

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 February 2010 (18.02.2010)

(10) International Publication Number
WO 2010/018584 A1

(51) International Patent Classification:

G06Q 30/00 (2006.01)

(21) International Application Number:

PCT/IL2009/000802

(22) International Filing Date:

13 August 2009 (13.08.2009)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/088,728 14 August 2008 (14.08.2008) US

(71) Applicant (for all designated States except US):
CHECKM8 INC. [US/US]; Suit 2003, 555 8th Avenue, New York, New York 10018 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BEN-DOR, Gilad** [IL/IL]; 71 Menachem Begin Road, 49732 Petach Tikva (IL). **NEUMAN, Ron** [IL/IL]; 23 Haetzet Street, 74014 Nes Ziona (IL).

(74) Agent: **APPELFELD ZER FISHER**; B.S.R. Tower 1, Floor 16, 2 Ben-Gurion Road, 52573 Ramat-Gan (IL).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

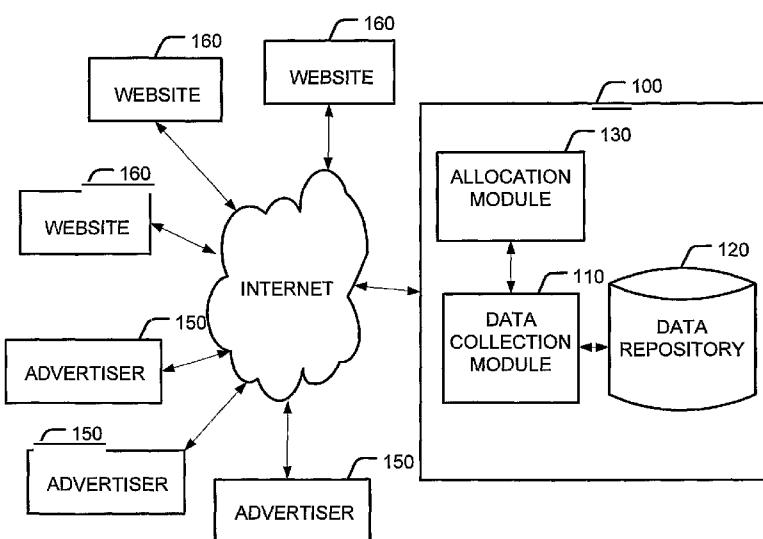
AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: INTERNET BASED ADVERTISEMENT INVENTORY FORECASTING AND ALLOCATION



(57) Abstract: An ad server and a method for forecasting and managing an Internet based advertisements inventory. Specifically, the data collection process required for forecasting the inventory utilizes a tabular table-like data structure upon which data is aggregated in real-time and stored periodically. Ads allocation in turn takes a "divide and conquer" approach in which campaign ordering derives from the categories and sizes of the ad units on the inventory, as well as their contextual affiliation.

FIG. 2

ALLOCATION

TECHNICAL FIELD

5 [0001] The present invention relates to the field of Internet advertisement, and more particularly, to the field of internet based advertisements inventory management.

BACKGROUND OF THE INVENTION

[0002] Internet advertising requires effective and optimized resource management.

10 The websites in which advertisements (ad units) appear have a limited number of items (ad types) and a limited quantity for each item. The total number of ad units and their associated ad type comprises an inventory of advertisements that necessitates optimal management. Optimal management is crucial due to timing constraints according to which some of the inventory items should be used before a

15 predetermined date or become obsolete.

[0003] FIG. 1 illustrates the concept of internet based advertisement inventory according to the prior art. A certain website **10** has a certain page **20** which is visited 100 times per day. Page **20** has a slot **30** for an advertisement, in the size of 728x90 pixels. Therefore, there are 100 units of the 728x90 item that can be sold every day, or

20 approximately 3000 per month. The number of units is often referred to as impressions.

[0004] In the case that website **10** adds a second ad unit **50** to page **20**, in the size of 160x600 pixels, website **10** doubled its inventory of ads. It now has 100 units of the 728x90 item and 100 units of the 160x600 item, or, as more commonly referred, 100 impressions of the 728x90 ad unit and 100 impressions of the 160x600 ad unit.

[0005] In the case that there is a second web page **40** in website **10** with merely a single 728x90 ad unit **60**, and second web page **40** receives 200 visits per day. In this case, second page **40** has 200 impressions of the 728x90 and first page **20** has 100 impressions of the 728x90 and 100 impressions of the 160x600. Altogether, website

30 **10** has 100 impressions of the 160x600 ad unit and 300 impressions of the 728x90 ad unit.

[0006] The challenge in managing an internet based advertisement inventory such as the one described above is to sell each and every ad unit of any as type on the corresponding websites preferably at the highest possible price. Clearly, if some ad

units are not sold on a certain website an opportunity loss occurs. On the other hand, if an ad unit is oversold (e.g., for more than one advertiser) there is a need to compensate one of the advertiser.

[0007] Therefore, there is a need in the Internet based advertisement industry to

5 forecast and manage each website's advertising inventory correctly. Forecasting and inventory management comprise two parts: forecasting the potential inventory based on historical site traffic and visitation patterns; and allocating the booked campaigns to the potential inventory.

[0008] Some of the challenges in regards with forecasting and allocating the

10 inventory are unique to Internet based inventory forecasting and management. For example, Inventory potential cannot be controlled. Unlike other types of inventory management environments such as TV advertising (based on TV shows historical rating) or airline inventory, the number of ad units or impressions that is expected in the future cannot be controlled since they are directly related to the number of visits to
15 the pages. The number of visits to the pages can be based on many objective factors. For example, a breaking news story can drive many visits to the page and cannot be anticipated.

[0009] Another example is the use of multiple advertising currencies. Internet

20 advertising may be sold in many forms - based on the number of impressions served, based on the number of days the advertisement runs, based on the number of clicks, based on the number of unique users that viewed the ad or based on other actions that are performed by users as a result of viewing the ad (such as signing up to the advertiser's website or requesting more info form the advertiser).

[0010] Another example for the unique nature of Internet based advertisements is that

25 it may be targeted to certain sections of the website (sports section, travel section) or certain user characteristics such as demographic information (the user's age, gender, and the like), the user's online behavior (e.g., visited shopping comparison sites and compared prices of new laptops), and others. Incidentally, audience types may overlap such that a certain user may have compared prices for new laptops, and is currently
30 reading the sports articles and is also a male. This information adds another level of complexity to the allocation of inventory.

[0011] Finally, Internet based advertising inventory forecasting and management is

characterized in that it deals with an extremely high amount of data. The amount of historical data is processed and used to forecast future traffic and visitation and as a

result inventory potential can be vast. Since inventory forecasting is a real-time application, with the sales representative speaking to the advertiser and trying to figure out how much inventory can be sold to this client, it is impractical for the inventory algorithms to process all the historical data. Prior art inventory management systems solve the computational complexity problem using a short history timeframe and by sampling the data within that timeframe. This in turn leads to large sampling errors that skew the forecast, or a small control group that may not be representative of the website's overall inventory.

10 BRIEF SUMMARY

[0012] Accordingly, it is a principal object of the present invention to overcome the disadvantages of the prior art. This is provided in the present invention by an ad server arranged to forecast and manage an Internet based advertisements inventory. Specifically, the data collection process required for forecasting the inventory utilizes a tabular table-like data structure upon which data is aggregated in real-time and stored periodically. Ads allocation in turn takes a "divide and conquer" approach in which the order of allocating ad units to websites derives from their size, number and contextual affiliation.

20 BRIEF DESCRIPTION OF THE DRAWINGS

[0013] For a better understanding of the invention and to show how the same may be carried into effect, reference will now be made, purely by way of example, to the accompanying drawings in which like numerals designate corresponding elements or sections throughout.

[0014] With specific reference now to the drawings in detail, it is stressed that the particulars shown are by way of example and for purposes of illustrative discussion of the preferred embodiments of the present invention only, and are presented in the cause of providing what is believed to be the most useful and readily understood description of the principles and conceptual aspects of the invention. In this regard, no attempt is made to show structural details of the invention in more detail than is necessary for a fundamental understanding of the invention, the description taken with the drawings making apparent to those skilled in the art how the several forms of the invention may be embodied in practice.

[0015] In the accompanying drawings:

[0016] FIG. 1 is a high level schematic block diagram of web pages within a websites exhibiting ad units according to the prior art;

[0017] FIG. 2 is a high level schematic block diagram showing the ad server according to the present invention;

5 [0018] FIG. 3 is a high level flow chart showing a method according to some embodiments of the invention;

[0019] FIG. 4 is a high level flow chart showing a method according to some embodiments of the invention; and

10 [0020] FIGS. 5A-5F are high level schematic block diagrams illustrating an exemplary implementation of the method according to the present invention.

