US009514248B1

(12) **United States Patent**
Guan et al.

(10) **Patent No.: US 9,514,248 B1**
(45) **Date of Patent: *Dec. 6, 2016**

(54) **SYSTEM TO GROUP INTERNET DEVICES BASED UPON DEVICE USAGE**

(71) Applicant: **Drawbridge, Inc.**, San Mateo, CA (US)

(72) Inventors: **Devin Guan**, Monte Sereno, CA (US); **Xiang Li**, San Mateo, CA (US); **Randy Cotta**, San Mateo, CA (US); **Rahul Bafna**, San Mateo, CA (US); **Obuli Venkatesan**, San Mateo, CA (US)

(73) Assignee: **Drawbridge, Inc.**, San Mateo, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **14/942,203**

(22) Filed: **Nov. 16, 2015**

Related U.S. Application Data

(63) Continuation of application No. 14/877,837, filed on Oct. 7, 2015.

(51) **Int. Cl.**
G06F 7/00 (2006.01)
G06F 17/30 (2006.01)

(52) **U.S. Cl.**
CPC ***G06F 17/30958*** (2013.01); ***G06F 17/3053*** (2013.01); ***G06F 17/30598*** (2013.01)

(58) **Field of Classification Search**
CPC G06F 21/316; G06F 2221/2117; G06F 2221/2151
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,423,255 B2 4/2013 Padmanabhan et al.
2014/0095320 A1* 4/2014 Sivaramakrishnan .. H04L 67/22 705/14.66

OTHER PUBLICATIONS

De Montjoye, Yves-Alexandre, et al., "Unique in the shopping mall: On the reidentifiability of credit card metadata", Science, 347(6221), (Jan. 30, 2015), 536-540.
Hotz, Robert Lee, "Even Nameless Data Can Reveal Identity, Study Finds", [online]. The Wall Street Journal. Retrieved from the Internet: <URL: http://www.wsj.com/articles/SB21515893924992633662404580429794002309442>, reprinted Jan. 29, 2015., Copyright 2014—Dow Jones & Company, Inc., (2014), 3 pgs.
Satuluri, Venu, et al., "Local Graph Sparsification for Scalable Clustering", Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data (SIGMOD '11), (2011), 721-732.
"U.S. Appl. No. 14/877,837, Non Final Office Action mailed Apr. 6, 2016", 37 pgs.
U.S. Appl. No. 14/877,837, filed Oct. 7, 2015, System to Group Internet Devices Based Upon Device Usage.

* cited by examiner

Primary Examiner — Taelor Kim
(74) *Attorney, Agent, or Firm* — Schwegman, Lundberg & Woessner, P.A.

(57) **ABSTRACT**

A system comprising: an index structure that associates devices with device feature information; a pairing engine to determine device pairs based upon device feature information; a feature vector generation engine to produce feature vectors corresponding to determined device pairs based upon feature values associated within the index structure with devices of the determined device pairs; a scoring engine to determine scores to associate with determined device pairs based upon produced feature vectors; a graph structure, wherein nodes within the graph structure represent devices of determined device pairs, and wherein edges between pairs of nodes within the graph structure indicate determined device pairs; a clustering engine to identify respective clusters of three or more nodes within the graph structure that represent respective groups of devices.

29 Claims, 16 Drawing Sheets

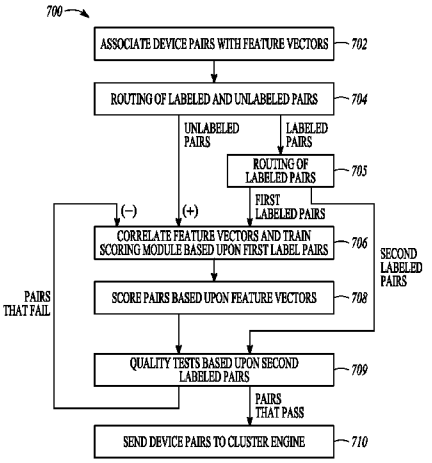

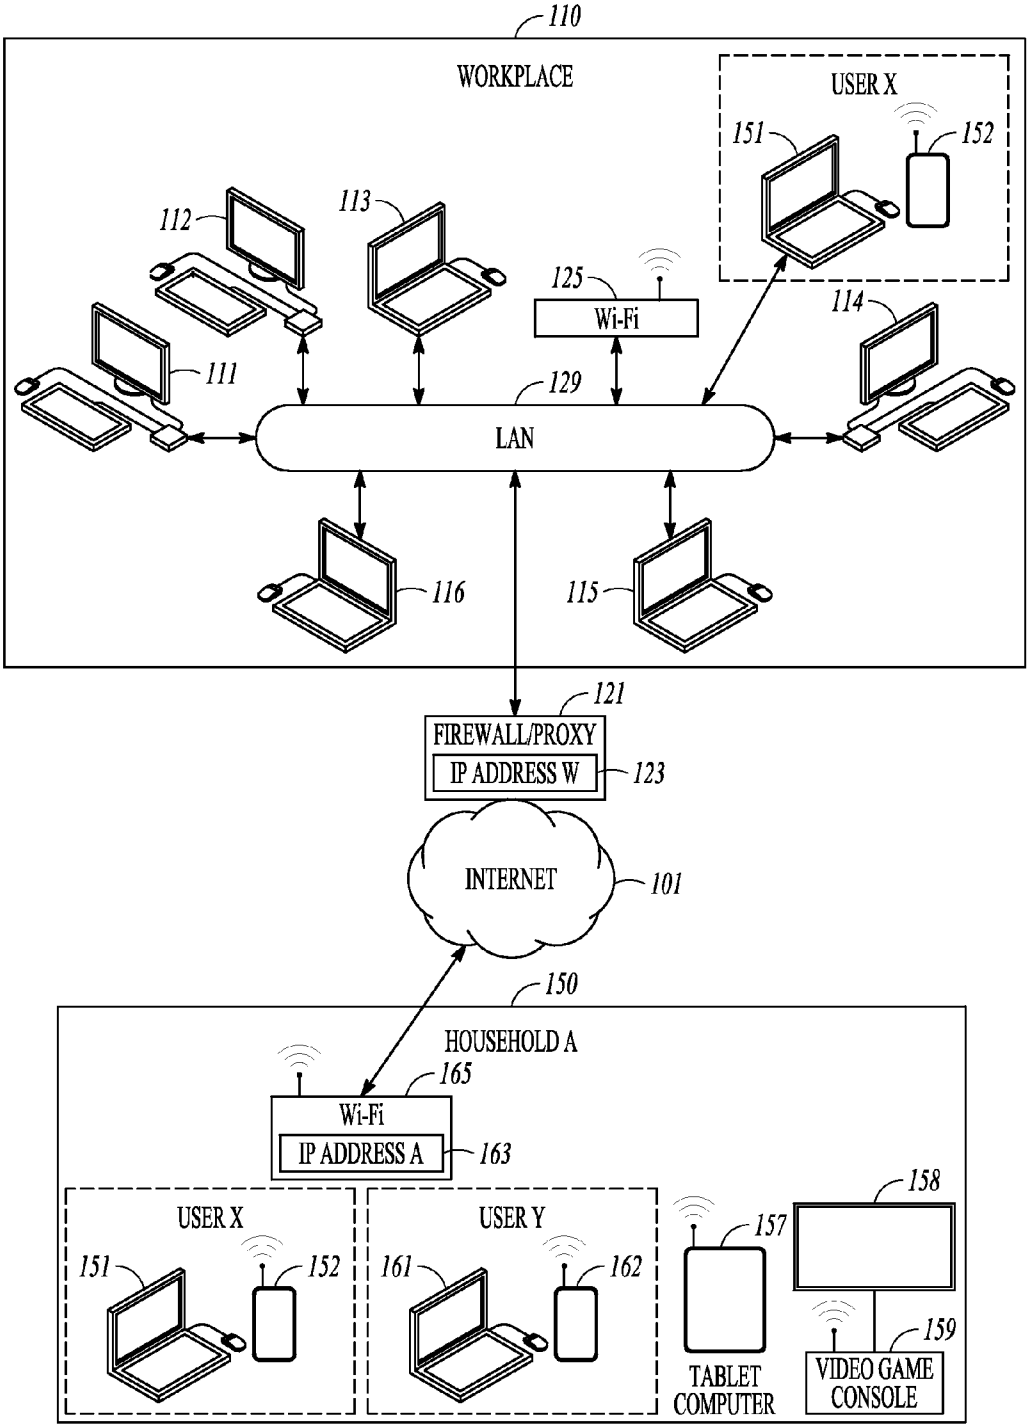
110
WORKPLACE
USER X
151
152
112
113
125
Wi-Fi
114
111
129
LAN
116
115
121
FIREWALL/PROXY
IP ADDRESS W
123
INTERNET
101
150
HOUSEHOLD A
165
Wi-Fi
IP ADDRESS A
163
158
USER X
USER Y
151
152
161
162
157
159
TABLET COMPUTER
VIDEO GAME CONSOLE

FIG. 1A

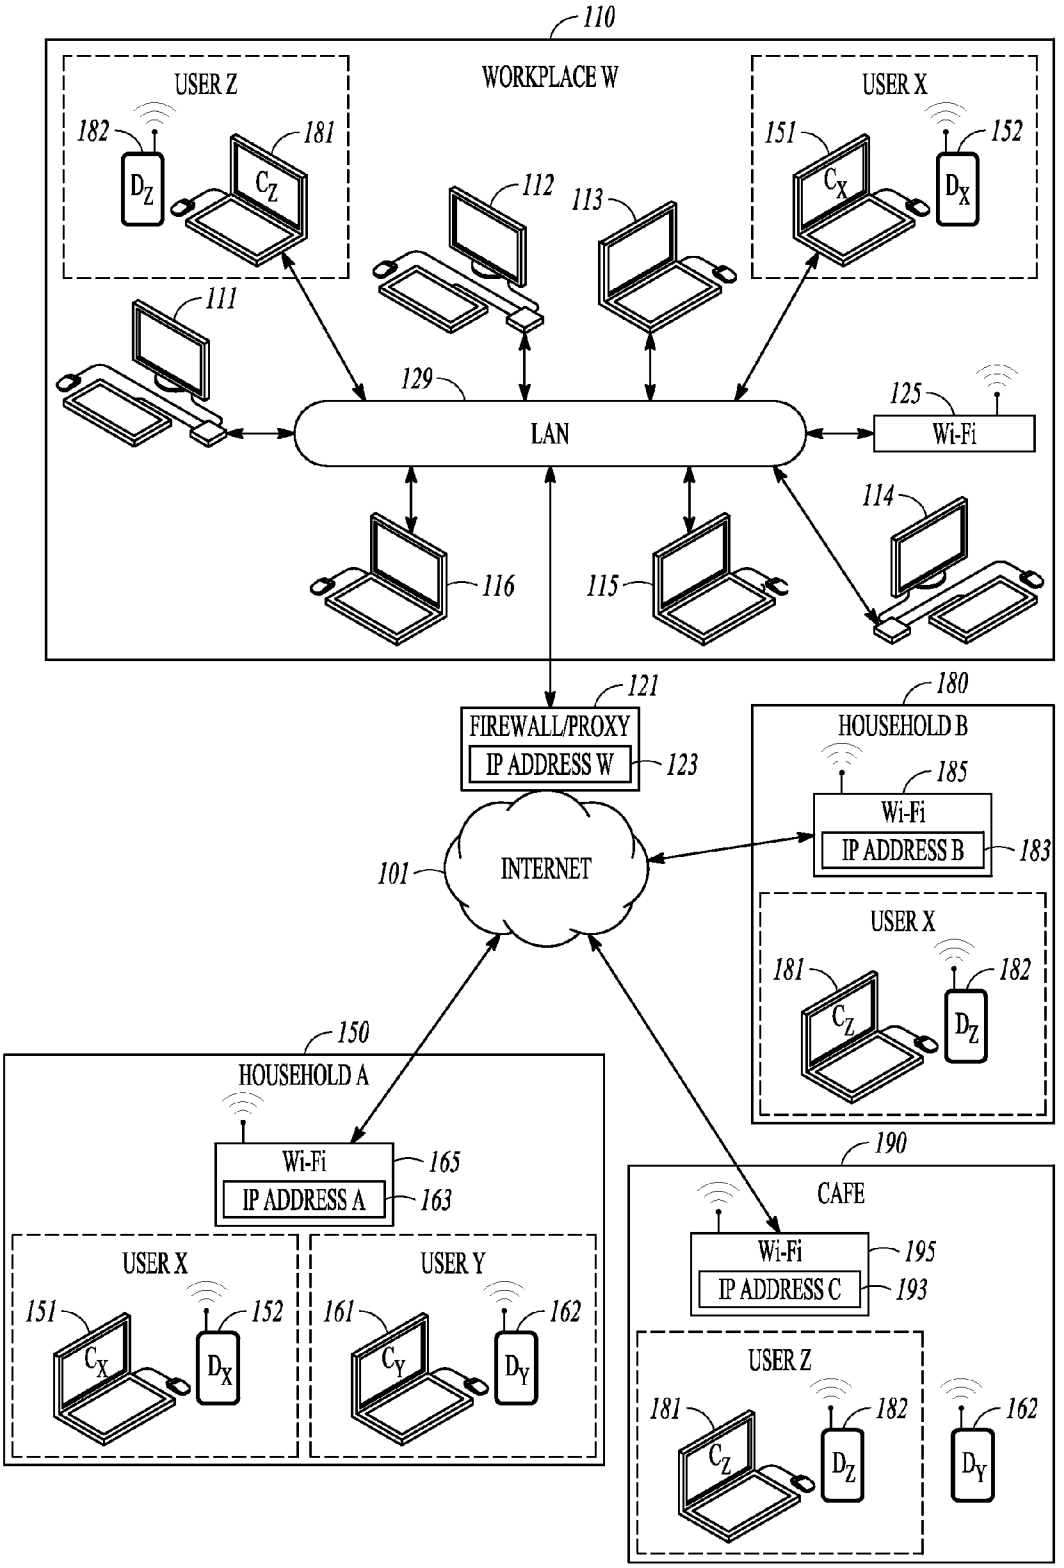
110
WORKPLACE W
USER Z
182
181
D_Z
C_Z
USER X
151
152
C_X
D_X
112
113
111
129
LAN
125
Wi-Fi
114
116
115
121
FIREWALL/PROXY
IP ADDRESS W
123
180
HOUSEHOLD B
185
Wi-Fi
IP ADDRESS B
183
101
INTERNET
USER X
181
182
C_Z
D_Z
150
HOUSEHOLD A
Wi-Fi
165
IP ADDRESS A
163
USER X
151
152
C_X
D_X
USER Y
161
162
C_Y
D_Y
190
CAFE
Wi-Fi
195
IP ADDRESS C
193
USER Z
181
182
C_Z
D_Z
162
D_Y

FIG. 1B

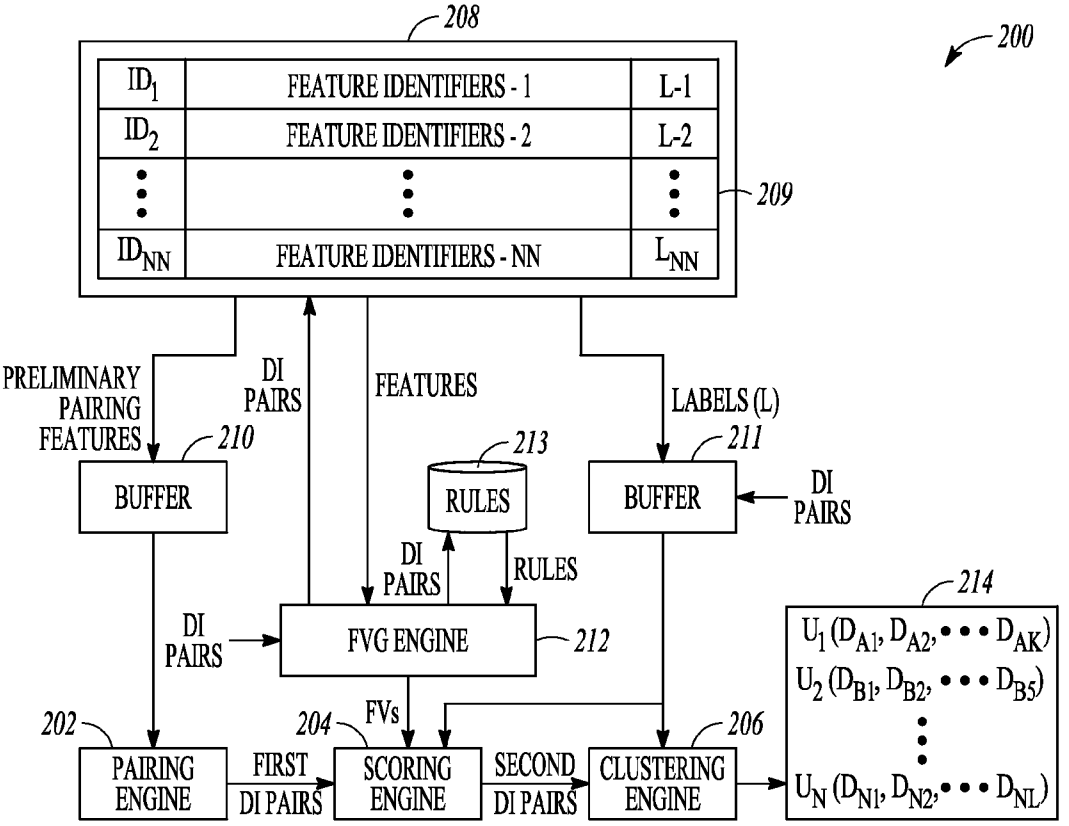
208
200
ID1
FEATURE IDENTIFIERS - 1
L-1
ID2
FEATURE IDENTIFIERS - 2
L-2
IDNN
FEATURE IDENTIFIERS - NN
LNN
209
PRELIMINARY PAIRING FEATURES
DI PAIRS
FEATURES
LABELS (L)
210
213
211
BUFFER
RULES
BUFFER
DI PAIRS
DI PAIRS
RULES
214
DI PAIRS
FVG ENGINE
212
U1 (DA1, DA2, • • • DAK)
U2 (DB1, DB2, • • • DB5)
UN (DN1, DN2, • • • DNL)
202
204
FVs
206
PAIRING ENGINE
FIRST DI PAIRS
SCORING ENGINE
SECOND DI PAIRS
CLUSTERING ENGINE

