IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF TEXAS
DALLAS DIVISION

VALTRUS INNOVATIONS LTD.,
KEY PATENT INNOVATIONS LIMITED,

Plaintiffs,

v.

GOOGLE LLC,

Defendant.

Case No. 3:24-cv-1795

JURY TRIAL DEMANDED

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Valtrus Innovations Limited ("Valtrus") and Plaintiff Key Patent Innovations Limited ("KPI"), bring this complaint because another defendant has challenged Valtrus's standing to assert patents in its own portfolio. There is a currently pending action against Defendant Google LLC ("Google"), Case No. 3:22-cv-00066-n ("Google I action"). Because that case is stayed, however, out of an abundance of caution and to preserve all rights, Plaintiffs are filing this lawsuit. The Plaintiffs will be filing a motion to amend and/or consolidate in the Google I Action, and will be filing a motion to stay this action pending the Court's decision on the motion to amend and/or consolidate in Google I.

Plaintiffs by and through their undersigned counsel, plead the following against Google and allege as follows:

THE PARTIES

1. Plaintiff Valtrus is the successor in interest to a substantial patent portfolio created by Hewlett Packard Enterprise and its predecessor companies, including Compaq, Verity, and Hewlett-Packard Development Company (collectively, "HPE"). Valtrus is an Irish entity duly organized and existing under the laws of the Republic of Ireland. The address of the registered office of Valtrus is: The Glasshouses GH2, 92 Georges Street Lower, Dun Laoghaire, Dublin A96 VR66, Ireland. HPE's worldwide corporate headquarters is located in Houston, Texas. One of HPE's primary US facilities is located in Plano, Texas.

2. Valtrus is the assignee and owns all right and title to U.S. Patent Nos. 6,728,704 ("the '704 Patent"), 6,816,809 ("the '809 Patent"), and 7,346,604 ("the '604 Patent") (collectively, the "Asserted Patents").

3. Plaintiff KPI is the beneficiary of a trust pursuant to which Valtrus owns, holds, and asserts the Asserted Patents. KPI is an Irish entity duly organized and existing under the laws

of the Republic of Ireland. The address of the registered office of KPI is: The Glasshouses GH2, 92 Georges Street Lower, Dun Laoghaire, Dublin A96 VR66, Ireland.

4. The Asserted Patents were developed by inventors working for HPE. HPE and its predecessors have been developing innovative search, computer processing, and server technology for decades.

5. On information and belief, Defendant Google is a limited liability company duly organized and existing under the laws of the State of Delaware, having a regular and established place of business in the Northern District of Texas, including at 3800 Railport Parkway, Midlothian, Texas 76065.

6. Google has become the dominant search engine in the United States and beyond, with a reported market share of around 90%. Google's strategy has included taking technology from other companies without regard to intellectual property rights and combining that technology to create and maintain market dominance. A core part of this strategy involved taking the innovations in the Asserted Patents. One need look no further than Google's employment records, which have listed as some of Google's most senior technologists a large number of inventors on the Asserted Patents, including a Senior Vice President at Google responsible for Search, the first Director of Google R&D Bangalore, and the Vice President at Google responsible for Google Apps and Cloud.

JURISDICTION AND VENUE

7. This is an action arising under the patent laws of the United States, 35 U.S.C. § 1 *et seq*. Accordingly, this Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

8. This Court has personal jurisdiction over Google because Google creates products and services that are and have been used, offered for sale, sold, and purchased in the Northern

District of Texas, and Google has committed, and continues to commit, acts of infringement in the Northern District of Texas, has conducted business in the Northern District of Texas, and/or has engaged in continuous and systematic activities in the Northern District of Texas.

9. Under 28 U.S.C. §§ 1391(b)-(d) and 1400(b), venue is proper in this judicial district because Google maintains a regular and established place of business in this district and has committed and regularly commits acts of infringement within this judicial district giving rise to this action. For example, Google operates a 260,000 square foot data center in Midlothian, Texas. Plaintiffs are informed and believe that this data center includes the infringing systems and practices the infringing methods described herein. This data center is one of only fourteen in all of North America. Google also maintains an office in the Dallas-Fort Worth area in Addison, Texas, and runs a content distribution network (CDN) node in the Dallas area. Google also provides, sells, and offers for sale infringing products and services to users in the Northern District of Texas.

10. Google also has significant operations in nearby cities, including Austin and Houston. For example, on information and belief, Google owns approximately 550,000 square feet of office space across three locations in downtown Austin, Texas. Google is also continuing to grow its presence in Texas. For example, on information and belief, Google is preparing to open an additional 750,000 square feet of offices in Austin at Block 185, a new office tower located at 601 West Second Street. Plaintiffs are informed and believe that these offices include employees responsible for the subject matter of this patent suit. For example, on information and belief, Google employees in Austin work on Google Cloud, finance, and engineering. On information and belief, Google also has an office in Houston, and employs more than 1700 people in Texas.

FIRST CLAIM

(Infringement of U.S. Patent No. 6,728,704)

11. Plainiffs re-allege and incorporate herein by reference Paragraphs 1-10 of their Complaint.

12. The '704 Patent, entitled "Method and apparatus for merging result lists from multiple search engines," was duly and lawfully issued on April 27, 2004. A true and correct copy of the '704 Patent is attached hereto as Exhibit 1.

13. The '704 Patent names Jianchang Mao, Rajat Mukherjee, Prabhakar Raghavan, and Panayiotis Tsaparas as co-inventors. The Patent represents the work of these data scientists at Verity, a leading competitor to Google at the time of the invention.

14. The '704 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '704 Patent, including the right to seek damages for any infringement thereof.

15. The '704 Patent is directed to addressing problems specific to and rooted in a major industrial process—the world wide web. There are a massive number of pages and hyperlinks on the web, and it is impossible for humans to manually search even a small fraction of those. At the time of the '704 Patent's invention, major industrial processes involving massive server arrays were used to perform machine searching. This problem, unique to the web, is made clear in the Patent's specification. *See, e.g.*, Ex. 1, 1:21-31 (describing the "explosion in the amount of information available" with the advent of the web, and the "paramount importance" of search engines in isolating useful information). The problem has only grown exponentially. Today, there are an estimated 130 trillion web pages across the Internet.[1]

[1] https://searchengineland.com/googles-search-indexes-hits-130-trillion-pages-documents-263378

16. The '704 Patent explains that, at the time of its invention, "the merging of multiple result lists into a single list [was] usually accomplished by examining and ranking every single entry of every list. . . . Thus, for large lists or large numbers of lists, the computation time required by the merging process can nullify any advantage gained by operating multiple search engines at the same time." *Id.* at 2:48-51, 53-56.

17. The '704 Patent states that the "invention allows for a reduction in computational overhead when merging and re-ranking multiple result lists. Ranking of results is accomplished by evaluating a subset of entries instead of every single one, thus reducing the number of calculations required" and allowing result lists to be merged quickly and efficiently without sacrificing the accuracy of the ranking process. *Id.* at 3:20-24.

18. The method described by the '704 Patent improves the function of a computer utilizing said method by reducing the considerable computational overhead associated with merging many results from numerous search engines. A reduction in computational overhead enables the computer to operate multiple search engines simultaneously without sacrificing response time, allowing the gathering, ranking, and presentation of millions of search results from across the web in a very short amount of time.

19. The '704 Patent provides detailed, specific steps for merging multiple result lists in this manner. For example, in one embodiment, each result list is assigned a representative scoring value based on a selected subset of entries from the list. *Id.* at 5:44-55; 7:12-14. In another embodiment, each list is instead "assigned a probability value equal to its average scoring value's percentage of the total of all average scoring values," and entries are selected from each list based on its probability value. *Id.* at 7:35-39.

20. The invention of the '704 Patent includes the following elements. First, it involves the unique environment of multiple search engines returning multiple results. *Id.* at 8:4-8 (claim 1). Second, rather than ranking every entry of every result list, the claimed method ranks only a subset of the result list from each search engine. *Id.* at 8:9-10 (claim 1); see also *id.* at 5:44-55 ("Lists are merged according to these subsets, rather than an evaluation of every single entry of every single list."). The method then uses the value of each entity in the subset to assign a rank to the subset. *Id.* at 8:10-13 (claim 1). Once the subset for each search engine is ranked based on the relevance of its results, a data structure is used to merge the results. *Id.* at 8:14-16 (claim 1). There are two specific data structures disclosed. In the highest average scoring value embodiment, the first choice will be the first unmerged entry in the result list with the highest average score. *Id.* at 6:6-28. In the probabilistic embodiment, results are merged in a pseudorandom manner in which the probability of selecting an entry from each list is determined by the average scoring value of the list. *Id.* at 7:15-54. Finally, the representative value of each search engine result based on the representative value of the subset in the result is not static, but changes in a predefined manner. *Id.* at 8:17-18 (claim 1). In the highest average scoring value embodiment, the unmerged entry from the result list with the highest average score is selected, and the representative value is then decremented. At some point, the result list will no longer have the highest average scoring value, and an entry from another result list with the new highest average scoring value will be selected instead. *Id.* at 7:15-34. In the probabilistic embodiment, once all entries in a result list have been selected using the pseudo-random model, that result list drops, and a new probabilistic model is calculated without it. *Id.* at 7:35-54.

21. Claim 1 of the '704 Patent recites a "method for merging result lists from multiple search engines," and describes concrete steps for achieving this result, as described in detail below.

Moreover, claim 1 includes the inventive concept of "assigning to each subset a representative value according to the scoring values assigned its entries," rather than investing the substantial time and computational resources necessary to evaluate every single entry of every single list. *Id.* at 8:12-13. The '704 Patent thus claims a technologically-based method which allows an Internet search engine to do something it could not do before: effectively merge results from several individual search engines based on representative values, rather than a computationally-intensive use of all entries in all result lists.

22. Moreover, the representative value of each result list is not static—instead, it changes in a "predetermined manner" based on the merging process used. *Id.* at 8:17-18. For example, in the average scoring value embodiment described above, an entry from the result list with the highest representative value is selected, and the representative value for that list is then decremented. *See id.* at 6:6-28. At a certain point, said result list will no longer have the highest representative value, and an entry from another list with the now-highest representative value is selected instead. *Id.* Thus, the ranking process changes as merging occurs.

23. Independent expert analysis and testing of the '704 Patent confirms that the claimed invention is not routine, well-understood, or conventional, and that it measurably improves the performance of a computer system as compared to previous approaches used for merging result lists from multiple search engines. For example, prior art meta-search engines "contemplated producing a merged, ranked list based solely on the individual scores of each entry from each search engine," whereas the '704 Patent recites using a "dynamically adjusted" representative value for each search engine result list to minimize computational overhead during the merging process. *See* Ex. 2 (Declaration of Dr. Amy N. Langville) at 17-18.

24. In independent expert testing, the method described by claim 1 of the '704 Patent returned a merged list of search results from seven separate search engines twice as fast as a method which simulated the type of merging process used in prior art meta-search engines. *See* Ex. 2 at 20 (showing that the method of the '704 Patent took an average of 20.29ms to return a merged list of ten results from seven search engines, while a prior art alternative approach took an average of 40.75ms).

25. The dependent claims of the '704 Patent identify additional inventive concepts as well, including unique data structures used to select the subset of entries from each search engine (claim 2-4) and embodiments for merging result lists and varying the representative value of claim 1 based on highest average values (claim 5) and probability values (claim 6). Ex. 1, 8:19-42. These embodiments are further described in the specification. *See id.* at 7:14-54.

26. During the prosecution of the '704 Patent, the Patent Office expressly acknowledged the novelty of the patented design. After the applicant pointed out that the prior art "neither teaches nor suggests 'selecting a subset of entries from each result list to form a set of selected entries,'" *see* Ex. 3 ('704 Patent prosecution history, September 2, 2003 Amendment) at 9, the patent examiner confirmed that the method of merging results specified in the independent claims was not in the prior art. *See* Ex. 4 ('704 Patent prosecution history, February 8, 2004 Notice of Allowability). In response to a rejection based on a prior art reference, the applicant had added language to describe that the rankings do not remain static—"in a predetermined manner" and "wherein the representative value varies in accordance with predetermined manner" to claim 1—and there was no contention by the examiner that the prior art contained these limitations. *See* Ex. 5 ('704 Patent prosecution history, January 29, 2004 Amendment) at 2.

27. Only a few years after the '704 Patent application issued, Google introduced its "Universal Search" feature, which corresponds to the features of the '704 Patent. Google, in its own words, "blend[s] results from more than just the web in order to provide the most relevant and useful results possible. In addition to web pages, for instance, the search results may include video, news, images, maps, and books."[2] As one contemporary article put it, "Google's hitting several of its vertical search services all at once, then bringing back those results to blend in with 'regular' results."[3] When this feature launched, commentators concluded that it transformed Google Search from "a painful process" to one in which "the guesswork was reduced" and results of all kinds could be shown together in a merged list.[4]

28. Plaintiffs are informed and believe, and thereon allege, that Google has infringed and unless enjoined will continue to infringe one or more claims of the '704 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed method of merging result lists in an infringing manner. Google practices every step of claim 1 of the '704 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

29. For example, the '704 accused product, Google Search, embodies every limitation of claim 1 of the '704 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

[2] https://developers.google.com/search/blog/2007/05/taking-advantage-of-universal-search
[3] https://searchengineland.com/google-20-google-universal-search-11232
[4] https://www.searchenginejournal.com/search-engines/universal-search/#:~:text=%E2%80%9CUniversal%20search%20is%20the%20ability%20to%20search%20all%20content%20across,integrated%20set%20of%20search%20results.%E2%80%9D

30. Google Search practices a method of merging result lists from multiple search engines, comprising the elements set forth below.