[0021] The drawings together with the description make apparent to those skilled in the art how the invention may be embodied in practice.

DETAILED DESCRIPTION OF THE INVENTION

15 [0022] Prior to setting forth embodiments of the invention in detail, it may be helpful to set forth definitions of certain terms that will be used hereinafter throughout the rest of the application.

[0023] The term "Ad Unit" or "Ad Slot" as used herein in this application, is defined as a particular unit size in a certain location in the page which is created to hold an ad.

20 [0024] The term "Impression" as used herein in this application, is defined as a single delivery of an ad to a web page to be seen by a user.

[0025] The term "Campaign" as used herein in this application, is defined as an order to serve advertisement(s). A campaign has a specific start and end date, one or more categories assigned to it, one or more formats assigned to it and an impression goal.

25 [0026] The term "Category" as used herein in this application, is defined as a collection of pages that can be grouped together based on a common denominator, usually by being all within context of a certain subject, such as *sports*, or more specifically *basketball*. When advertising campaigns are scheduled, they are scheduled to particular categories and not to individual pages.

30 [0027] The term "Category Cluster" as used herein in this application, is defined as a collection of one or more categories that all share the exact same campaigns that are scheduled to them within a certain time period.

[0028] Embodiments of the present invention provide an ad server arranged to forecast and manage an Internet based advertisements inventory. Specifically, the data

collection process required for forecasting the inventory utilizes a tabular table-like data structure upon which data is aggregated in real-time and stored periodically. Ads allocation in turn takes a "divide and conquer" approach in which campaign ordering derives from the categories and sizes of the ad units on the inventory, as well as their 5 contextual affiliation.

[0029] Before explaining at least one embodiment of the invention in detail, it is to be understood that the invention is not limited in its application to the details of construction and the arrangement of the components set forth in the following description or illustrated in the drawings. The invention is applicable to other 10 embodiments or of being practiced or carried out in various ways. Also, it is to be understood that the phraseology and terminology employed herein is for the purpose of description and should not be regarded as limiting.

[0030] FIG. 2 is a schematic block diagram of the ad server according to the present invention. Ad server 100 is operatively associated with advertisers 150 and websites 15 160 via the Internet 140. Ad server 100 comprises a data collection module 110, a data repository 120 and an ad allocating module 130 all operatively associated therebetween.

[0031] In operation, data collection module 110 in cooperation with data repository 120 are arranged to analyze the historical inventory of each website ad server 100 is 20 associated therewith and further apply predefined operations on the historical data in order to forecast the future inventory potential.

[0032] following are Internet based advertisements inventory properties that are considered by data collection module 110 while analyzing the data from the websites:

[0033] Extent of historical data to be used - in order for the forecast to be accurate, it 25 is imperative to analyze as much historical data as possible, which could result in terabytes of historical log files that need to be analyzed.

[0034] Sampling errors - if sampling is used in order to reduce the amount of historical data that needs to be analyzed, the result set becomes prone to sampling errors which can become very significant in some cases.

[0035] Real-time availability - inventory potential estimation is part of the website's 30 sales process in which sales representatives need to be able to respond in real-time to inventory availability inquiries coming from clients. This increases the complexity of the problem as the results of the forecast need to be immediate and up-to-date at all times.

[0036] Seasonality - many websites experience seasonal attributes that affect their traffic patterns (for example, e-commerce websites always get significantly more user traffic on the weeks before Christmas). The algorithm needs to be able to deal with seasonality effects.

5 [0037] Trends - in addition to seasonality, websites experience growth trends year over year and quarter over quarter that need to be taken into consideration when creating longer term forecasts.

[0038] Special Events - there are special events that can be planned for in advance and affect the traffic on top of seasonality and trends. For example, the Olympics

10 10 games that take place every four years will increase traffic to sports sites, therefore increasing potential inventory.

[0039] Spikes and valleys - these are abnormal historical traffic trends that do not repeat themselves in an expected manner. For example, a breaking news item could drive lots of traffic to a site on a certain date and create an abnormal traffic spike. This

15 15 is not a seasonality event (it does not repeat itself in an expected manner) neither a trend. However, if these results are included in the historical results that are used to forecast potential inventory, this abnormal traffic can skew the forecast. Therefore, an inventory potential calculation must know how to clean out abnormal traffic spikes from the historical data used for forecasting.

20 [0040] According to some embodiments of the invention, data collection module 110 in cooperation with data repository 120 are arranged to process a large amount of historical data, thus increasing the accuracy of the calculations significantly. The data collected is arranged within data repository 120 in a form of a tabular table-like data structure.

25 [0041] While collecting data for inventory potential calculation purposes, at least some of the following parameters are used in constructing the data repository 120:

[0042] Page views - the total number of times that a page belonging to the category was viewed by the user within the time period.

[0043] Impressions (optional) - the total number of times that an ad was delivered on

30 30 a page belonging to the category within the time period. This is collected separately for each individual ad unit type. Collecting impression data is optional since they may also be calculated later as will be described later on in this document.

[0044] Unique users - the total numbers of users that viewed a page belonging to the category within the time period (unique users are page views if and only if each user visited only once within the time period).

5 [0045] Attributes - includes many different types of targeting attributes such as browser type, bandwidth level, age, country, state, city, and the like.

[0046] According to some embodiments of the invention, page views, impressions and unique users may be collected and stored on an hourly basis for every day of the week and for every category in the category hierarchy, including leaf categories, parent categories and root categories (other resolutions may apply).

10 [0047] The weekly data may be collected and stored in one of two methods: The first method is a "Model Week" (or any other time period) - always keep a week's worth of data. The new data is always averaged with the old one in a form of a weighted average. For example, when data for the most recent Monday is collected, the Monday of the model week is recalculated based on weighted average. The other 15 method is called "detailed historic model" - in this model, many weeks of data are kept separately to allow more precision and to factor seasonality effects into the forecasting.

[0048] Following are suggested considerations for applying the various aforementioned parameters while implementing the aforementioned data structure:

20 [0049] Attributes - these may be collected and stored on an hourly basis for every day of the week and for every category in the category hierarchy as well. However, from precision/performance considerations, it is usually recommended to store attributes with less granularity. For example, each attribute may have a single daily value (not broken down by hour), a single weekly value (not broken down by days and hours), or even kept just for the root category of the category tree through a weighted average of sub-categories. This will depend on the implementation and the 25 precision and performance requirements of the application.

[0050] Trends - in a similar manner to attributes, trends may be collected and stored for every category in the category hierarchy, and hourly/weekly. However, from 30 precision/performance considerations, it may be enough to store a global value per category.

[0051] Seasonality - seasonality can be stored on a daily or weekly basis factor per category.

[0052] FIG. 3 is a high level flow chart showing an inventory forecasting method according to some embodiments of the invention. Calculating the potential inventory is performed in the following manner which, in a non limiting manner, includes the essential steps required. The user initiates the method by creating an inventory query 5 that exhibiting start date, end date, channel/section, and ad format (The query may be translated into the practical question: "how much potential inventory do we have on this section, for this ad format, on these days"); Responsive to the query, extracting historical data- such that the query dates are resolved into days, and the historical page view data for these dates on that specific category are extracted from a 10 predefined model week 310 (For example, if the dates resolve to Monday and Tuesday, the historical page view data for the average Monday and average Tuesday is extracted); based on the expected dates, a seasonality factor may be applied to each day individually 320; applying trend factor based on the expected dates 330 (For example, if the start date is 3 months ahead and the monthly growth trend is 3%, a 3% 15 growth per each month is applied to each individual day - results is multiplies by 1.03 to the power of 3; applying special event factor in the case that the user had entered a special event that falls on the dates 340.

[0053] According to some embodiments of the invention, the product of the aforementioned method is an inventory potential calculation, broken down by day and 20 category. Results may be added up to show potential inventory across the entire time period on each category, daily on all categories, or for the entire time period across all categories combined.

[0054] According to some embodiments of the invention, a further parameter may be used - "Targeting attributes" Targeting attributes shall be used when querying 25 potential inventory for certain audience groups, for example, only users that are female or only users from the UK. The attribute factor will be applied to the potential calculation to yield the result. In case of multiple attributes that may intersect with each other, more advanced implementation methods exist.

