



(19) **United States**

(12) **Patent Application Publication**
Paraschivescu

(10) **Pub. No.: US 2022/0171638 A1**

(43) **Pub. Date: Jun. 2, 2022**

(54) **DATA STRUCTURES FOR VISUALIZATION OF HIERARCHICAL DATA**

Publication Classification

(71) Applicant: **Quick Eye Technologies Inc.**, Chicago, IL (US)

(51) **Int. Cl.**
G06F 9/451 (2006.01)
G06F 3/0482 (2006.01)
G06F 3/04842 (2006.01)
H04L 67/10 (2006.01)

(72) Inventor: **Andrei Paraschivescu**, Chicago, IL (US)

(52) **U.S. Cl.**
CPC *G06F 9/451* (2018.02); *H04L 67/10* (2013.01); *G06F 3/04842* (2013.01); *G06F 3/0482* (2013.01)

(21) Appl. No.: **17/675,883**

(22) Filed: **Feb. 18, 2022**

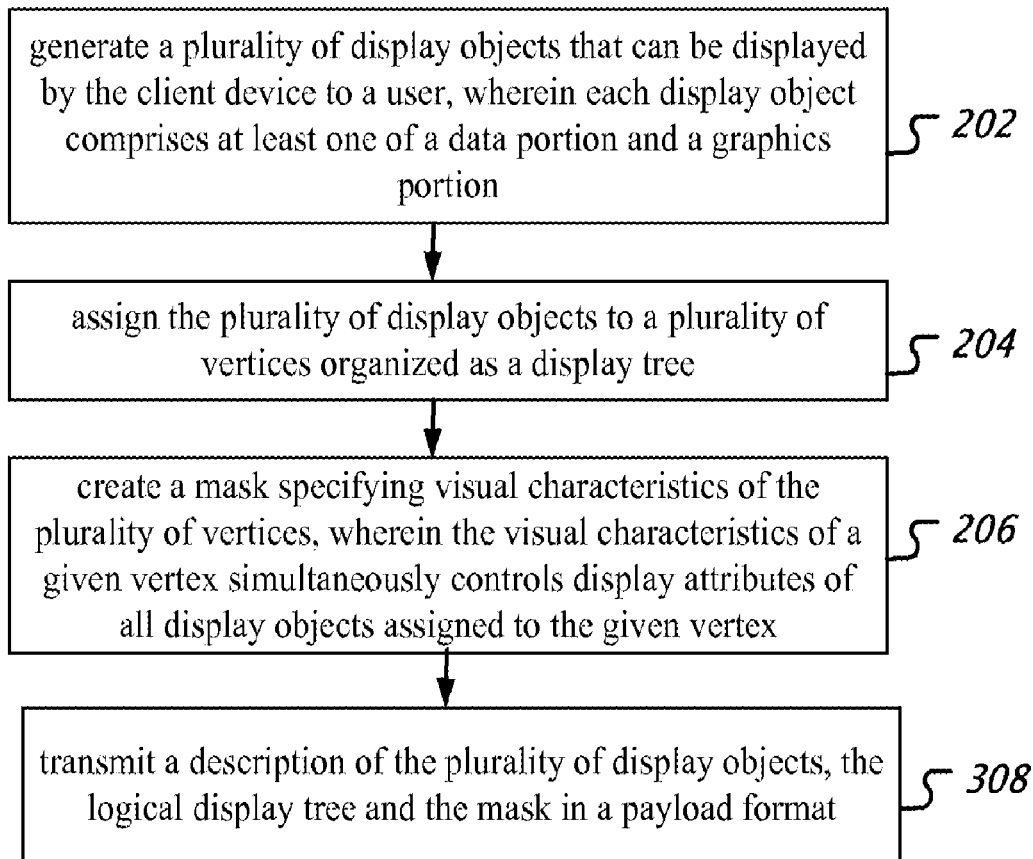

(57) **ABSTRACT**

Related U.S. Application Data

(63) Continuation of application No. 14/769,041, filed on Aug. 19, 2015, now Pat. No. 11,256,524, filed as application No. PCT/US2014/017209 on Feb. 19, 2014.

To provide visualization data to a client device, a server generates a plurality of display objects for selectively displaying at the client device to a user. Each display object includes at least one of a data portion and a graphics portion. The plurality of display objects is assigned to a plurality of vertices organized as a logical display tree. A mask specifying visual characteristics of the plurality of vertices is created. The visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex. The server transmits a description of the plurality of display objects, the logical display tree and the mask in a payload format.