31. For example, a query transmitted to Google Search returns results from multiple search engines displayed to the user as a single merged list of results. *See, e.g.*, Google Search Results for "in-n-out burger" Query 1 (red arrows added):

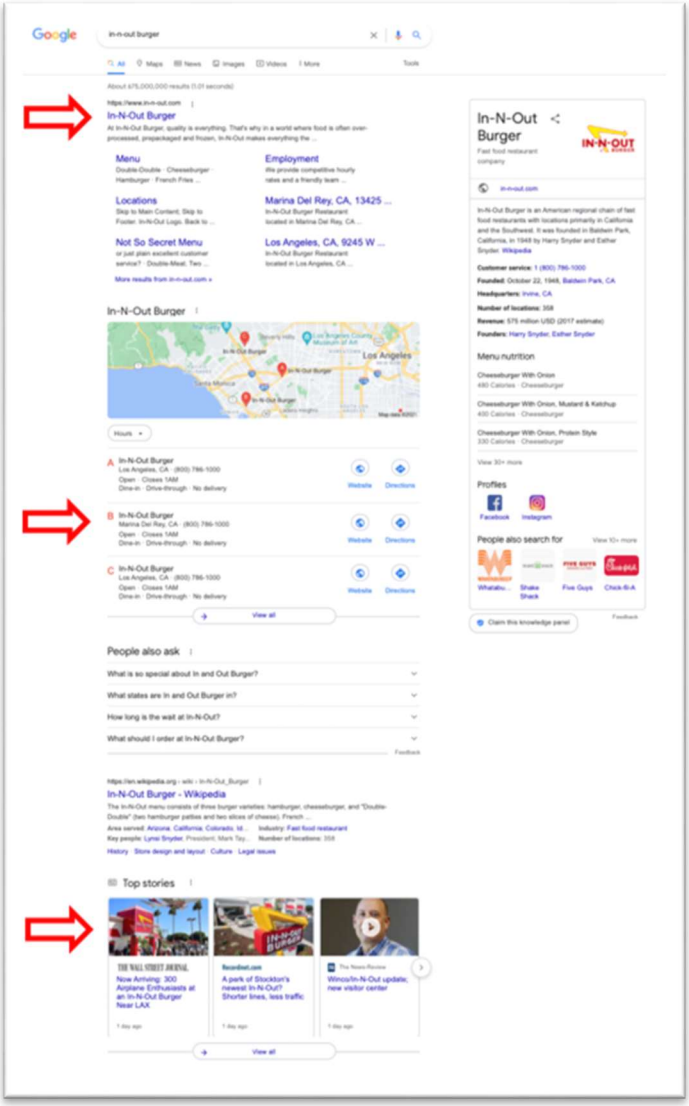

32. Google Search practices a method comprising transmitting a query to a set of search engines.

33. For example, a query transmitted to Google Search returns results from multiple search engines. This process begins by transmitting that query to said search engines. *See, e.g., id.* (red arrows added):

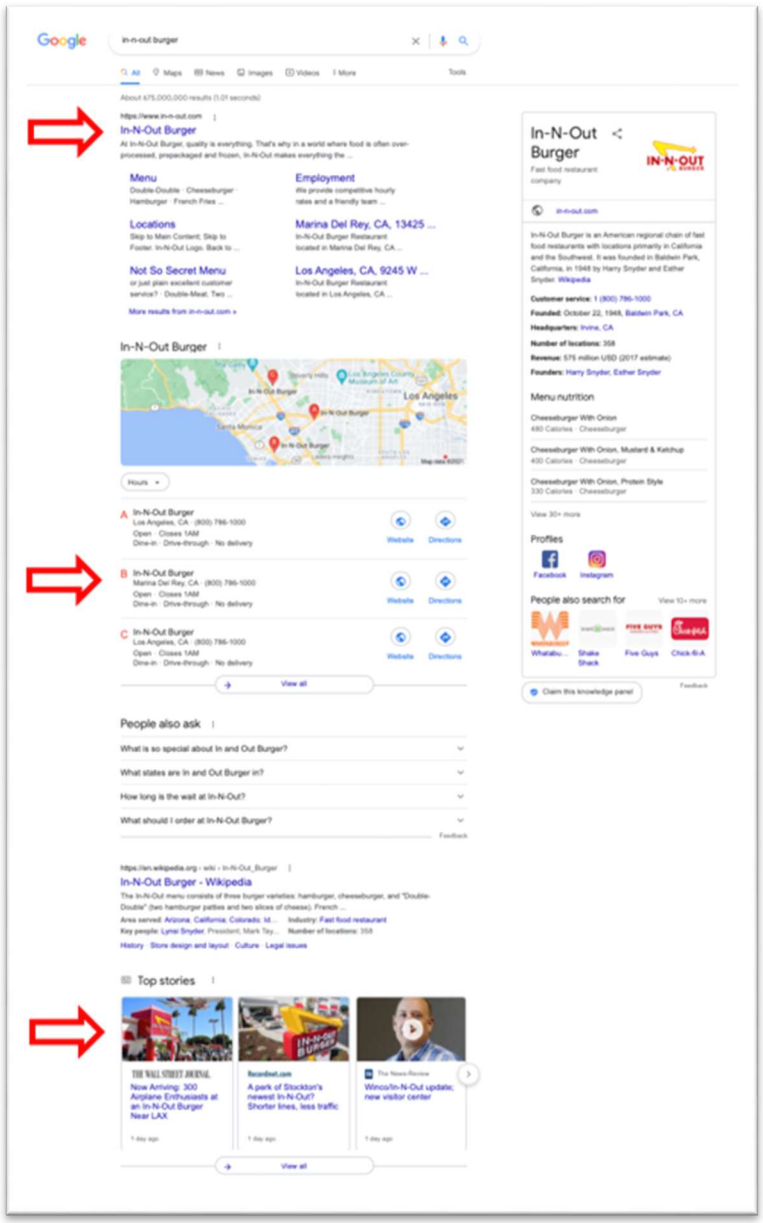

34. Google Search practices a method comprising receiving in response to said query a result list from each search engine of said set of search engines, each result list including one or more entries.

35. For example, a query transmitted to Google Search returns one or more results from each of several search engines. Said results are grouped into result lists associated with each of said search engines. In order to display said result lists, Google Search first receives a result list from each search engine. *See, e.g., id.* (red arrows added):

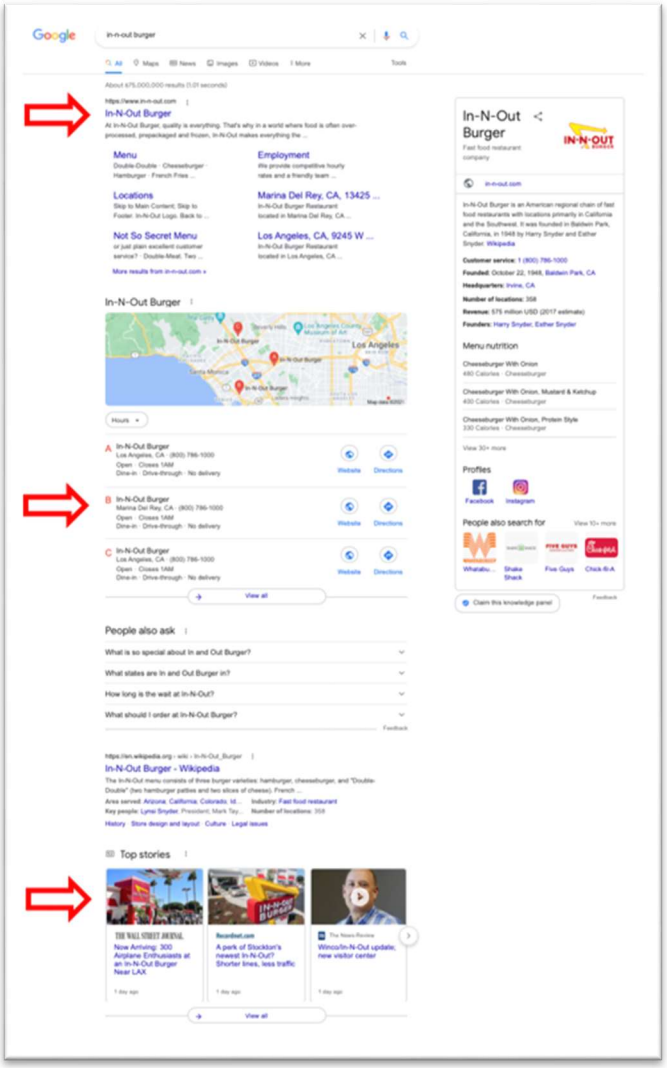

36. Google Search practices a method comprising selecting a subset of entries from each result list to form a set of selected entries.

37. For example, a query transmitted to Google Search may return many millions of results across a set of multiple search engines. Only a subset of these results, or entries, is selected for display to the user near the top of the Google Search result list. *See, e.g.*, Google Search Results for "in-n-out burger" Query 2 (callout added):

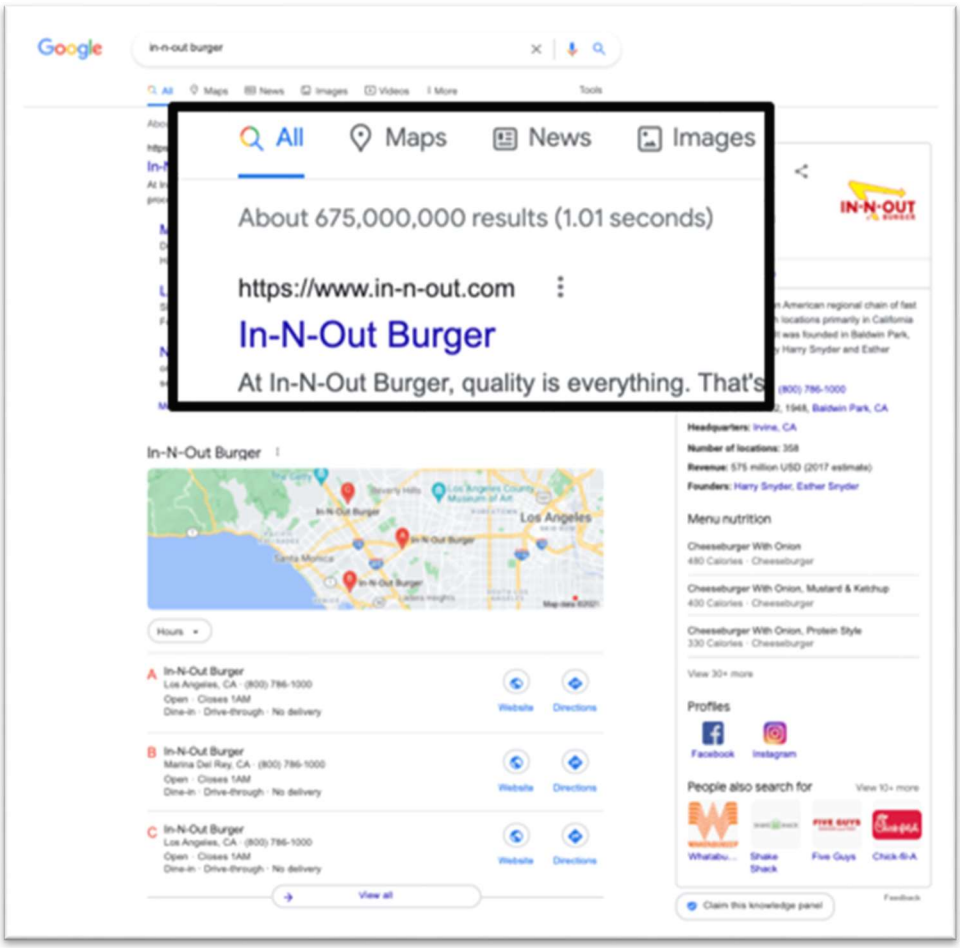

38. Google Search practices a method comprising assigning to each selected entry of said set of selected entries a scoring value according to a scoring function.