[0055] According to some embodiments of the invention, in order to further control 30 the performance of the forecasting process, additional factors may be considered for either improving the performance of the method and its precision. There is a tradeoff between performance and precision, such that if these methods will be used to improve the performance of the algorithm, it will in fact degrade from its precision and vice versa, if the methods are employed to improve the precision of the algorithm,

there will be performance degradation. The use of these methods for performance or precision should be considered based on performance and precision requirements of the user based on the specific use scenarios of it.

[0056] Following are manners in which precision and performance may be controlled:

5 [0057] Resolution of data stored - data can be stored in various resolutions, such as minute, hour, day, week and month. Depending on the type of data, the impact can be large or small.

[0058] Page views, impressions and unique users - to increase performance, this information can be collected on a daily resolution (instead of hourly) or only on top-10 level categories (instead of every category). Increasing precision can be done by collecting the daily on a single minute resolution; however in most implementations the impact on precision will be negligible.

[0059] Attributes - to increase precision, attributes can be collected on a daily or even hourly basis, although we believe impact on precision will be minimal. Alternatively, 15 data may be collected on each individual category or just top-level categories.

[0060] Trends, Special Days and Seasonality - it is recommended to collect these on an individual category level, however based on precision requirements level they can be collected on a top-level category only.

[0061] FIG. 4 is a high level flow chart showing an inventory allocation method 20 according to some embodiments of the invention. The inventory allocation method depends on the aforementioned forecasting method and comprises the following concepts:

[0062] Buckets - Viewing each category as a bucket that can contain a certain number of impressions. The size of the bucket is directly correlated to the number of potential 25 impressions on the category;

[0063] Cups - Breaking each campaign down to chunks of impressions (or cups). Each cup can hold a certain number of impressions; and

[0064] Filling the buckets with cups - optimal inventory allocation is achieved by the filling the buckets (categories) with impressions using the cups (impression chunks).

30 [0065] The allocation method comprising: (optional) reading the inventory database 410 (all the relevant campaign and category data can be read from the database and kept in the processing computer's memory cache); (Optional) creating cluster categories 420 (categories can be grouped into clusters, and the algorithm processing time can be reduced as the algorithm will run based on the category clusters and not

based on each individual category); for each hour - allocating inventory separately for every day and for every hour within the day 430 (in order to allocate inventory for a specific week, we will run the algorithm separately for each hour of each day of the week); for each priority class —allocating ad units based on campaign priority class, 5 starting with the highest priority class first (each campaign has a priority class assigned to it when it is inserted into the system, which gives it an absolute preference compared to other campaigns); calculating category weights 440 - every category is assigned a weight which is calculated for every category i (l..n categories) category weight (i) = category impressions (i) / sum (category impressions[l..n]); iterating over 10 categories until all campaign impressions are allocated to the categories in full 450 wherein the allocation method goes through all categories one-time randomly, and then goes through all categories one-time randomly again and again until it finishes allocating all impressions. Randomization is based on the weights previously assigned to the categories (category weights), in a manner in which the weight of the category 15 is the respective probability for it to be randomly chosen. Thus, if the first category has a weight of 50 and the second category has a weight of 100, the probability that the second category will be selected is twice as much as the probability that the first category will be selected. The next stage is calculating campaign weights 460 such that every campaign is assigned an absolute weight which is a combination of the 20 weight assigned to it by the user, its level of starvation (a campaign is considered in starvation if it has served less than impressions that expected) and its unit price, and possibly additional parameters. The average campaign weight is then calculated as follows: for every campaign i (l..n campaigns) , average campaign weight (campaign i) = campaign weight (i) / sum (campaign weights[l..n]). The next stage is selecting 25 single campaign 470 so that a single campaign is selected from all the campaigns targeted to the category. The campaign is selected randomly, but is selected based on the campaign weights assigned previously (campaign weights), in a manner in which the weight of the campaign is the respective probability for it to be randomly chosen. So if the first campaign has a weight of 50 and the second campaign has a weight of 30 100, the probability that the second campaign will be selected is twice as much as the probability that the first campaign will be selected. The next stage is allocating impression chunk so that impressions of the selected campaign are allocated to the 35 selected category 480 in the following manner: a chunk of the campaign's impressions is allocated to the category based on the *effective chunk size*, and the campaign's ..

impressions and category's available impressions are reduced accordingly. The *effective chunk size* shall be equal to the *defined chunk size* (we will describe an implementation later in which the *effective chunk size* can be equal or smaller to the *defined chunk size*). For example, if the chunk size if 100 impressions, the campaign

5 hourly impressions are 500 and the category available hourly impressions are 1800, then the campaign hourly impressions shall be reduced to 400, and the category available hourly impressions shall be reduced to 1700. In the method's next iteration, the campaign weights and category weights are recalculated to account for the impression allocation performed in the previous iteration **490**.

10 [0066] FIGS. **5A-5F** are high level block diagram illustrating categories and campaigns as manipulated by the allocation method. In a non limiting example, FIG. 5A shows the status prior to the inventory allocation, before any impressions have been allocated from the campaigns to the categories. There are provided three categories **510**, 520, 530, having 100, 50 and 150 available impressions each, respectively. There are further provided two campaigns, **540** and **550** having 50 and 15
15 impressions each, respectively (all within a 1 hour allotment). The chunk size is 25 impressions.

[0067] FIG. 5B shows the status after a first iteration of the method, in the case that **530** is randomly selected to be filled (since it has the highest probability) and **550** is randomly selected to fill. 25 impressions of **550** are then allocated to **530**.

[0068] FIG. 5C shows the status after the second iteration of the method, in which **530** is randomly selected again and **550** is selected.

[0069] FIG. 5D shows the status after the third iteration of the method. The odds of selecting **510** and **530** are now equal. So are the odds of selecting **540** or **550**. In this case **510** was selected and **540** was selected.

[0070] FIG. SE shows the status after the fourth iteration of the method. Assuming **530** is selected again and **540** is selected. Now **540** is fully allocated.

[0071] FIG. 5F shows the status after the fifth iteration, assuming **510** is selected.

[0072] Finally, **530** is once again chosen as it still has the highest probability. All campaign impressions are now fully allocated and the algorithm terminates.

[0073] According to some embodiments of the invention, a category may have multiple ad units (ad sizes and location) assigned to it. For example, a category may have a first 728x90 ad unit on the top of the category's pages, a second 160x600 ad unit on the right of the category's pages and a third 728x90 ad unit on the bottom of

the category's pages. In that case, the inventory allocation shall be done separately for each ad unit. The inventory has inherent support for multiple ad units per category in step 2 of the algorithm. In the case of multiple ad units, step 2 of the algorithm shall be revised so that it loops through every combination of category and ad unit.

5 [0074] According to some embodiments of the invention, in more complex real-life scenarios, the allocation of the chunk to the categories described in step 5 of the basic implementation of the algorithm, can take a more complex implementation when certain limitations are imposed on the campaign or the category. These limitations include but are not limited to targeting criteria such as demographic, geographic and
10 other. In the event that one or more of these limitations exist, the chunk size is reduced to an effective chunk size by applying the limitations to the chunk size in various different ways, depending on the limitation.

[0075] According to some embodiments of the invention, the method defined above is based on all campaigns being ordered on an impression basis, i.e. a certain number of
15 impressions are requested for the campaign. In more complex real-life scenarios, multiple campaigns can exist in multiple advertising currencies - impressions, clicks, days, acquisitions, etc. In order to apply the above method in this scenario, an additional setup step is required to convert all campaigns into corresponding impression-based campaigns, which shall be called virtual campaigns. The algorithm
20 then uses the virtual campaigns instead of the actual campaigns. The conversion can be done in various manners, depending on the currency of the campaign:

[0076] Campaigns ordered by clicks (CPC) - to convert CPC campaigns to impression-based campaigns (CPM), an "exchange rate" is determined between a click and an impression. The exchange rate is determined by dividing the number of
25 clicks generated in the campaign to date by the number of impressions delivered in the campaign to date. The exchange rate is then applied to the number of clicks to yield the number of impressions. For example, if a certain campaign had been ordered with 25,000 clicks, and to this date has served 100,000 impressions and generated 200 clicks, then the CPC to CPM exchange rate shall equal $200 / 100,000 = 500$. The
30 exchange rate is then multiplied by the number of clicks to yield the number of impressions in virtual campaign in the following manner: $25,000 \times 500 = 12,500,000$ impressions.

[0077] Campaigns ordered by CPA or other action-based models such as CPL - an exchange rate is applied as defined above between the model and CPM to yield the respective number of impressions for the virtual campaign.

