1 Pruning: A method of speeding up queries by skipping data unnecessary to the query when executing a query.
*2 Graph-structured data: A structure using nodes and edges to represent data. Nodes represent the objects of data, while edges represent the relationships or connections between nodes.
*3 Recursive query processing: In a database or program, a method of solving a problem by dividing it into small parts and successively solving each part. This method is used to trace graph-structured data, particularly for obtaining data with parent-child relationships or a hierarchical structure.
*4 Compared to Hitachi's existing technology.
*5 Traceability: The state of having recorded the manufacturer, supplier, vendor, and other information about a product at each process from procurement of raw materials and parts to production, distribution, sales, and maintenance, and enabling this history to be traced.
*6 Norifumi Nishikawa, Akira Shimizu, Akira Ito, Shinji Fujiwara, Yuto Hayamizu, Masaru Kitsuregawa, and Kazuo Goda. 2025. Dynamic Pruning for Recursive Joins. In Companion of the 2025 International Conference on Management of Data (SIGMOD-Companion '25), June 22–27, 2025, Berlin, Germany. ACM, 15 pages.

Background and issues

The rapid growth in AI and Big Data utilization is heightening the importance of data processing technology, particularly in demand for more efficient data retrieval, which is crucial for enhancing AI performance and addressing societal problems. In a database storing huge amounts of data, it is necessary to represent the complex interconnections across data efficiently. The graph-structured data that makes this possible is used in a wide variety of analysis operations, from discovering transportation routes and product recommendations in e-commerce to product quality control, medical care data analysis, and the analysis of unauthorized access, for example. A problem, however, is that as the amount of data in graph-structured data increases and the data layers become deeper, retrieval speed slows down, hindering fast data analysis and decision-making. To help address this very real problem faced by industry, Hitachi and the University of Tokyo have been conducting research and development to establish new data platform technology.

Dynamic pruning technology for faster recursive query processing

One of the fruits of the joint research between Hitachi and the University of Tokyo is the development of dynamic pruning technology, which significantly enhances the retrieval speed of graph-structured data in a database. Up to now, successively tracing through graph-structured data has been done by a method called recursive query processing, which requires repeatedly reading unnecessary data. The newly developed technology, based on intermediate results[*7] obtained during recursive query processing, is able to estimate in real-time the scope of the next data to be read and accurately identify the required data in each recursion. Even when the amount of data increases or the number of data layers becomes deeper, this approach is able to reduce the amount of data to be read, thereby significantly improving retrieval speed (Figure 2).

*7 Intermediate results: Retrieval results data obtained in the preliminary stage of a recursive query

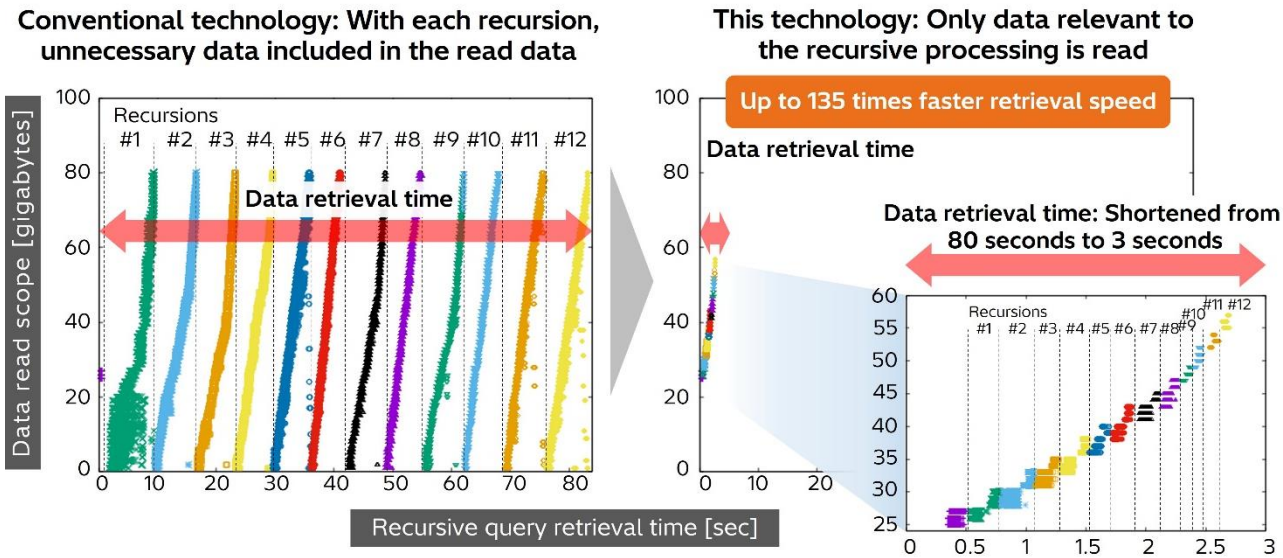

Figure 2. Comparison of data read scope and data retrieval time in recursive query processing