39. Entries are displayed to a user of Google Search in order of relevance. On information and belief, these entries are assigned a scoring value according to a scoring function in order to be sorted and displayed in this way. For example, a list of results from the Maps subset of entries is sorted according to a scoring function which considers metrics such as the distance of each result from the user at the time of the query. *See, e.g.,* Google Maps Result List (red arrow added):

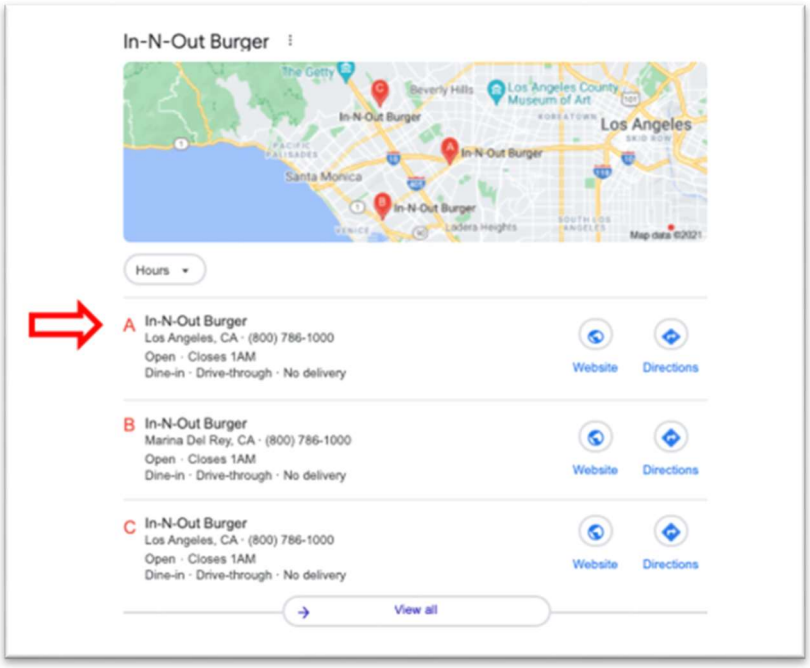

40. Google Search practices a method comprising assigning to each subset a representative value according to the scoring values assigned to its entries.

41. For example, Google Search presents results in subsets that appear in different orders for different queries. On information and belief, Google Search orders each subset based on a representative value according to the scoring values assigned to its entries. *See, e.g.*, Google Search Results for "in-n-out burger" Query 1 (red arrows added):

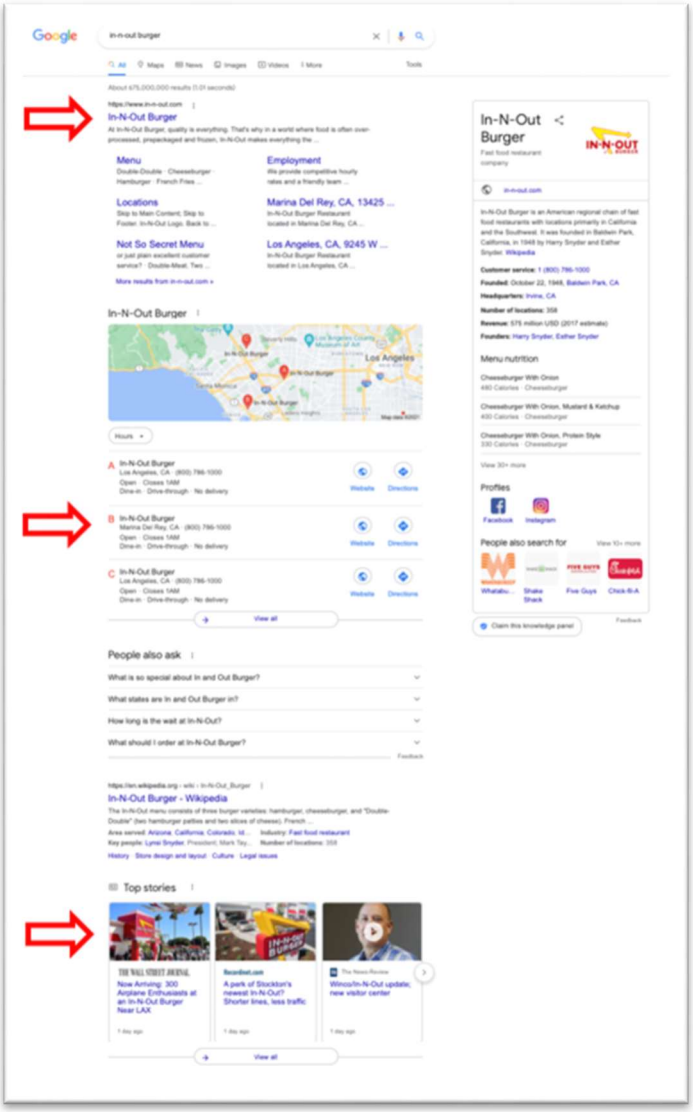

42. Google Search practices a method comprising producing a merged list of entries in a predetermined manner based on the representative value assigned to each result list.

43. For example, when a user transmits a query involving a term associated with a physical location, Google Search produces a merged list of entries where in Maps results are generally displayed higher in the merged list than News results because, on information and belief, the Maps result list has been assigned a higher representative value. *See, e.g., id.* (red arrows added):

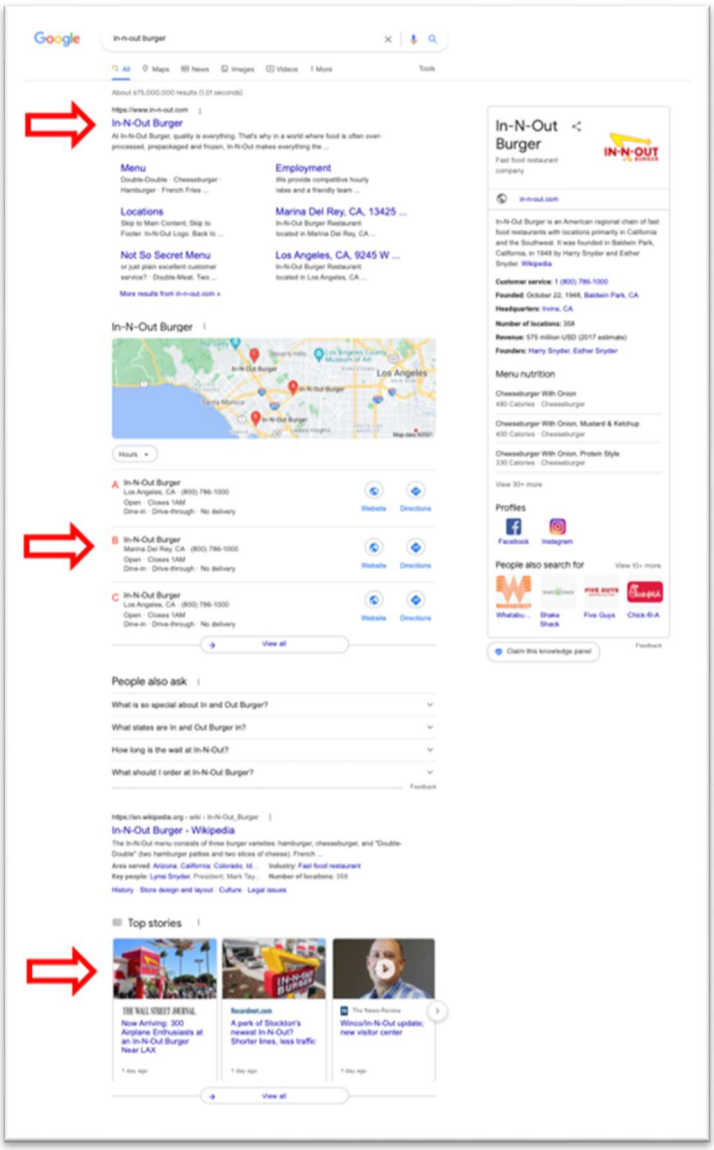

44. Google Search practices a method comprising producing a merged list of entries in a predetermined manner based on the representative value assigned to each result list, wherein the representative value varies in accordance with predetermined manner.

45. For example, result lists from each individual search engine are generally displayed as individual groups of results in the merged list presented to the user. On information and belief, the representative value of each result list, which determines the order of said groups, varies in

accordance with the nature of a user's query. *See, e.g.,* Google Search Result Groupings (red arrows added):

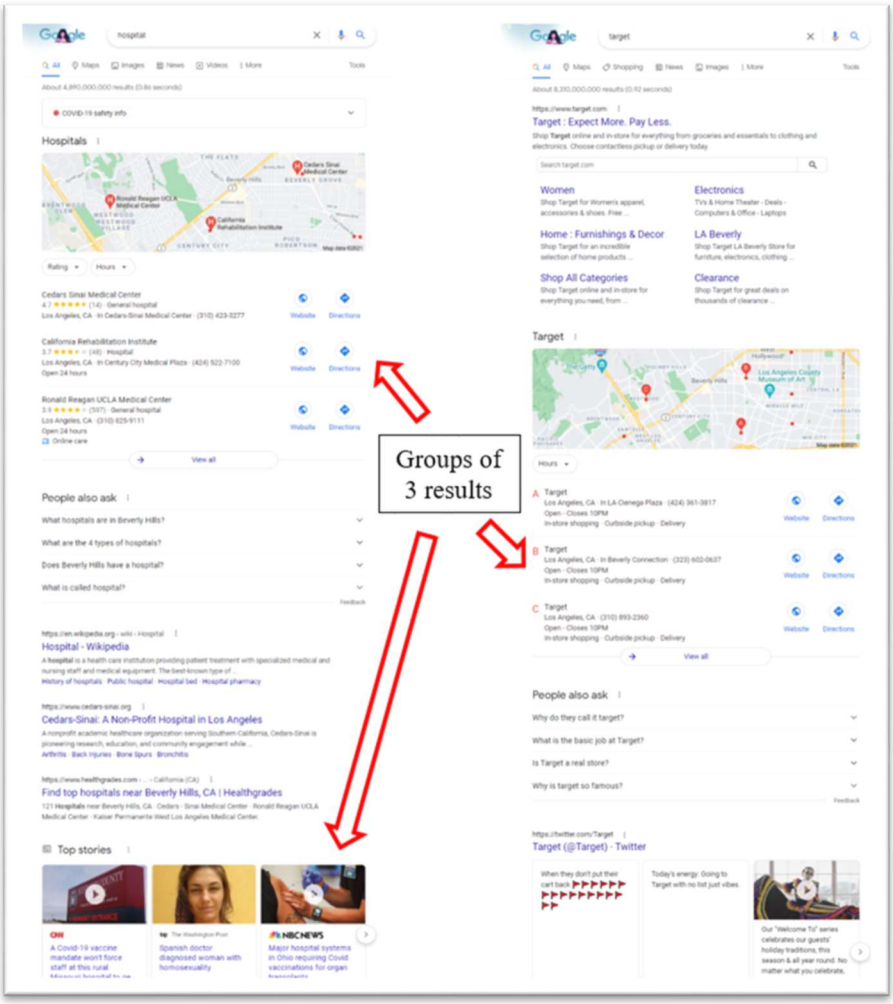

46. Furthermore, the '704 accused product, Google Search, embodies every limitation of claim 2 of the '704 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

47. Google Search practices the method of claim 1 wherein said selecting includes selecting a consecutive number of entries from each result list, including the first entry from each result list.

48. For example, as described above with respect to claim 1, when a user transmits a query for a search term with a physical location, a list of entries from, for example, the Maps result list is displayed. *See, e.g.,* para. 42, *supra*. This list includes the first entry from the Maps result list (*i.e.*, the In-N-Out Burger location closest to the user at the time of the query), as well as two other consecutive entries from said result list (*i.e.*, the second- and third-closest locations).

49. Furthermore, the '704 accused product, Google Search, embodies every limitation of claim 12 of the '704 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

50. Google Search practices a method of merging result lists from multiple search engines, comprising the elements set forth below.

51. For example, a query transmitted to Google Search returns results from multiple search engines displayed to the user as a single merged list of results. *See, e.g.*, Google Search Results for "in-n-out burger" Query 1 (red arrows added):

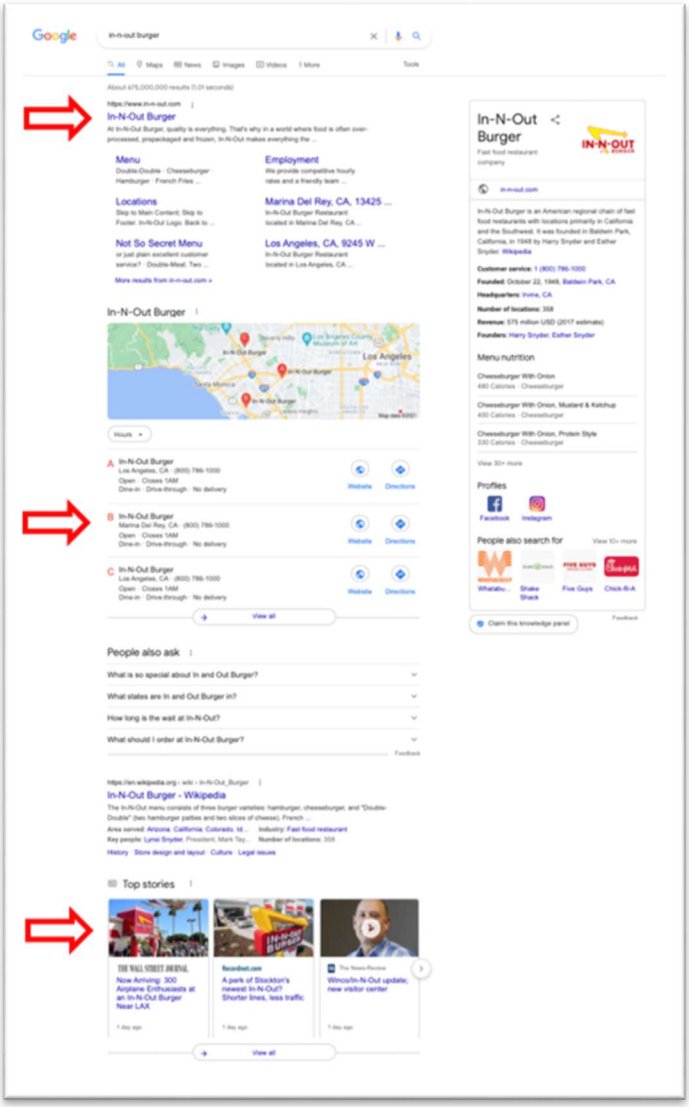

52. Google Search practices a method comprising transmitting a query to a set of search engines.

53. For example, a query transmitted to Google Search returns results from a set of multiple search engines. This process begins by transmitting that query to said search engines. *See, e.g., id.* (red arrows added):