(60) Provisional application No. 61/766,531, filed on Feb. 19, 2013.






FIG. 1

dir	type	user	last
4038			
NA			
			2
docs			53
econBoost			58
econSetup			34
econWrap			60
ico			24
fx			205
GUIin			503
localPY			30
ngen			69
objectMert			256
objectMertc			248
om0_0			606
om0_1			730
omex			37
om.b			58
omRiskMan			68
physical			62
physical2			57
riskMan			91
rmSetup			47
sandbox			96
setup1			130
setup1c			118
setup2			16
solver			40
solModel			40
switch			15
Transport			78
twolayer			83
twolayer2			80
utils			39

FIG. 2

	dir	type	user	last
4038	N/A	-	exe	docs
N/A	-	-	-	media
	1	-	1	-
docs	29	5	4	5
econEcost	23	14	6	-
econSetup	12	7	-	-
econWrap	18	12	2	-
foo	14	4	-	-
lx	58	13	64	-
GUIkit	18	59	-	-
localPY	14	7	-	-
ngen	39	9	1	-
objectMan	23	33	-	-
objectManrc	16	32	-	-
om0_0	21	70	-	-
om0_1	20	72	1	-
cmex	15	9	1	-
cmLib	19	12	-	-
omRiskMan	21	13	-	1
physical	20	15	-	-
physical2	16	14	-	-
riskMan	32	11	21	-
rmSetup	17	9	-	1
sandbox	19	19	-	-
setup1	19	14	-	-
setup1c	14	13	-	-
setup2	6	4	-	-
solver	22	4	6	-
ssModel	12	7	-	-
switch	11	-	2	-
Transport	16	13	2	-
twolayer	16	16	-	-
twolayer2	10	13	-	-
utils	15	9	-	-

FIG. 3

	dir	type	user	last			
4038	N/A	-	exe	docs	media	devel	
N/A	-	-	-	-	-	-	-
	1	-	1	-	-	-	-
docs	29	5	4	5	-	-	10
econBoost	23	14	6	-	-	-	16
econSetup	12	7	-	-	-	-	15
econWrap	18	12	2	-	-	-	28
foo	14	4	-	-	-	-	6
ix	58	13	64	-	-	-	70
GUIkit	18	59	-	-	-	-	426
localPY	14	7	-	-	-	-	9
ngen	39	9	1	-	-	-	20
objectMan	23	33	-	-	-	-	200
objectManr	16	32	-	-	-	-	198
om0_0	21	70	-	-	-	-	515
om0_1	20	72	1	-	-	-	637
omex	15	9	1	-	-	-	12
omlib	19	12	-	-	-	-	27
omRiskMan	21	13	-	1	-	-	31
physical	20	15	-	-	-	-	27
physical2	16	14	-	-	-	-	27
riskMan	32	11	21	-	-	-	27
rmSetup	17	9	-	1	-	-	20
sandbox	19	19	-	-	-	-	58
setup1	19	14	-	-	-	-	97
setup1c	14	13	-	-	-	-	91
setup2	6	4	-	-	-	-	6
solver	22	4	6	-	-	-	8
ssModel	12	7	-	-	-	-	21
switch	11	-	2	-	-	-	-
Transport	16	13	2	-	-	-	47
twolayer	16	16	-	-	-	-	61
twolayer2	10	13	-	-	-	-	57
utils	15	9	-	-	-	-	15

FIG. 4


	dir	type	user	last
0	N/A	7	exe	docs
	media			devel
N/A	-	-	-	-
	1	-	1	-
docs	29	5	4	5
econBoost	23	14	6	-
econSetup	12	7	-	-
econWrap	18	12	2	-
lca	14	4	-	-
lx	58	13	64	-
GLIM	18	59	-	-
localPY	14	7	-	-
agen	39	9	1	-
objecMan	23	33	-	-
objectfanz	16	32	-	-
omd g	21	70	-	-
omd f	20	72	1	-
cmx	15	9	1	-
cmib	19	12	-	-
omRiscMan	21	13	-	1
physical	20	15	-	-
phys celz	16	14	-	-
radMan	32	11	21	-
rmSetup	17	9	-	1
sandbox	19	19	-	-
setup1	19	14	-	-
setup1a	14	13	-	-
setup2	6	4	-	-
solver	22	4	6	-
ssModel	12	7	-	-
switch	11	-	2	-
Transport	16	13	2	-
twoplayer	16	16	-	-
twoplayer2	10	13	-	-
utils	15	9	-	-