FIG. 2

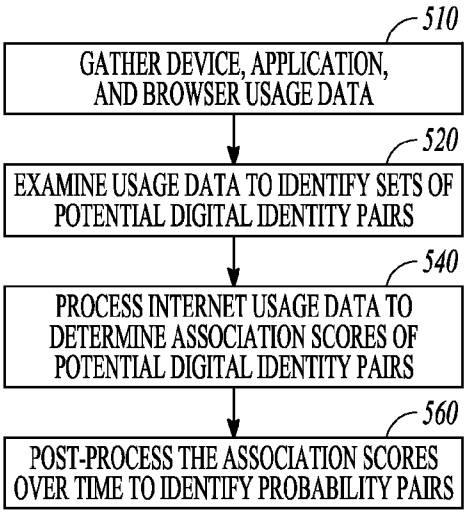
510
GATHER DEVICE, APPLICATION, AND BROWSER USAGE DATA
520
EXAMINE USAGE DATA TO IDENTIFY SETS OF POTENTIAL DIGITAL IDENTITY PAIRS
540
PROCESS INTERNET USAGE DATA TO DETERMINE ASSOCIATION SCORES OF POTENTIAL DIGITAL IDENTITY PAIRS
560
POST-PROCESS THE ASSOCIATION SCORES OVER TIME TO IDENTIFY PROBABILITY PAIRS

FIG. 3

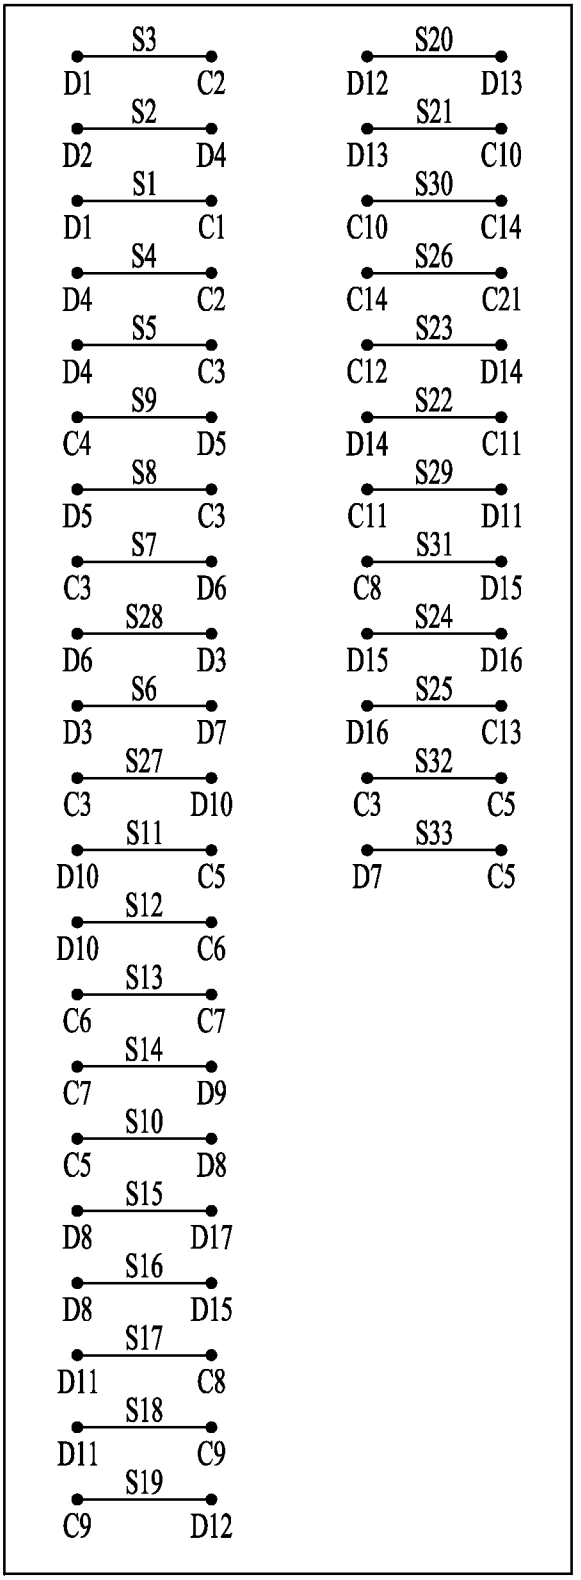
S3
D1 C2
S2
D2 D4
S1
D1 C1
S4
D4 C2
S5
D4 C3
S9
C4 D5
S8
D5 C3
S7
C3 D6
S28
D6 D3
S6
D3 D7
S27
C3 D10
S11
D10 C5
S12
D10 C6
S13
C6 C7
S14
C7 D9
S10
C5 D8
S15
D8 D17
S16
D8 D15
S17
D11 C8
S18
D11 C9
S19
C9 D12
S20
D12 D13
S21
D13 C10
S30
C10 C14
S26
C14 C21
S23
C12 D14
S22
D14 C11
S29
C11 D11
S31
C8 D15
S24
D15 D16
S25
D16 C13
S32
C3 C5
S33
D7 C5

FIG. 4A

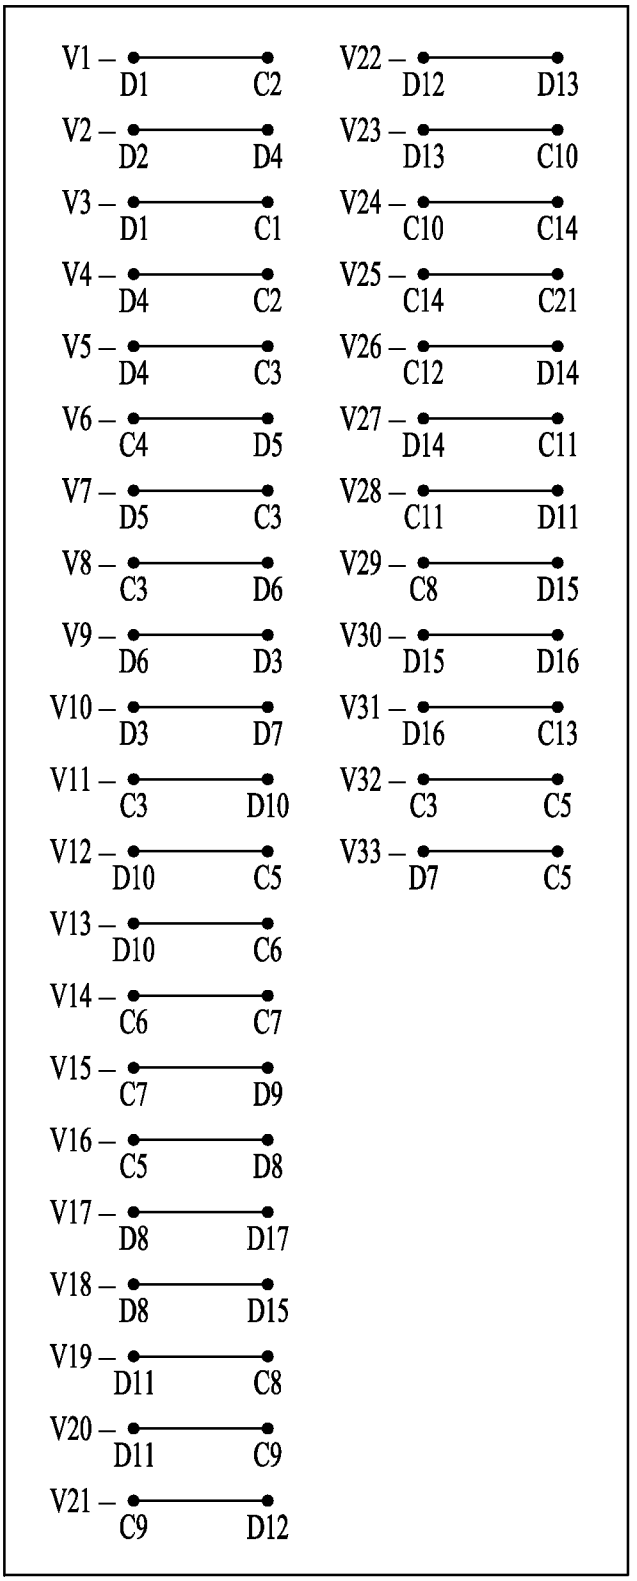
V1 – D1 C2
V2 – D2 D4
V3 – D1 C1
V4 – D4 C2
V5 – D4 C3
V6 – C4 D5
V7 – D5 C3
V8 – C3 D6
V9 – D6 D3
V10 – D3 D7
V11 – C3 D10
V12 – D10 C5
V13 – D10 C6
V14 – C6 C7
V15 – C7 D9
V16 – C5 D8
V17 – D8 D17
V18 – D8 D15
V19 – D11 C8
V20 – D11 C9
V21 – C9 D12
V22 – D12 D13
V23 – D13 C10
V24 – C10 C14
V25 – C14 C21
V26 – C12 D14
V27 – D14 C11
V28 – C11 D11
V29 – C8 D15
V30 – D15 D16
V31 – D16 C13
V32 – C3 C5
V33 – D7 C5

FIG. 4B

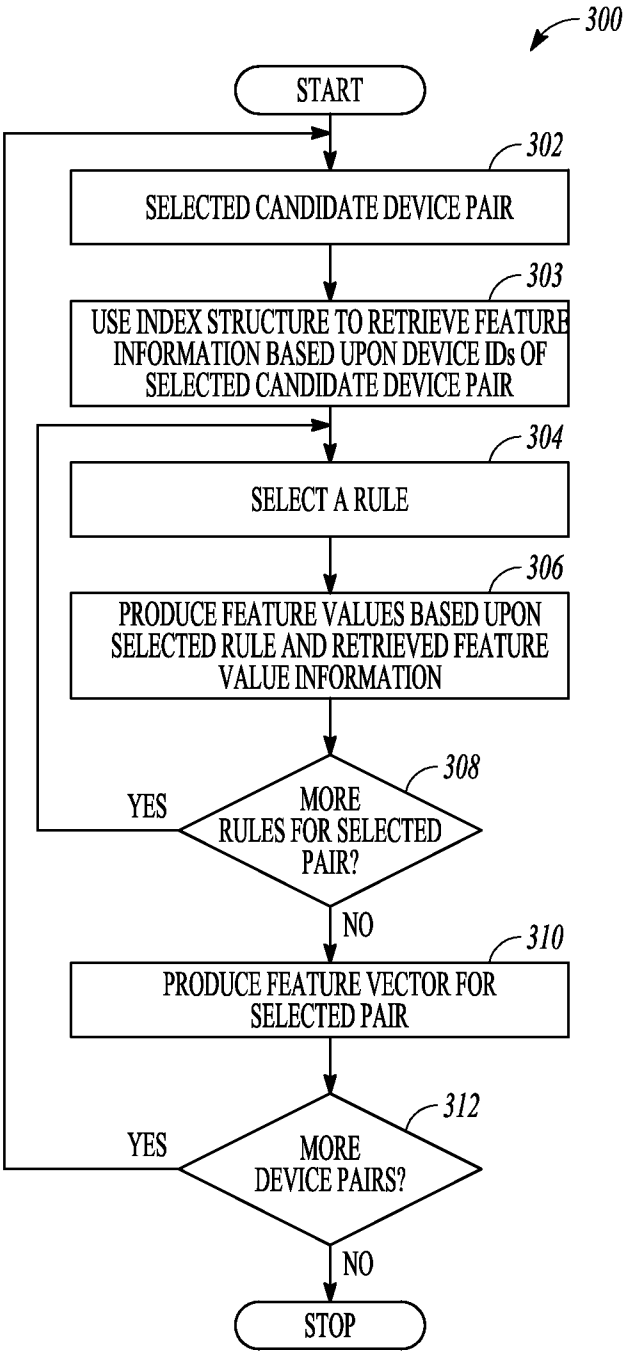
300
START
302
SELECTED CANDIDATE DEVICE PAIR
303
USE INDEX STRUCTURE TO RETRIEVE FEATURE INFORMATION BASED UPON DEVICE IDs OF SELECTED CANDIDATE DEVICE PAIR
304
SELECT A RULE
306
PRODUCE FEATURE VALUES BASED UPON SELECTED RULE AND RETRIEVED FEATURE VALUE INFORMATION
308
MORE RULES FOR SELECTED PAIR?
YES
NO
310
PRODUCE FEATURE VECTOR FOR SELECTED PAIR
312
MORE DEVICE PAIRS?
YES
NO
STOP

FIG. 5

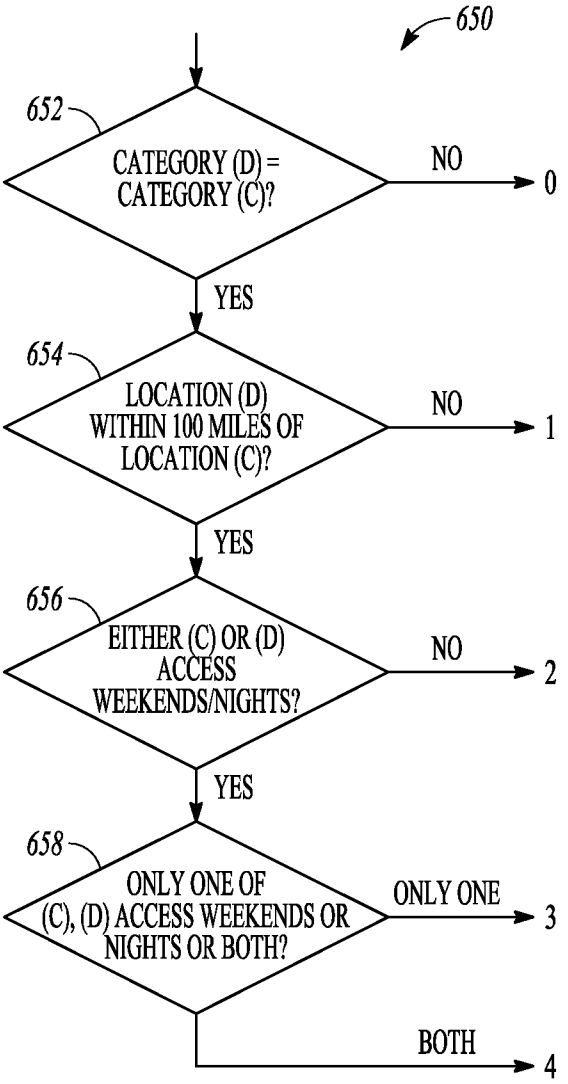
650
652
CATEGORY (D) = CATEGORY (C)?
NO
0
YES
654
LOCATION (D) WITHIN 100 MILES OF LOCATION (C)?
NO
1
YES
656
EITHER (C) OR (D) ACCESS WEEKENDS/NIGHTS?
NO
2
YES
658
ONLY ONE OF (C), (D) ACCESS WEEKENDS OR NIGHTS OR BOTH?
ONLY ONE
3
BOTH
4

FIG. 6A

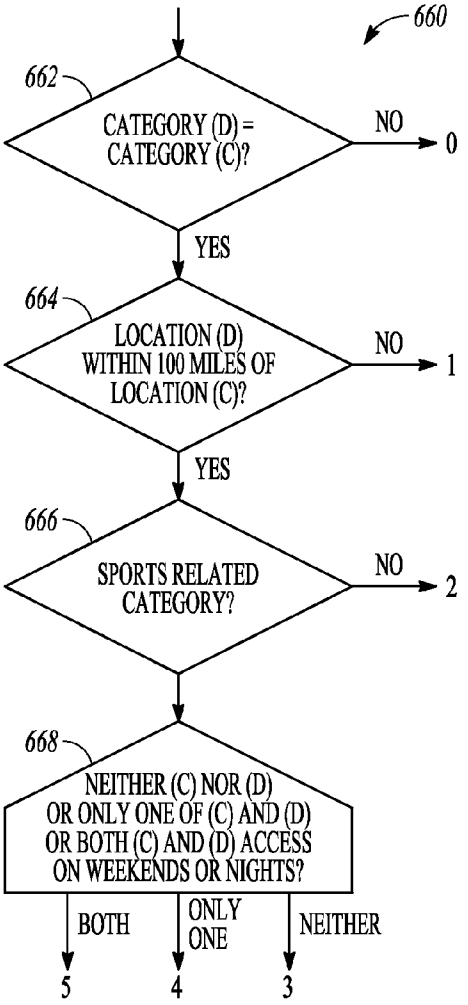

FIG. 6B

FEATURE NAME	DESCRIPTION
L	("LABEL") CAN WEE VERIFY THAT CD BELONG TO UNIQUE USER? OR THAT THEY DO NOT?
NT	WHAT TYPE OF DEVICE ID IS D? (IDFA, RAW, ANDROID, APPLE, ...)
01	HOW MANY TOTAL OBSERVATIONS DO WE HAVE FOR D?
07	HOW MANY TOTAL OBSERVATIONS DO WE HAVE FOR C?
CO	HOW MANY UNIQUE IPs WHERE ONLY C OR D OCCUR? (IP SET SYMMETRIC DIFFERENCE)
GEN	DO CD HAVE THE SAME GENDER? OPPOSITE GENDER? NO GENDER INFO?
AGE	DO CD HAVE THE SAME AGE? HOW DIFFERENT ARE THEIR AGES?
GEO	DO CD HAVE THE SAME HOME LOCATION? HOW CLOSELY MATCHED OR BADLY MISMATCHED ARE THEIR INFERRED HOME LOCATIONS