This technology underwent verification using a database modeled on graph-structured data analysis for product shipment decisions in the manufacturing industry. The testing confirmed a significant reduction in data read amounts in recursive query processing, and data retrieval speed was up to 135 times faster than with the conventional methodology.

The research was partially funded by grants under the “Big Data Value Co-creation Platform Engineering” Social Cooperation Program*8 established by Hitachi and the University of Tokyo in the university’s Institute of Industrial Science in 2022, and under the Integrated Healthcare System project of the Strategic Innovation Promotion Program (SIP).

*8 “Institute of Industrial Science, The University of Tokyo and Hitachi set up "Big Data Value Co-creation Platform Engineering" social cooperation program” (completed March 31, 2025).

Looking ahead

Hitachi has already incorporated this technology into its ultra-high-speed database engine, “Hitachi Advanced Database (HADB),” and has begun provisioning. Moreover, the HADB applying this technology can be used together with IoT-Compas, which replicates in a digital twin the linkage between data and operations in production processes, as part of the “Hitachi Intelligent Platform”*9 services supporting IoT and data utilization.

Hitachi and the University of Tokyo will seek to further advance dynamic pruning technology and its AI integration, aiming to apply this technology to fields beyond manufacturing, such as social insurance and finance, thereby promoting technical innovation to address societal problems.

*9 The Hitachi Intelligent Platform has been registered with the Amazon Web Services (AWS) Industrial Data Fabric Solutions.

About the Hitachi Advanced Database (HADB)

HADB is a made-in-Japan relational database that supports standard SQL, enabling fast and timely analysis of large amounts of data. Based on the principle of “out-of-order execution”*10 for maximizing the use of server and storage capability by executing multiple tasks in parallel, starting with the data available for processing, this product is able to perform ultra-fast search processing.
Website: https://www.hitachi.com/products/it/software/prod/hadb/index.html

*10 Out-of-order execution is a principle devised by University Professor Kitsuregawa (President, Research Organization of Information and Systems) of the University of Tokyo and Kazuo Goda.

About Hitachi, Ltd.
Through its Social Innovation Business (SIB) that brings together IT, OT (Operational Technology), and products, Hitachi contributes to a harmonized society where the environment, wellbeing, and economic growth are in balance. Hitachi operates globally in four sectors – Digital Systems & Services, Energy, Mobility, and Connective Industries – and the Strategic SIB Business Unit for new growth businesses. With Lumada at its core, Hitachi generates value from integrating data, technology and domain knowledge to solve challenges of customers and society. Revenues for FY2024 (ended March 31, 2025) totaled 9,783.3 billion yen, with 618 consolidated subsidiaries and approximately 280,000 employees worldwide. Visit us at www.hitachi.com

About the Institute of Industrial Science, The University of Tokyo
The Institute of Industrial Science, The University of Tokyo (UTokyo-IIS) is one of the largest university-attached research institutes in Japan. UTokyo-IIS is comprised of over 120 research laboratories—each headed by a faculty member—and has over 1,200 members (approximately 400 staff and 800 students) actively engaged in education and research. Its activities cover almost all areas of engineering. Since its foundation in 1949, UTokyo-IIS has worked to bridge the huge gaps that exist between academic disciplines and real-world applications.

Inquiries about technologies
Research & Development Group
Hitachi, Ltd.

Contact form:
https://www8.hitachi.co.jp/inquiry/hitachi-ltd/hqrd/news/en/form.jsp

Institute of Industrial Science,
The University of Tokyo.
Kazuo Goda

+81-3-5452-6594
kgoda@tkl.iis.u-tokyo.ac.jp

Inquiries about products and services
Managed & Platform Services Business Division, AI & Software Services Business Unit
Hitachi, Ltd.

Contact form:
https://www8.hitachi.co.jp/inquiry/hitachi-ltd/it/soft/en/form.jsp?UM_Key=global

---

Information contained in this news release is current as of the date of the press announcement, but may be subject to change without prior notice.