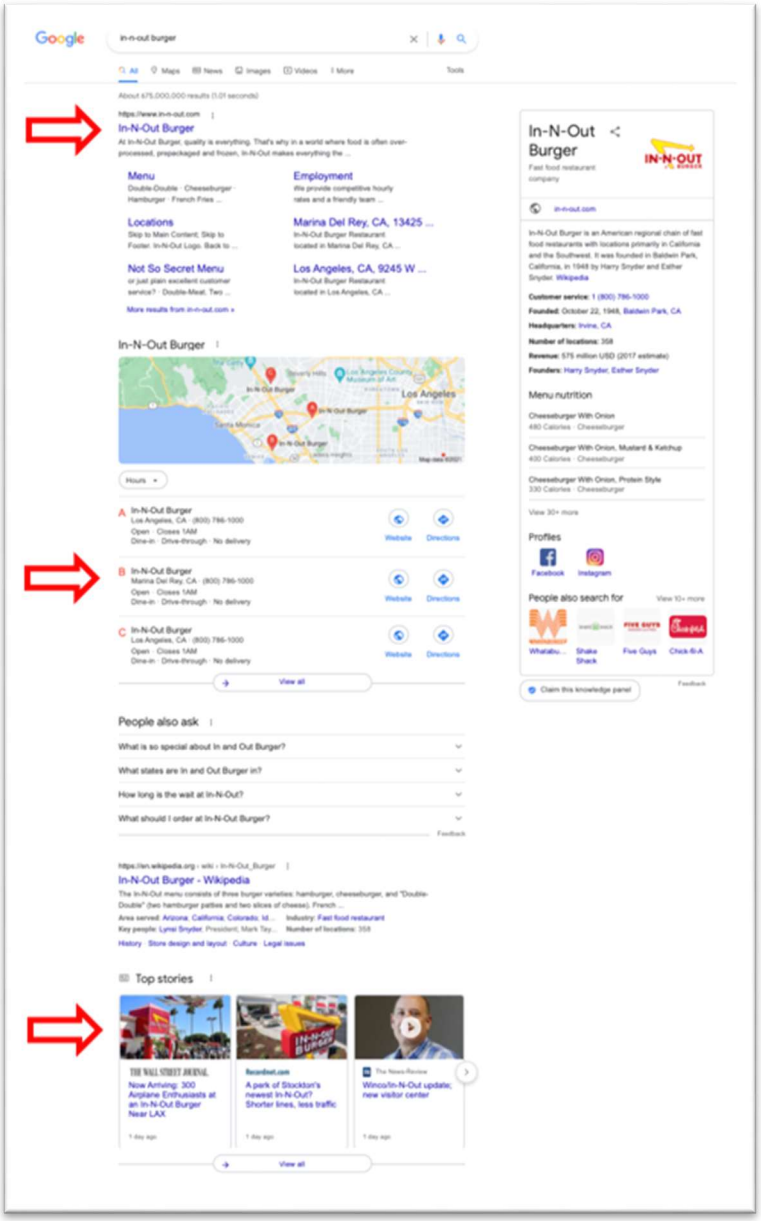

54. Google Search practices a method comprising receiving in response to said query a set of result lists, each result list in said set of result lists including one or more entries.

55. For example, a query transmitted to Google Search returns one or more results from each of several search engines. Said results are grouped into result lists associated with each of said search engines. In order to display said result lists, Google Search first receives a result list from each search engine. Each result list is generally displayed as a group of entries from search engines including, but not limited to, Web, Maps, News, Videos, Images, Shopping, Books, Flights, and Finance. *See, e.g., id.* (red arrows added):

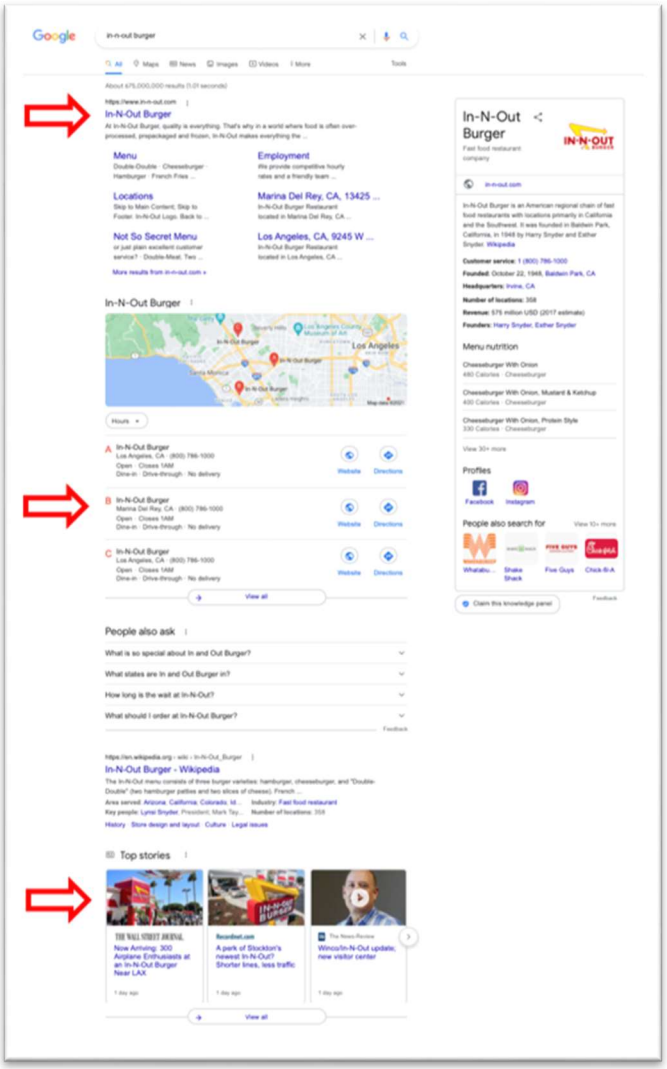

56. Google Search practices a method comprising selecting a subset of entries from each result list.

57. For example, a query transmitted to Google Search may return many millions of results across a set of multiple search engines. Only a subset of these results, or entries, is selected for display to the user near the top of the Google Search result list. *See, e.g.*, Google Search Results for "in-n-out burger" Query 2 (callout added):

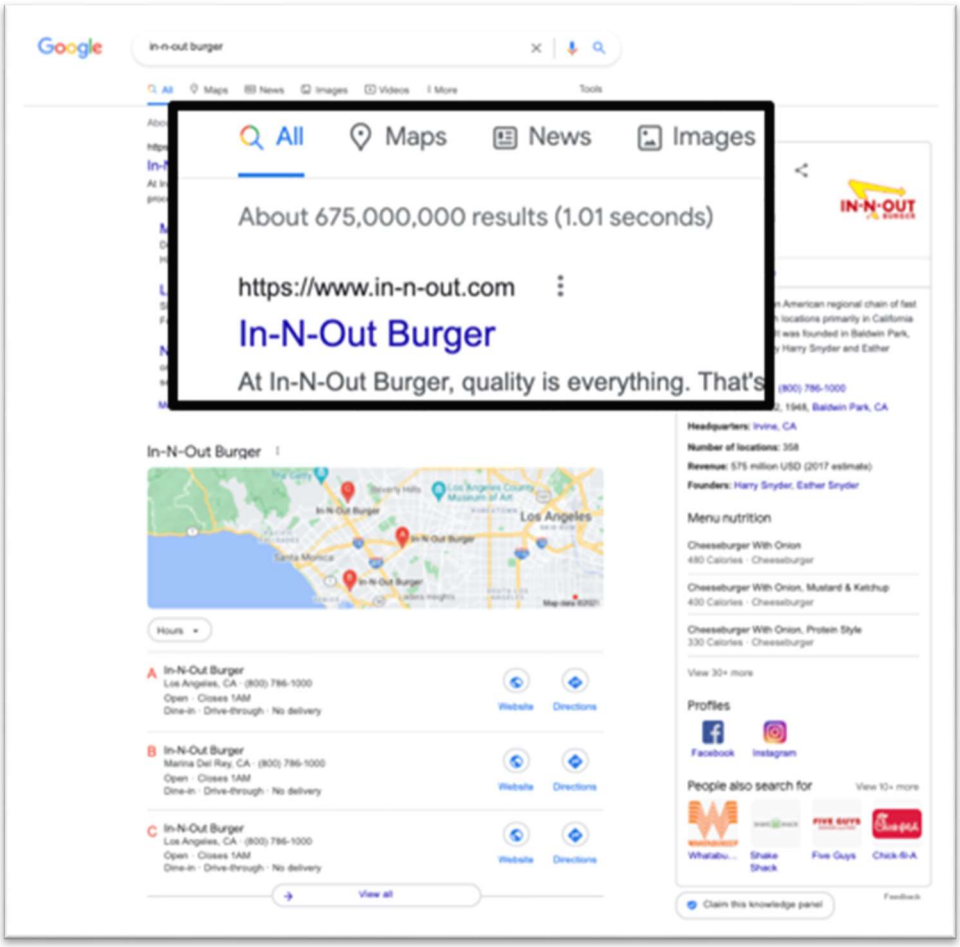

58. Google Search practices a method comprising determining a scoring value for each entry of said subset of entries.

59. For example, on information and belief, entries selected by Google Search have a scoring value determined for them based on their relevance to the user's search. For example, Maps results (which represent physical locations) are scored based on their distance from the user at the time of the query, a metric that is not used (or possible) for news articles and web pages.

Entries with a higher scoring value are placed higher on the list within that subset. *See, e.g.,* Google Maps Result List (red arrow added):

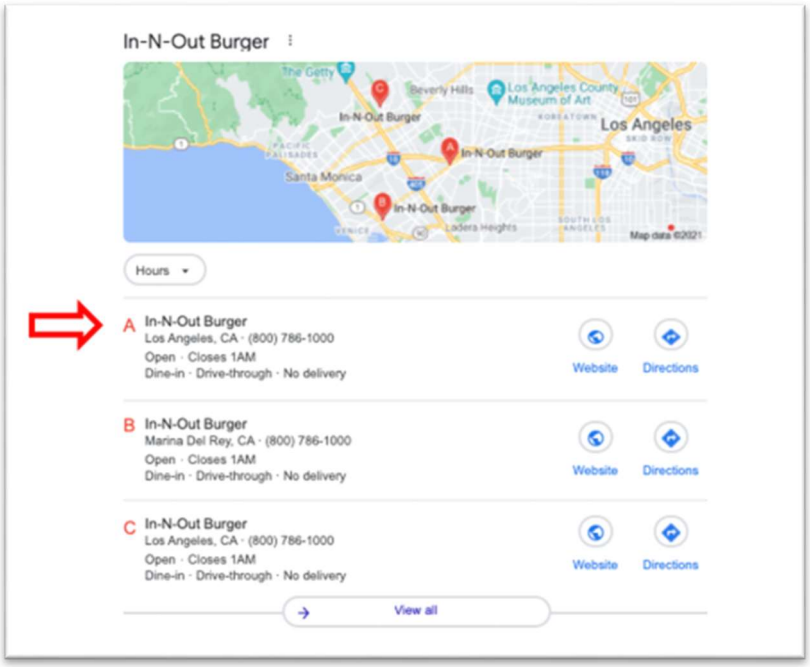

60. Google Search practices a method comprising characterizing said subset in accordance with a representative value.

61. For example, on information and belief, the order of different search engine result subsets in response to a query is dependent on the importance of each subset to a user's particular query. Google Search characterizes each subset by assigning a higher representative value to subsets that are more important to a query, with importance measured by the scoring values assigned to entries within each subset. For example, when a user transmits a query for a search term with a physical location (e.g., "in-n-out burger"), Google Search typically assigns a higher representative value to the Maps search engine subset than it assigns to the News search engine subset.

62. Google Search practices a method comprising merging entries in a predetermined manner into a single list based on said representative value.