FIG. 7

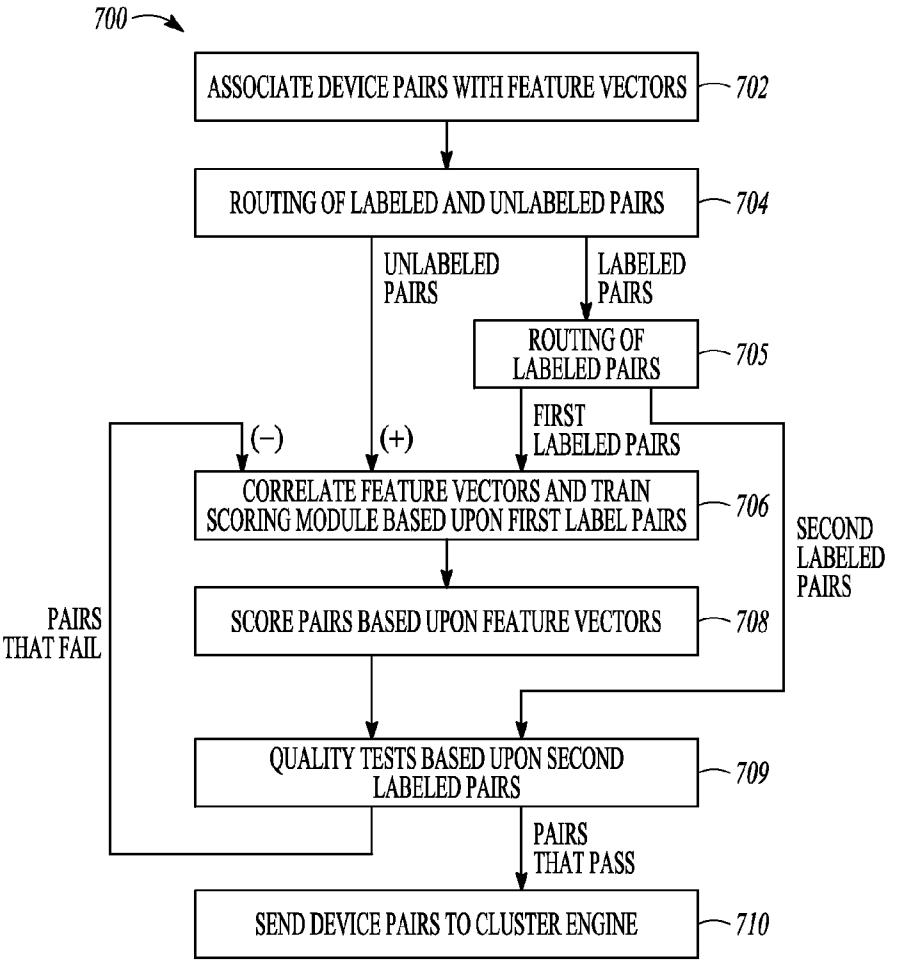
700
ASSOCIATE DEVICE PAIRS WITH FEATURE VECTORS
702
ROUTING OF LABELED AND UNLABELED PAIRS
704
UNLABELED PAIRS
LABELED PAIRS
ROUTING OF LABELED PAIRS
705
(−)
(+)
FIRST LABELED PAIRS
CORRELATE FEATURE VECTORS AND TRAIN SCORING MODULE BASED UPON FIRST LABEL PAIRS
706
SECOND LABELED PAIRS
PAIRS THAT FAIL
SCORE PAIRS BASED UPON FEATURE VECTORS
708
QUALITY TESTS BASED UPON SECOND LABELED PAIRS
709
PAIRS THAT PASS
SEND DEVICE PAIRS TO CLUSTER ENGINE
710

FIG. 8

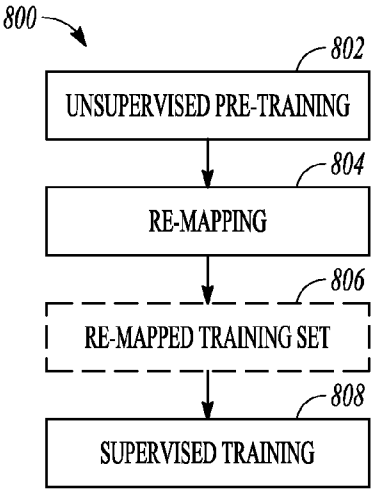
800
802
UNSUPERVISED PRE-TRAINING
804
RE-MAPPING
806
RE-MAPPED TRAINING SET
808
SUPERVISED TRAINING

FIG. 9

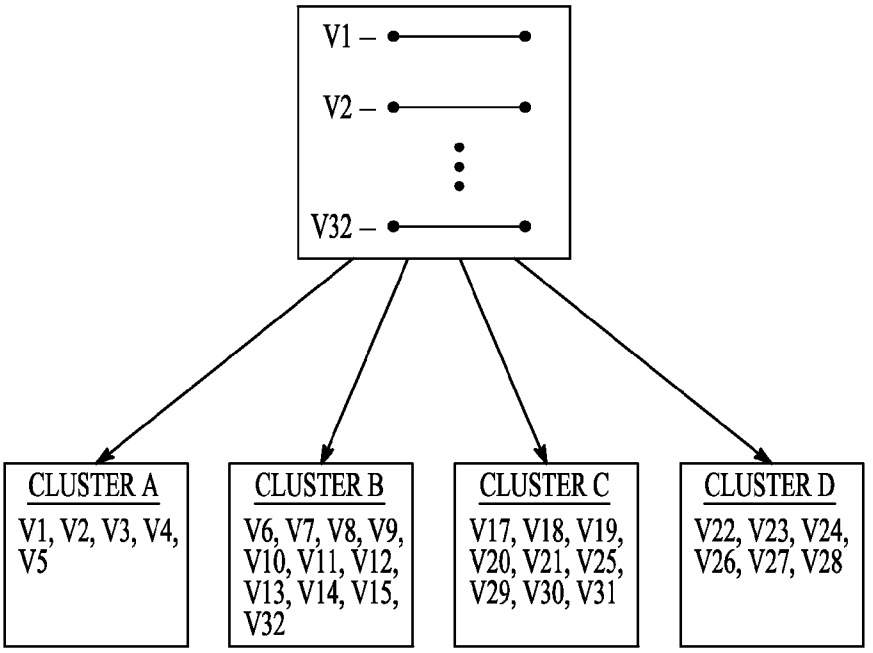

FIG. 10A

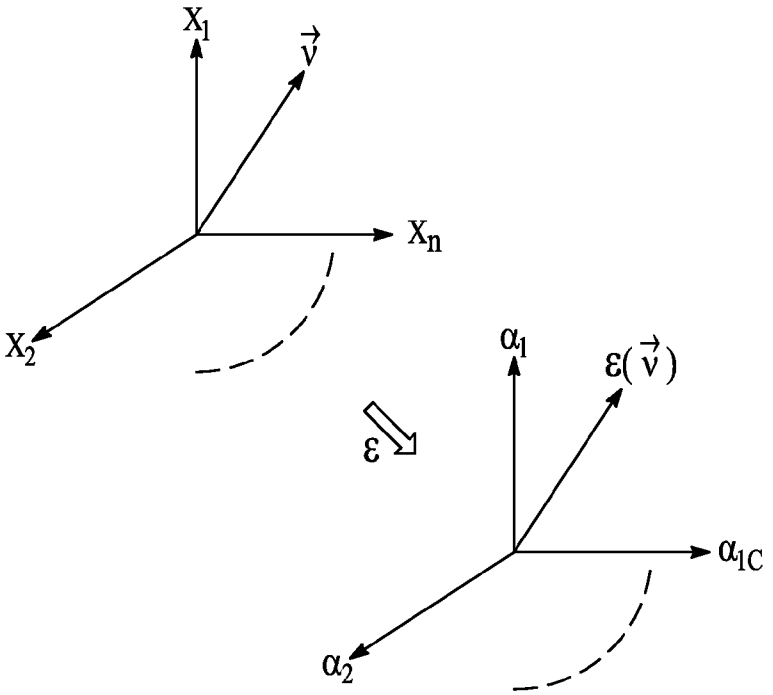

FIG. 10B

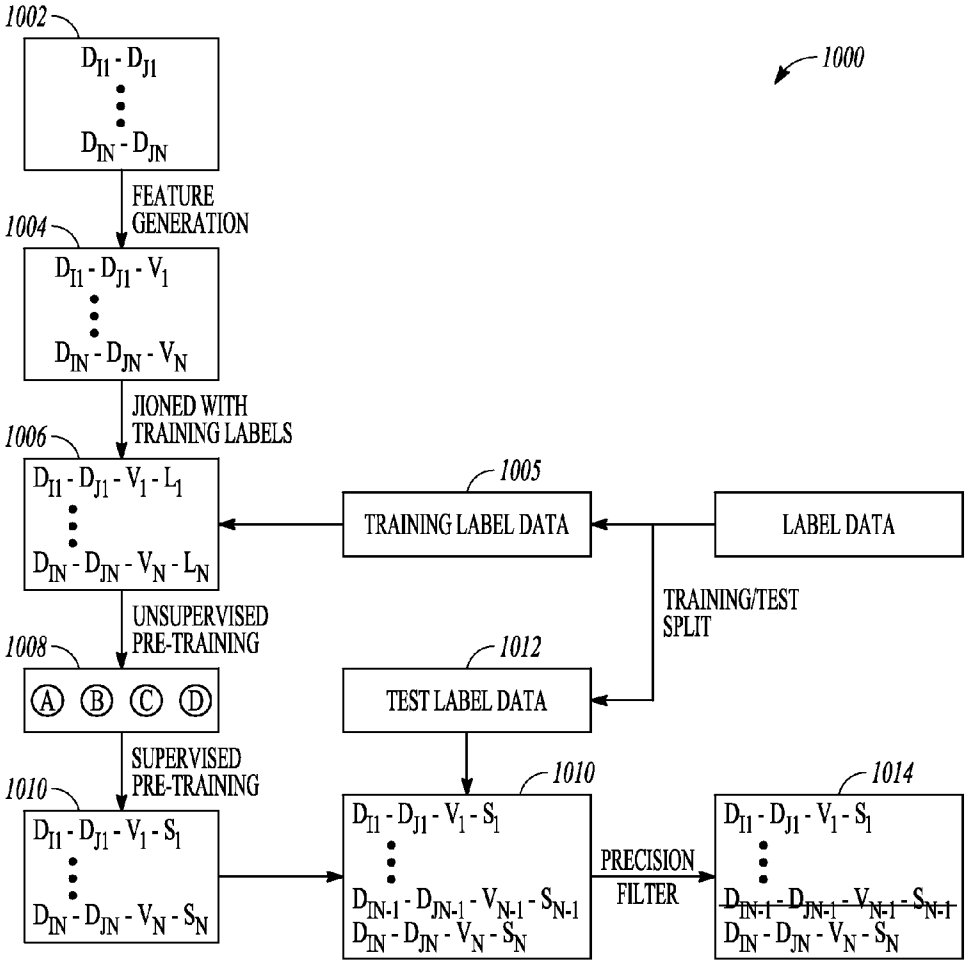
1000
1002
D_{I1} - D_{J1}
D_{IN} - D_{JN}
FEATURE GENERATION
1004
D_{I1} - D_{J1} - V_1
D_{IN} - D_{JN} - V_N
JIONED WITH TRAINING LABELS
1006
D_{I1} - D_{J1} - V_1 - L_1
D_{IN} - D_{JN} - V_N - L_N
1005
TRAINING LABEL DATA
LABEL DATA
TRAINING/TEST SPLIT
UNSUPERVISED PRE-TRAINING
1008
A B C D
1012
TEST LABEL DATA
SUPERVISED PRE-TRAINING
1010
D_{I1} - D_{J1} - V_1 - S_1
D_{IN} - D_{JN} - V_N - S_N
1010
D_{I1} - D_{J1} - V_1 - S_1
D_{IN-1} - D_{JN-1} - V_{N-1} - S_{N-1}
D_{IN} - D_{JN} - V_N - S_N
PRECISION FILTER
1014
D_{I1} - D_{J1} - V_1 - S_1
~~D_{IN-1} - D_{JN-1} - V_{N-1} - S_{N-1}~~
D_{IN} - D_{JN} - V_N - S_N

FIG. 11

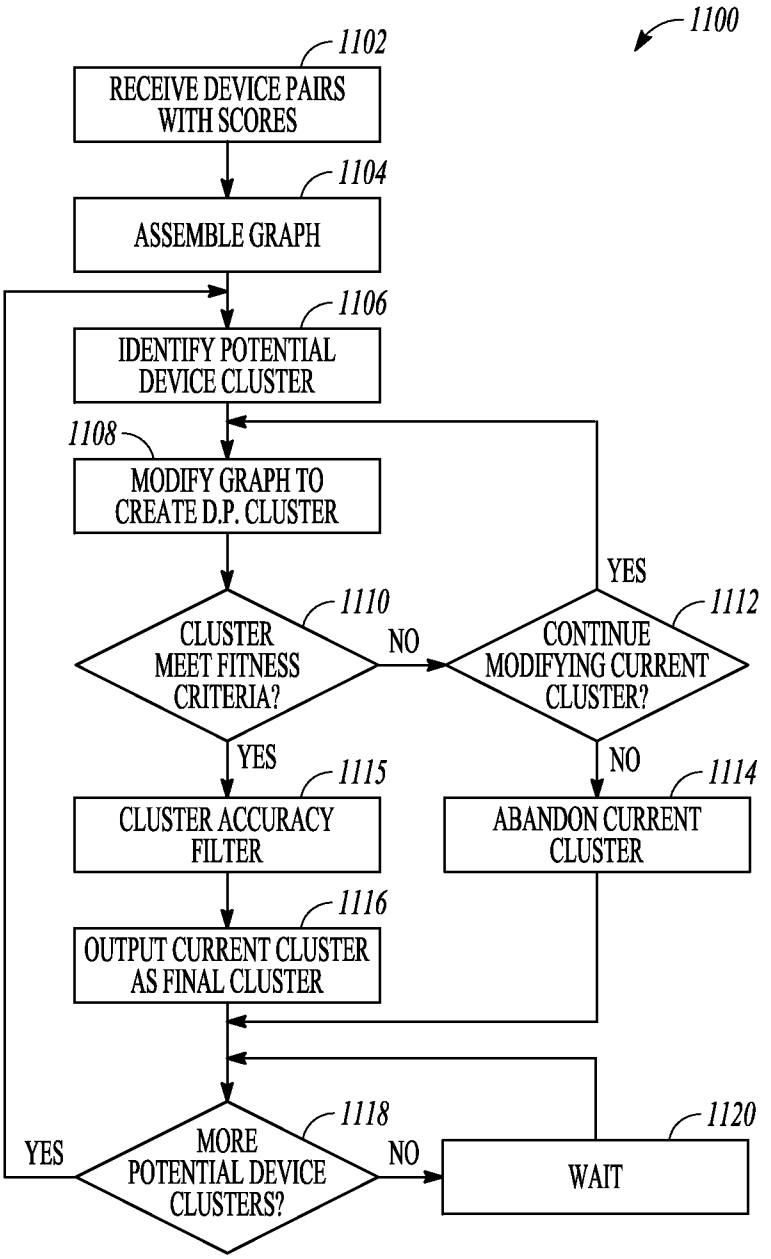
1100
1102
RECEIVE DEVICE PAIRS WITH SCORES
1104
ASSEMBLE GRAPH
1106
IDENTIFY POTENTIAL DEVICE CLUSTER
1108
MODIFY GRAPH TO CREATE D.P. CLUSTER
1110
CLUSTER MEET FITNESS CRITERIA?
NO
YES
1112
CONTINUE MODIFYING CURRENT CLUSTER?
YES
NO
1115
CLUSTER ACCURACY FILTER
1114
ABANDON CURRENT CLUSTER
1116
OUTPUT CURRENT CLUSTER AS FINAL CLUSTER
1118
MORE POTENTIAL DEVICE CLUSTERS?
YES
NO
1120
WAIT

FIG. 12

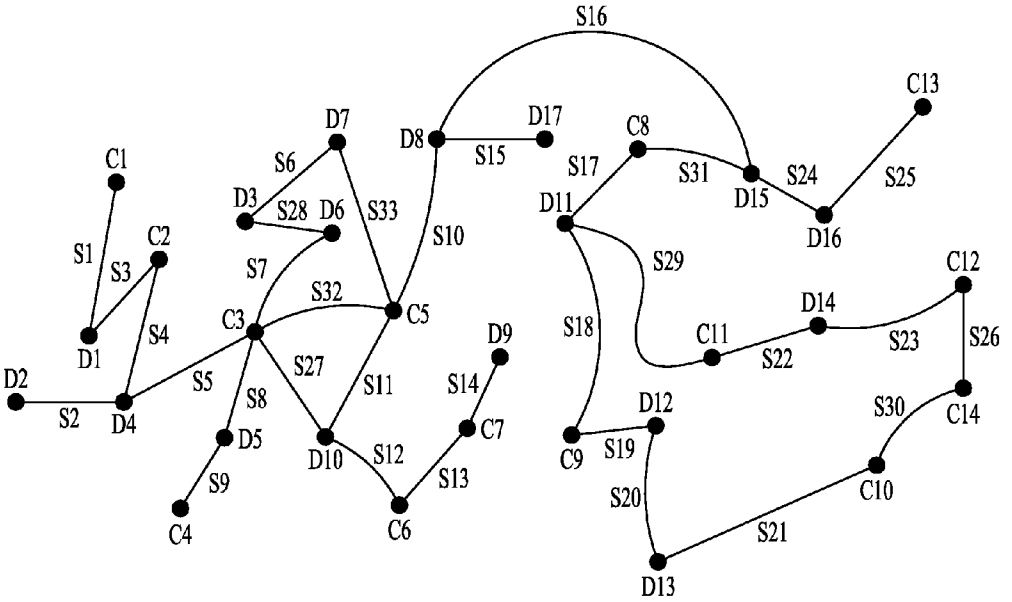
S16
C13
D7
D8
D17
C8
C1
S6
S15
S17
S31
S24
S25
D3
S28
D6
S33
D11
D15
D16
S1
C2
S10
S3
S7
S29
C12
S32
C5
D14
C3
S4
D9
S18
C11
S23
S26
D1
S27
S22
D2
S5
S11
S14
S8
D12
S30
C14
S2
D4
C7
D5
D10
S12
C9
S19
S9
S13
S20
C10
C4
C6
S21
D13