5 [0078] Following is a more detailed description of the method of allocation inventory according to some embodiments of the invention:

[0079] Method setup according to some embodiments of the invention comprises:

[0080] Reading database and creating data structure (optional) - to improve algorithm performance, all the relevant campaign and category data can be read from the database and kept in the processing computer's memory cache. This step is optional

10 but is recommended for better performance.

[0081] Creating cluster categories (optional) - to improve algorithm performance, an optional process is recommended, in which categories can be grouped into clusters, and the algorithm processing time can be reduced as the algorithm will run based on the category clusters and not based on each individual category. This is an optional

15 step and will be described in detail in a more advanced implementation of the algorithm.

[0082] Chunk size - the chunk size can be a constant size. However, more advanced implementations of the chunk size can assign a variable size to it. The advanced implementation will be described later in this document.

20 [0083] Generating Virtual campaigns - streamlining all campaigns under the same currency (impressions). Impression-based campaigns are translated as is to virtual campaigns. Non-impression based campaigns are translated to virtual campaigns based on the exchange rates described above.

25 [0084] The outer loop of the method according to some embodiments of the invention comprises:

[0085] Loop 1: For each hour - To resolve the inventory allocation problem, we first take the divide and conquer approach and break down the problem to daily and hourly basis, meaning we allocate inventory separately for every day and for every hour within the day. To allocate inventory for a specific week, we will run the algorithm

30 separately for each hour of each day of the week.

[0086] Loop 2: For each priority class - As a next step, we start the allocations based on campaign priority class, starting with the highest priority class first. Each campaign has a priority class assigned to it when it is inserted into the system, which gives it an absolute preference compared to other campaigns. In other words, if the

first campaign belongs to the "standard" class and the second campaign belongs to the "high" class, then the second campaign will always have a preference over the first.

[0087] The main loop of the method according to some embodiments of the invention comprises:

5 [0088] Step 1: Calculate category weights - every category is assigned a weight which is calculated as follows: for every category i (1.. n categories) category weight $(i) = \text{category impressions } (i) / \text{sum (category impressions}[1..n]).$

[0089] Step 2: Loop over categories x ad units - The main part of the algorithm loops over each an every category and each and every ad unit assigned to the category until 10 all virtual campaign impressions are allocated to the categories and ad units in full.

The algorithm goes through all categories x ad units one-time randomly, and then goes through all categories x ad units one-time randomly again and again until it finishes allocating all impressions. The randomization is based on the weights previously assigned to the categories (category weights), in a manner in which the 15 weight of the category is the respective probability for it to be randomly chosen. So if the first category has a weight of 50 and the second category has a weight of 100, the probability that the second category will be selected is twice as much as the probability that the first category will be selected.

[0090] Step 3: Calculate virtual campaign weights - every virtual campaign is 20 assigned an absolute weight which is a combination of the weight assigned to it by the user, its level of starvation (a campaign is considered in starvation if it has served less than impressions that expected) and its unit price, and possibly additional parameters. The average virtual campaign weight is then calculated as follows: for every virtual campaign i (1.. n campaigns) average virtual campaign weight (campaign i) = virtual campaign weight $(i) / \text{sum (virtual campaign weights}[1 ..n]).$

[0091] Step 4: Select single virtual campaign - As a next step, a single virtual campaign is selected from all the campaigns targeted to the category. The campaign is selected randomly, but is selected based on the campaign weights assigned previously 30 (campaign weights), in a manner in which the weight of the campaign is the respective probability for it to be randomly chosen. So if the first campaign has a weight of 50 and the second campaign has a weight of 100, the probability that the second campaign will be selected is twice as much as the probability that the first campaign will be selected.

[0092] Step 5: Allocate impression chunk - As a final step in the algorithm, impressions of the selected virtual campaign are allocated to the selected category in the following manner: a chunk of the campaign's impressions are allocated to the category based on the *effective chunk size*, and the campaign's impressions and category's available impressions are reduced accordingly. For example, if the chunk size is 100 and the effective chunk size is 60 impressions, the campaign hourly impressions are 500 and the category available hourly impressions are 1800, then the campaign hourly impressions shall be reduced to 440, and the category available hourly impressions shall be reduced to 1740.

10 [0093] In the algorithm's next iteration, the campaign weights and category weights are recalculated to account for the impression allocation performed in the previous iteration.

[0094] When applying the algorithm either in its basic form or its advanced form, additional factors can be considered that can improve the performance of the 15 algorithm or improve its precision. There is a tradeoff between performance and precision, such that if these methods will be used to improve the performance of the algorithm, it will in fact degrade from its precision and vice versa, if the methods are employed to improve the precision of the algorithm, there will be performance degradation. The use of these methods for performance or precision should be 20 considered based on performance and precision requirements of the user based on the specific use scenarios of it.

[0095] According to some embodiments of the invention, categories may be grouped into clusters such as a cluster will contain one or more categories. The grouping is done in a way in which all categories that belong to the same cluster have the exact 25 same campaigns targeted to them, such as if campaign A is targeted to category 1 and category 1 and category 2 belong to the same category cluster, then campaign A is also targeted to category 2. When grouping into clusters, it reduces the number of iterations of the algorithm and therefore improves the speed/performance of the algorithm. Another version of this clustering can be that the clustering is applied only 30 to some of the categories, while the remaining categories remain un-clustered. This can be employed when precision and drill down capabilities are required on some categories but not on all of them. .

[0096] According to some embodiments of the invention, instead of using fixed chunk sizes as described above, the size of the chunks can be variable and determined

by the performance requirements versus precision. A larger chunk size will imply less iterations of the algorithm and thus speed up the performance of the algorithm. However, large chunk size can mean less optimized precision of the algorithm. The chunk size should be determined based on the average and min-max category sizes in

5 impressions and based on the user's precision/performance requirements.

[0097] In addition to the algorithm-based optimization described above (category clusters and variable chunk sizes), the following system-level optimizations and considerations exist:

[0098] Double buffering - this method allows the inventory to complete calculations 10 and hold the data structure in a first buffer in the computer's memory cache, and the snapshot is then ready to be queried by the user. At the same time, a new calculation of inventory begins in a second buffer in the computer memory. When the second calculation is complete, the buffers are swapped so that the user queries the freshest data, and then the calculation begins in the other buffer again.

15 [0099] Multiple computers - the inventory allocation task can be distributed between multiple computers such that each computer will have specific categories assigned to it and that computer will be the sole computer assigned to those categories.

[00100] Extent of allocation - the extent to which the allocation is performed 20 (allocation dates) is a major parameter in performance. If the inventory is required to be allocated for a 2-months period, it will require double the amount of iterations compared to a 1-month period. The extent depends on the user's particular implementation of the inventory algorithm.

[00101] The invention can be implemented advantageously in one or more 25 computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. A computer program is a set of instructions that can be used, directly or indirectly, in a computer to perform a certain activity or bring about a certain result. A computer program can be written in any form of programming 30 language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment.

[00102] Suitable processors for the execution of a program of instructions include, by way of example, digital signal processors (DSPs) but also general purpose

microprocessors, and the sole processor or one of multiple processors of any kind of computer. Generally, a processor will receive instructions and data from a read-only memory or a random access memory or both. The essential elements of a computer are a processor for executing instructions and one or more memories for storing instructions and data. Generally, a computer will also include, or be operatively coupled to communicate with, one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

5 [00103] In the above description, an embodiment is an example or implementation of the inventions. The various appearances of "one embodiment," "an embodiment" or "some embodiments" do not necessarily all refer to the same embodiments.

[00104] Although various features of the invention may be described in the context of a single embodiment, the features may also be provided separately or in any 10 suitable combination. Conversely, although the invention may be described herein in the context of separate embodiments for clarity, the invention may also be implemented in a single embodiment.

15 [00105] Reference in the specification to "some embodiments", "an embodiment", "one embodiment" or "other embodiments" means that a particular feature, structure, or characteristic described in connection with the embodiments is included in at least 20 some embodiments, but not necessarily all embodiments, of the inventions.

[00106] It is to be understood that the phraseology and terminology employed herein is not to be construed as limiting and are for descriptive purpose only.

25 [00107] The principles and uses of the teachings of the present invention may be better understood with reference to the accompanying description, figures and 30 examples.

[00108] It is to be understood that the details set forth herein do not construe a limitation to an application of the invention.

[00109] Furthermore, it is to be understood that the invention can be carried out or practiced in various ways and that the invention can be implemented in embodiments other than the ones outlined in the description above.