FIG. 5

The image shows a window titled "Organizer" displaying a directory tree. The table below represents the data shown in the window.

dir	type	user	last	docs	media	devel
342	N/A	-	-	-	-	-
N/A	-	-	-	-	-	-
	3	-	1	-	-	-
back	29	5	4	5	-	10
econdev	23	14	6	-	-	16
econdev0	12	7	-	-	-	15
econdev1	16	12	2	-	-	26
log	14	4	-	-	-	6
ix	58	13	64	-	-	70
Clash	18	59	-	-	-	426
clash	14	7	-	-	-	9
nger	39	9	1	-	-	20
nbnet	23	33	-	-	-	200
nbnet	16	32	-	-	-	198
onc_0	21	70	-	-	-	519
onc_1	29	72	1	-	-	637
onc	15	9	1	-	-	12
onc	19	12	-	-	-	27
onprod	21	13	-	1	-	31
physnet	20	15	-	-	-	27
physnet2	16	14	-	-	-	27
rdn	32	11	21	-	-	27
rdnet	17	9	-	1	-	20
sandbox	19	18	-	-	-	98
snup1	19	14	-	-	-	07
snup1c	14	13	-	-	-	01
snup2	8	4	-	-	-	6
snup	22	4	6	-	-	8
snup	12	7	-	-	-	21
snup	13	-	2	-	-	-
transport	16	13	2	-	-	47
trclaver	16	16	-	-	-	61
trclaver2	19	13	-	-	-	57
trcl	15	9	-	-	-	16


FIG. 6



The image shows a screenshot of a file organizer window titled "Organizer". The window contains a table with columns for file names, sizes, types, users, and last modified dates. The table lists several folders and files, including "econBoost", "econSetup", "econWrap", "GUkit", "ngan", "Transport", and "utils".

	dir	type	user	last			
842	N/A	7	eve		docs	media	devel
econBoost	23	14	6		-	-	16
econSetup	12	7	-		-	-	15
econWrap	18	12	2		-	-	28
GUkit	18	59	-		-	-	426
ngan	39	9	1		-	-	20
Transport	16	13	2		-	-	47
utils	15	9	-		-	-	15


FIG. 7



The image shows a screenshot of a file organizer window titled "Organizer". The window displays a table with columns for file names, sizes, types, users, and last modified dates. The table is organized into a tree view with a root folder named "701".


	dir	type	user	last
701	N/A	7	eve	
econBoost	23	14	6	
econSetup	12	7	-	
econWrap	18	12	2	
GUIKit	18	59	-	
ngan	39	9	1	
Transport	16	13	2	
utils	15	9	-	

FIG. 8



	dir	type	user	last
	701		exe	devel
econBoost		14	6	16
econSetup		7	-	15
econWrap		12	2	28
GUIka		59	-	426
ngen		9	1	20
Transport		13	2	47
utils		9	-	15

FIG. 9



The image shows a terminal window titled "Organizer" displaying a table of project statistics. The table has columns for project name, count, type, user, and last. The "GLib" project has a sub-table for languages: C/C++ (426), python (103), and hg (323).

		dir	type	user	last			
701	7	exe		devel				
econBoost	14	6		16				
econSetup	7	-		15				
econWrap	12	2		28				
GLib	59	-		426	C/C++	python	hg	
						103	323	
ngen	9	1		20				
Transport	13	2		47				
utils	9	-		15				

FIG. 10

The image shows a screenshot of a file organizer application window titled "Organizer". The window displays a directory tree with columns for "dir", "type", "user", and "last". The tree structure is as follows:

dir	type	user	last
/	7	0	0
econDocset	14	8	16
econSetup	7	-	15
econWfrep	12	2	26
GLUED	59	0	428
	N/A	N/A	N/A
	2012	59	2012
	2011	-	2011
	2010	-	2010
	2009	-	2009
	2008	-	2008
	2007	-	2007
	2006	-	2006
	2005	-	2005
	2004	-	2004
	2003	-	2003
	2002	-	2002
	2001	-	2001
	2000	-	2000
	1999	-	1999
	1998	-	1998
	1997	-	1997
	1996	-	1996
	1995	-	1995
ngen	9	1	20
Transport	13	2	47
utils	9	-	15