FIG. 13

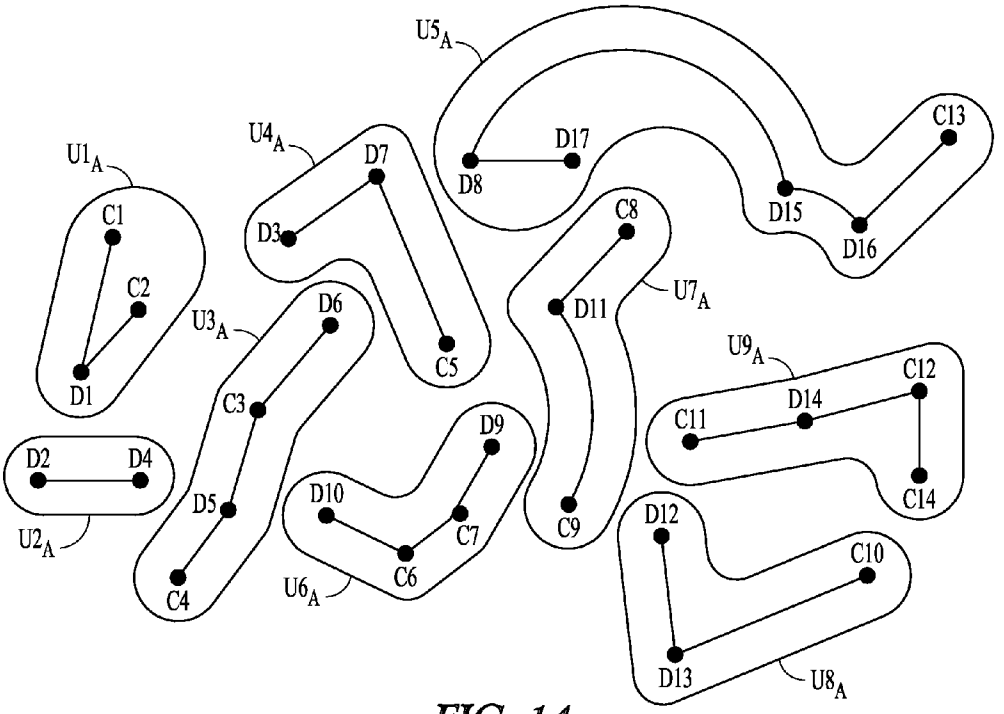

FIG. 14

$U1_B$ (C1, D1, C2)
$U1_B$ (D2, D4)
$U3_B$ (C4, D5, C3, D6)
$U4_B$ (D3, D7, C5)
$U5_B$ (D8, D17, D15, D16)
$U6_B$ (D10, C6, C7, D9)
$U7_B$ (C8, D11, C9)
$U8_B$ (D12, D13, C10)
$U9_B$ (C11, D14, C12, C14)

FIG. 15

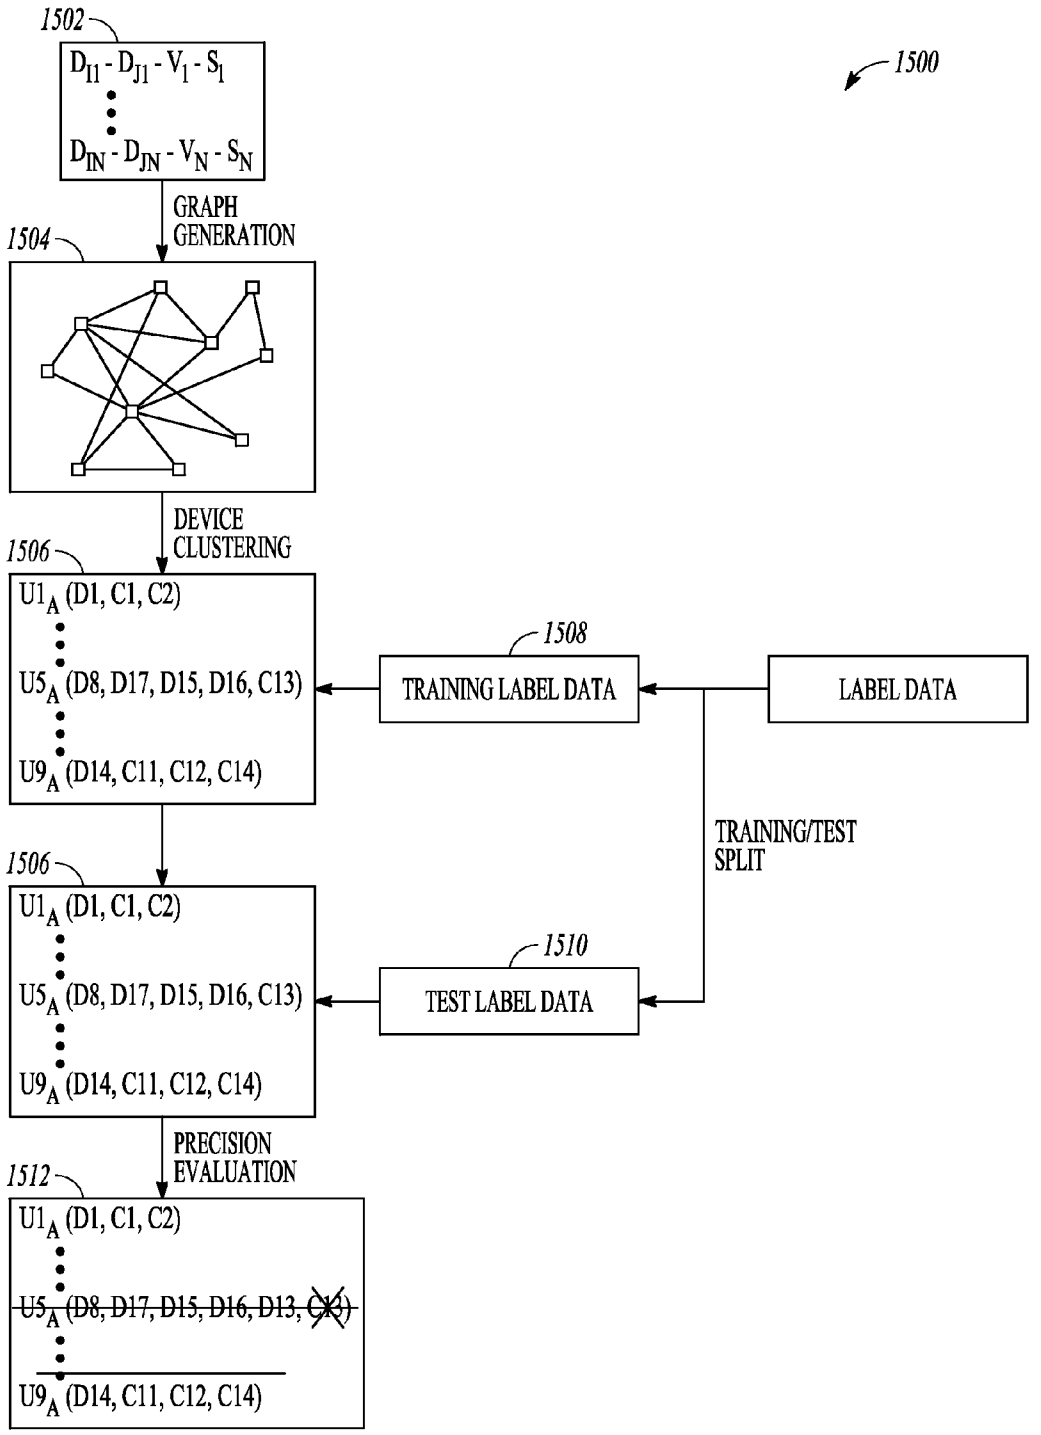
1500
1502
D_{I1} - D_{J1} - V_1 - S_1
D_{IN} - D_{JN} - V_N - S_N
GRAPH GENERATION
1504
DEVICE CLUSTERING
1506
$U1_A$ (D1, C1, C2)
$U5_A$ (D8, D17, D15, D16, C13)
$U9_A$ (D14, C11, C12, C14)
1508
TRAINING LABEL DATA
LABEL DATA
TRAINING/TEST SPLIT
1506
$U1_A$ (D1, C1, C2)
$U5_A$ (D8, D17, D15, D16, C13)
$U9_A$ (D14, C11, C12, C14)
1510
TEST LABEL DATA
PRECISION EVALUATION
1512
$U1_A$ (D1, C1, C2)
~~$U5_A$ (D8, D17, D15, D16, D13, C13)~~
$U9_A$ (D14, C11, C12, C14)

FIG. 16

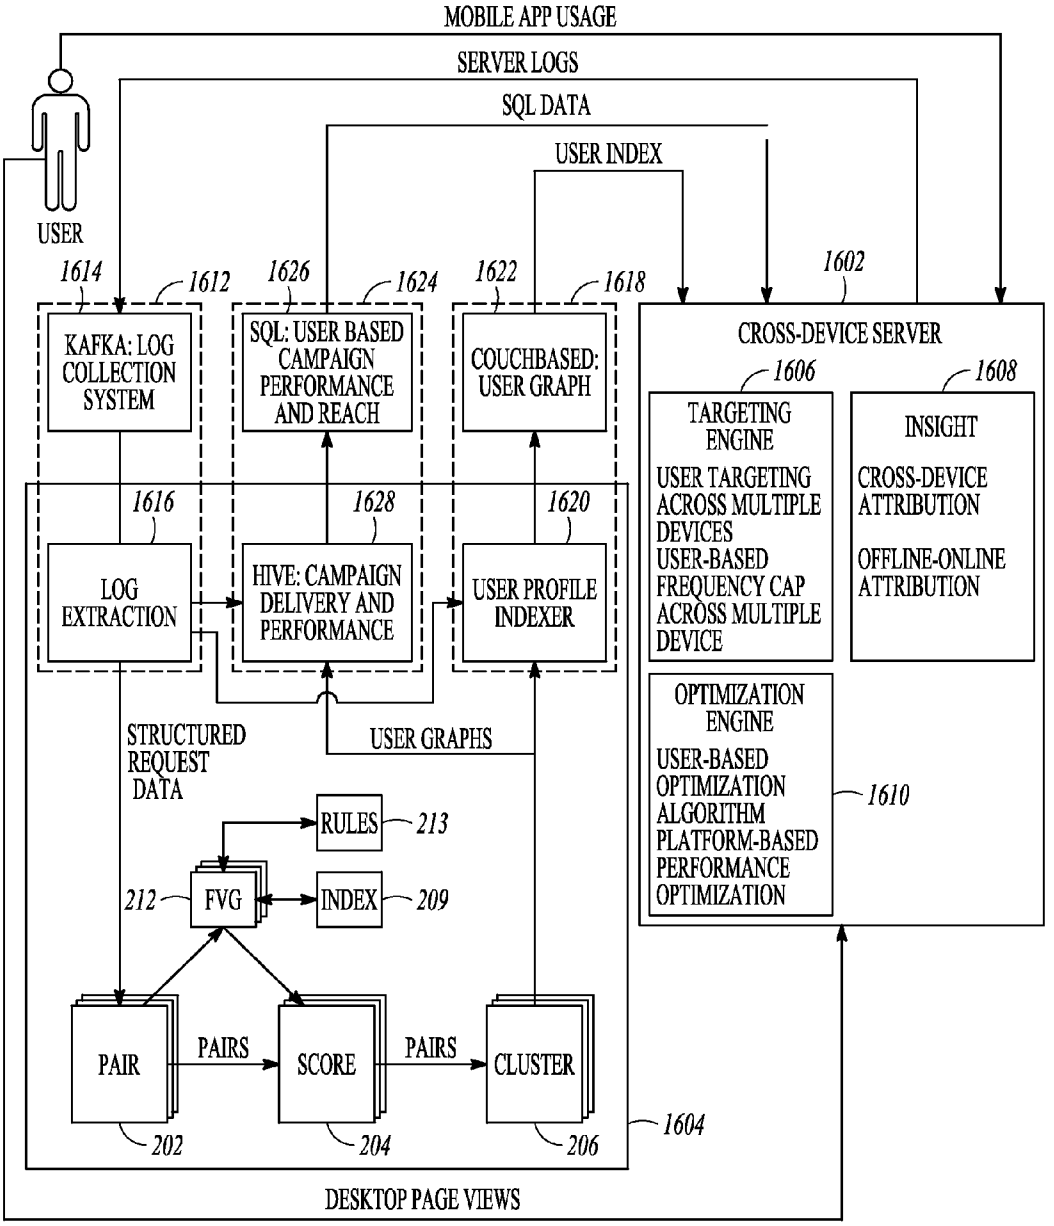
MOBILE APP USAGE
SERVER LOGS
SQL DATA
USER INDEX
USER
1614
1612
1626
1624
1622
1618
1602
KAFKA: LOG COLLECTION SYSTEM
SQL: USER BASED CAMPAIGN PERFORMANCE AND REACH
COUCHBASED: USER GRAPH
CROSS-DEVICE SERVER
1606
1608
TARGETING ENGINE
USER TARGETING ACROSS MULTIPLE DEVICES
USER-BASED FREQUENCY CAP ACROSS MULTIPLE DEVICE
INSIGHT
CROSS-DEVICE ATTRIBUTION
OFFLINE-ONLINE ATTRIBUTION
1616
1628
1620
LOG EXTRACTION
HIVE: CAMPAIGN DELIVERY AND PERFORMANCE
USER PROFILE INDEXER
OPTIMIZATION ENGINE
USER-BASED OPTIMIZATION ALGORITHM
PLATFORM-BASED PERFORMANCE OPTIMIZATION
1610
STRUCTURED REQUEST DATA
USER GRAPHS
RULES
213
212
FVG
INDEX
209
PAIR
PAIRS
SCORE
PAIRS
CLUSTER
202
204
206
1604
DESKTOP PAGE VIEWS

FIG. 17

SYSTEM TO GROUP INTERNET DEVICES BASED UPON DEVICE USAGE

PRIORITY APPLICATION

This application is a Continuation of U.S. patent application Ser. No. 14/877,837, filed Oct. 7, 2015, which application is incorporated herein by reference in its entirety.

BACKGROUND

The Internet provides a communication network over which persons can exchange information using a wide variety of different devices. For example, a user may own a smartphone, a mobile tablet, a laptop computer, and a connected TV. As users work, socialize, research, and buy products across different Internet connectable devices, companies will continue to shift focus to reaching users more effectively across their multiple devices. Although a person may own and use different devices to communicate over the Internet, the relationship among different devices and users of the different devices is not readily apparent to outsiders such as companies seeking to reach the person across his or her multiple devices.

A person may use different devices with different device identifiers to communicate over the Internet. A person may communicate anonymously over the Internet without disclosing a personal identifier. User device connections to the Internet may be transitory and dynamic. Devices typically connect to the Internet through a connection point associated with an Internet Protocol (IP) address. User devices may use different network addresses at different times. In the course of communication over the Internet, user device identifying information may be exchanged such as a device identifier or a user identifier. However, the identifying information used during Internet communication by a user using one device may be different from the identifying information used during Internet communication by the same user using a different device. A person may use different IP addresses at different times when communicating over the Internet with different devices. A user may have multiple different email accounts and may participate in use of social media under different pseudonyms. Thus, there is no readily available solution to identify users using different devices accessing the Internet.

SUMMARY

A system is provided that includes an index structure that associates devices with device feature information. A pairing engine determines device pairs based at least in part upon at least a portion of the device feature information. A feature vector generation engine produces feature vectors corresponding to determined device pairs, based at least in part upon feature values associated within the index structure with devices of the determined device pairs. A scoring engine determines scores to associate with determined device pairs based at least in part upon produced feature vectors associated with the determined device pairs. A graph structure includes nodes that represent devices of determined device pairs and includes edges between pairs of nodes that indicate determined device pairs. A clustering engine identifies respective clusters of nodes within the graph structure that represent respective groups of devices.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGS. **1**A-**1**B are illustrative drawings representative of network connected device usage at different network

addresses associated with different physical locations in accordance with some embodiments.

FIG. **2** is an illustrative schematic diagram of a system in accordance with some embodiments.

FIG. **3** is an illustrative process flow diagram representing configuration of the digital identity pairing engine of the system of FIG. **2** in accordance with in some embodiments.

FIG. **4**A is illustrative drawing representing a plurality of example device pairs identified using the pairing engine of FIG. **2** in accordance with some embodiments.

FIG. **4**B is an illustrative drawing representing the example device pairs of FIG. **4**A associated with example feature vectors in accordance with some embodiments.

FIG. **5** is an illustrative process flow diagram representing configuration of a feature vector generation engine of the system of FIG. **2** in accordance with some embodiments.

FIG. **6**A is an illustrative flow diagram of a process representing example details of a first example rule applied according to the process of FIG. **5**, to produce an example feature value as a function of feature information associated with a candidate device pair, in accordance with some embodiments.

FIG. **6**B is an illustrative flow diagram of a process representing example details of a second example rule applied according to the process of FIG. **5**, to produce an example feature value as a function of feature information associated with the candidate device pair, in accordance with some embodiments

FIG. **7** is an illustrative chart providing examples of lower level feature information within an input storage system and rules within a rule storage of the system of FIG. **2** that are used to produced feature values in accordance with some embodiments.

FIG. **8** is an illustrative flow diagram representing configuration of a scoring engine of the system of FIG. **2** in accordance with some embodiments.

FIG. **9** is an illustrative flow diagram representing details of a training module of the scoring engine of FIG. **8** in accordance with some embodiments.

FIG. **10**A is an illustrative drawing representing example in which an unsupervised pre-training module remaps the device pairs of FIG. **7**B to device pair clusters based upon the sample device pair feature vectors in accordance with some embodiments.

FIG. **10**B is an illustrative drawing representing use of a remapping module of the training module of FIG. **8** in accordance with some embodiments.

FIG. **11** is an illustrative diagram representing a data flow process including data transformations that occur within computer readable storage devices within the scoring engine of the system of FIG. **2** in accordance with in some embodiments.

FIG. **12** is an illustrative flow diagram representing configuration of a cluster engine of the system of FIG. **2** in accordance with some embodiments.

FIG. **13** is an illustrative drawing representing an example graph produced using the cluster engine of FIG. **12** for the device pairs of FIGS. **4**A-**4**B in accordance with some embodiments.

FIG. **14** is an illustrative drawing representing a set of example proposed user device clusters user device clusters identified using the cluster engine of FIG. **12** within the graph of FIG. **13** in accordance with some embodiments.

FIG. **15** is an illustrative drawing representing an example set of final device clusters produced using the cluster engine of FIG. **12** based upon the proposed user device clusters of FIG. **14** in accordance with some embodiments.

FIG. **16** is an illustrative diagram representing a data flow process diagram including data transformations that occur during user device clustering within the cluster engine of the system of FIG. **2** in some embodiments.

FIG. **17** is an illustrative drawing representing the system of FIG. **2** configured to operate in a network environment accordance with some embodiments.

