Recent Patents for Leading Computer Architecture Companies

Joshua J. Yi, *The Law Office of Joshua J. Yi, PLLC, Austin, TX, 78750, USA*

One way to reveal what technologies or areas companies are investing their research and development resources is to analyze their recently issued patents. Toward that end, this article examines patents issued to the companies that have the most issued patents in the "Electric digital data processing" subcategory in the first quarter of 2021. This subcategory appears to be the subcategory in the International Patent Classification system that is closest to computer architecture.

Table 1 shows the number of patents that were issued to the top seven companies in descending order of this subcategory (labeled as the "Category 2" column as Table 1) during the first quarter of 2021.

The second column shows the total number of patents that were issued to each company in the first quarter of 2021. For the purposes of this article, a single "company" includes all related companies, e.g., Samsung Display, Samsung Electronics, etc., are grouped together under as "Samsung."

The third and fourth columns show how many patents were classified in the "Computing; calculating or counting" ("Category 1") and "Electric digital data processing" ("Category 2"), respectively, categories. While the "Electric digital data processing" subcategory appears to be the closest International Patent Classification subcategory to computer architecture, including its parent category ("Computing; calculating or counting") in Table 1 also includes sister subcategories that may be relevant to computer architecture. For example, the "Computing; calculating or counting" category (G06) includes the "Computer Systems Based on Specific Computational Models" subcategory (G06N), which includes biological models, mathematical models, quantum-mechanical phenomena, and machine learning.

The fourth and fifth columns show the percentage of issued patents (in the second column) that were classified under the "Computing; calculating or counting" ("Category 1") and "Electric digital data processing" ("Category 2"), respectively, categories.

The results in Table 1 show interesting results for each company. First, IBM has the most Category 2 patents, as well as the most issued patents overall. More specifically, IBM has almost three ***times*** the number of issued Category 2 patents (1429) than the second-place company (Samsung, 540 Category 2 patents) and more than Dell/EMC, Intel, Amazon, and Apple combined (1334).

Second, although Samsung had slightly more issued patents in the first quarter of 2021 than IBM, it had only approximately one-third the number of Category 1 and Category 2 patents. In addition, though Samsung had almost four times more issued patents than Microsoft, Samsung only had 22.3% and 14.4% more Category 1 and Category 2, respectively, patents than Microsoft. This particular result may not be particularly surprising given that Samsung produces a wide range of products from cellular phones to semiconductor products to durable goods such as refrigerators.

Third, although Dell/EMC had the fewest number of issued patents of the companies in Table 1 with 493, because 81.5% of its issued patents are Category 2 patents, Dell/EMC has the fourth-highest number of Category 2 patents with 402. Similarly, 68.4% of Microsoft's issued patents are Category 2 patents, which places Microsoft in third-place despite having the fifth highest number of issued patents of the companies in Table 1.

Fourth, although Intel has the third highest number of issued patents of the companies in Table 1, because it had the third-lowest percentage of Category 1 and Category 2 patents, Intel only had 452 and 367 Category 1 and Category 2, respectively, patents issued in the first quarter of 2021. Although beyond the scope of this particular article, a potentially interesting avenue of study would be to analyze which subcategories each company's patents are classified under, in addition to those within Category 2. The percentage of Category 1 and Category 2 patents for Amazon are similar to those for Intel.

0272-1732 © 2021 IEEE
Digital Object Identifier 10.1109/MM.2021.3087210
Date of current version 2 July 2021.

TABLE 1. Patents issued in the first quarter of 2021, listed in descending order of Category 2.

Company	Total	Category 1[a]	Category 2[b]	% Category 1	% Category 2
IBM	2,437	1,731	1,429	71.0%	58.6%
Samsung	2,465	679	540	27.5%	21.9%
Microsoft	690	555	472	80.4%	68.4%
Dell / EMC[c]	493	420	402	85.2%	81.5%
Intel	793	452	367	57.0%	46.3%
Amazon	603	383	286	63.5%	47.4%
Apple	751	320	279	42.6%	37.2%

[1]Category 1 is the "Computing; calculating or counting" category in the International Patent Classification system. The code for this category is G06.
[2]Category 2 is the "Electric digital data processing" subcategory of the "Computing; calculating or counting" category in the International Patent Classification system. The code for this category is G06F.
[3]Because Dell and EMC are currently a single company, this article presents the combined results for the two companies.

Finally, although Apple has the fourth highest number of issued patents in the first quarter of 2021, it has the lowest numbers of Category 1 and Category 2 patents. Concomitantly, it has the second-lowest percentages of Category 1 (42.6%) and Category 2 (37.2%) patents, only ahead of Samsung.