FIG. 11




FIG. 12




FIG. 13




FIG. 14

dir	type	user	last
163			
Jan			9
Feb			-
Mar			27
Apr			-
May			-
Jun			-
Jul			-
Aug			-
Sep			-
Oct			-
Nov			19
Dec			48

FIG. 15

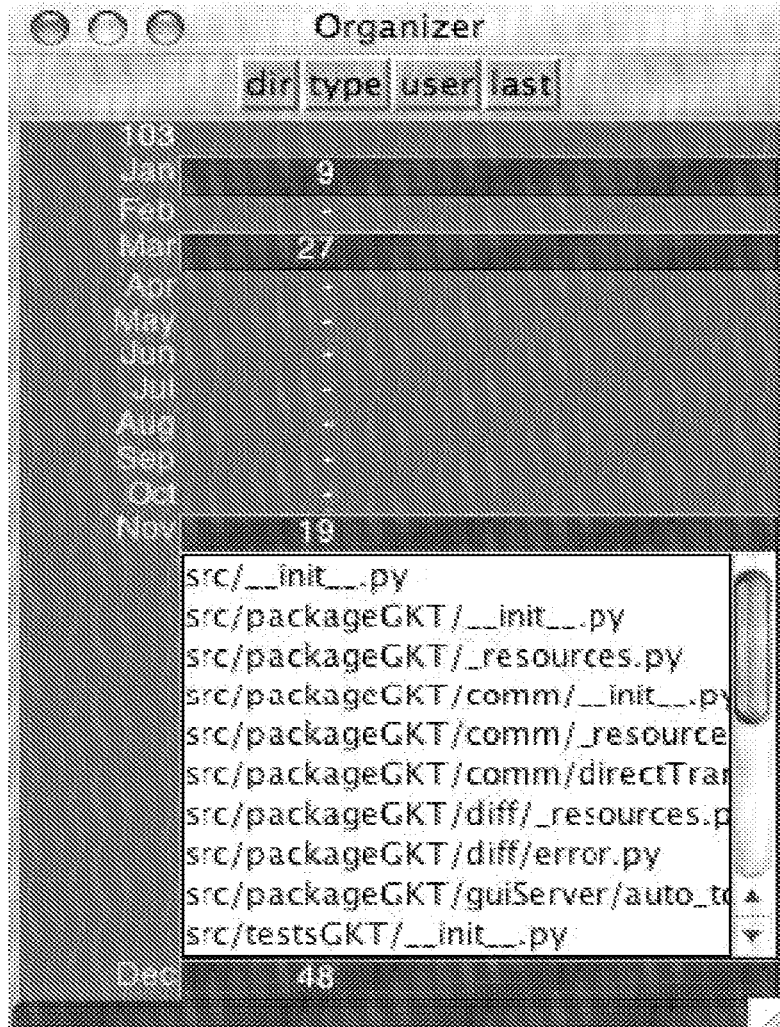


FIG. 16




FIG. 17




FIG. 18

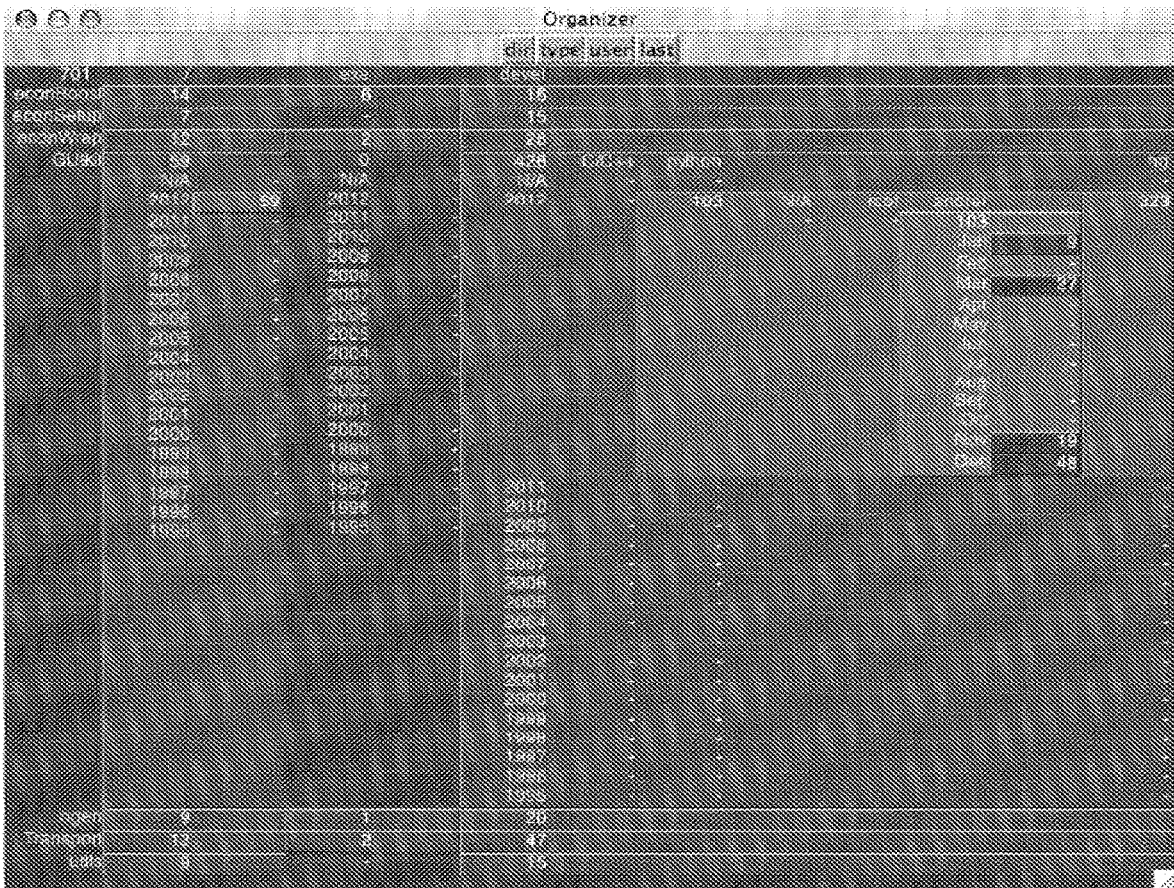


FIG. 19





FIG. 20

Organizer

	dir	type	user	last
root	0	0	0	0
bin	4	0	0	0
boot	2	0	0	0
dev	4	0	0	0
etc	4	0	0	0
home	64	0	0	0
lib	8	0	0	0
lib64	8	0	0	0
media	64	0	0	0
mnt	4	0	0	0
opt	4	0	0	0
proc	4	0	0	0
run	4	0	0	0
src	4	0	0	0
tmp	4	0	0	0
usr	64	0	0	0
var	64	0	0	0
video	64	0	0	0
work	64	0	0	0
com	4	0	0	0
com1	4	0	0	0
com2	4	0	0	0
com3	4	0	0	0
com4	4	0	0	0
com5	4	0	0	0
com6	4	0	0	0
com7	4	0	0	0
com8	4	0	0	0
com9	4	0	0	0
com10	4	0	0	0
com11	4	0	0	0
com12	4	0	0	0
com13	4	0	0	0
com14	4	0	0	0
com15	4	0	0	0
com16	4	0	0	0
com17	4	0	0	0
com18	4	0	0	0
com19	4	0	0	0
com20	4	0	0	0
com21	4	0	0	0
com22	4	0	0	0
com23	4	0	0	0
com24	4	0	0	0
com25	4	0	0	0
com26	4	0	0	0
com27	4	0	0	0
com28	4	0	0	0
com29	4	0	0	0