63. For example, Google Search merges entries from each of a set of multiple independent search engines into a single list, with the single, merged list being produced in a predetermined manner based on the representative value of each independent search engine's result list. For example, when a user transmits a query involving a term associated with a physical location, Google Search produces a merged list of entries where in Maps results are generally displayed higher in the merged list than News results because, on information and belief, the Maps result list has been assigned a higher representative value. *See, e.g., id.* (red arrows added):

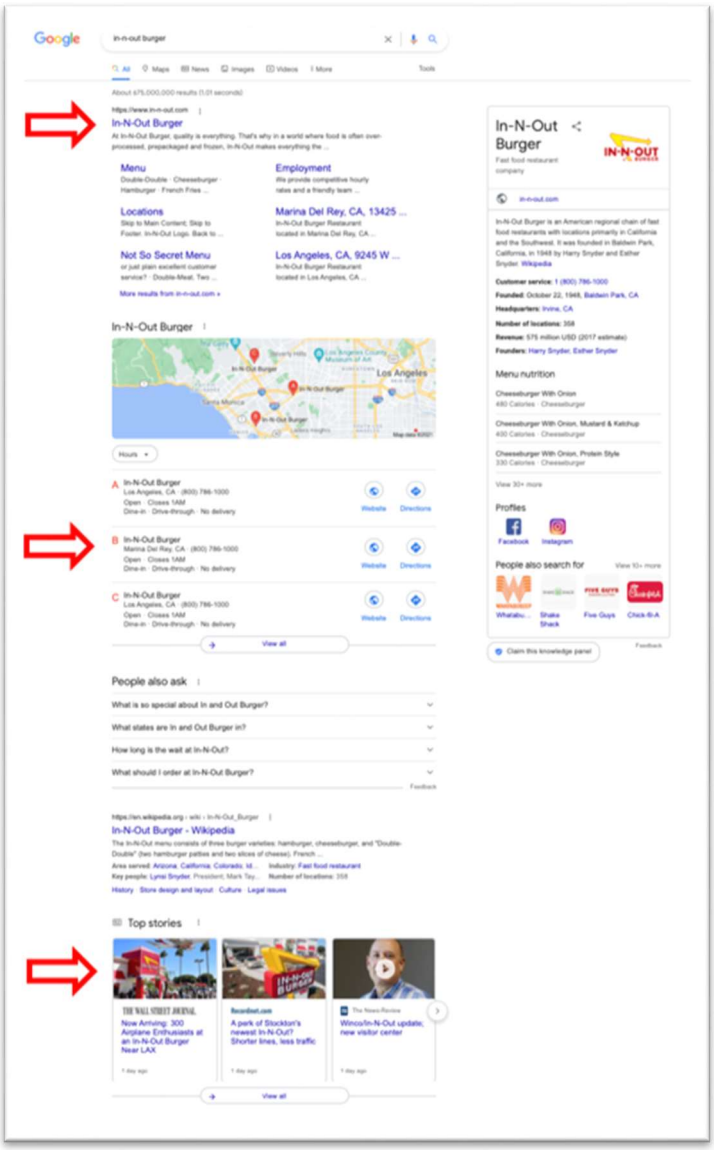

64. Google Search practices a method comprising merging entries in a predetermined manner into a single list based on said representative value, wherein said representative value varies in accordance with said predetermined manner.

65. For example, as described above, Google Search assigns a representative value to each subset of search engine results based on the importance of the search engine to a user's query. The representative value varies in accordance with a predetermined manner. Result lists from each individual search engine are typically displayed as a single group in the merged list. *See, e.g.*, *id.* The order of individual search engine result list groupings is based on the representative value assigned to each individual result list. On information and belief, result lists with higher representative values appear to be shown higher on the merged list than those with lower representative values, and the representative value of each list varies depending on the nature of the user's query.

66. Furthermore, the '704 accused product, Google Search, embodies every limitation of claim 13 of the '704 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

67. Google Search practices the method of claim 12 wherein said selecting includes selecting a consecutive number of entries from each result list, including the first entry from each result list.

68. For example, as described above with respect to claim 12, when a user transmits a query for a search term with a physical location (e.g., "in-n-out burger"), a list of entries from the Maps result list is typically selected for display in the merged list. *See, e.g.*, para. 54, *supra*. This list includes the first entry from the Maps result list (i.e., the In-N-Out Burger location closest to

the user at the time of the query), as well as two other consecutive entries from said result list (i.e., the second- and third-closest locations to the user at the time of the query).

69. Google has had knowledge of the '704 Patent since at least January 1, 2013. On this date, two patents assigned to Google Inc. were issued, both of which cite the '704 Patent as a reference: U.S. Patent Nos. 8,346,791 and 8,346,792. The '704 Patent is also cited as a reference by two later patents, U.S. Patent Nos. 8,521,725 (issued August 27, 2013) and 9,152,714 (issued October 6, 2015), both of which name Google Inc. as assignee. Furthermore, two of the co-inventors of the '704 Patent, Rajat Mukherjee and Prabhakar Raghavan, are listed as inventors on at least seven other patents or publications which are assigned to Google: U.S. Patent Nos. 8,782,071 (issued July 15, 2014), 9,251,168 (issued February 2, 2016), 10,079,785 (issued September 18, 2018), 10,402,889 (issued September 3, 2019), and 10,917,371 (issued February 9, 2021); and U.S. Patent Application Publication Nos. 2014/0258466 (published September 11, 2014) and 2016/0371425 (published December 22, 2016). Two of these patents, U.S. Patent Nos. 8,782,071 and 10,402,889, relate directly to search technology. Furthermore, Prabhakar Raghavan now holds the position of Senior Vice President at Google, where he "is responsible for Google Search," among other products.[5] Given that four of its own patents cite the '704 Patent as a reference, two of the co-inventors of the '704 Patent are listed as inventors on numerous later patents assigned to Google, and one of the co-inventors of the '704 Patent is now in a senior leadership role at Google, Google knew or was willfully blind to the fact that its conduct was infringing by, at the very latest, July 15, 2014.

70. Plaintiffs are informed and believe, and thereon allege, that Google actively, knowingly, and intentionally has induced infringement of the '704 Patent by, for example, offering

[5] https://research.google/people/PrabhakarRaghavan/

for public use Google Search with the intent to encourage and facilitate infringing uses of that service in the Northern District of Texas, in the United States, and throughout the world.

71. As a result of Google's infringement of the '704 Patent, Plaintiffs have been damaged. Plaintiffs are entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

72. In addition, Google's infringing acts and practices have caused and are causing immediate and irreparable harm to Plaintiffs.

73. Plaintiffs are informed and believe, and thereon allege, that Google's infringement of the '704 Patent has been and continues to be willful. As noted above, Google has long had knowledge of the '704 Patent and its infringement of the '704 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Plaintiffs' patent rights. Thus, Google's infringing actions have been and continue to be consciously wrongful.

SECOND CLAIM

(Infringement of U.S. Patent No. 6,816,809)

74. Plaintiffs re-allege and incorporate herein by reference Paragraphs 1-73 of their Complaint.

75. The '809 Patent, entitled "Hardware based utilization metering," was duly and lawfully issued on November 9, 2004. A true and correct copy of the '809 Patent is attached hereto as Exhibit 6.

76. The '809 Patent names Edgar Circenis as inventor.

77. The '809 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '809 Patent, including the right to seek damages for any infringement thereof.

78. The '809 Patent relates to novel circuitry in a processor and the use of this novel circuitry, in particular, "methods that use central processor metering to determine processor utilization for billing and other purposes." Ex. 6 at 1:6-8. At the time of the '809 Patent's invention, colocation servers existed, in which different customers would pay to have their processes run on common partitioned processors. These systems "allow computer users to acquire a given computing capacity that may be tailored to the user's specific need by, among other things, charging the user only for actual utilization of the system central processor units (CPUs)." *Id.* at 2:44-48. At the time of the invention, such metering was "accomplished by *software* running within the computer system's operating system." *Id.* at 1:16-18 (emphasis added).

79. "In a computer system having hardware that may be partitioned, gathering processor utilization data from a hardware system requires communications between the metering application and all operating systems running within the hardware. The need for communication with different operating systems poses significant challenges because operating systems by their design are separated from other operating systems and do not have visibility to utilization data from other operating systems." *Id.* at 1:18-27. Additional difficulties could arise, for example, "when network access is restricted, agent software is not installed, and operating systems are temporarily out of service." *Id.* at 1:34-36.

80. The '809 Patent overcomes these problems in the prior art by proposing "a *hardware* based method for measuring processor utilization." *See id.* at 7:30-8:65 (emphasis added). More specifically, the invention uses "inputs from [a] system clock and [an] idle indicator" to measure CPU cycles "where the CPU is not in an idle state, but instead is performing a service for the user of the system." *Id.* at 4:1-4. Claim 1 recites three hardware elements: "an idle indicator," "a counter coupled to the hardware indicator," and a "data usage provider." *See id.* at

7:30-43. The idle indicator is a specific hardware structure, as described in the specification. *See id.* at 4:36-56. Likewise the data usage provider is a specific hardware structure that can, for example, allow the data generated by the counter to persist after removal of the processor or its loss of power. *See id.* at 4:8-23.

81. Independent expert analysis of the '809 Patent confirms that "this specific arrangement of specialized hardware to accomplish utilization metering was not a routine practice at the time of the '809 Patent's invention . . . [and] avoid[s] the many potential pitfalls of software-based metering." Ex. 7 (Declaration of Dr. Thomas M. Conte) at 8.

82. Moreover, the specific apparatus and method of the '809 Patent create quantifiable improvements to the functioning of a computer system utilizing them. For example, independent expert testing on the '809 Patent invention found that using one particular arrangement of hardware and the method claimed by the '809 Patent resulted in an improvement of over 78% per watt as compared to a prior art approach. *See id.* at 12.

83. The dependent claims of the '809 Patent add additional inventive structure. For example, claim 3 specifies an even more detailed implementation of the hardware counter which involves the use of a pin on the processor. Ex. 6 at 7:46-51. Claim 4 specifies a design in which the idle indicator is visible through an external register. *Id.* at 7:52-56. Claim 5 describes how the counter can increment, and claim 6 describes how the count can persist when a processor is turned back on. *Id.* at 7:57-62. Claims 7 and 8 describe how the counter can be modified. *Id.* at 7:63-66. Claim 9 describes an interface to an external network associated with the counter. *Id.* at 7:67-68. And claims 10 through 12 describe a hardware structure for combining multiple counters. *Id.* at 8:3-13. The method and apparatus claims 13 through 20 are analogous. *See id.* at 8:14-65. All of these claims describe specialized circuitry.

84. Unlike prior software-based metering approaches, the method and apparatus claimed by the '809 Patent allow for accurate measurement of processor utilization across multiple processors and operating systems, cutting out the need for operating system-specific software agents and avoiding the difficulties caused by network or operating system failures. During prosecution of the '809 Patent, the applicant emphasized that the prior art did not disclose hardware-based circuits for tracking the recited information: "[The prior art of record] are software implementations, and the [cited art] teaches away from hardware based methods." Ex. 8 ('809 Patent prosecution history, January 6, 2004 Amendment and Response) at 7. The Patent Office agreed, finding that "the prior arts do not disclose a plurality [of] processors and the idle indicator of processors in a busy state." Ex. 9 ('809 Patent prosecution history, January 8, 2004 Interview Summary) at 1. There is no concession in the record that these hardware structures are routine or conventional. To the contrary, they are repeatedly emphasized as novel during prosecution, a position which was accepted by the Patent Office. *See* Ex. 10 ('809 Patent prosecution history, April 5, 2004 Notice of Allowability).

85. Plaintiffs are informed and believe, and thereon allege, that Google has infringed and unless enjoined will continue to infringe at least one or more of the method claims of the '809 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the claimed hardware based utilization metering method in an infringing manner. Google practices every step of claim 1 of the '809 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

86. For example, the '809 accused product, Google Cloud, embodies every limitation of claim 1 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The

further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

87. Google Cloud uses a hardware based utilization metering device comprising the elements below.

88. For example, Google Cloud offerings are powered by central processing units ("CPUs"), which include a hardware based utilization metering device comprising the following elements.

89. Google Cloud uses a hardware based utilization metering device comprising an idle indicator coupled to a processor, wherein the idle indicator receives an indication when the processor is in a first state.

90. For example, the hardware based utilization metering device used by Google Cloud includes a processor having hardware performance counters. Certain of these counters increment only while the processor core is in a C0 (busy) state. These counters, including MPERF and APERF counters, are stopped and started by a hardware idle indicator.

91. Google Cloud uses a hardware based utilization metering device comprising a counter coupled to the idle indicator and coupled to a system clock, wherein the counter receives a measure of system time from the system clock and receives data related to the indication when the processor is in the first state, and generates a counter value indicative of time the processor is in the first state.

92. For example, the hardware based utilization metering device used by Google Cloud includes a processor having hardware performance counters. Certain of these counters increment only while the processor core is in a C0 (busy) state. These counters, including MPERF and

APERF counters, are stopped and started by a hardware idle indicator, and generate a counter value indicative of time the processor is in a first state (such as a C0 busy state).

93. Google Cloud uses a hardware based utilization metering device comprising a data usage provider coupled to the counter, wherein the data usage provider is capable of providing the counter value.