SPECIFIC OBJECTIVE CHARACTERISTICS MAY INDICATE WHETHER A PATENT IS MORE INTERESTING THAN AVERAGE.

INTERESTING PATENTS

Specific objective characteristics may indicate whether a patent is more interesting than average. As a first example, a large number of claims may indicate that the patent is broad. As a second example, patents that are part of a large patent family could be more interesting for the same reason that a large number of claims could be more interesting, namely, the family collectively could have a broad scope. Third, a patent with a large number of cited prior art references may indicate that the patent is more likely to be valid. Fourth, the amount of time needed to prosecute a patent could indicate how interesting the patent is. On one hand, patents that are litigated in court spend significantly longer time in prosecution than patents that are not litigated.[a] On the other hand, patents that have very short prosecution times could indicate the inventor is trying to get the patent issued as quickly as possible, which could indicate that the inventor thinks the patent is interesting. Patents with shorter prosecution times also provide earlier public notice as to what inventions a company may be working on. Fifth, reissue[b] patents could be more interesting because it could indicate that the patentee apparently believed that the patent was interesting enough to be reissued despite having to surrender any damages prior the reissue date, as well as taking the time and effort to get the patent reissued.

PATENTS WITH SHORTER PROSECUTION TIMES ALSO PROVIDE EARLIER PUBLIC NOTICE AS TO WHAT INVENTIONS A COMPANY MAY BE WORKING ON.

The rest of this article presents a recently issued patents from each of the above companies that may be interesting based on having a few of the above characteristics and/or are directed toward interesting subject matter, e.g., artificial intelligence or machine learning.

Assignee: IBM
Number: 10,902,352
Filed: January 6, 2020
Issued: January 26, 2021

[a]John A. Allison, Mark A. Lemley, Kimberly A. Moore, & R. Derek Trunkey, *Valuable Patents*, 92 Geo. L.J. 435 (2004).

[b]A reissue patent fixes an error in the issued patent or changes the scope of the claims. 35 U.S.C. § 251(a) ("Whenever any patent is, through error, deemed wholly or partly inoperative or invalid, by reason of a defective specification or drawing, or by reason of the patentee claiming more or less than he had a right to claim in the patent, the Director shall. . .reissue the patent.").

Title: Labeling of data for machine learning

Inventors: P. Ghosh (Bangalore, India), S. Godbole (Bangalore, India), S. Joshi (Gurgaon, India), S. Merugu (Bangalore, India), A. Verma (New Delhi, India)

Abstract: A computer generates labels for machine learning algorithms by retrieving, from a data storage circuit, multiple label sets that contain labels that each classifies data points in a corpus of data. A graph is generated that includes a plurality of edges, each edge between two respective labels from different label sets of the multiple label sets. Weights are determined for the plurality of edges based on a consistency between data points classified by two labels connected by the edges. An algorithm is applied that groups labels from the multiple label sets based on the weights for the plurality of edges. Data points are identified from the corpus of data that represent conflicts within the grouped labels. An electronic message is transmitted in order to present the identified data points to entities for further classification. A new label set is generated using the further classification received from the entities.

Assignee: Samsung

Number: 10,950,009

Filed: January 15, 2020

Issued: March 16, 2021

Title: AI encoding apparatus and operation method of the same, and AI decoding apparatus and operation method of the same

Inventors: Q. Dinh (Suwon-si, South Korea), M. Choi (Suwon-si, South Korea), K. Choi (Suwon-si, South Korea)

Abstract: Provided is an artificial intelligence (AI) decoding apparatus including a memory storing one or more instructions; and a processor configured to execute the one or more instructions to, when an image is input to a second DNN including a plurality of layers, obtain first result values based on an operation between the image and a first filter kernel; obtain second result values based on an operation between the image and a second filter kernel, from a first layer including the first and second filter kernels from among the plurality of layers; perform normalization by transforming the first result values into first values by using a first scale factor; perform normalization by transforming the second result values into second values by using a second scale factor; and transform the first values and the second values into integer values included in a preset range.