DESCRIPTION OF EMBODIMENTS

The following description is presented to enable any person skilled in the art to create and use a system to group Internet connectable devices belonging to the same user based upon anonymous Internet user data associated with the devices. Various modifications to the embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the invention. Moreover, in the following description, numerous details are set forth for the purpose of explanation. However, one of ordinary skill in the art will realize that the invention might be practiced without the use of these specific details. In other instances, well-known data structures and processes are shown in block diagram form in order not to obscure the description of the invention with unnecessary detail. Identical reference numerals may be used to represent different views of the same item in different drawings. Flow diagrams in drawings referenced below are used to represent processes. A machine such as a computer system that includes one or more processors and storage devices is configured to perform these processes. The flow diagrams represent configuration of one or more processors and/or storage devices of a computer system configured using computer program code to perform the acts described with reference to the flow diagrams. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

Digital Identity Pairing Overview

Referring to FIG. **1**A, the people who live at household A **150** use several mobile internet-connected digital devices to communicate with the internet **101**. Specifically, the laptop computer systems (**151** and **161**), the cellular phones (**152** and **162**), and the tablet computer system **157** are all mobile internet-connected digital devices, each having a unique device identifier. These mobile devices can be and often are taken with the user to other locations that provide wireless (or wired) internet access. For example, user X in the household may use laptop computer system **151** and cellular phone **152** both at household **150** and at a workplace W **110**.

When user X brings laptop computer system **151** and cellular phone **152** to workplace W **110** those two devices can no longer use Wi-Fi router **165** to access the internet **101**. Instead, user X connects laptop computer system **151** to the local area network **129** at work and configures cellular phone **152** to use a local Wi-Fi network provided by wireless access point **151**. With these two internet connections, user X then will be able to access the internet **101** using laptop computer system **151** and cellular phone **152** through the firewall/proxy server **121** at the workplace **110**.

When user X is at home A **150**, the laptop computer system **151** and the cellular phone **152** will both use a single Internet Protocol (IP) address A **163** that is on Wi-Fi router **165**. Similarly, when that same user X is at workplace **110**, laptop computer system **151** and cellular phone **152** will both use the single IP address W **123** that is on firewall/proxy server **121**. Of course, X can selectively send and receive information using one or the other or both of the laptop computer system **151** and the cellular phone **152**. However, a third party who communicates with X over the computer system **151** may be unaware that X also can be reached over cellular phone **152**, and a third party who communicates with X over the cellular phone **152** may be unaware that X also can be reached over the laptop computer **151**.

Nevertheless, given the specific internet usage pattern data described above with reference to FIG. **1**A, an astute observer who knows nothing about user X and who is unaware that user X owns both the laptop computer system **151** and the cellular phone **152** could make the rational inference that laptop computer system **151** and cellular phone **152** are very likely used by the same person since laptop computer system **151** and cellular phone **152** are used together at both household **150** and at workplace **110**. After having made such an inference, a machine may link together a digital identifier associated with laptop computer system **151** and a digital identifier associated with cellular phone **152**. Such a pairing of distinct digital identities to a single user is referred to as digital identity pairing.

There are many different internet-connected devices (**151**, **152**, **161**, **162**, **157**, and **159**) at household A **150** that may access the internet **101**. All of these internet-connected devices at household A **150** will use the same IP address A **163** that is on Wi-Fi router **165**. Similarly, when the user of laptop computer system **151** and cellular phone **152** is at workplace **110** those two devices will have the same IP address W **123** as all of the other internet connected devices (**111**, **112**, **113**, **114**, **115**, and **116**) that share the connection to the internet **101** through Wi-Fi router **125** and firewall/proxy device **121**. Therefore, with a large number of internet connected devices in use, the task of specifically identifying laptop computer system **151** and cellular phone **152** as a pair of related digital devices used by the same human user is a difficult task. In accordance with some embodiments, Internet web servers are used to send information to user devices and to receive information from user devices in response to user device requests. Server logs stored in computer readable storage devices associated with the web servers store information concerning these send/receive requests. In accordance with some embodiments, analysis of the server logs is performed to determine digital identity candidate digital identity pairings.

System Overview

FIG. **2** is an illustrative block diagram representing a system **200** in accordance with some embodiments. The system includes a computer readable storage device **208** to store device identifiers and indicia of features associated with the identified devices, a pairing engine **202**, a feature vector generation (FVG) engine **212**, a scoring engine **204** and a clustering engine **206**. The pairing engine **202** identifies candidate device pairs based upon lower level preliminary pairing-identifying features indicated within the storage device **208** as associated with the devices of the candidate device pairs. It will be understood that the term "candidate" is used herein to indicate that the pairing determinations are preliminary and may undergo further evaluation in the scoring engine **204** and the clustering engine **206**. The FVG engine **212** produces feature vectors in response to candidate device pairs as a function of additional feature information associated with individual devices of the candidate device pairs. More particularly, the pairing engine provides to the FVG **212** candidate device pair identifiers. In response to the received candidate device pair identifiers, the FVG engine **212** uses individual identi-

ties of the devices of the candidate device pair to identify additional feature information, associated within the storage device **208** with those individual devices of that candidate device pair. The FVG engine **212** generates feature vector values to associate with the identified candidate device pair as a function of the additional feature information identified as associated with individual devices of the identified candidate device pair. The scoring engine **204** determines scores to associate with identified candidate device pairs based at least in part upon feature vectors produced using the FVG engine **212**. The clustering engine **206** produces a graph structure within a computer readable storage device that represents at least a portion of the identified candidate device pairs and the determined scores. The clustering engine **206** identifies digital identity clusters $U_1, U_2, \ldots U_N$ within the graph structure indicative of clusters or groups of devices that are associated with the same user.

In accordance with some embodiments, each individual cluster of digital identities stored within a storage device acts as a digital identity group communication structure $U_1, U_2, \ldots U_N$, comprising a group of devices that are associated with the same user and that can be used to communicate with the user who owns the grouped devices. Thus, for example, digital identity group communication structure U_1 includes devices $D_{A1}, D_{A2}, \ldots D_{AK}$. More specifically, digital identity groupings reveal the multiple devices that are grouped that can be used to communicate with the owner of the grouped devices. In contrast, in the absence of digital identity grouping, a third party who wishes to communicate with a particular device user cannot as readily identify the different devices over which to communicate with that particular device user. Thus, a digital identity group communication structure identifies a group of devices associated with the same user so as to expand the available avenues of communication with that user, thereby enhancing opportunities for communication with the user. It will be appreciated that device identity groups $U_1, U_2, \ldots U_N$ that previously were unknown to some third party wishing to communicate with users associated with the groups are determined based upon digital device pairings determined using the pairing engine **202**, scoring using the scoring engine **204** based upon feature vectors produced using the FVG engine **212**, and clustering using the clustering engine **206**.

More particularly, in some embodiments, the computer readable feature storage device **208** includes one or more storage devices that store an index structure **209** that associates device identifier information with lower level feature information and label data. The lower level feature information can be stored within the storage device **208** itself or can be stored in a different storage location (not shown) indicated within the storage device **208**. The low level data can be obtained from server logs of internet connected device requests, for example. The system includes a computer readable rules storage device **213** that stores rules used to generate higher level feature values based upon lower level feature information indicated within the storage device **208**. More specifically, the rules storage device **213** stores instructions to configure the FVG engine **212** to identify additional higher level feature information to associate with a candidate device pair, in response to the pairing engine **202**, based upon lower level features associated with individual devices of candidate device pairs and to use the additional feature information, together with low level feature information, to produce feature values used within feature vectors. The system **200** includes buffer circuitry **210**, **211** configured to receive as input certain lower level feature information and the label data, respectively. In some embodiments, the lower level feature information received within the buffer circuitry **210** includes deviceIDs, common source/destination identifiers (e.g. IP address), and timestamps.

The pairing engine **202** uses the preliminary pairing-identifying feature information such as a deviceID, a common source/destination identifier (e.g. IP address), and a timestamp to determine preliminary pairings of devices. As used herein, the term "device pair" refers to an information structure that is stored in a storage device and that indicates a pairing of distinct digital identities to a single user. The pairing engine **202** acts as a pre-filter that produces device pairs that represent a preliminary "candidate" pairing of devices. The pairing engine **202** also associates a preliminary pair score with each device pair, which provides an indication of the likelihood that the devices of the pair actually are associated with the same user.

The FVG engine **212** is configured to use a set of rules from the rules storage device **213** to produce feature vectors to associate with candidate device pairs as a function of lower level feature information associated within the storage device **212** with the different individual devices of the candidate device pairs. The lower level feature information indicated within the storage device **212** includes user profile information associated with individual devices. The lower level feature information indicated within the storage device **208** may be gathered over the Internet and may include information such as demographic data such as age, gender or marital status and/or behavior data such as user intent to purchase some item or personal interests such as likes or dislikes.

The scoring engine **204** receives as input, candidate device pair identifiers (such as deviceIDs of the devices within the pairs), associated preliminary pair scores produced by the pairing engine **202**, associated feature vectors produced by the FVG engine **212**, and associated label information. In some embodiments, the preliminary pair scores associated with candidate device pairs that are produced using the pairing engine **202** are included as feature values of features vectors associated with candidate device pairs. The scoring engine **204** produces precision pair scores for the received device pairs as a function, at least in part, of their associated feature vectors. The precision pair scores produced using the scoring engine **204** are indicative of the likelihood that the device pairs in fact are associated with the same user. Pair scores produced by the scoring engine **204** preferably supplant the preliminary pair scores produced using the pairing engine **202**.

The cluster engine **206** receives as input, device pairs and associated precision pair scores produced using the scoring engine **204**. The cluster engine **206** produces a graph structure (not shown) within a computer readable storage device based upon the device pairs. Graph nodes represent devices associated with device identifiers. Graph edges indicate potential pairings of devices of candidate device pairs. Precision pair scores are associated with the graph edges. The cluster engine **206** identifies user device clusters within the graph based upon pair scores associated edge connections within the graph. The cluster engine **206** identifies groups of user devices that are associated with a common user. The cluster engine **206** stores in computer readable storage device **214** information that identifies respective groups of devices associated with respective users.

Example user device clusters $U_1, U_2, \ldots U_N$, which acts as digital identity group communication structures for corresponding owners of the devices in the clusters, that are identified using the clustering engine **206** are shown stored

within the output storage system **214** associated with corresponding user identifier information. Different groups of device identifiers that are associated with the different respective users. Each cluster associates deviceIDs with a unique user, although the actual identity of the unique user is unknown. A first user device cluster U_1 indicated by device IDs ($D_{A1}, D_{A2}, \ldots, D_{AK}$) is associated with a unique user U_1, for example. Clusters of device IDs may be referred to herein as "device clusters" or "user device clusters".

In some embodiments, the cluster engine **206** provides feedback **216** to the scoring engine **204**. One example feature value considered by the scoring engine **204**, in accordance with some embodiments, is an indication of the number of additional devices indicated in the graph as potentially paired to a device of a candidate device pair. This example indication provides a view of the larger picture that is not necessarily present in the triad information alone. This additional potential pairing information is an example of "feedback" **216** through which the cluster engine **206** can affect precision scores produced using the scoring engine **204**, which in turn, can influence precision scores associated with edges of the graph structure.

Also, for example, the cluster engine **206** can provide a previous precision pair score as feedback **216** to the scoring engine **204** for determination by the scoring engine **204** of a later updated precision pair score. A previous pair score produced using the score engine **204** can allow the scoring engine **204** to determine a new updated pair score to improve upon a previous pair score that was generated when more limited information about a candidate device pair was available, e. g., when less feature information was available.

Devices and Device Identifiers

As used herein the terms "device ID" and "device identifier" and "digital identity" refer to an identifier for a user device, a user device program stored in a computer readable storage device, a user device information structure (e.g., a cookie) stored in a computer readable storage device or a device user identity, for example. As used herein, "computer readable" storage device refers to a non-transitory storage device from which stored information can be extracted using a programmable computer or device, such as a smartphone, for example. Examples of deviceIDs include web browser cookies, cellular telephone device identifiers, MAC addresses, userids, and other identifiers that are linked to a specific client device, client program, or device user. As used herein, the term "device" or "user device" shall be used generally to refer to digital entities such as laptop computer systems, desktop computer systems, the cellular phones, tablet computer system, smart watches, smart devices such as internet connected appliances, and web browsers, for example. The teachings of the present disclosure may be used with a wide variety of different deviceIDs. In example digital identity pairings that will be disclosed, web browser cookies on laptop computer systems, desktop computers, and device identifiers on cellular phone devices are used as deviceIDs. However, the disclosed techniques may be used with any other suitable deviceIDs that can be used to identify specific client devices, web browsers, users, or other digital identities. As used herein, the letter "D" may be used to refer to a physical device such as a mobile phone and the letter "C" may be used to refer to a cookie that is indicative of physical device such as a laptop computer, for example. As noted above, both a mobile phone and a browser cookie encoded in a computer readable storage device may be referred to herein as "devices".

Features and Feature Information

As used herein, the term "features" refers to attributes associated with a device. As used herein, the term "feature information" refers to information indicative of one or more features. In accordance with some embodiments, feature information is collected through observing Internet communications. In accordance with some embodiments, feature information is used to produce, feature values which are constituent components of feature vectors.

For example, lower level feature information includes information such a source device attributes (e.g., iPad, ID for installed applications (e.g., the application named, "Words with Friends" is installed), device user's ender, device ID, geographic location, dates, etc.), destination device attributes (e.g., MacOSX installed, Three other devices associated, etc.), contextual attributes (e.g., web site topics, etc.), and IP attributes (e.g., how many unique devices have been observed on this IP address, how many days have has this IP address been observed). In general, information relating to devices of device pairs comes from several broad categories: spatial and temporal observations, demographics and satellite pairs. Low level demographic information includes information such as: What are the age, gender, home location and consumer interests of the C and D sides of the pair? Do the C and D descriptions appear consistent with one unique user? Low level satellite pairs information includes information such as: Are there other pairs related to C or to D? How do these pairs compare with the current pair? The satellite pairs information is an injection of information from the perspective of a device graph (described more fully below) back into a pairing prediction process.

The FVG engine **212** produces feature vectors that represent higher level feature information about a candidate device pair as a function of lower level feature information. More particularly, in some embodiments, the FVG engine **212** is configured to use rules to produce feature vectors as a function of device usage information. In accordance with some embodiments, many higher level features are developed based upon spatio-temporal information. In some embodiments, the feature vector generation engine **212** evaluates lower level features using spatio-temporal heuristics (e.g., collect all observations for device A on IP address C, how often does device A appear on IP address C at nighttime or during the weekends.

Pairing Engine

The following explanation of the pairing engine **202** in accordance with some embodiments includes information from commonly owned co-pending U.S. Patent Application Publication No. US 1014/0095320, filed May 10, 1013, published Apr. 3, 2014.

FIG. **3** is an illustrative process flow diagram representing configuration of the digital identity pairing engine of the system of FIG. **2** in accordance with in some embodiments. Initially, module **510** collects Internet usage data from the feature input storage system **208** for many different digital identities. In some embodiments, the Internet usage data that is collected for analysis include data triads that include a client device deviceID, a network source/destination address identifier (e.g., an IP address), and a timestamp.

The common source/destination identifier is the identity of some source or destination that client devices (as identified by their deviceIDs) will likely have in common if the two client devices are related. In the situation depicted in FIG. **1**A, the IP addresses are the common source/destination identifier that may be used to link related client devices.

Specifically, the digital identity pairing system may use the fact that laptop computer system **151** and cellular phone **152** both share the IP address W **123** when at work and the IP address A **163** when at home to deduce that laptop computer system **151** and cellular phone **152** are related client devices.

The timestamps in each data triad may be used to ensure that the data used is relevant. The ownership of internet connected devices may change over time such that very old internet usage data should not be used. In some embodiments, the network source/destination addresses are implemented using IP addresses. Many IP addresses are "dynamic addresses" that may be used by different entities at different times. Thus, internet usage data observations should have relatively close temporal relations in order to provide accurate digital identity pairing results. In addition to ensuring that internet usage observations are temporally proximate, certain embodiments of the disclosed system use the timestamps of internet usage data triads in a more sophisticated manner as will be disclosed in a later section of this document.

The triads of internet usage data (deviceIDs, common source/destination identifier, and timestamp) may be collected by internet web servers that track each internet server request received. Alternatively, for example, in some embodiments individual application programs (such as games, media aggregators, utilities, etc.) that run on client devices and report usage information to servers on the internet also can be sources of usage data.

Referring again to FIG. **3**, after collecting internet usage data the next step in digital identity pairing, in accordance with some embodiments, is to determine a set of candidate digital identity pairs. Generally, in an Internet environment, there are many millions of different digital identities involved in internet activities every day. Attempting to analyze every possible permutation of digital identities as a potential digital identity pair would be an extremely difficult and probably futile task. Thus, to reduce the size of the digital identity pairing problem, module **520** analyzes the gathered internet usage data to identify a much smaller number of candidate digital identity pairs that have a decent probability of being related.

In embodiments that use IP addresses as common source/destination identifiers, two different techniques have been used to select potential digital identity pairs for further analysis. A first strategy is to examine the number of different digital identities known to use the same IP address. Specifically, if less than a threshold number of digital identities are known to use a specific IP address then all of the different logical pairings of digital identities from that single IP address may be viewed as potential digital identity pairs. The reasoning is that if there are just a few different digital identities related to a single common IP then there is a good probability that some of those different digital identities are associated with the same person and that one may be able to statistically link the digital identities belonging to that same person. For example, a family household that shares a single internet account will likely have family members that use more than one digital identity that can be statistically linked.

In some embodiments, for example, a threshold value is set to six such that if there are six or less digital identities seen at a particular IP address then various logical combinations of those six or less digital identities may be considered potential digital identity pairs. For example, in FIG. **1**A household A **150** has only six different digital devices (**151**, **152**, **161**, **162**, **157**, and **159**) that couple to the internet **101** through a single IP address A **163** on Wi-Fi router **165** such that the various digital devices in household A **150** may be considered as candidate digital identity pairs. In contrast, a very large number of digital devices couple to the internet **101** though the single IP address W **123** at workplace W **110** such that the digital identity pairing system does not immediately consider all of the combinations of digital devices at workplace **110** as potential digital identity pairs. In effect, the system identifies family households (which often have less than six internet connected devices) and then attempts to pair up digital devices from the family household that are used by the same user.