com30	4	0	0	0
com31	4	0	0	0
com32	4	0	0	0
com33	4	0	0	0
com34	4	0	0	0
com35	4	0	0	0
com36	4	0	0	0
com37	4	0	0	0
com38	4	0	0	0
com39	4	0	0	0
com40	4	0	0	0
com41	4	0	0	0
com42	4	0	0	0
com43	4	0	0	0
com44	4	0	0	0
com45	4	0	0	0
com46	4	0	0	0
com47	4	0	0	0
com48	4	0	0	0
com49	4	0	0	0
com50	4	0	0	0
com51	4	0	0	0
com52	4	0	0	0
com53	4	0	0	0
com54	4	0	0	0
com55	4	0	0	0
com56	4	0	0	0
com57	4	0	0	0
com58	4	0	0	0
com59	4	0	0	0
com60	4	0	0	0
com61	4	0	0	0
com62	4	0	0	0
com63	4	0	0	0
com64	4	0	0	0
com65	4	0	0	0
com66	4	0	0	0
com67	4	0	0	0
com68	4	0	0	0
com69	4	0	0	0
com70	4	0	0	0
com71	4	0	0	0
com72	4	0	0	0
com73	4	0	0	0
com74	4	0	0	0
com75	4	0	0	0
com76	4	0	0	0
com77	4	0	0	0
com78	4	0	0	0
com79	4	0	0	0
com80	4	0	0	0
com81	4	0	0	0
com82	4	0	0	0
com83	4	0	0	0
com84	4	0	0	0
com85	4	0	0	0
com86	4	0	0	0
com87	4	0	0	0
com88	4	0	0	0
com89	4	0	0	0
com90	4	0	0	0
com91	4	0	0	0
com92	4	0	0	0
com93	4	0	0	0
com94	4	0	0	0
com95	4	0	0	0
com96	4	0	0	0
com97	4	0	0	0
com98	4	0	0	0
com99	4	0	0	0
com100	4	0	0	0