Assignee: Microsoft

Number: 10,956,162

Filed: June 28, 2019

Issued: March 23, 2021

Title: Operand-based reach explicit dataflow processors, and related methods and computer-readable media

Inventors: R. Clancy (Cary, NC, USA), M. Brown (Raleigh, NC, USA), Y. Tekmen (Raleigh, NC, USA), B. Stempel (Raleigh, NC, USA), M. Mcilvaine (Raleigh, NC, USA), T. Speier (Wake Forest, NC, USA), R. Smith (Raleigh, NC, USA), G. Gupta (Bellvue, WA, USA), D. Harper, III (Seattle, WA, USA)

Abstract: Operand-based reach explicit dataflow processors, and related methods and computer-readable media are disclosed. The operand-based reach explicit dataflow processors support execution of a producer instruction that explicitly names a target consumer operand of a consumer instruction in a consumer operand encoding namespace of the producer instruction. The produced value from execution of the producer instruction is provided or otherwise made available as an input to the named target consumer operand of the consumer instruction as a result of processing the producer instruction. The target consumer operand is encoded in the producer instruction as an operand target distance relative to the producer instruction. Instructions in an instruction stream between the producer instruction and the targeted consumer instruction that have no operands do not consume an operand reach namespace in the producer instructions. This provides for a deeper explicit consumer naming reach for a given bit size of the operand reach namespace.

Assignee: Dell/EMC

Number: 10,936,376

Filed: December 12, 2019

Issued: March 2, 2021

Title: Methods, systems, and computer readable mediums for workload clustering

Inventors: C. Gong (Sherborn, MA, USA), V. Korolyov (Cork, Ireland)

Abstract: Methods, systems, and computer readable mediums for optimizing a system configuration are disclosed. In some examples, a method includes determining whether a system configuration for executing a workload using a distributed computer system is optimizable and in response to determining that the system configuration is optimizable, modifying the system configuration such that at least one storage resource for storing workload data is located at a server node that is executing the workload in the distributed computer system.

Assignee: Intel

Number: 10,917,415

Filed: January 10, 2018

Issued: February 9, 2021

Title: Machine learning-based determination of program code characteristics

Inventors: L. Chen (Hillsboro, OR, USA)

Abstract: A technique includes processing a plurality of sets of program code to extract call graphs; determining similarities between the call graphs; applying unsupervised machine learning to an input formed from the determined similarities to determine latent features of the input; clustering the determined latent features; and determining a characteristic of a given program code set of the plurality of program code sets based on a result of the clustering.

Assignee: Amazon

Number: 10,896,459

Filed: April 7, 2020

Issued: January 19, 2021

Title: Recommendation system using improved neural network

Inventors: R. Joseph (Seattle, WA, USA), O. Rybakov (Seattle, WA, USA)

Abstract: Some aspects of the present disclosure relate to generating and training a neural network by separating historical item interaction data into both inputs and outputs. This may be done, for example, based on date. For example, a neural network machine learning technique may be used to generate a prediction model using a set of inputs that includes both a number of items purchased by a number of users before a certain date as well as some or all attributes of those items, and a set of outputs that includes the items purchased after that date. The items purchased before that date and the associated attributes can be subjected to a time-decay function.

Assignee: Apple

Number: 10,884,811

Filed: January 12, 2018

Issued: January 5, 2021

Title: Scheduler for AMP architecture with closed loop performance controller using static and dynamic thread grouping

Inventors: J. Andrus (Sunnyvale, CA, USA), J. Dorsey (San Francisco, CA, USA), J. Magee (Orlando, FL, USA), D. Chimene (San Francisco, CA, USA), C. de la Cropte de Chanterac (San Francisco, CA, USA), B. Hinch (Mountain View, CA, USA), A. Venkataraman (Sunnyvale, CA, USA), A. Dorofeev (San Jose, CA, USA), N. Gamble (San Francisco, CA, USA), R. Blaine (San Carlos, CA, USA), C. Pistol (Cupertino, CA, USA)

Abstract: Systems and methods are disclosed for scheduling threads on a processor that has at least two different core types, such as an asymmetric multiprocessing system. Each core type can run at a plurality of selectable voltage and frequency scaling (DVFS) states. Threads from a plurality of processes can be grouped into thread groups. Execution metrics are accumulated for threads of a thread group and fed into a plurality of tunable controllers for the thread group. A closed-loop performance control (CLPC) system determines a control effort for the thread group and maps the control effort to a recommended core type and DVFS state. A closed-loop thermal and power management system can limit the control effort determined by the CLPC for a thread group, and limit the power, core type, and DVFS states for the system. Deferred interrupts can be used to increase performance.

JOSHUA J. YI is a solo practitioner who serves as the court-appointed technical advisor for the Honorable Alan D Albright, United States District Judge for the Western District of Texas, Waco Division. His research interests include microarchitecture and performance methodology. Yi received a Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, and a JD from the University of Texas at Austin. Contact him at josh@joshuayipatentlaw.com.