In another embodiment, the digital identity pairing system considers the specific IP address origin and determines whether that IP address is an address where paired digital identities are likely to be found (such as household residences as set forth above). The static IP addresses on the internet are commonly allocated by the Internet Corporation for Assigned Names and Numbers (ICANN). By examining who owns a particular IP address, one may make a determination as to whether it will be easy to identify related digital identities may be located at that IP address. Thus, for example, IP addresses that are used by an internet service provider (ISP) to provide residential internet service may be good IP addresses to use when identifying potential digital identity pairs. Various other systems of identifying residential household IP addresses may also be used. In addition, other techniques of identifying likely digital identity pairs may also be used in module **520** in addition to or instead of the systems for identifying residential households.

After selection of sets of potential digital identity pairs using module **520**, module **540** then processes the gathered internet usage data, i.e. the gathered featured information, specifically the triad information, to determine preliminary pair scores for the candidate digital identity pairings. Since the observed internet usage data will vary over time and certain chance activities may cause false digital identity associations to be detected, the preliminary pair scores produced using the pairing engine **202** may be post-processed to remove noise. For example, association scores may be smoothed out over time using various techniques. Thus, at stage **560**, the association score data that has been generated over time may be Module **560** post-processes pair scores produced using module **540** such that outlier data points are largely filtered out. Module **560** produces a set of high probability digital identity pairings.

It will be appreciated that in some embodiments, the preliminary pair scores produced using the pairing engine **202** provide a preliminary pair score that is more coarse, i.e. less precise, than the pair score produced using the scoring engine **204**. In some embodiments, in determining a precision pair score, the score engine **204** takes into account the preliminary pair score produced using the pairing engine **202**. The pairing engine **202** identifies these candidate device pairs, together with their associated preliminary pair scores, to the scoring engine **204**. Moreover, as more fully explained below, the device IDs of identified candidate device pairs are used with the index structure **209** within storage device **208** to retrieve additional feature information for use by the FVG engine **212** to produce feature values corresponding to the candidate device pairs.

FIG. **1**B illustrates the household A **150** and the workplace W **110** of FIG. **1**A with an additional household B **180** and a cyber cafe C **190**. In household A **150** there are two digital device users: User X and User Y. User X regularly uses laptop computer system **151** and cellular phone **152** such that the digital identities of laptop computer system **151** and cellular phone **152** should be identified as paired digital

identities. User X's laptop computer system **151** is digitally identified using a cookie (C) on a web browser such that it is indicated as C_X and user X's cellular phone **152** is digitally identified with a deviceID (D) such that it is indicated as D_X. This digital device referencing indication nomenclature will be used with the other laptop and cellular phones as well. User Y of household A **150** regularly uses laptop computer system C_Y **161** and cellular phone D_Y **162** such that C_Y **161** and D_Y **162** should also be identified as a related pair of digital identities. Recall that all of the digital devices (C_X **151**, D_X **152**, C_Y **161**, and D_Y **162**) in household A **150** will use the same IP address A **163** that is assigned to Wi-Fi router **165** when those digital devices are used at household A **150**. Note that although this embodiment is based upon using web browser cookies and mobile device identifiers, any other similar identifiers that can be associated with the digital devices may be used.

User Z resides at household B **180** regularly uses laptop computer system C_Z **181181** and cellular phone D_Z **182182182**. While at household B **180**, both C_Z **181181** and D_Z **182182182** will use IP address B **183** that is assigned to Wi-Fi router **185** in use at household B **180**. Both user X and user Z work together at workplace W **110** such that C_X **151**, D_X **152**, C_Z **181181**, and D_Z **182182182** are regularly used at workplace W **110**. While at workplace W **110** those digital devices will all use IP address W **123** that is assigned to firewall/proxy **121** at workplace W **110**. Many other digital devices (**111**, **112**, **113**, **114**, **115**, and **116**) will also use IP address W **123** at workplace W **110**.

Finally, FIG. **1**B also illustrates a cyber cafe **190** that offers free Wi-Fi service to customers of cyber cafe **190**. User Z and user Y frequent cyber cafe C **190** such that C_Z **181181**, D_Z **182182182**, and D_Y **162** are illustrated at cyber cafe **190** where IP address C **193** is used on Wi-Fi router **195**. Note that many other visitors (not shown) will also frequent cyber cafe **190**. However, the various digital devices only seen together at cyber cafe **190** will not be considered potential digital identity pairs since there are too many digital identity pairings seen together at cyber cafe **190**.

After collecting Internet usage data (as explained with reference to module **510**), the next step in identifying digital identity pairs is to select a set of potential digital identity pairs as explained with reference to **520** of FIG. **3**. As set forth in the previous section, the various combinations of digital identities that are associated with IP addresses having six or less digital identities may be selected as potential digital identity pairings. In the example of FIG. **1**B, the various combinations of the digital devices at households A **150** and B **180** are therefore considered candidate digital identity pairs. (The workplace W **110** and cyber cafe C **190** have too many different digital identities associated with them and thus do not provide good candidates for digital identity pairs.) For simplicity, only the possible digital identity pairings (C_X, D_X), (C_X, D_Y), (C_Y, D_X), and (C_Y, D_Y) from household A **110** will be analyzed in this example.

In one particular embodiment, for example, the pairing engine **202** uses a variation of Bayesian probability analysis to calculate a preliminary pair score, also referred to as an "association" score, for each of the potential cookie and deviceID digital identity pairs. In some embodiments, a "support" score and "confidence" score also are determined. The support, confidence, and association scores may be defined as follows:

Support=P(cookie, deviceID)

Confidence=P(cookie\deviceID)

Association (cookie→deviceID)=P(cookie|deviceID)/P (cookie)

These three scores may be used to identify digital identity pairings and to rate the confidence in a digital identity pairing that has been made. The support score gives an indication of how much data support there is for the analysis of this particular cookie and deviceID pair. The confidence score gives an indication of how much confidence there is in the preliminary pair score, also referred to as an association score. The preliminary pair score produced using the pairing engine **202** provide rating of how closely the cookie and deviceID are associated. In some embodiments, the preliminary pair score produced using the pairing engine **202** is used as a feature in score engine **204**.

In some embodiments, the support score, confidence score, and preliminary pair (association) score are calculated using the internet usage, triad feature information, described above. The following relationships describe how the internet usage information is used in accordance with some embodiments to calculate the support, confidence, and preliminary pair (association) scores.

co-occurrences(cookie, deviceID)=number of times both cookie and deviceID at the same location (same source identifier IP address).

P(cookie, deviceID)=co-occurrences(cookie, deviceID)/total sample size

P(cookie\deviceID)=co-occurrences (cookie, deviceID)/occurrences(deviceID)

P(cookie)=number of occurrence (cookie)/total sample size

FIG. **4**A is an illustrative drawing representing a plurality of example candidate device pairs and preliminary pair scores identified using the pairing engine **202** of FIG. **2**. For example, a device pair data structure associates a preliminary pair score S3 with a candidate device pair having device identifier D1 and having cookie identifier C2. It will be understood that in accordance with some embodiments, preliminary pair score S3 is a value between 0 and 1. The candidate device pairs and their corresponding preliminary scores are provided as input to the score engine **204**. In some embodiments, the preliminary scores are provided as feature values within feature vectors.

Feature Vectors

As used herein a feature vector refers to a set including multiple feature values that are associated with a candidate device pair. Feature vectors are information structures stored in a computer readable storage. Each feature value represents one or more features associated with a candidate device pair.

In accordance with some embodiments, feature vectors are used to evaluate correlations between identified device pairs during training of the scoring engine **204**. The FVG engine **212** produces feature values that are included within a feature vector as a function of feature information that may be associated with the individual devices of a candidate device pair or with the pair of devices and with one or more rules. Thus, in some embodiments, for example, each feature vector represents N features (e.g., F_1, F_2, F_3, F_4, . . . , F_N). In some embodiments, each candidate device pair is associated with a feature vector (V) in which each feature has an integer numerical value (F_{value1}, F_{value2}, . . . F_{valueN}) that indicates the state of that feature as relates to the associated device pair. Thus, for example, the FVG engine **212** uses a rule to produce feature vector value F_{value1} based upon feature information associated with F_1 features associated with the devices of the pair; uses a different rule to produce feature vector value F_{value2} based upon different feature information associated with F_2 features associated with the devices of the pair; etc., for example.

Feature Vector Generation Engine

FIG. **5** is an illustrative process flow diagram **300** representing configuration of a FVG engine **212** in accordance with some embodiments. Modules of the diagram correspond to computer program code (not shown) that configures a computer system that includes one or more processors and/or control components to cause performance of the specific acts represented by the different modules. Module **302** selects a candidate device pair. Module **303** uses the index structure **209** stored within the storage device **208** to retrieve stored feature information associated with individual device IDs of the selected pair. Module **304** selects a rule from the rules storage device **213**. Module **306** produces a feature value, for use in a feature vector, based upon the retrieved feature information and the currently selected rule.

Decision module **308** determines whether there are additional rules to apply to the currently feature information associated with the devices of the selected candidate device pair. In response to a decision by module **308** that additional rules are to be applied to the currently selected candidate device pair, control flows back to module **304** and another rule is selected. Which rules to apply to a candidate device pair, for example, may depend upon what feature information is available for the pair. On the other hand, in response to a determination by decision module **308** that no additional rules are to be applied to the currently selected candidate device pair, module **310** provides to the scoring engine **204** a feature vector that includes feature values produced based upon individual feature information associated with the individual devices of the currently selected candidate device pair, and decision module **312** determines whether the feature input storage system **208** includes additional device pairs to evaluate. In response to a determination by decision module **312** that additional device pairs are to be evaluated, control flows to module **302** and another candidate device pair is selected. On the other hand, in response to a determination by decision module **312** that no additional device pairs are to be evaluated, the process ends and awaits identification of additional candidate device pairs.

It will be appreciated that in accordance with some embodiments, the pairing engine **202** provides a preliminary determination of candidate device pairs. The initial candidate pair determination is used as an index for use to search the feature input storage system **208** for separate and possibly different feature information associated with the separate different devices of a currently selected candidate device pairs. The FVG engine **212** applies one or more rules from the rules storage device **213** to the separate and possibly different information identified using the pairing-engine-provided index information so as to produce feature values used to produce a feature vector associated with selected candidate device pair.

For example, a candidate device pair may include a laptop computer and a smartphone. Separate internet usage information, also referred to herein as low level feature information, may be collected separately for the laptop and for the smartphone. The index structure **209** is used to identify usage information collected for the laptop and to identify separately collected usage information collected for the smartphone. In module **306**, FVG engine **212** applies one or more rules to produce a feature vector value for the device pair based upon the usage information collected separately for the two devices of the pair.

It will be further understood, as explained more fully below, that the produced feature vectors of numerous candidate device pairs are used to produce precision pair scores for numerous corresponding candidate device pairs. FIG. **4**B is an illustrative drawing representing the example device pairs of FIG. **4**A associated with example feature vectors V1 to V33 in accordance with some embodiments. In some embodiments, the weights shown in FIG. **4**A may be subsumed as individual feature values within feature vectors. More particularly, through supervised and unsupervised training, the scoring engine **204** produces precision pair scores to associate with the candidate device pairs identified by the pairing engine **202**. The precision pair scores produced using the scoring engine **204** are used, in turn, by the clustering engine **206** to facilitate clustering of user devices.

FIG. **6**A is an illustrative flow diagram of a process **650** representing example details of an example first rule applied by module **306** to produce an example higher level feature value as a function of the lower level feature information associated in the index structure **209** with two devices of an example candidate device pair identified using the pairing engine **202**, in accordance with some embodiments. It will be appreciated that, that module **306** configures the FVG engine **212** according to details of the rule currently selected using module **304**, to produce a feature value based upon lower level feature information. Module **306** configures the FVG engine **212** differently to apply different rules. More particularly, in accordance with some embodiments, module **306** uses instructions stored in rules storage **213** to configure the FVG engine **212** according to a rule selected using module **304** to determine a feature value for a higher level feature based upon lower level feature information selected using module **303** associated with devices of a candidate device pair selected using module **302**. As used herein, a "higher" level feature is a feature that has a value that is dependent upon the values of one or more other features. A produced feature value indicates an attribute of the relationship between device C and device D.

The example first rule process **650** of FIG. **6**A uses multiple pieces of lower level feature information indicated in the index structure **209** to produce a feature value for a feature indicative of the strength of the relationship between a candidate device pair includes device C and device D. As explained herein, a vector includes multiple feature values, each produced using a different rule. The example rule produces one of five possible feature values for the example feature: 0, 1, 2, 3, and 4. The higher level feature value produced using the example rule is included in a feature vector as explained more fully below. The example rule is hierarchical. It includes a hierarchy of branching determinations in which some determinations are predicates of other dependent determinations.

Referring to FIG. **6**A, first decision module **652** determines whether both devices of the example candidate device pair C and D access information from the same IAB category, such as an Interactive Advertising Bureau (IAB) category. The first decision module **652** produces a feature value of 0 in response to a determination that C and D do not access the content with the same category. The first process **650** ends with production of feature value 0 if C and D do not access the content with the same category.

In response to the first decision module **652** determining that both devices of the example candidate device pair access information from the same category, a second decision module **654** determines whether the category is sports. The second decision module **654** produces a feature value of 1 in response to a determination that that although both C and D access the content with the same category, the estimated geographic locations of devices C and D are more than one-hundred miles apart. The first process **650** ends with production of feature value 1 if C and D access the

content with the same category but devices C and D are located more than one-hundred miles apart.

In response to the second decision module **654** determining that both devices access the sports category, a third decision module **656** determines whether either of the two devices accesses the sports content during times on weekends or during times in the evenings. The third decision module **656** produces a feature value of 2 in response to a determination that neither device C nor device D accesses sports content during weekends or evenings. The first process **650** ends with production of feature value 2 if both device C and device D access sports content but neither of them accesses sports content during weekends or evenings.

In response to the third decision module **656** determining that at least one of the devices accesses the sports category on the weekend or in the evening, a fourth decision module **658** determines whether only one or both of the two devices accesses the sports category on the weekend or in the evening. The fourth decision module **658** produces a feature value of 3 in response to a determination that only on or the other of the two devices accesses the sports category on the weekend or in the evening. The fourth decision module **658** produces a feature value of 4 in response to a determination that both of the two devices access the sports category on the weekend or in the evening.

FIG. **6**B is an illustrative flow diagram of a process **660** representing example details of an example second rule applied by module **306**, in accordance with some embodiments. The example second rule process **660** of FIG. **6**B is applied to the same candidate pair, C and D, as in FIG. **6**A, to produce one of six possible feature values for the example feature: 0, 1, 2, 3, 4, and 5. A first decision module **662** determines whether both devices of the example candidate device pair C and D access information from the same category. The first decision module **662** produces a feature value of 0 in response to a determination that C and D do not access the content with the same category. The process **660** ends with production of feature value 0 if C and D do not access the content with the same category.

In response to the first decision module **662** determining that both devices of the example candidate device pair access information from the same category, a second decision module **664** determines whether the devices C and D are geographically located within one-hundred miles of each other. The second decision module **664** produces a feature value of 1 in response to a determination that that although both C and D access the content with the same category, they are not located within one-hundred miles of each other. The second process **660** ends with production of feature value 1 if C and D access the content with the same category but are not located within one-hundred miles of each other.

In response to the second decision module **664** determining that the devices are located within one-hundred miles of each other, a third decision module **666** determines whether the category is sports related. The third decision module **666** produces a feature value of 2 in response to a determination that the category is not sports related. The second process **660** ends with production of feature value 2 if both device C and device D access the same category are located within one-hundred miles of each other, but that the category is not sports related.

In response to the third decision module **666** determining that the accessed content is sports related, a fourth decision module **668** determines whether neither, only one, or both of the two devices accesses the soccer sports category on the weekend or in the evening. The fourth decision module **668** produces a feature value of 3 in response to a determination that neither devices accesses soccer sports related content on the weekends or evenings. The fourth decision module **668** produces a feature value of 4 in response to a determination that only on or the other of the two devices accesses the sports category on the weekend or in the evening. The fourth decision module **668** produces a feature value of 5 in response to a determination that both of the two devices access the sports category on the weekend or in the evening.

It will be appreciated that the rules processes of FIGS. **6**A-**6**B provide two example rules applied by the FVG engine **212** to produce example feature values for higher level features produced based upon lower level feature information. It will be understood that the feature vector generation engine **212** is configured to apply numerous rules to determine numerous corresponding features values to include in a feature vector based upon observed lower level feature information.

Thus, the FVG engine **212** produces high level feature values (e.g., 0, 1, 2, 3, 4, 5) using low level feature information and a set of pre-defined rules. In some embodiments, generating higher level feature values involves using heterogeneous lower level features such as features indicative of categories (e.g., IAB categories), geographical location (e.g., within one-hundred miles) and time (e.g., on weekends or evenings). Moreover, in some embodiments, generating a high level feature value involves applying multiple rules in a sequential order (e.g., decision modules **662**, **664**, **664**, **666** and **668** each implements pre-defined rules in a pre-defined order.)

Each rule in the example sequence produces a higher level feature. As the FVG engine **212** applies the sequence of example rules, lower level features are blended with higher level features. In particular, for example, the first decision module **662** can produce a feature value 0 based upon only one low level feature, an category. The second decision module **664** can produce a value of 1 to represent a higher level feature that represents a combination of a lower level category feature and a geographic proximity feature. The third decision module **666** can produce a feature value of 2 based upon a higher level feature that represents a combination of a lower level category feature, a geographic proximity feature, and a sports category feature. The fourth decision module **668** can produce a feature value of 3, 4, or 5 based upon a combination of a lower level category feature, a geographic proximity feature, a sports category feature, and a time-frame feature.

FIG. **7** is an illustrative chart providing examples of low level feature information within an input storage system **208** and rules within a rule storage **213** of the system of FIG. **2** that are used to produced feature values in accordance with some embodiments. Many of the example rules in the table are relatively rough descriptors of the observation history for a device C, a device D and an IP address. It will be understood that not all rules will be applicable to every potential device pair and that only some portion of the rules may be applicable to any given device pair.

Labels

Label information stored in the storage device **208** is used to train the scoring engine **204** to generate device pair scores as a function of device pair feature vectors. More particularly, feature vector includes a label value associated with a corresponding device identifier. Labeled data are used to train the scoring engine **204** to determine the likelihood that both devices of an unlabeled device pair are associated with the same user. The feature vectors associated with labeled device pairs are commonly referred to as "ground truths".