FIG. 21



The image shows a window titled "Organizer" with a table of files and folders. The table has columns for "dir", "type", "user", and "last". The data is as follows:

dir	type	user	last
216		evr	10ve
econBoss	14	6	16
econSetup	7	-	15
econWres	12	2	28
nger	9	1	20
Transport	13	2	47
utils	9	-	11

FIG. 22

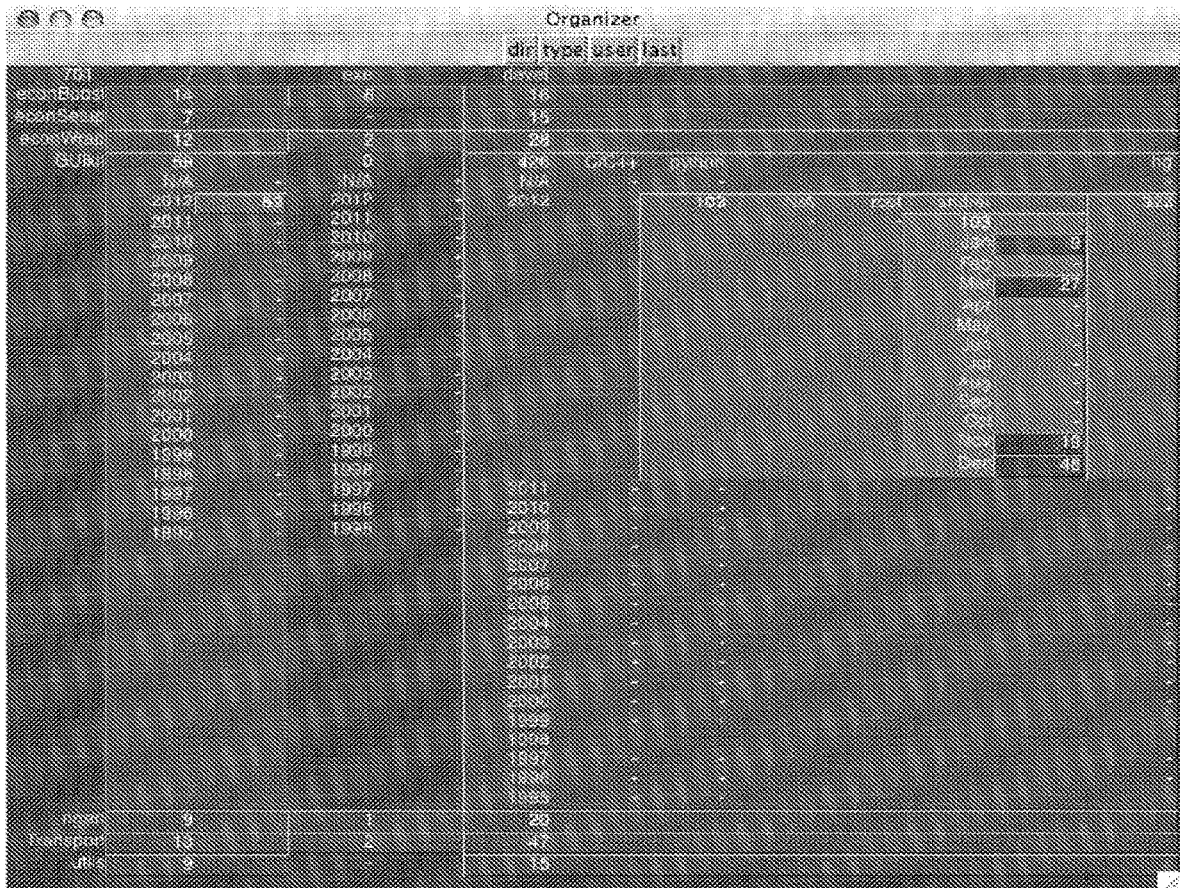




FIG. 24



The image shows a screenshot of a file manager window titled "Organizer". The window displays a table of files and directories. The table has columns for "dir", "type", "user", and "last". The data is as follows:

dir	type	user	last
701	-	-	-
anonBoast	14	6	16
anonSetup	7	-	13