[00110] It is to be understood that the terms "including", "comprising",

5 "consisting" and grammatical variants thereof do not preclude the addition of one or more components, features, steps, or integers or groups thereof and that the terms are to be construed as specifying components, features, steps or integers.

[00111] If the specification or claims refer to "an additional" element, that does not preclude there being more than one of the additional element.

10 [00112] It is to be understood that where the claims or specification refer to "a" or "an" element, such reference is not to be construed that there is only one of that element.

[00113] It is to be understood that where the specification states that a component, feature, structure, or characteristic "may", "might", "can" or "could" be included, that particular component, feature, structure, or characteristic is not required to be included.

15 [00114] Where applicable, although state diagrams, flow diagrams or both may be used to describe embodiments, the invention is not limited to those diagrams or to the corresponding descriptions. For example, flow need not move through each illustrated box or state, or in exactly the same order as illustrated and described.

20 [00115] Methods of the present invention may be implemented by performing or completing manually, automatically, or a combination thereof, selected steps or tasks.

[00116] The term "method" may refer to manners, means, techniques and procedures for accomplishing a given task including, but not limited to, those manners, means, techniques and procedures either known to, or readily developed

25 from known manners, means, techniques and procedures by practitioners of the art to which the invention belongs.

[00117] The descriptions, examples, methods and materials presented in the claims and the specification are not to be construed as limiting but rather as illustrative only.

30 [00118] Meanings of technical and scientific terms used herein are to be commonly understood as by one of ordinary skill in the art to which the invention belongs, unless otherwise defined.

[00119] The present invention may be implemented in the testing or practice with methods and materials equivalent or similar to those described herein.

[00120] Any publications, including patents, patent applications and articles, referenced or mentioned in this specification are herein incorporated in their entirety into the specification, to the same extent as if each individual publication was specifically and individually indicated to be incorporated herein. In addition, citation

5 or identification of any reference in the description of some embodiments of the invention shall not be construed as an admission that such reference is available as prior art to the present invention.

[00121] While the invention has been described with respect to a limited number of embodiments, these should not be construed as limitations on the scope of the 10 invention, but rather as exemplifications of some of the preferred embodiments. Other possible variations, modifications, and applications are also within the scope of the invention. Accordingly, the scope of the invention should not be limited by what has thus far been described, but by the appended claims and their legal equivalents.

15

20

25

30

CLAIMS

1. a computer implemented method of allocating Internet based advertisements responsive of a plurality of campaigns exhibiting priorities and weights, the inventory exhibiting a data repository comprising ad units associated with impressions, arranged in a day by hour tabular form and further associated with categories exhibiting weights, the method comprising:
 - 5 for each hour, allocating inventory separately for every day and for every hour within the day;
 - 10 for each priority class, allocating ad units based on campaign priority class, starting with the highest priority class first;
 - calculating category weights;
 - iterating over categories until all campaign impressions are allocated to the categories in full in view of category weights;
 - 15 calculating campaign weights such that every campaign is assigned an absolute weight which is a combination of the weight assigned to it by the user, its level of starvation and its unit price;
 - selecting single campaign such that a single campaign is selected from all the campaigns targeted to the category in view of campaign weights;
 - 20 allocating impression chunk so that impressions of the selected campaign are allocated to the selected category;
 - recalculating the campaign weights and category weights to account for the impression allocation performed in the previous iteration.
2. The method of claim 1, wherein allocating inventory separately for every day and for every hour within the day is preceded by reading the inventory data repository.
- 25 3. The method of claim 1, wherein allocating inventory separately for every day and for every hour within the day is preceded by creating cluster categories.
3. The method of claim 1, wherein the data repository is created by a method comprising:
 - 30 responsive of a user query exhibiting start date, end date, channel/section, and ad format, extracting historical data such that the query dates are resolved into days, and the historical page view data for these dates on that specific category are extracted from a predefined weekly data model;
 - applying a seasonality factor based on the expected dates, on each day

applying trend factor based on the expected dates;

applying special event factor in the case that the user had entered a special event that falls on the dates.

5. The method of claim 4, wherein the weekly data model is a model week in which a week's worth of data is kept wherein new data is averaged with the old one in a form of a weighted average.

6. The method of claim 4, wherein the weekly data model is a detailed historic model in which a plurality of weeks of data is kept separately.

10 7. The method of claim 1, wherein allocating inventory separately for every day and for every hour within the day is preceded by converting all campaigns into impression based virtual campaigns using a predefined conversion rate.

8. An ad server for forecasting and managing an Internet based advertisements inventory, the server operatively associated with advertisers and websites exhibiting

15 ad slots for accommodating ad units, the server comprising:

a data collection module;

a data repository; and

an ad allocating module,

wherein the data collection module is arranged, responsive of a user query

20 exhibiting start date, end date, channel/section, and ad format and in cooperation with the data repository to:

extract historical data such that the query dates are resolved into days, and the historical page view data for the dates on that specific category are extracted from a predefined weekly data model;

25 apply a seasonality factor based on the expected dates, on each day individually;

apply trend factor based on the expected dates;

apply special event factor in the case that the user had entered a special event that falls on the dates;

and wherein the ad allocating module is arranged to:

30 for each hour, allocate inventory separately for every day and for every hour within the day;

for each priority class, allocate ad units based on campaign priority class, starting with the highest priority class first

calculating category weights;

iterate over categories until all campaign impressions are allocated to the categories in full;

calculate campaign weights such that every campaign is assigned an absolute weight which is a combination of the weight assigned to it by the user, its level of
5 starvation and its unit price;

select single campaign such that a single campaign is selected from all the campaigns targeted to the category;

allocate impression chunk so that impressions of the selected campaign are allocated to the selected category,

10 recalculate the campaign weights and category weights to account for the impression allocation performed in the previous iteration.

9. The server of claim 8, further arranged to allocate inventory separately for every day and for every hour within the day is preceded by reading the inventory data repository.

15 10. The server of claim 8, further arranged to allocate inventory separately for every day and for every hour within the day is preceded by creating cluster categories.

20

25

30

1 / 6

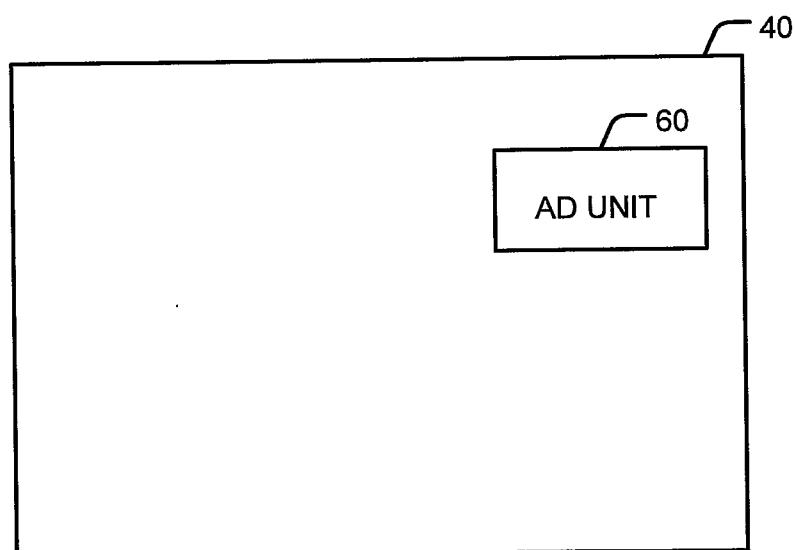
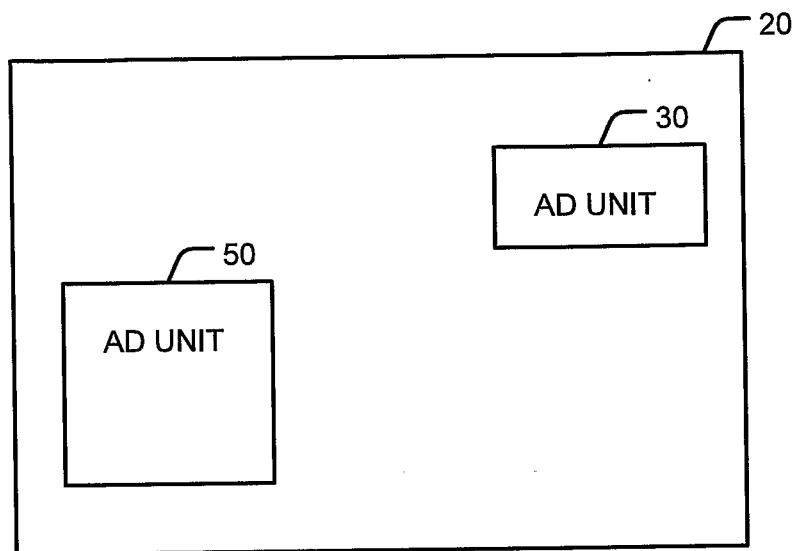