Specifically, label data are provided to indicate feature vectors that are verifiably known to be associated with the same user, and label data also are provided that indicate feature vectors that are verifiably known not to be associated with the same user. These known device pair relationships are referred to herein as "labeled". In accordance with some embodiments, each labeled device pair is associated with either label −1 or label +1. All other device pairs, referred to herein as "unlabeled", are associated with a label 0. As explained above with reference to feature information within the input storage system **208** and rules within the rules storage device **213**, respective candidate device pairs are associated with respective feature vectors. Labeled device pairs with label=−1 are associated with feature vectors known to be indicative of a pair of devices that is not associated with the same user. Labeled device pairs with label=+1 are associated with feature vectors known to be indicative of a pair of devices that is associated with the same user. The labeled device pairs are used to learn during a training stage to assess whether or not feature vectors of respective unlabeled candidate device pairs are indicative of the respective candidate device pairs being associated with the same user.

Scoring Engine

FIG. **8** is an illustrative flow diagram **700** representing configuration of the score engine **204** in accordance with some embodiments. Modules of the diagram correspond to computer program code (not shown) that configures a computer system that includes one or more processors and/or control components to cause performance of the specific acts represented by the different modules. Module **702** associates candidate device pairs identified using the pairing engine **202** with feature vectors produced using the feature vector generation (FVG) engine **212**. In some embodiments, preliminary pair scores produced using the pairing engine **202** are incorporated into the feature vectors used to determine precision pair scores for associated candidate pairs. Module **702** also receives label information which it associates with candidate pairs and feature vectors. Each candidate device pair also is associated with a preliminary pair score. For example, the candidate device pair having device identifier D1 and a cookie identifier C2 is associated with feature vector V1 and with and with preliminary pair score S3.

Referring again to FIG. **8**, module **704** identifies unlabeled and labeled candidate device pairs from among the candidate device pairs received from the pairing engine **202**. Module **704** provides the labeled device pairs to module **705**. Module **704** provides the unlabeled candidate device pairs to training module **706**. Label data selection module **705** designates first labeled candidate device pairs for provision to training module **706** for use in training a scoring model implemented using a scoring module **708** and designates second labeled candidate device pairs for provision to testing module **709** for use in testing the quality of scores produced using the scoring model. In some embodiments, the label data selection module **705** designates about eighty percent of the received labeled candidate device pairs for use in training and designates about twenty percent of the received labeled candidate device pairs for use in testing.

Training module **706** evaluates correlations between feature vectors of the unlabeled candidate device pairs and feature vectors of the first labeled candidate device pairs that have been designated for use in training and determines pair scores to associate with the first candidate device pairs based upon the correlation evaluations. As explained above, feature vectors associated with unlabeled candidate pairs have a label=0, and feature vectors associated with labeled candidate device pairs have label=−1 or a label=+1. Training module **706** produces a scoring model used by scoring module **708**. The training module **706** produces a scoring model that associates model parameters with vector features. The training module **706** uses training data comprising an array of feature vectors associated with candidate device pairs and corresponding labels to determine model parameters. In accordance with some embodiments, the determined model parameters are indicative of the importance of the features within a vector to a determination of whether the devices of a candidate device pair actually are associated with the same user. A goal of training is to produce a training model that minimizes a residual between model predictions and observed values in training label data.

The model parameters produced using training module **706**, used by the scoring model, are used to configure the feature-based scoring module **708**. Scoring module **708** uses model parameters in the scoring model, combined with the feature vectors associated with candidate device pairs, to produce precision pair scores between 0 and 1 for the unlabeled candidate device pairs. A precision pair score for a candidate device pair represents a more accurate estimate than a corresponding preliminary pair score of likelihood that the devices of the candidate device pair actually are associated with the same user. In accordance with some embodiments, a precision pair score is determined as a function of feature vectors (which may include a feature value indicative of a preliminary pair score) associated with each device pair and model parameters within the trained scoring model.

Score testing module **709** evaluates quality of precision pair scores determined using the scoring model. In particular, score testing module **709** determines precision pair score quality for candidate device pairs as a function of correlation between feature vectors of candidate device pairs and the second labeled candidate device pairs designated by module **705** for use in testing. Candidate device pairs having precision pair scores that pass a score quality threshold are passed to module **710**, which transmits the candidate device pairs and their precision pair scores to the cluster module **206**. Candidate device pairs having precision pair scores that do not pass the score quality threshold are removed from use in the feature based scoring module **708**.

FIG. **9** is an illustrative process flow diagram **800** representing details of configuration of the scoring engine **204** to implement the training module **706** of FIG. **8** in accordance with some embodiments. The training module **706** implements both a semi-supervised learning and a supervised learning process. An unsupervised pre-training module **802** performs pre-training based upon feature vectors of labeled and unlabeled device pairs. In accordance with some embodiments, the unsupervised pre-training module **802** produces the mapping that a remapping module **804** will use to remap features in feature vectors of labeled device pairs to the device remap pair clusters determined by module **802**. The remapping module **804** maps candidate device pair feature vectors to remap clusters that are stored in a computer readable storage device and that are used to facilitate supervised learning. The remapping module **804** maps candidate device pair feature vectors to remap clusters so as to cluster together device pairs having greater feature vector similarity as determined using the unsupervised pre-training process. A supervised training module **808** determines scores to associate with device pair feature vectors based upon features of labeled device pairs. The determined scores are indicative of likelihood that both devices of pair are associated with the same user. The device pair remap

clustering performed using module **804** facilitates supervised learning by improved targeting of labeled feature vectors to feature vector remap clusters with which they have the greatest affinity. More specifically, the unsupervised pre-training module **802** and remapping module **804** act to map feature vectors that correspond to candidate device pairs to remap clusters so as to cluster candidate device pairs based upon their representation in feature space. As explained above, labeled device pairs have a label value of −1 or +1 and unlabeled device pairs have label value 0. It will be appreciated that in accordance with some embodiments, the unsupervised pre-training restructures the feature vector data to reduce dimensions with very sparse data.

A supervised training module **808** determines precision pair scores to associate with candidate device pair feature vectors based upon features of labeled device pairs. The determined scores are indicative of likelihood that both devices of a candidate device pairs are associated with the same user. The device pair remap clustering performed using module **804** facilitates supervised learning by improved targeting of labeled feature vectors to feature vector remap clusters with which they have the greatest affinity.

FIG. **10**A is an illustrative drawing representing example in which the unsupervised pre-training module **804** remaps the device pairs of FIG. **7**B to device pair remap clusters based upon the candidate device pair feature vectors in accordance with some embodiments. Each candidate device pair remap cluster is associated with a set of feature vectors that is distinct from a set of feature vectors of other device pair remap clusters, as determined using an unsupervised training process. This intermediate-mapping to device pair remap clusters is commonly known as a "feature-space embedding" and the process of using the unlabeled examples is often termed the "unsupervised pre-training" step of a semi-supervised learning process.

In some embodiments, the unsupervised pre-training module **802** uses Bayes Networks to train the remapping module **804** to produce the device pair remap clusters. In alternative embodiments, the unsupervised pre-training process can use PCA, Tree-Embedding, Auto-encoders or RBMs, for example. More particularly, in accordance with some embodiments, the supervised learning process produces a remapped training data set **806**. FIG. **10**B is an illustrative drawing representing use of the remapping module **804** in remapping of an example feature vector Vin an original feature space spanned by $\{x_1, x_2, \ldots, x_n\}$ according to a generic nonlinear embedding function E into a latent space spanned by $\{l_1, l_2, \ldots, l_k\}$, where k is not necessarily >n or <n. The latent space representation E(V) allows more efficient learning.

The supervised learning module **808** infers correlation functions from labeled training data, which can be used to estimate likelihood of unlabeled device pairs belonging to same users. In some embodiments, the supervised learning process uses Bayesian Network or Tree Ensembles to produce the device pair scores and to produce the scoring model. In alternative embodiments, the supervised learning process can use Regression Models, or Neural Networks. A goal of this process is to minimize residual between model estimation of likely device pairs and labeled data. The supervised learning module **808** produces a mathematical scoring model, which consists of model parameters in a feature space that is defined in the remapping module **804**. The scoring model is used to configure the scoring module **708** described above to score all other candidate device pairs.

Data Transformations During Scoring

FIG. **11** is an illustrative diagram representing a data flow process **1000** including data transformations that occur within computer readable storage devices within the scoring engine **204** in accordance with in some embodiments. The scoring engine **204** receives candidate device pairs data structures **1002** as input from the pairing engine **202**. It will be understood that during the training stage, candidate device pairs are used to train the scoring engine **204**. The candidate device pair data structures include deviceIDs identifying each device of a device pair and include corresponding preliminary pair scores for the paired devices. In a first data transformation, feature vectors are associated with the candidate device pairs resulting in data structures that associated candidate device pairs, their preliminary pair scores and their associated feature vectors **1004**. In some embodiments, the preliminary pair scores are incorporated within the feature vectors. In a second data transformation, training label data **1005** are associated with the candidate device pairs resulting in candidate device pairs, their preliminary pair scores, their associated feature vectors and their associated training labels **1006**. Recall that labeled feature pairs have a label value of −1 or +1 and that unlabeled device pairs have a label value of 0. In a third data transformation, unsupervised pre-training involving modules **802** and **804** uses device pair feature vectors and their associated training labels to produce affinity scores of candidate device pair clusters (A, B, C, D) **1008** in which clusters are created, which then appended to feature vectors. In a fourth data transformation, supervised training involving supervised training module **808** to produce precision pair scores to associate with the candidate device pairs **1010** and to determine parameters used for the scoring model.

Test label data **1012** is used by testing module **709** to perform a precision evaluation of the scores assigned to the device pairs **1008** during the fourth data transformation. Test label data **1012** used during precision evaluation is disjoint from the training label data **1005**. In other words, the labeled device pairs, used as test label data **1012**, are disjoint from the labeled device pairs used for training label data **1002**. Device pairs are put into different quantized pair score buckets. In a fifth data transformation, by measuring precision and recall in each score bucket, we can transform device pair model scores to precision scores. Device pairs in precision pair score bucket below quality threshold will be filtered out. Final device pairs **1014** having invalidated pair scores are removed from the corpus of training results used to configure the feature based scoring module **708** of FIG. **8**. By way of example, the device pair D_{i1}-D_{j1} is shown to be removed by the precision filter.

Clustering Engine

FIG. **12** is an illustrative process flow diagram **1100** representing configuration of the clustering engine **206** in accordance with some embodiments. Modules of the diagram correspond to computer program code (not shown) that configures a computer system that includes one or more processors and/or control components to cause performance of the specific acts represented by the different modules. Module **1102** receives candidate device pairs and associated precision pair scores from the scoring engine **204**. It will be understood that some of the first candidate pairs identified using the pairing engine **202** may be filtered out by the scoring engine **204** such that a different second set of candidate device pairs may be received by clustering engine **206**. Module **1104** assembles a graph that incorporates the received device pairs. In some embodiments, module **1104** configures a computer system to act as a graph generator.

FIG. **13** is an illustrative drawing representing an example graph produced using the cluster engine of FIG. **12** for the candidate device pairs of FIGS. **4**A-**4**B in accordance with some embodiments. It will be understood that the graph uses precision pair scores produced using the scoring module **708** of the scoring engine **206** rather than the preliminary pair scores produced using the pairing engine **202**. Graph nodes correspond to device identifiers. Graph edges that connect graph nodes indicate device pairs. Precision pair scores associated with the graph edges indicate the likelihood that the associated devices of the device pairs actually are associated with the same user.

In the course of generating a graph, module **1104** prunes some graph edges from the device graph. More particularly, in accordance with some embodiments, module **1104** implements a local graph sparsification process to clean the graph before graph clustering is performed. Inputs to the local graph sparsification process include device pairs (sets of two device IDs) with associated precision pair scores provided using the scoring module **708**. Outputs from the graph sparsification process include fewer pairs, judiciously chosen. One simple way of choosing pairs to remove before doing clustering is to impose a flat cut on the score. For example, all pairs with precision scores <0.1, for example, would be thrown away. Although the flat cut method is simple it sometimes is not optimal since it tends to be overly harsh on less dense portions of the graph than it is on over-dense portions of the graph, destroying our ability to keep many good clusters. Alternatively, local graph sparsification improves upon the flat cut approach. A judiciously pruned graph provides a better indication of relationships between device pairs. Specifically, for example, devices associated with different device pairs that share a deviceID in common also share a node in common in the graph. Moreover, the edge scores indicate the relative strengths of different device pairs.

Module **1106** selects potential user device clusters within the graph for evaluation as to whether they actually include one or more user device clusters. FIG. **14** is an illustrative drawing representing a set of example proposed user device clusters user device clusters $U1_A$-$U9_A$ identified using the graph of FIG. **13** in accordance with some embodiments. It can be seen that a number of graph edges have been removed to produce the proposed user device clusters. For example, the graph edge between device D4 and Cookie C2, which is shown in FIG. **13**, is absent from FIG. **14**.

Module **1108** modifies the proposed user device cluster to attempt to identify a valid user device cluster based upon one or more cluster fitness requirements. The user device cluster modification may involve adding, deleting or combining edges within the potential device cluster to attempt to meet the cluster fitness requirements. Decision module **1110** determines whether the fitness requirements have been achieved through the modifications by module **1108**. In response to a determination by decision module **1110** that the fitness requirement has not yet been achieved, decision module **1112** determines whether to continue to modify the potential cluster to meet the fitness requirements. In response to decision module **1112** determining to continue to modify, control flows back to module **1108**. In response to decision module **1112** determining to not continue to modify, module **1114** abandons the proposed cluster.

In response to decision module **1110** determining that the cluster fitness requirements have been achieved, a cluster accuracy filter module **1115** determines which device identifiers of a proposed user device cluster are to be associated with a final user device cluster and which devices of the user proposed device cluster are to be removed. Module **1116** outputs the cluster as a final cluster. Following module **1114** or following module **1116**, depending upon control flow for a given potential user device cluster, control flows to decision module **1118**, which determines whether there are more potential device clusters in the graph to be evaluated. In response to a determination that that are more potential device clusters in the graph to be evaluated, control flows to back to module **1106** and another potential user device cluster is identified. In response to a determination that there are no additional device clusters to be evaluated, control flows to module **1120**, which causes decision module to wait for new potential device clusters.

FIG. **15** is an illustrative drawing representing an example set of final device clusters $U1_B$-$U9_B$ produced using the cluster engine of FIG. **12** based upon the proposed user device clusters of FIG. **14** in accordance with some embodiments. One or more of the final device clusters of FIG. **14** are modified relative to the proposed user device clusters to remove device indicators as determined by the second model training process. In particular, for example, device identifiers for cookie C13 was present in proposed user device cluster $U5_A$ but is missing from user device cluster $U5_B$. The final device clusters $U1_B$-$U9_B$ can act as digital identity group communication systems for communication with users who own, or are associated with, devices identified within the clusters. In other words, the clusters act as digital identity group communication systems for use to identify multiple devices that can be used to communicate with associated users.

Cluster Fitness Determination with Label Propagation

In accordance with some embodiments, modules **1108**-**1116** perform user device clustering via label propagation to arrive final user device clusters that include sets of nodes assigned to unique user identifiers. Label propagation in accordance with some embodiments, involves initially assigning each vertex in the pairing graph a unique label, and then propagating the label to other vertices through the connected edges, and updating the labels associated with each vertex in the graph. Label propagation and label updating are performed iteratively, until it can no longer update the label for every single node in the graph (a node will not update its label if no incoming label has a better aggregated score than the current label's score). Once label propagation stops, whatever the nodes/vertices that have the same label will be treated as belonging to the same user.

To propagate the label, in accordance with some embodiments, each node of the graph sends its label to its neighboring nodes, and it will also receive the label sent over from its neighbor nodes at the same time. In addition to sending and receiving labels, each vertex also updates its label based on its current label and all the labels that it receives from its neighbor. The whole process is carried out at each individual node and in parallel. Within each node, there are three parameters that control the behavior of sending the label to other nodes. First is a percentage number that controls what fraction of neighbor nodes will receive the label from current node, second is an integer number that limits the maximum number of nodes that will receive the label from current node and the third is an integer number that controls how far a label can travel through the graph.

The details regarding the percentage parameters that controls the fraction of neighbor nodes to receive label are the following: Each node first computes maximum value of the pair scores associated with all the edges that connect the

node to its neighbors, it then sends out label only to those neighbors whose scores are above this percentage of the maximum value.

Regarding the second parameter, Even if there are a lot of neighbors that are qualified to receive the label from current node, this parameter is used to restrict it to top few nodes. We will sort all the qualifying edges in descending order and propagate labels only through the top edges determined by this parameter.

The third parameter controls how far a label can traverse throughout the graph. We will track how many edges a label has traveled so far, and we will not allow a label to travel any further once the number of edges it traveled exceeds this integer limit.

As label travels through the graph, it collects pair scores along the path it travels. Each label starts with initial score of 1. For a label that originates from one node, say node A, for example, to another node B, for example, the score of that label is defined as the score of the label at node A, multiplied by the score of the edge that connects node A and B and divided by the number of edges the label has traveled to reach this node. Each node will aggregate the score for all the labels that it receives, pick the label with maximum score and assign that label to itself if the new score is greater than the score of the current label. The score of the new label will also be retained for the next iteration. If the label changes the current node will propagate it to its neighbors for them to consume in next iteration.

It will be appreciated that the final output of the label propagation module depends on the joint behavior of all the above parameters. While it is difficult to predict which combination of those parameters will provide the best performance, we can actually define a user clustering performance metrics (e.g. precision of the user clusters), allocate a separate set of the pairing input as training data and do a grid search to find the combination of all the parameters that yields the best performance on the output user clusters. We then apply the optimal value of those parameters to the future pairing input and generate optimal user clusters.

Cluster Fitness Determination with Simulated Annealing

Alternatively, in accordance with some embodiments, modules **1108-1116** perform user device clustering via label simulated annealing. Simulated annealing is an iterative, probabilistic, cluster fitness-driven algorithm Annealing is performed once per node (device) in the graph as each node acts as a "seed" for an eventual user device cluster. The cluster fitness function is maximized when both (i) the members of the cluster are strongly interconnected, meaning most of the possible pairs between the member nodes are observed to exist and they have very high scores, and (ii) the cluster is very weakly connected to devices outside of the members of the cluster. The simulated annealing process proposes to add or subtract edges (pairs) between devices at random and accepts these proposals with probability related to the fitness of the resulting cluster. The simulated annealing process iterates and the acceptance probabilities evolve according to a simulated temperature schedule (starting in the more randomized "high-temperature" phase evolving into the more optimization-focused "Low-Temperature" phase). The process stops when temperature is below some value and the value of the cluster fitness function has stabilized.