94. For example, the hardware based utilization metering device used by Google Cloud includes a processor having a data usage provider coupled to the counter. This provider is capable of providing the counter value. *See* Google Cloud metrics[6]:

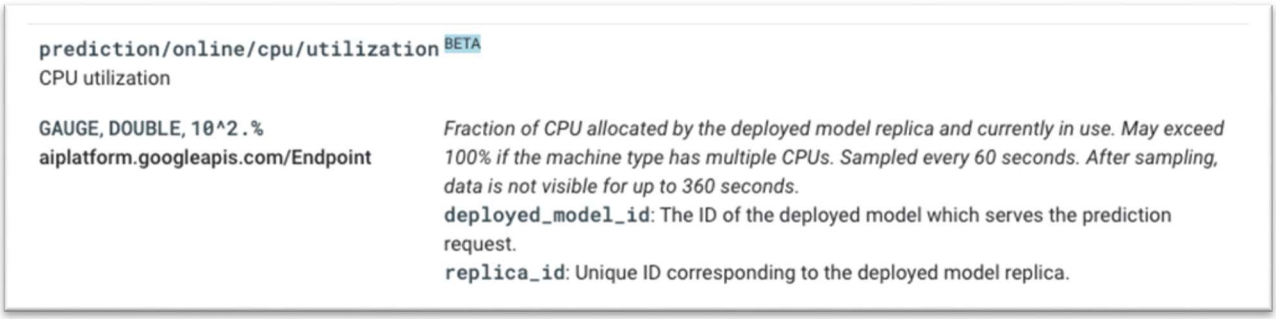

`prediction/online/cpu/utilization` BETA
CPU utilization

GAUGE, DOUBLE, `10^2.%` aiplatform.googleapis.com/Endpoint	*Fraction of CPU allocated by the deployed model replica and currently in use. May exceed 100% if the machine type has multiple CPUs. Sampled every 60 seconds. After sampling, data is not visible for up to 360 seconds.* `deployed_model_id`: The ID of the deployed model which serves the prediction request. `replica_id`: Unique ID corresponding to the deployed model replica.

95. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 2 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

96. Google Cloud uses the hardware based utilization metering device of claim 1, wherein the first state is a busy state.

97. For example, each core has performance counters including MPERF and APERF. These counters increment only while the processor core is in a C0 busy state.[7]

[6] https://cloud.google.com/monitoring/api/metrics_gcp
[7] *See, e.g.*, https://github.com/torvalds/linux/blob/464fddbba1dfbc219f1e9145127a482d2159dee5/arch/x86/kernel/acpi/cstate.c; *see also, e.g.*,

98. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 3 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

99. Google Cloud uses the hardware based utilization metering device of claim 2, wherein the processor comprises an operating system that halts the processor when the processor is not processing, and wherein the idle indicator is a hardware device coupled to a pin of the processor, wherein the hardware device reads a signal asserted on the pin to receive the idle indication.

100. For example, processors used by Google Cloud include an operating system that halts the processor using, for example, the 'MWAIT' instruction when the processor is not processing. *See, e.g.*, *id.* The hardware based utilization metering device used by Google Cloud and described above in claims 1 and 2 includes a processor having hardware performance counters. Certain of these counters increment only while the processor core is in a C0 (busy) state. These counters, including MPERF and APERF counters, are stopped and started by a hardware idle indicator.

101. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 5 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

---

https://github.com/torvalds/linux/blob/5bfc75d92efd494db37f5c4c173d3639d4772966/arch/x86/include/asm/mwait.h

102. Google Cloud uses the hardware based utilization metering device of claim 2, wherein the counter value increments based on the measure of system time when the processor is in the busy state.

103. For example, each core of the processors used by Google Cloud has performance counters including MPERF and APERF. Certain of these counters increment only while the processor core is in a C0 (busy) state, such that the counter value increments based on the measure of system time when the processor is in the C0 (busy) state.

104. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 6 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

105. Google Cloud uses the hardware based utilization metering device of claim 1, wherein when the processor is powered on, the counter receives a current value of the counter value.

106. For example, a core of a processor of the type included in Google Cloud can gate off power to its internal logic when it enters any non-C0 state, such as C6. As part of the sequence that occurs when a core enters C6, the core's internal state—which includes the counter values described in claim 1—is saved to cache, DRAM, or other memory. This state is used to reinitialize the core, including the counter value described above with respect to claim 1, to the saved core state when the processor is powered on again (i.e., reenters C0), such that the counter receives a current value of the counter value when the processor is powered on.

107. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 9 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The

further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

108. Google Cloud uses the hardware based utilization metering device of claim 1, further comprising an interface coupling the data usage provider to an external network, the interface used for providing processor utilization data to the external network.

109. For example, the hardware based utilization metering device used by Google Cloud, as described above in claim 1, includes a processor having a data usage provider coupled to the counter. As shown below, this provider is capable of providing processor utilization data (i.e., the counter value) to an external network. *See* Google Cloud metrics:

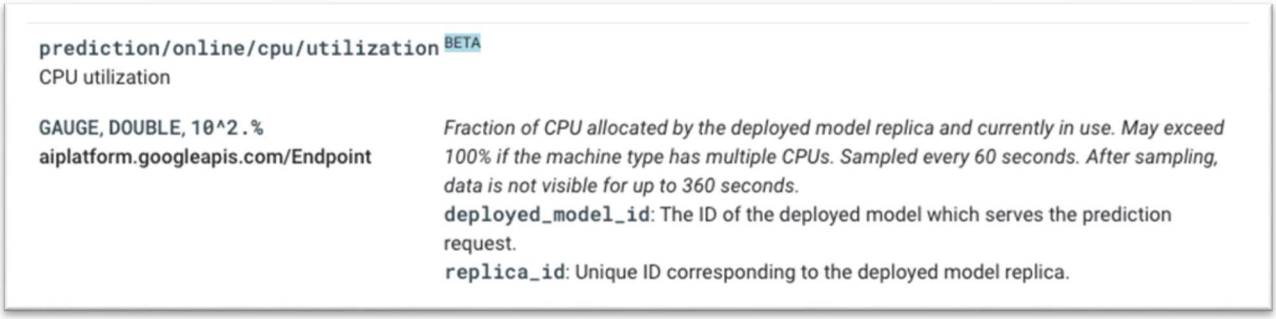

`prediction/online/cpu/utilization` BETA
CPU utilization

GAUGE, DOUBLE, `10^2.%` aiplatform.googleapis.com/Endpoint	*Fraction of CPU allocated by the deployed model replica and currently in use. May exceed 100% if the machine type has multiple CPUs. Sampled every 60 seconds. After sampling, data is not visible for up to 360 seconds.* `deployed_model_id`: The ID of the deployed model which serves the prediction request. `replica_id`: Unique ID corresponding to the deployed model replica.

110. Furthermore, the '809 accused product, Google Cloud, embodies every limitation of claim 13 of the '809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

111. Google Cloud practices a hardware based method for measuring processor utilization in a computer system comprising a plurality of processors. The method comprises the elements described below.

112. For example, Google Cloud offerings are powered by a plurality of processors. *See, e.g.*, Confidential Computing[8]:

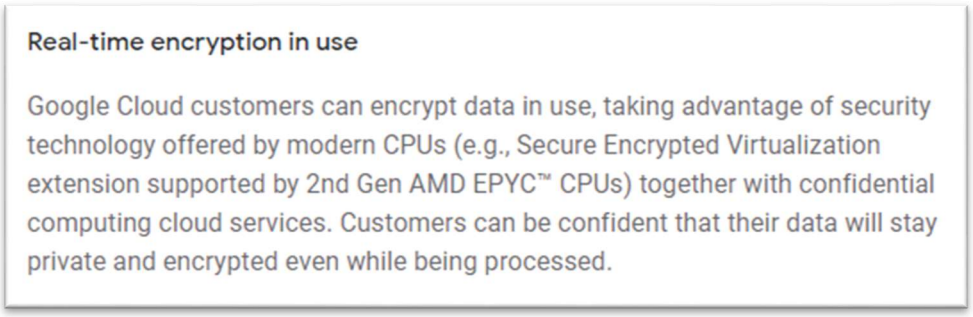

113. Each CPU contains a plurality of cores, each of which is a processor. These processors use a hardware-based method of incrementing performance counters when a specified event occurs or to count the number of processor clock cycles required to complete a specific hardware function. This method can be used to measure processor utilization.

114. Google Cloud practices a method comprising determining when any of the plurality of processors is busy.

115. For example, processors used by Google Cloud have performance counters including MPERF and APERF. An exemplary performance counter in these processors increments only while the processor core is in the C0 state, meaning the core is executing instructions, and is used to determine when any of a plurality of such processors is busy.

116. Google Cloud practices a method comprising providing a busy indication to a counter associated with a busy processor.

117. For example, the processor cores used by Google Cloud provide a busy indication to a counter while the core is in the C0 state, which controls when the counter will increment.

118. Google Cloud practices a method comprising receiving at the counter a measure of computer system time.

[8] https://cloud.google.com/confidential-computing

119. For example, the APERF counter in processors used by Google Cloud receives at the counter a core clock as a pulse train, a measure of computer system time, while the core is in the C0 state.

120. Google Cloud practices a method comprising incrementing a counter value in the counter based on the provided busy indication and an amount of computer system time that the processor is determined to be busy.

121. For example, counters in the processors used by Google Cloud are incremented by hardware based on busy indications and the core clock. A counter increments in proportion to the actual number of core clocks cycles while the core is in C0, a measurement of the amount of computer system time that the processor is determined to be busy. For example, each core has performance counters including MPERF and APERF. The MPERF and APERF counters are incremented by hardware based on busy indications and clock cycles. The APERF counter increments in proportion to the actual number of core clocks cycles while the core is in C0.

122. Google Cloud practices a method comprising maintaining the counter value.

123. For example, counters in the processors used by Google Cloud are read/write registers, meaning that their values are preserved for access by software. Processors used by Google Cloud maintain certain of these counter values as they increment to allow the counter to retain values from clock cycle to clock cycle. The MPERF and APERF counter values are incremented when hardware detects an occurrence of a specified event or when it counts the number of processor clock cycles required to complete a specific hardware function. Processors maintain this counter value as they increment to allow values to retain in the counter from clock cycle to clock cycle.

124. Furthermore, the ’809 accused product, Google Cloud, embodies every limitation of claim 14 of the ’809 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

125. Google Cloud practices the method of claim 13, further comprising reinitializing the counter value in the counter, using the maintained value of the counter value, when the processor is powered on.

126. For example, a core of a processor can gate off power to its internal logic when it enters any non-C0 state, such as C6. As part of the sequence that occurs when a core enters C6, the core’s internal state—which includes the counter values described in claim 13—is saved to cache, DRAM, or other memory. This state is used to reinitialize the core, including the counter value described above with respect to claim 13, to the saved core state when the processor is powered on again (i.e., re-enters C0).

127. Google has had knowledge of the ’809 Patent at least since April 14, 2021. On that date, Plaintiffs sent to Google a notice letter with an exemplary list of Plaintiffs’ patents, which included the ’809 Patent, that were infringed by specific exemplary Google products. Therefore, Google either knew of or was willfully blind to the ’809 Patent and its infringement of the ’809 Patent no later than this date.

128. Plaintiffs are informed and believe, and thereon allege, that Google actively, knowingly, and intentionally has induced infringement of the ’809 Patent by, for example, using, selling, and offering for sale Google Cloud products which rely on processors having the claimed counters, which in turn use the process claimed by the ’809 Patent. Google offers for sale and

sells said Google Cloud products with the intent to encourage and facilitate infringing uses of those products in the Northern District of Texas, in the United States, and throughout the world.

129. As a result of Google's infringement of the '809 Patent, Plaintiffs have been damaged. Plaintiffs are entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

130. In addition, Google's infringing acts and practices have caused and are causing immediate and irreparable harm to Plaintiffs.

131. Plaintiffs are informed and believes, and thereon alleges, that Google's infringement of the '809 Patent has been and continues to be willful. As noted above, Google had knowledge of the '809 Patent and its infringement of the '809 Patent no later than April 14, 2021. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Plaintiffs' patent rights. Thus, Google's infringing actions have been and continue to be consciously wrongful.

THIRD CLAIM

(Infringement of U.S. Patent No. 7,346,604)

132. Plaintiffs re-allege and incorporate herein by reference Paragraphs 1-131 of their Complaint.

133. The '604 Patent, entitled "Method for ranking hypertext search results by analysis of hyperlinks from expert documents and keyword scope," was duly and lawfully issued on March 18, 2008. A true and correct copy of the '604 Patent is attached hereto as Exhibit 11.

134. The '604 Patent names Krishna A. Bharat and George A. Mihaila as co-inventors.

135. The '604 Patent has been in full force and effect since its issuance. Valtrus owns by assignment the entire right and title in and to the '604 Patent, including the right to seek damages for any infringement thereof.