FIG. 1
(Prior Art)

2 / 6

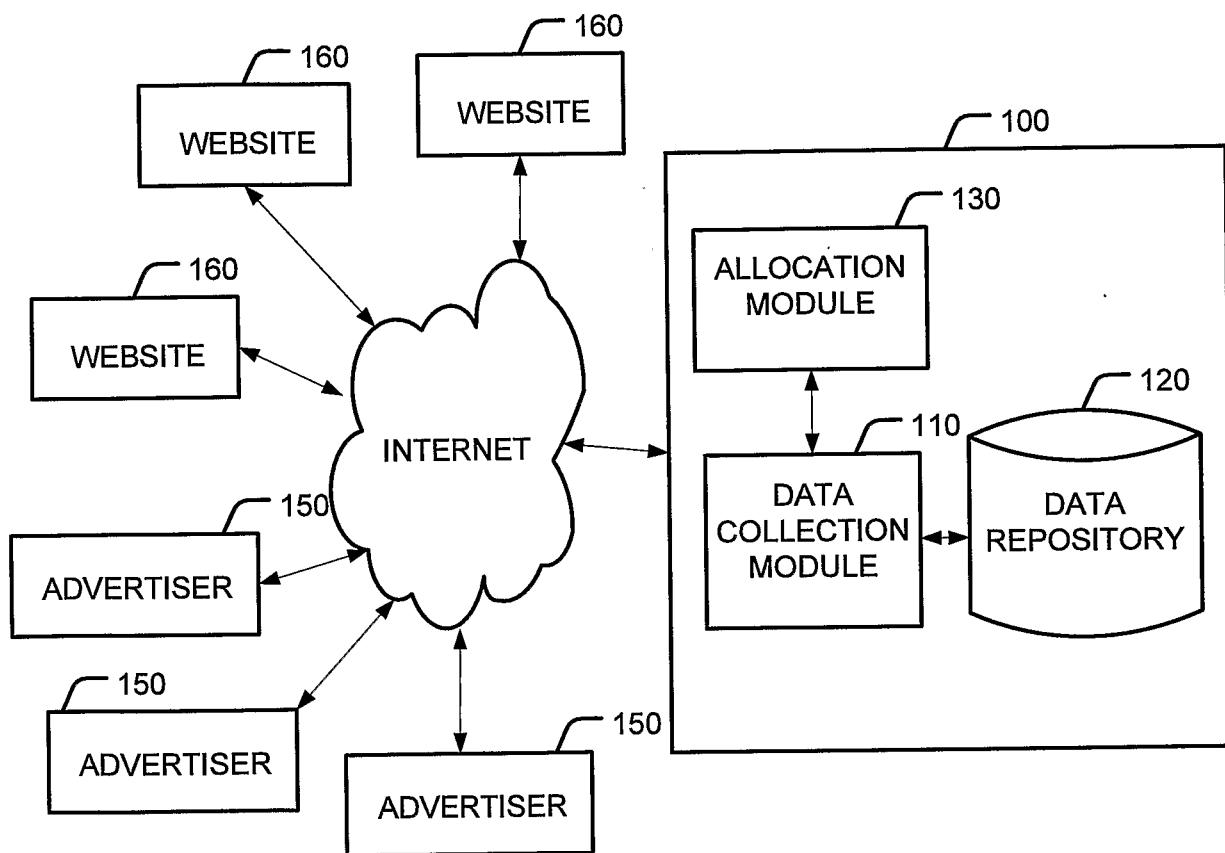


FIG. 2

3 / 6

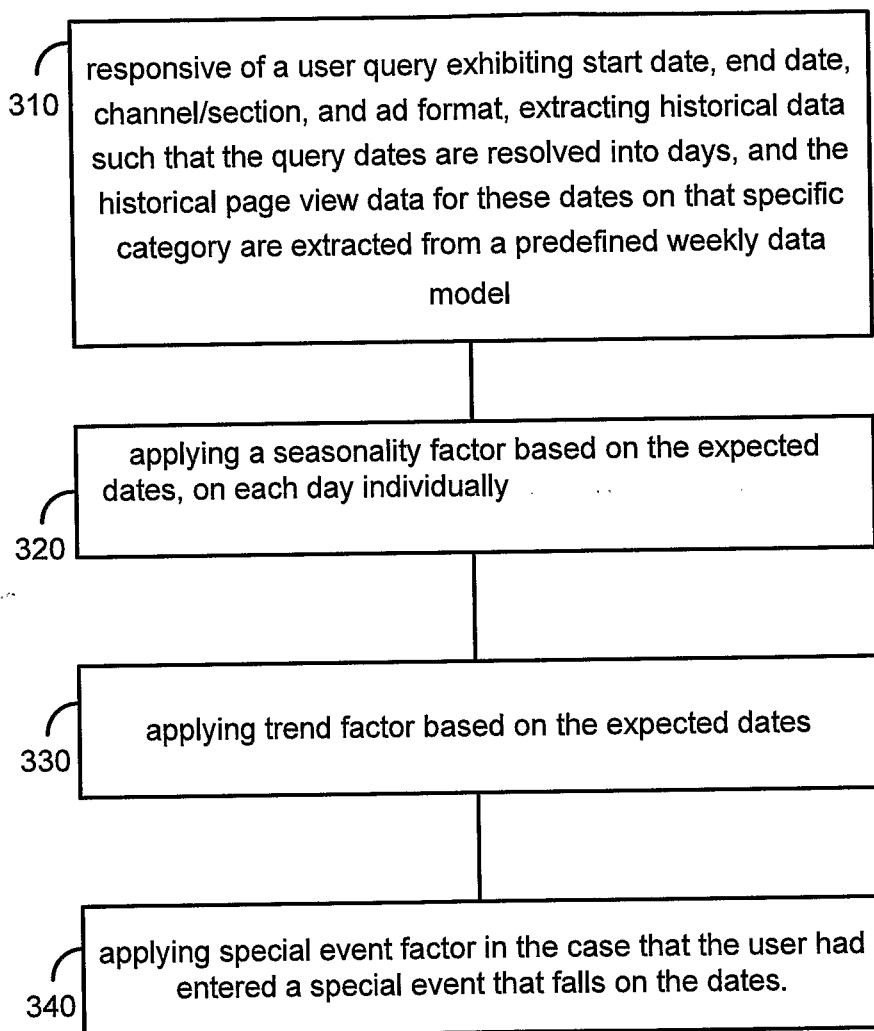


FIG. 3

4 / 6

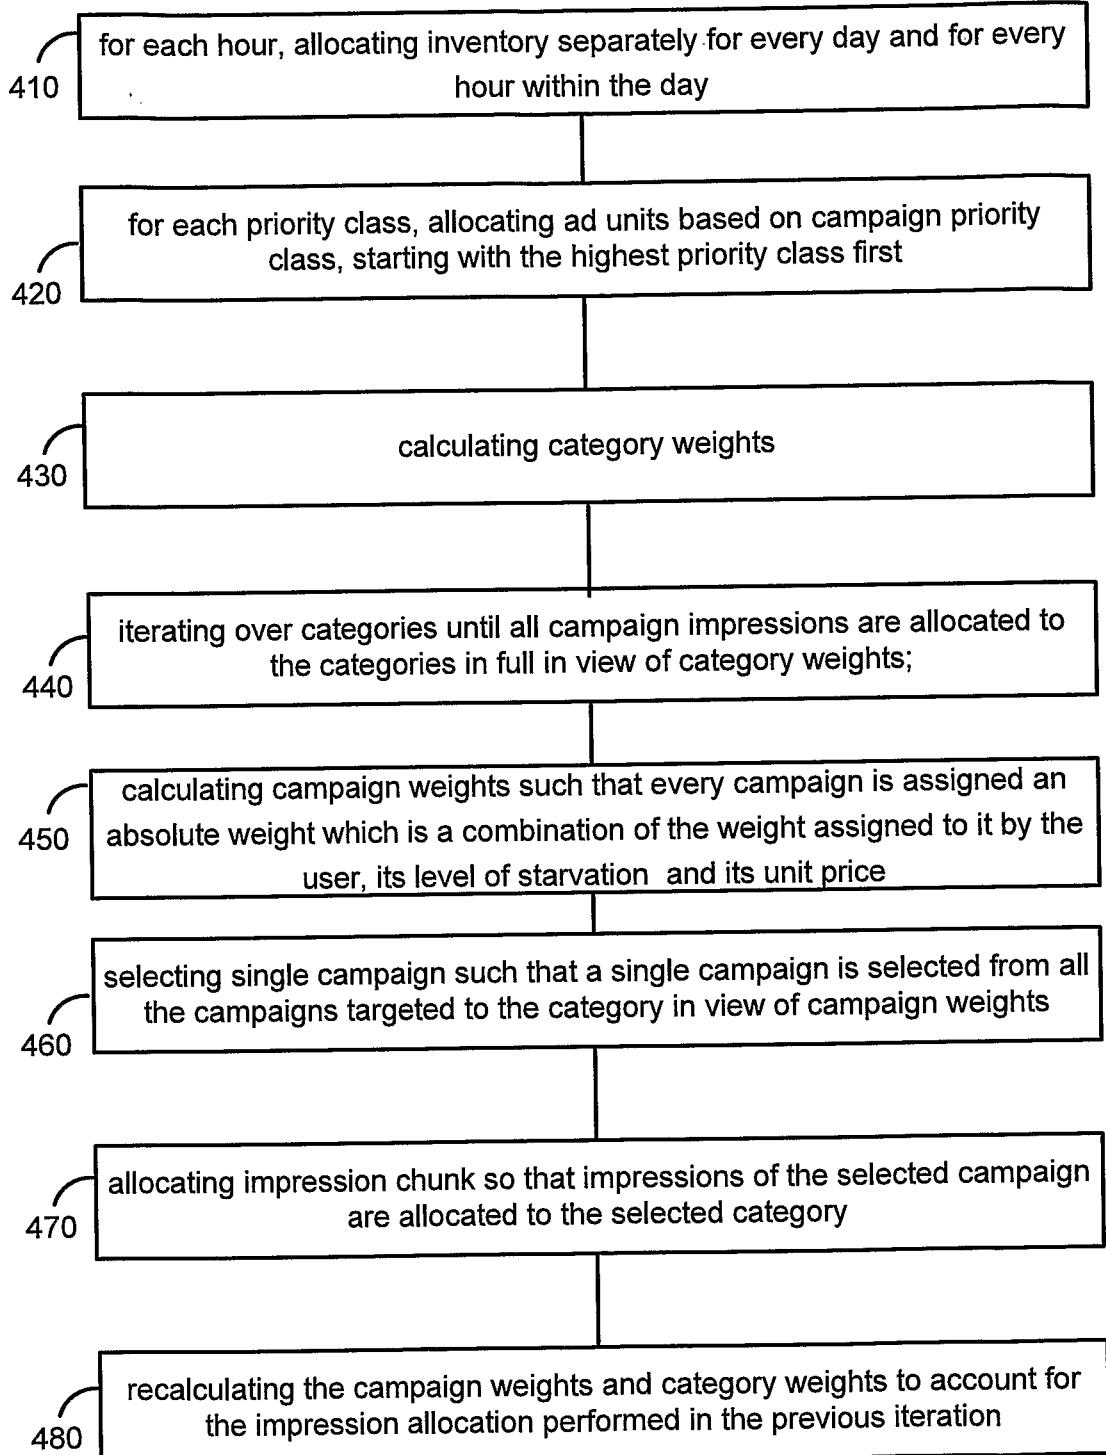


FIG. 4

5 / 6

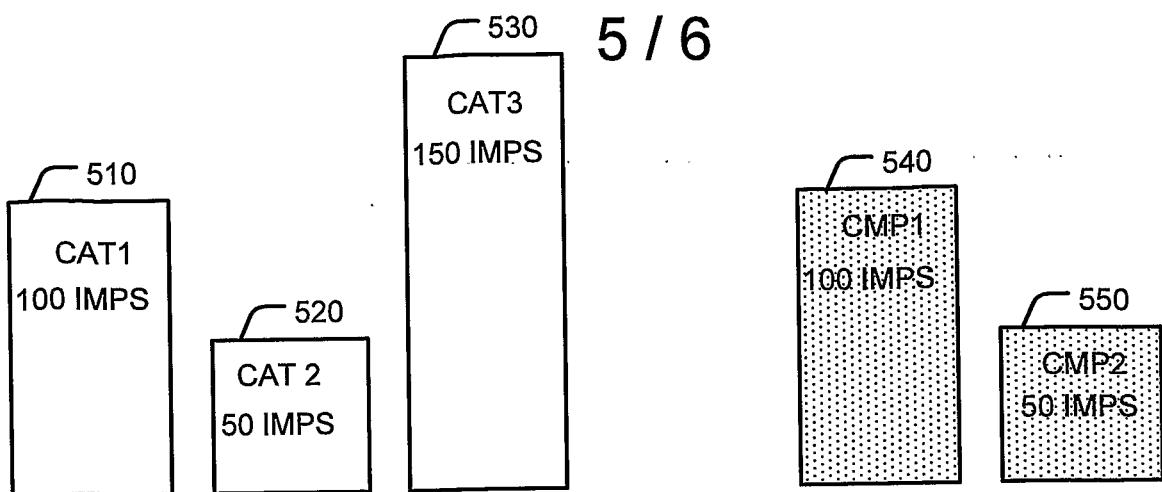


FIG. 5A

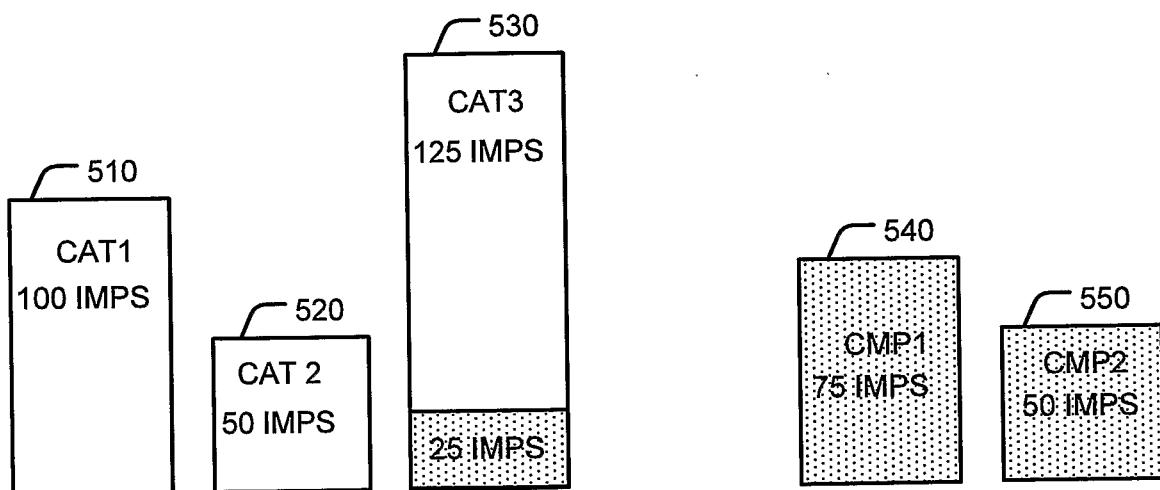


FIG. 5B

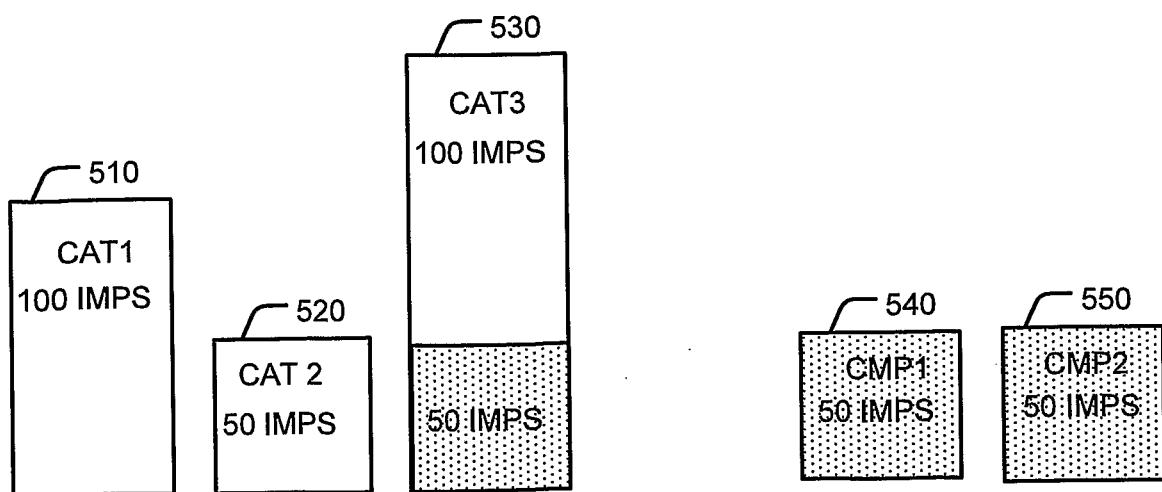


FIG. 5C

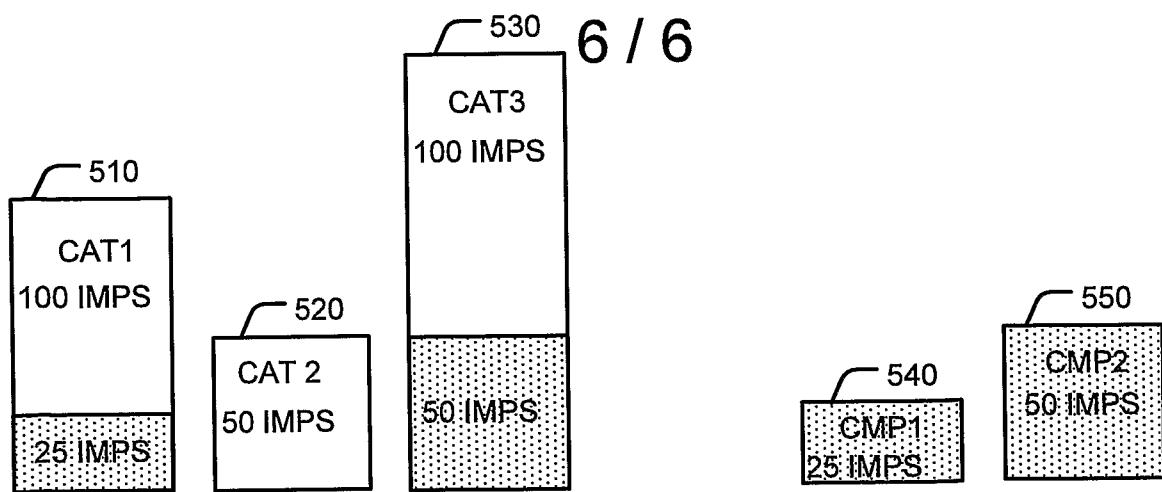


FIG. 5D

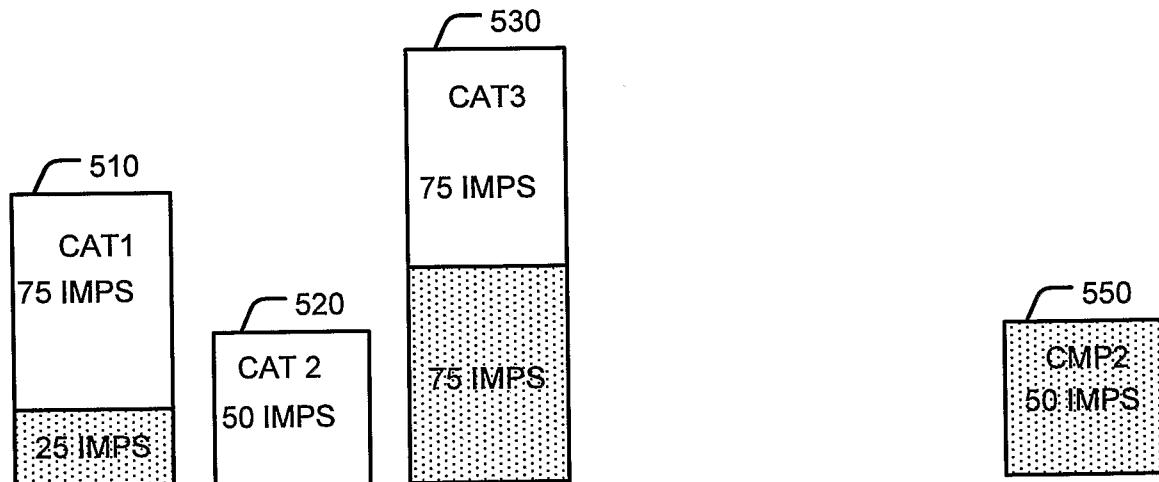


FIG. 5E

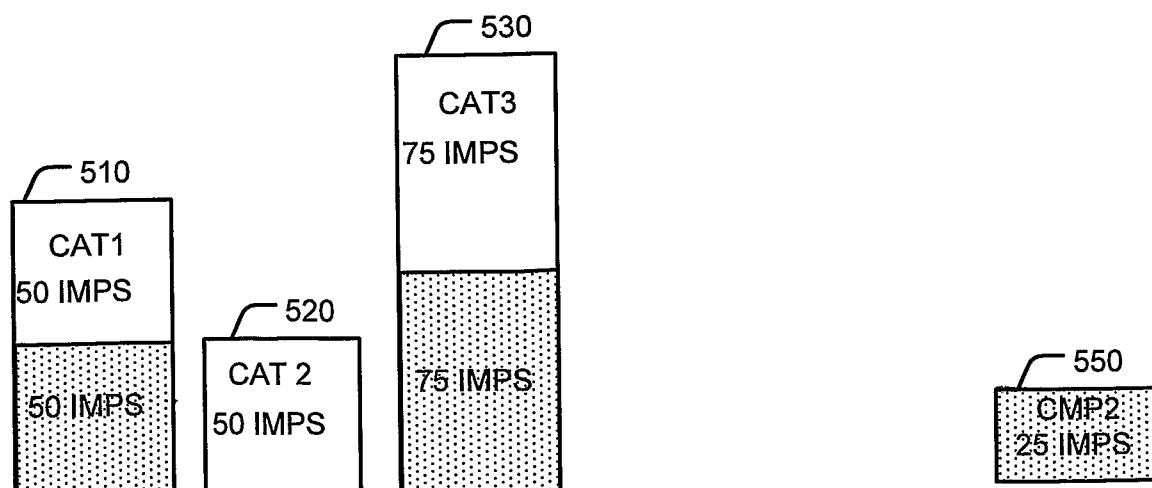


FIG. 5F

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IL 09/00802

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(8) - G06Q 30/00 (2009.01)
 USPC - 705/14.42

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 USPC: 705/14.42

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
 USPC: 705/1, 8, 14.1, 14.2, 14.24, 14.4, 4.49, 14.52, 22, 28,500; 700/90, 91, 92, 99, 100; 718/100, 102, 103, 104; 707/1, 3, 6, 10

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