Simulated Annealing is an example of the class of "Fitness-Driven" Clustering methods. "Fitness-Driven" methods are one of the triad taxonomy of clustering algorithms as "divisive", "agglomerative" or "fitness-driven", where of course hybrids can be formed by combining examples of any of these. The process and end-result of annealing requires the specification of the cluster fitness function wherein, contrary to intuition from popular examples, there is much freedom to customize the algorithm. As noted above, typically-used cluster fitness function is the product of two factors which is maximized when (i) the devices within the found clusters are maximally inter-connected and (ii) the devices within a given cluster are maximally isolated from those outside of the cluster. Given this definition of cluster fitness, there are nearly zero free-parameters in the annealing algorithm (what is left is the temperature schedule), so it would appear that there is nothing to tune and the clusters you get from applying annealing are "The annealing clusters".

However, in specific problem domains it is very often beneficial to add, by hand, additional terms to the cluster fitness function to penalize or encourage various effects/metrics that are seen in the end clusters. Since these terms are added by hand it is not known what their individual strengths should be compared to each other, or to the original baseline fitness described above, in order to obtain an optimal clustering. We typically introduce one such free-parameter that needs to be tuned for each additional term in the fitness function.

For example, in the example graphs described above, there can be several distinct "types" of candidate device pairs (e.g., desktop-cookie to desktop-cookie, or mobile-device to mobile device, or mobile-web-cookie to mobile device, etc.), each with its own pair scoring model. Between pair types, the performance as a function of score is quite different, owing to the many differences in systematic effects that go into collecting data regarding the various "devices" (differing rates of activity on laptops vs. smartphones for example). Thus, it makes sense to add additional factors to the cluster fitness function to account for these differences (for instance, to trust clusters built on a good mix of cookie-device pairs more than those built primarily on cookie-cookie pairs if your cookie-device model is the more performant of the two). In this case the full fitness function can be written as:

$$\text{Fitness}(C) = \sum_a \lambda_a \left[\left(\frac{C_{in}^a}{2} \right)^{-1} \sum_i^{C_{in}^a} s_i \right] \cdot \left[\frac{\sum_i^{C_{in}^a} s_i}{\sum_i^{C_{in}^a} s_i + \sum_j^{C_{ex}^a} s_j} \right]$$

where the sum over 'a' is a sum over pair types (cookie-device, cookie-cookie, etc.) the various C's denote the number of internal or external pairs of type "a", and the lambdas are the free parameters that we need to tune to find optimal performance.

Tuning these parameters is done as separate step, optimizing a combination of cluster precision (using labeled data) and heuristic metrics using a subgraph of the entire device graph. In practice the cluster precision is the dominant factor in our choice of optimal parameters but we may supplement this with simple metrics like average cluster fitness, average number of devices/cluster, etc. Finding the optimal parameters given a desired quantity to be optimized is a routine optimization problem. Some of the simplest techniques to solve this include grid search and Monte Carlo. For fitness functions with much more than four free parameters the computation typically is prohibitively expensive and one generally must do something more intelligent

to find a good optimum (for example Markov chain, Monte Carlo or Bayesian variants thereof).

The partitioning of data in this optimization phase is analogous to that of basic supervised learning: we select a random sub-graph of the entire graph to be our "test" data set and another to be our "training set" we optimize metrics using the clusters in the training set and verify this performance on the testing set before going on to apply the optimal clustering to the entire device graph.

Data Transformations During Device Clustering

FIG. **16** is an illustrative diagram representing a data flow process **1500** including data transformations that occur during user device clustering within the cluster engine **206** in some embodiments. The clustering engine **206** receives as input from the scoring engine **204**, candidate device pairs associated with feature vectors and precision pair scores **1502**. In a first data transformation, a graph **1504** is created in which user devices are represented by graph nodes, device pairs are represented by graph edges and device pair scores are associated with corresponding graph edges. In a second data transformation, graph modification is performed to produce proposed user device clusters **1506**. The example proposed user device clusters of FIG. **14** are shown. In a third data transformation, test label data **1510** is used to determine whether device pairs within user the proposed user device clusters meet a second threshold level of precision. Final user device clusters **1512** are produced in which user device pairs have removed if they do not meet the second precision threshold level. The final device clusters **1512** of FIG. **15** are shown in which, for example, cluster U_{5B} is removed from the final set of clusters.

Finding a Balance Between Pairing and Clustering

The approach pairing of candidate devices using the pairing engine **202** and in the score engine **204** is designed to produce pair scores that come as close as possible to "Oracle Precision", meaning all pairs whose devices are truly owned by the same person have score=1 and all of those that do not have score=0. In reality of course the scoring model is not perfect, and scores for the vast majority of pairs lie between these extremes. Part of this imperfection is irreducible (pure noise that cannot be predicted, typically termed "variance"), but much of it is reducible (typically termed "bias")

There are many other useful pieces of information that are not included in the pairing engine. In principle we can featurize any of this information and input it directly to the pairing engine to improve our predictions. Much of this information has to do with information about devices that are not included in the pair (say A-B) itself, questions like: How many of A's friends are also B's friends? How many of A's friends have been seen on the same IP as B? How many friends do A and B have, respectively? How many of A's good friends are also friends of B? Some of this information has actually been fed into the pairing engine already (what we were calling "device graph feedback to pairing" or "satellite features") . . . .

But a more efficient way of thinking about this information is apparent: we are trying to give the pairing engine a view of the local properties of the device graph that can be built around this pair. From the device graph point of view all of these questions can be answered at once, and with relative ease, whereas the amount of work necessary to provide a comprehensive encoding of the information in the device graph as features to the Pairing Engine is prohibitive. Thus, as a matter of efficiency and practicality, we draw a line, figuratively speaking, as to how much information we will add to the pairing engine **202** and/or to the score engine **204** to generate good pair scores and how much will only be taken into account in the clustering engine **206** after constructing the device graph.

Taking this modular structure "Pairing Engine=>Clustering Engine", rather than trying to make one big model goes beyond convenience, but actually improves the end result. The fact is that, while the pairing candidate devices using the pairing engine **202** and/or the scoring engine **204** has the theoretical capability to entirely eliminate the reducible error (bias) in our predictions, the amount of work necessary to improve the performance of the pairing engine **202** and/or the scoring engine **204** typically grows exponentially with the desired lift in performance. Stopping short and sending the pair scores of an approximate pairing model on to the separate clustering engine **206** allows to recoup the majority of the possible bias reduction with much less computational effort.

System in Network Environment

FIG. **17** is an illustrative drawing representing the system **200** of FIG. **2** configured to operate in a network environment accordance with some embodiments. The system in a network environment **1600** includes a first server system **1602** to receive mobile application usage information and page view information from users. The first server system **1602** includes a targeting engine **1606** to target information delivery to a user across multiple devices and to provide a user-based frequency-cap for information delivered to a user across multiple devices, an insight engine **1608** to provide cross-device attribution across multiple devices of a user, and to provide offline-online attribution to user, an optimization engine **1610** to produce user-based optimization and platform-based optimization.

The system includes a second server system **1604** to produce user device cluster indicating groups of devices associated with users. The second server system **1604** includes the pairing engine **202**, the scoring engine **204**, the cluster engine **206** the FVG engine **212**, the index structure **209** and the rules storage **213**. In accordance with some embodiments, the second server is configured to act as a distributed massively parallel processing and distributed storage system in accordance with a Hadoop framework. Thus, in effect, a multiplicity of instances (as shown) of the pairing engine **202**, scoring engine **204** and cluster engine **206** engines run in parallel.

The system **1600** includes a first interface (indicated within dashed lines) **1612** between the first server **1602** and Hadoop cluster **1604**. The first interface **1612** includes a third server **1614** that collects a log containing information from the first server **1602** relating to users' interaction with devices such as mobile application usage and page views. The log includes an append-only, ordered sequence of records, ordered by time. Typically, each entry is assigned a unique sequential log entry number. Since log entries are time-ordered, a log entry number can act as a "timestamp" of the entry. The third server **1614** is configured to handle high-throughput, low-latency real-time data feeds. In accordance with some embodiments, the third server **1614** is configured to act as a message broker that mediates communications between different applications, such as pushing bidding model to server **1602** and pushing log files from server **1610** to Hadoop cluster **1604**. In accordance with some embodiments, the third server **1614** is implemented according to Apache Kafka system documentation. The first interface **1612** also includes a log extractor module **1616** within the second server **1604** that extracts logs from the third server **1614** and provides the logs to the pairing engine **202**. The log extractor **1616** acts as a data normalizer that

converts unstructured data received within the third server **1614** to structured data such as key value pairs, for example, suitable for input to the pairing engine **202**. It will appreciated that information within the log may be received from different user device platforms that structure data differently.

The system **1600** includes a second interface (indicated within dashed lines) **1618** between the first server **1602** and the second server **1604**. The second interface **1618** includes a user profile indexer module **1620** within the second server **1604** that receives and stores user device cluster information produced by the cluster engine **206** that indicates groups of devices associated with users. As explained above, the cluster engine **206** produces user device clusters structures that associate clusters of user device IDs with unique users. The user profile indexer **1620** stores index information that associates user deviceID clusters with unique users. In some embodiments, the user profile indexer module also receives information from the log extraction module **1616**, such as user demographic information including one or more of gender, age, income, location and behavioral patterns (e.g., search patterns), for example. The second interface **1618** also includes a fourth server **1622** that obtains the user device cluster structures from the user profile indexer and provides it to the first server. The fourth server **1622** is configured to serve many concurrent users. In accordance with some embodiments, the fourth server **1622** is configured to store, retrieve, and manage document-oriented information, sometimes referred to as semi-structured data. Document oriented databases are a category of NoSQL databases, which are designed around an abstract notion of a "Document". In accordance with some embodiments, the fourth server **1622** is implemented using Couchbase servers, which can be clustered from a single machine to very large scale deployments spanning many machines, and that are configured to provide scalable key-value or document access with low latency and high sustained throughput.

The system includes a third interface (indicated within dashed lines) **1624** between the first server and the second server. The third interface **1624** includes a campaign delivery and performance module **1626** within the second server **1604** that receives and stores the count of user device cluster structures that indicate groups of devices associated with users. In some embodiments, the campaign delivery and performance module **1626** also receives information from the log extraction module **1616** such as, the number of ad impressions on a user's desktop devices, number of impressions received on a user's mobile device and number of impressions received on multiple devices of a user, for example. In accordance with some embodiments the campaign delivery and performance module is configured according to the Hive data warehouse infrastructure, which is based upon Hadoop. A Hive infrastructure is best suited for batch jobs over large sets of data such as ad campaign data, for example. The third interface **1624** also includes a fifth server **1628** that extracts campaign performance information from the campaign delivery and performance module **1626** and provides it to the first server **1602**. In accordance with some embodiments, the fifth server **1628** includes a SQL server that provides access by the second server to information relating to user based campaign and performance reach.

The foregoing description and drawings of embodiments are merely illustrative of the principles of the invention. Various modifications can be made to the embodiments by those skilled in the art without departing from the spirit and scope of the invention, which is set forth in the appended claims.

The invention claimed is:

1. A system comprising:

a computer readable storage device storing an index structure that associates device identifiers with device feature information;

a pairing engine configured to determine candidate device pairs based at least in part upon at least a portion of the device feature information;

a feature vector generation engine configured to produce multiple feature value feature vectors corresponding to determined candidate device pairs, based at least in part upon feature information associated within the index structure with devices of the determined candidate device pairs;

wherein the feature vector engine is configured according to at least one rule to produce a feature value for at least one feature represented within a feature vector corresponding to at least one determined candidate device pair based at least in part upon both device feature information associated within the index structure with a first device identifier of the at least one determined candidate device pair and device feature information associated within the index structure with a second device identifier of the at least one determined candidate device pair;

a scoring engine configured to determine scores to associate with determined candidate device pairs, including the at least one determined candidate device pair, based at least in part upon produced feature vectors associated with the determined candidate device pairs;

a computer readable storage device storing a graph structure, wherein nodes within the graph structure represent device identifiers, including the first device identifier and the second device identifier of the at least one determined candidate device pair, and wherein edges between pairs of nodes within the graph structure indicate determined candidate device pairs;

a clustering engine configured to identify respective clusters of two or more nodes within the graph structure that represent respective groups of devices.

2. The system of claim **1** further including:

a computer readable storage device storing information structures that identify multiple different groups of device identifiers, indicated within the graph structure, that are associated with the different respective users.

3. The system of claim **1** further including:

a targeting engine to use the groups of device identifiers to target information over the internet to users associated with the groups of device identifiers.

4. The system of claim **1**,

wherein the index structure associates device identifiers with one or more network source/destination address identifiers.

5. The system of claim **1**,

wherein the index structure associates device identifiers with one or more IP addresses.

6. The system of claim **1**,

wherein the index structure associates device identifiers with one or more network source/destination address identifiers and one or more timestamps.

7. The system of claim **1**,

wherein the index structure associates device identifiers with one or more network source/destination address identifiers; and

wherein the pairing engine configured to identify device pairs based at least in part upon device identifiers and associated network source/destination address identifiers.

8. The system of claim **1**,

wherein the index structure associates device identifiers with one or more IP addresses; and

wherein the pairing engine configured to identify device pairs based at least in part upon device identifiers and associated one or more IP addresses.

9. The system of claim **1**,

wherein the index structure associates device identifiers with one or more network source/destination address identifiers and one or more timestamps; and

wherein the pairing engine configured to identify device pairs based at least in part upon device identifiers and associated one or more network source/destination address identifiers and one or more timestamps.

10. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs.

11. The system of claim **1**,

wherein the feature vector generation engine determines feature values to associate with respective determined device pairs based at least in part upon feature information associated within the index structure with device identifiers of devices of respective determined device pairs.

12. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs based at least in part upon feature information associated within the index structure with device identifiers of devices of respective determined device pairs.

13. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs; and

wherein the scoring engine is configured to evaluate correlations between feature vectors associated with different respective determined device pairs.

14. The system of claim **13**,

wherein the scoring engine produces a scoring model based upon the evaluation of the correlations between feature vectors associated with different respective determined device pairs.

15. The system of claim **14**,

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs.

16. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs;

wherein the scoring engine is configured to produce multiple remap clusters of determined device pairs based at least in part upon evaluation of affinities between feature vectors associated with the determined device pairs;

wherein the scoring engine is configured to evaluate correlations between feature vectors associated with different respective determined device pairs within the same remap clusters;

wherein the scoring engine produces a scoring model in a computer readable storage device based upon the evaluation of the correlations between feature vectors associated with different respective determined device pairs; and

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs.

17. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs;

wherein the scoring engine is configured to evaluate correlations between feature vectors associated with different respective determined device pairs;

wherein the scoring engine produces a scoring model stored in a computer readable storage device based upon evaluation of correlations between vectors of labeled device pairs and vectors of unlabeled device pairs within the produced clusters; and

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs.

18. The system of claim **1**,

wherein the scoring engine is configured to produce multiple remap clusters of determined device pairs in which to evaluate correlations between vectors of labeled device pairs and vectors of unlabeled device pairs;

wherein the scoring engine produces a scoring model based upon evaluation of correlations between vectors of labeled device pairs and vectors of unlabeled device pairs within the produced remapped clusters; and

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs.

19. The system of claim **18**,

wherein the index structure associates respective device identifiers with label information indicative of whether respective devices associated with respective device identifiers are labeled or unlabeled.

20. The system of claim **1**,

wherein the index structure associates respective device identifiers with label information indicative of whether respective devices associated with respective device identifiers are labeled or unlabeled;

wherein the scoring engine is configured to evaluate correlations between vectors of labeled device pairs and vectors of unlabeled device pairs;

wherein the scoring engine produces a scoring model stored in a computer readable storage device based upon the evaluation of correlations between vectors of labeled device pairs and vectors of unlabeled device pairs within the produced clusters; and

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs.

21. The system of claim **1**,

wherein the clustering engine is configured to identify clusters of nodes within the graph structure based at least in part upon one or more cluster fitness requirements.

22. The system of claim **1**,

wherein the clustering engine is configured to use a label propagation process to identify clusters of nodes within the graph structure.

23. The system of claim **1**,

wherein the clustering engine is configured to use a simulated annealing process to identify clusters of nodes within the graph structure.

24. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs;

wherein the scoring engine is configured to evaluate correlations between feature vectors associated with different respective determined device pairs;

wherein the scoring engine produces a scoring model based at least in part upon the evaluation of the correlations between feature vectors associated with different respective determined device pairs;

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs; and

wherein the clustering engine is configured to identify clusters of nodes within the graph structure based at least in part upon the produced scores.

25. The system of claim **24**,

wherein the clustering engine is configured to use the produced scores to identify clusters of nodes within the graph structure based at least in part upon one or more cluster fitness requirements.

26. The system of claim **1**,

wherein the scoring engine is configured to produce multiple remap clusters of determined device pairs in which to evaluate correlations between vectors of labeled device pairs and vectors of unlabeled device pairs;

wherein the scoring engine produces a scoring model based upon evaluation of correlations between vectors of labeled device pairs and vectors of unlabeled device pairs within the produced remapped clusters;

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs; and

wherein the clustering engine is configured to identify clusters of nodes within the graph structure based at least in part upon the produced scores.

27. The system of claim **26**,

wherein the clustering engine is configured to use the produced scores to identify clusters of nodes within the graph structure based at least in part upon one or more cluster fitness requirements.

28. The system of claim **1**,

wherein the feature vector generation engine produces feature vectors that include multiple feature values to associate with respective determined device pairs;

wherein the scoring engine is configured to evaluate correlations between feature vectors associated with different respective determined device pairs;

wherein the scoring engine produces a scoring model stored in a computer readable storage device based upon evaluation of correlations between vectors of labeled device pairs and vectors of unlabeled device pairs within the produced clusters; and

wherein the scoring engine uses the scoring model to produce scores for at least a portion of the determined device pairs; and

wherein the clustering engine is configured to identify clusters of nodes within the graph structure based at least in part upon the produced scores.

29. The system of claim **28**,

wherein the clustering engine is configured to use the produced scores to identify clusters of nodes within the graph structure based at least in part upon one or more cluster fitness requirements.

* * * * *