136. In 1999, the co-inventors of the '604 Patent, Krishna Bharat and George A. Mihaila, published a white paper describing a new computer search engine called "Hilltop." *See* Hilltop Paper.[9] In Bharat and Mihaila's testing, Hilltop was either as accurate as or more accurate than the best commercial search engines then available: "Thus, for about 87% of the queries, *Hilltop* returned the desired page as the first result, comparable with *Google* at 80% of the queries, while *DirectHit* and *AltaVista* succeeded at rank 1 only in 43% and 20% of the cases, respectively." *Id.* The lead author of the Hilltop Paper was Krishna Bharat, then a researcher at Compaq, who was hired by Google in 1999, the same year the Hilltop Paper was published. Google introduced the Hilltop search engine in its "Florida" update, with one industry commentator stating that it "conquered a fatal flaw in PageRank" by "adding topical relevance to PageRank's one-dimensional vision of authority."[10]

137. The application of Hilltop to Google Search fundamentally changed the performance and nature of Google's business. As one commentator put it, "Krishna Bharat rewrote how Google worked in the early 2000s by applying the Hilltop Algorithm."[11] Moreover, the '604 Patent's ability to objectively distinguish between authoritative experts and spam pages was specifically acknowledged as transforming search optimization:

> Sites that relied on link farms, stuffed meta tags, superfluous keyword density, hidden tags, and invisible text all plummeted in SERPs [(search results pages)]. Those exchanging links with off-topic sites were punished as well. Google had reinforced its point: It wanted the most relevant, not the most optimized, links at the top of SERPs.[12]

---

[9] Available at http://ftp.cs.toronto.edu/pub/reports/csrg/405/hilltop.html.
[10] https://web.archive.org/web/20210227213440/https://stafflessseo.com/google-search-ranking-brief-history-known-knowns-known-unknowns-unknown-unknowns/
[11] https://www.seobythesea.com/2014/03/incomplete-google-ranking-signals-1/
[12] https://web.archive.org/web/20210227213440/https://stafflessseo.com/google-search-ranking-brief-history-known-knowns-known-unknowns-unknown-unknowns/

138. The ’604 Patent covers, among other things, the Hilltop search engine, and is directed to creating a ranked list of search results based on an index of expert documents formed during a preprocessing step. It “relates generally to software programs and, more specifically, to search engines that search large numbers of hypertext documents.” Ex. 11 at 1:9-11.

139. The ’604 Patent states that “[a] broad search query can produce a huge result set. This set is hard to rank based on content alone, since the quality and ‘authoritativeness’ (namely, a measure of how authoritative the page is on the subject) of pages cannot be assessed solely by analyzing their content.” *Id.* at 1:31-35. Because prior art approaches could not “distinguish between authoritative and non-authoritative pages,” *id.* at 1:41-42, they failed to detect and screen out so-called spam pages, which are “created for the purpose of misleading search engines and may contain spurious words that do not pertain to the topic of the page.” *Id.* at 1:36-38.

140. The ’604 Patent explains, “[b]y combining relevant out-going links from many experts on the query topic, it is possible to find the pages that are most highly regarded by the community of pages related to the query topic.” *Id.* at 3:4-7.

141. The ’604 Patent’s method for ranking hypertext search results provides an innovative approach to overcoming deficiencies in existing methods of improving the authoritativeness of ranked search results. The ’604 Patent describes three such existing methods: “ranking based on human classification,” “ranking based on usage information,” and “ranking based on connectivity,” *id.* at 1:46, 56, 66, and notes that none of these are able to provide authoritative, relevant search results for a large number and variety of queries. *See id.* at 1:44-2:15. The result of the techniques described by the ’604 Patent, by contrast, “is to generate a listing of pages that are highly relevant to the user’s query and of high quality,” an innovation which “provides a novel and advantageous system and method of searching large numbers of hypertext

documents, such as the hypertext documents of the world wide web." *Id.* at 8:27-29, 39-42. The '604 Patent provides specific steps for carrying out this method. For example, in one embodiment, an "expert reverse index" is created by identifying pages which are "about a certain topic and [have] links to many non-affiliated pages on that topic." *Id.* at 4:47-48. Detailed instructions are given for determining whether two hosts are "affiliated" and the degree of a page's expertise, and for "adding keywords in the expert pages to an expert reverse index." *Id.* at 4:57-6:36. This invention improves the performance and efficiency of computers by returning highly relevant and authoritative search results to users while avoiding the high computational overhead and slow response times of other methods. *See id.* at 1:44-2:45.

142. The '604 Patent's claimed method for producing a ranked list of highly relevant and authoritative results "involves two broad phases: (a) expert lookup and (b) target ranking." *Id.* at 2:53-54. The Patent explains that an expert page on the Internet is "a page that is about a certain topic and has links to many 'non-affiliated' pages on that topic. Two pages are non-affiliated conceptually if they are authored by authors from non-affiliated organizations." *Id.* at 2:55-58. These expert pages are identified before and independent of any given query in order to build an expert reverse index. *See id.* at 5:62-6:36. Building the expert index involves searching a corpus of hypertext documents available to the engine for certain defined characteristics, which are typically detected via machine learning. *See id.* at 4:58-5:40. When a user submits a query, "a lookup is done on the expert reverse index to find and rank matching expert pages." *Id.* at 2:65-66. "Next, target ranking looks at the out-going links from identified expert pages," making it "possible to find the pages that are most highly regarded by the community of pages related to the query topic. This is the basis of the high relevance that the described embodiment of the invention delivers." *Id.* at 3:3-9. By identifying expert pages in a pre-processing step and then ranking pages

pointed to by those experts, the inventive method of the '604 Patent screens out spam pages more effectively than prior approaches and allows for faster retrieval of highly relevant results. The claims recite this ranking process, which involves ranking the expert documents in accordance with a user's query, then ranking target documents pointed to by those ranked expert documents for presentation to the user. *See id.* at 8:50-55.

143. Independent expert analysis and testing of the '604 Patent further confirms that the claimed invention is not routine, well-understood, or conventional, and that the prior art of record "fails to disclose, among other things, forming a set of expert documents without reference to a query, and ranking expert documents in accordance with a search query." Ex. 2 at 12-13.

144. Furthermore, the '604 Patent "also results in a quantifiably faster ranking process as compared to alternatives which attempt to evaluate authoritativeness at the time a search query is submitted to the search engine." *Id.* at 13. Because "the '604 Patent forms the set of expert documents in a pre-processing step which occurs prior to and without reference to any particular query . . . this method avoids the need to compute a new subgraph measuring connectivity between expert and subject pages for every possible query, a process which can be computationally intensive to generate, requiring the consideration of millions of pages." *Id.* In independent expert testing, the method described by claim 1 of the '604 Patent returned a list formed by, on average, ten expert results nearly twice as fast as a method which simulated one aspect of an alternative used in prior art meta-search engines. *See id.* at 15 (demonstrating that the method of the '604 Patent took an average of 12.97ms to return a list of results, while a prior art alternative approach took an average of 24.27ms).

145. The dependent claims of the '604 Patent contain additional inventive elements, emphasizing that the invention can operate in the vast, technologically-complex environment of

the world wide web (claims 2 and 3), and providing detailed instructions for implementing the invention. *See* Ex. 11 at 8:57-62. For example, claims 5 through 10 and 19 recite instructions on defining key words and phrases within an expert document in order to qualify what links are appropriate targets in the second phase, as well as how to define what constitutes an expert document. *See id.* at 8:66-9:20. Still other claims provide instructions for defining expert documents (claims 11 and 12), scoring the target documents identified in expert documents (claims 13 through 17), and determining that two sources are affiliated (claim 18). *See id.* at 9:21-10:9. Each of these elements are objective and allow for a machine learning process to define the expert pages, rank the expert pages and targets within the expert pages, and perform the search.

146. The ’604 Patent issued after a determination by the Patent Trial and Appeal Board that the invention it claimed was novel and non-obvious. *See generally* Ex. 12 (’604 Patent prosecution history, June 27, 2007 Decision on Appeal). In its brief to the PTAB, the applicant emphasized the objective improvement in search engine processes provided by the invention: “The invention is a solution to a problem caused by the fact that the [web] contains an extremely large number of documents. Conventional search engines . . . may require excessive amounts of time and computer resources to actually search the entire content of the web . . . [and] have difficulty in filtering out documents which are of low quality.” *See* Ex. 13 (’604 Patent prosecution history, October 13, 2004 Appeal Brief) at 7. “The present invention reduces the time and effort required to provide a ranked list of documents after a query is submitted by breaking the process into two broad phases,” which include expert lookup and target ranking. *Id.*

147. Plaintiffs are informed and believe, and thereon allege, that Google has infringed one or more claims of the ’604 Patent, in violation of 35 U.S.C. § 271, by, among other things, using, selling, and offering for sale, without authority or license, Google products that use the

claimed method of analysis of hyperlinks from expert documents in an infringing manner. Google practices every step of claim 1 of the '604 Patent in the United States, including one or more steps that it practices in the Northern District of Texas.

148. For example, the '604 accused product, Google Search, embodies every limitation of claim 1 of the '604 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

149. Google Search practices a computer-implemented method for searching a large number of hypertext documents in accordance with a search query.

150. For example, Google Search "looks through hundreds of billions of webpages" when a user submits a query. *See, e.g.*, How Search works – Organizing information[13]:

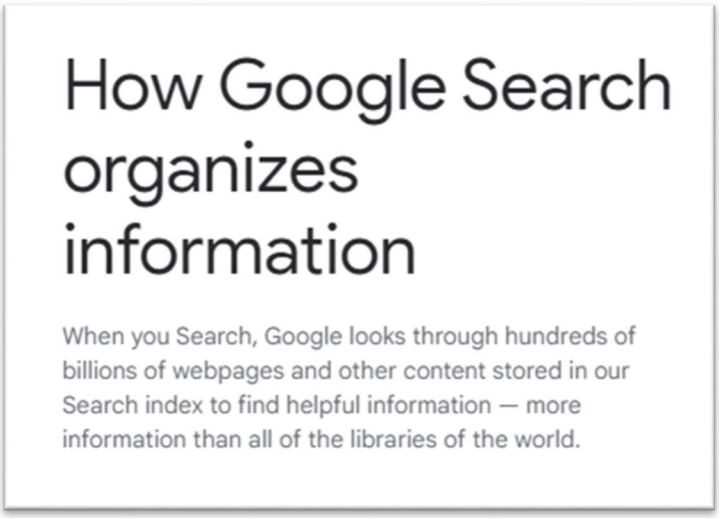

151. On information and belief, Google Search practices a method comprising forming a set of expert documents from the set of all hypertext documents crawled without reference to the search query.

[13] https://www.google.com/search/howsearchworks/how-search-works/organizing-information/

152. For example, publicly available information indicates that Google Search uses an algorithm known as "Hilltop" to build a special index of expert documents from the set of all webpages that Google Search crawlers have indexed. *See, e.g.*, Aaron Aders, *Build Trust and Authority in Google Search*, Inc.com[14]:

> To offer an understanding of authority, the Hilltop algorithm uses "a special index of expert documents" as starting points for authority signals. These "expert documents" are simply Web pages that Google has deemed an authoritative Hilltop. When a Hilltop links to a website outside of this special club, that website earns trust and authority. This increased brand authority earns higher search engine rankings.

153. Google Search uses an algorithm to build a special index of expert documents from the set of all webpages that crawlers have indexed. This process occurs separately from any particular query submitted by a Google Search user. Google processes a search engine's database pages and identifies a subset of pages as good sources of unknown links for specific topics. For example, considering all pages with out-degree greater than a threshold X, Google tests to determine if those URLs point to X distinct hosts that are not affiliated. Each is considered an expert page.

154. An inverted index then maps keywords to the experts in which they appear and a list is structured of positions that match within the experts. The matches map to the appearance of keywords within a key phrase of an expert page.

[14] https://www.inc.com/aaron-aders/build-trust-and-authority-in-google-search.html

155. Google Search "systems" form said index of expert documents by determining "which content demonstrates expertise, authoritativeness, and trustworthiness." This determination is made based at least in part on "if other prominent websites link or refer to the content." On information and belief, at least part of this determination occurs without reference to a particular query. *See* How Search works – Ranking results – Quality of content[15] (red rectangular annotations added):

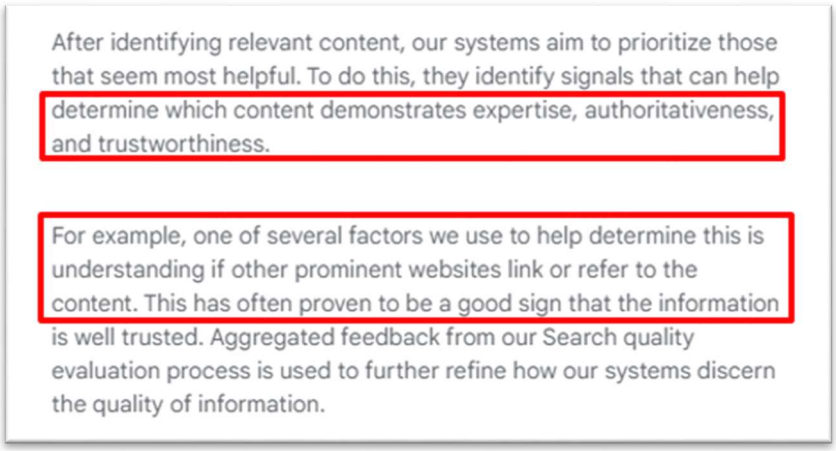
After identifying relevant content, our systems aim to prioritize those that seem most helpful. To do this, they identify signals that can help determine which content demonstrates expertise, authoritativeness, and trustworthiness.

For example, one of several factors we use to help determine this is understanding if other prominent websites link or refer to the content. This has often proven to be a good sign that the information is well trusted. Aggregated feedback from our Search quality evaluation process is used to further refine how our systems discern the quality of information.