 USPTO WEST (PGPB, USPT, EPAB, JPAB); Google Scholar
 Search Terms Used: advertising, campaign, inventory, resources, managing, limited, constraint, unit, product, good, unavailable, available, impression, chunk, allocate, weight, factor, seasonal, holiday, day, hour, week, model, paradigm, history, prediction etc.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
Y	US 2006/0080171 A1 (JARDINS et al.) 13 April 2006 (13.04.2006), entire document, especially para [0007], [0011], [0012], [0030], [0033], [0116]-[0122] [0138], [0142], [0150], [0162] and Fig 7	1-10
Y	US 2007/0271 145 A1 (VEST) 22 November 2007 (22.11.2007), entire document, especially para [0022], [0048], [0055]-[0057] and Figs 5-6	1-10
Y	US 2007/0005420 A1 (ROY et al.) 04 January 2007 (04.01.2007), entire document, especially para [0054]-[0056]	4-6 and 8-10
A	US 2003/01 10171 A1 (OZER et al.) 12 June 2003 (12.06.2003)	1-10
A	US 2008/0077504 A1 (GAUSEBECK et al.) 27 March 2008 (27.03.2008)	1-10
A	US 2005/0149396 A1 (HOROWITZ et al.) 07 July 2005 (07.07.2005)	1-10
A	US 2005/0222865 A1 (FOX) 06 October 2005 (06.10.2005)	1-10



Further documents are listed in the continuation of Box C.

D

* Special categories of cited documents	τ	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance		
"E" earlier application or patent but published on or after the international filing date	"X"	document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on novelty of claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"γ"	document of particular relevance, the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&"	document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search
 15 December 2009 (15.12.2009)

Date of mailing of the international search report

23 DEC 20Q9

Name and mailing address of the ISA/US
 Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, Virginia 22313-1450
 Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

 PCT Helpdesk: 571-272-4300
 PCTOSP 571-272-7774