156. Google Search practices a method comprising ranking the expert documents in accordance with the search query.

157. For example, Google Search "ranking systems" use algorithms to evaluate results received in response to a user's query in order "to present the most relevant, useful results." *See* How Search works – Ranking results[16] (red rectangular annotations added):

[15] https://www.google.com/search/howsearchworks/how-search-works/ranking-results/#quality
[16] https://www.google.com/search/howsearchworks/how-search-works/ranking-results/

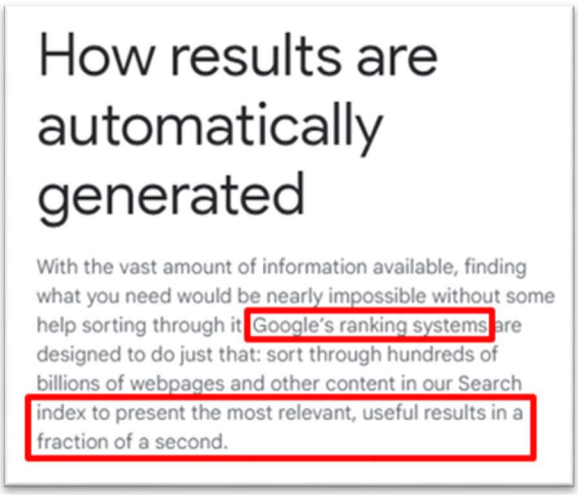

158. Google Search algorithms consider the "expertise of sources," among other factors, in order to "prioritize those that seem most helpful." *Id.* (red rectangular annotations added):

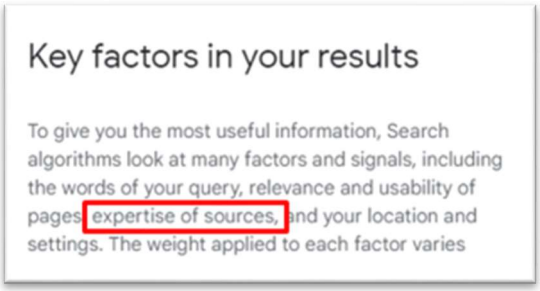

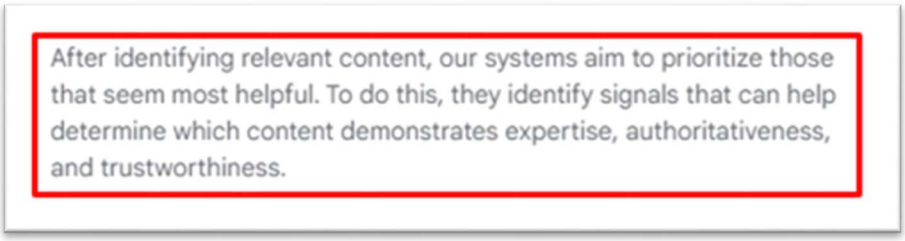

159. In order to prioritize those documents that "seem most helpful" or relevant to a user's query, Google Search algorithms rank the expert documents described above in accordance with a user's search query. As described in further detail below, Google Search prioritizes target documents pointed to by the ranked expert documents and presents them in an order based in part on the ranked expert documents.

160. Google Search practices a method comprising ranking target documents pointed to by the ranked expert documents.

161. For example, Google Search crawlers automatically visit webpages as well as hyperlinks on those pages. Thus, when an expert document is crawled and ranked, Google Search algorithms also crawl and rank target documents to which the expert document links. *See, e.g.*, How Search works – Organizing information (red rectangular annotation added):

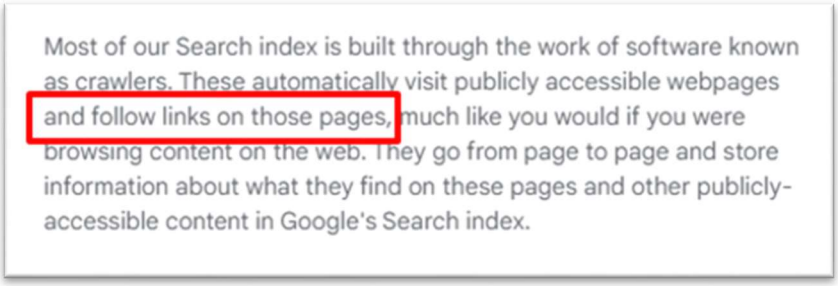

162. Google Search practices a method comprising returning a results list based on the ranked target documents.

163. For example, when a user submits a query to Google Search, it returns a results list. This list presents what Google Search has determined to be the most relevant or useful results based at least in part on whether its algorithms have identified a webpage as being a target document linked to by an expert document. Target documents that are linked to by authoritative expert webpages are ranked higher than other webpages, and thus appear in a better position on the results list shown to the user. For example, target sites linked to by highly relevant websites, top directories, or .edu pages are ranked higher than those without such links.

164. Furthermore, the '604 accused product, Google Search, embodies every limitation of claim 2 of the '604 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

165. Google Search practices the computer-implemented method of claim 1, wherein the hypertext documents are pages in the world wide web.

166. For example, the hypertext documents crawled by Google Search (from which a set of expert documents is formed) are webpages, which are pages in the world wide web. *See* How Search works – Organizing information (red rectangular annotations added):

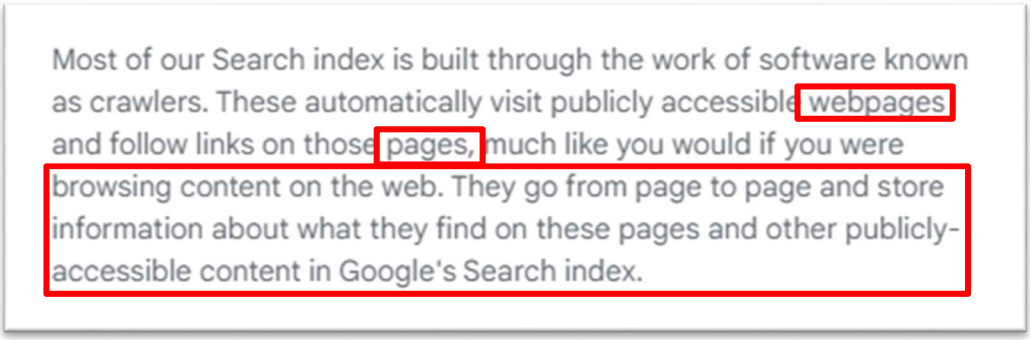

167. Furthermore, the '604 accused product, Google Search, embodies every limitation of claim 3 of the '604 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

168. Google Search practices the computer-implemented method of claim 1, wherein the hypertext documents are sites in the world wide web.

169. For example, when a hypertext document is crawled as described above with respect to claim 1, Google Search crawlers also crawl and index other websites which are linked to the hypertext document. Websites are sites in the world wide web. *See, e.g.*, Quality of content (red rectangular annotation added)[17]:

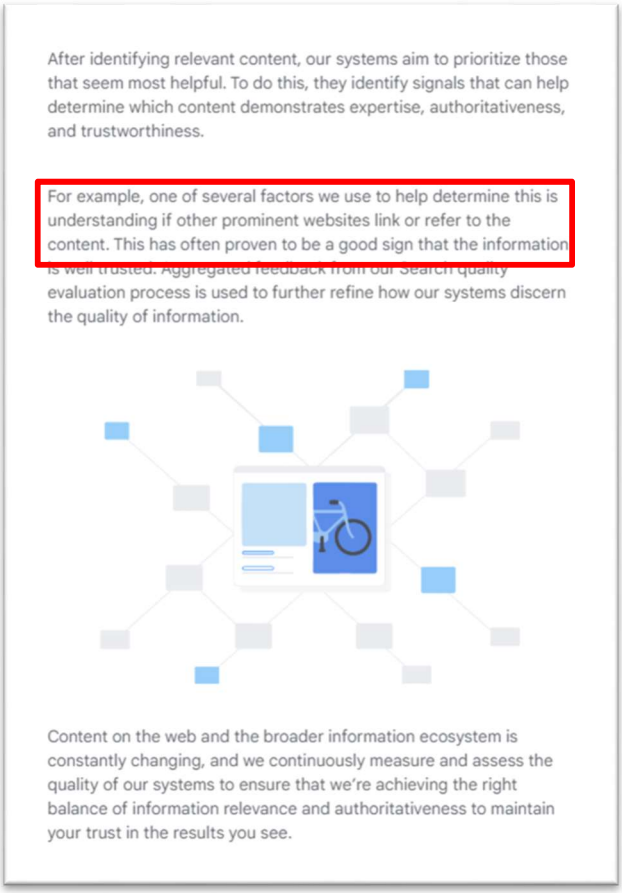

170. Furthermore, the '604 accused product, Google Search, embodies every limitation of claim 20 of the '604 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

171. Google Search practices the computer-implemented method of claim 1, forming a set of expert documents occurs before a search query is received.

[17] https://www.google.com/search/howsearchworks/how-search-works/ranking-results/#quality

172. For example, Google Search uses an algorithm to build a special index of expert documents from the set of all webpages that crawlers have indexed. This process occurs separately from any particular query submitted by a Google Search user. Google processes a search engine's database pages and identifies a subset of pages as good sources of unknown links for specific topics. For example, considering all pages with out-degree greater than a threshold X, Google tests to determine if those URLs point to X distinct hosts that are not affiliated. Each is considered an expert page. An inverted index then maps keywords to the experts in which they appear and a list is structured of positions that match within the experts. The matches map to the appearance of keywords within a key phrase of an expert page. Google algorithms first select the expert pages and identify the highest quality sites to rank. This occurs in a separate step, before a query is received.

173. Google has had knowledge of the '604 Patent since at least June 26, 2012. On this date, U.S. Patent No. 8,209,339, which cites the '604 Patent as a reference, was issued with Google Inc. as assignee. The '604 Patent is also cited as a reference by U.S. Patent No. 8,849,817 (issued September 30, 2014), which names Google Inc. as assignee. Furthermore, one of the co-inventors of the '604 Patent, Krishna A. Bharat, is listed as an inventor on two other patents which are assigned to Google: U.S. Patent Nos. 7,359,894 (issued April 15, 2008) and 9,037,575 (issued May 19, 2015). U.S. Patent No. 9,037,575 relates to the ranking of news articles. Bharat, who has been described as a "Google veteran" and "the inventor of Google News," has spent nearly two decades working for Google during the time since the filing of the '604 Patent application.[18] Given that two of its own patents cite the '604 Patent as a reference, and that one of the co-inventors of the

[18] https://www.cnbc.com/2019/08/22/google-news-inventor-krishna-bharat-returns-after-four-year-hiatus.html

'604 Patent has since been extensively involved with the development of Google products, Google knew of or was willfully blind to the fact that its conduct was infringing by, at the very latest, June 2012.

174. Plaintiffs are informed and believe, and thereon allege, that Google actively, knowingly, and intentionally has induced infringement of the '604 Patent by, for example, offering for public use Google Search with the intent to encourage and facilitate infringing uses of that service in the Northern District of Texas, in the United States, and throughout the world.

175. As a result of Google's infringement of the '604 Patent, Plaintiffs have been damaged. Plaintiffs are entitled to recover damages sustained as a result of Google's wrongful acts in an amount subject to proof at trial.

176. In addition, Google's infringing acts and practices have caused immediate and irreparable harm to Plaintiffs.

177. Plaintiffs are informed and believe, and thereon allege, that Google's infringement of the '604 Patent has been willful. As noted above, Google has had knowledge of the '604 Patent and its infringement of the '604 Patent. Google has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for Plaintiffs' patent rights. Thus, Google's infringing actions have been consciously wrongful.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs pray for judgment against Google as follows:

A. That Google has infringed each of the Asserted Patents, and unless enjoined will continue to infringe one or more of the Asserted Patents;

B. That Google has willfully infringed one or more of the Asserted Patents;

C. That Google pay Plaintiffs damages adequate to compensate Plaintiffs for Google's past infringement of each of the Asserted Patents, and present and future infringement of applicable Asserted Patents, together with interest and costs under 35 U.S.C. § 284;

D. That Google be ordered to pay prejudgment and post-judgment interest on the damages assessed;

E. That Google pay Plaintiffs enhanced damages pursuant to 35 U.S.C. § 284;

F. That Google be ordered to pay supplemental damages to Plaintiffs, including interest, with an accounting, as needed;

G. That Google be enjoined from infringing the applicable Asserted Patents, or if its infringement is not enjoined, that Google be ordered to pay ongoing royalties to Plaintiffs for any post-judgment infringement of the Asserted Patents;

H. That this is an exceptional case under 35 U.S.C. § 285, and that Google pay Plaintiffs' attorneys' fees and costs in this action; and

I. That Plaintiffs be awarded such other and further relief, including equitable relief, as this Court deems just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), Plaintiffs hereby demand a trial by jury on all issues triable to a jury.

July 12, 2024

Respectfully submitted,

By: */s/ Jon B. Hyland*

Jon B. Hyland (State Bar No. 24046131)
HILGERS GRABEN, PLLC
7859 Walnut Hill Lane, Suite 335
Dallas, Texas 75230
Phone: (972) 645-7728
Fax: (402) 413-1880
E-mail: jhyland@hilgersgraben.com

Jason Sheasby (*pro hac vice* forthcoming)
Amy E. Proctor (*pro hac vice* forthcoming)
Connor He-Schaefer (*pro hac vice* forthcoming)
IRELL & MANELLA LLP
1800 Avenue of the Stars, Suite 900
Los Angeles, California 90067
Telephone: (310) 277-1010
E-mail: jsheasby@irell.com
E-mail: aproctor@irell.com
E-mail: che-schaefer@irell.com

COUNSEL FOR PLAINTIFFS