anonWrap	12	5	20
GLib2	50	-	420
ngen	9	5	20
Transport	15	2	47
ubs	9	-	19

FIG. 25



The image shows a screenshot of a window titled "Organizer". Inside the window, there is a table with four columns: "dir", "type", "user", and "last". The table contains two rows of data. The first row has values "3456", "/", "eve", and "devel". The second row has values "563", "111", and "2782". The "type" column is empty for both rows. The table is displayed in a dark-themed interface.

dir	type	user	last
3456	/	eve	devel
	563	111	2782

FIG. 26




FIG. 28A




FIG. 28B




FIG. 28C




FIG. 29

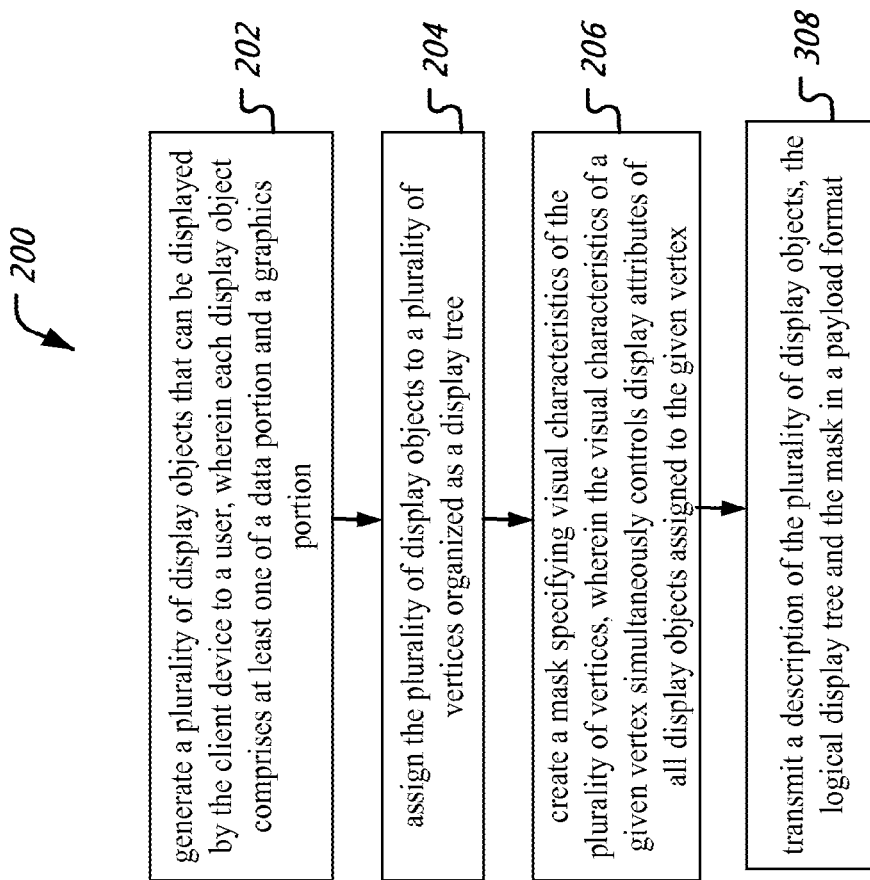


FIG. 30

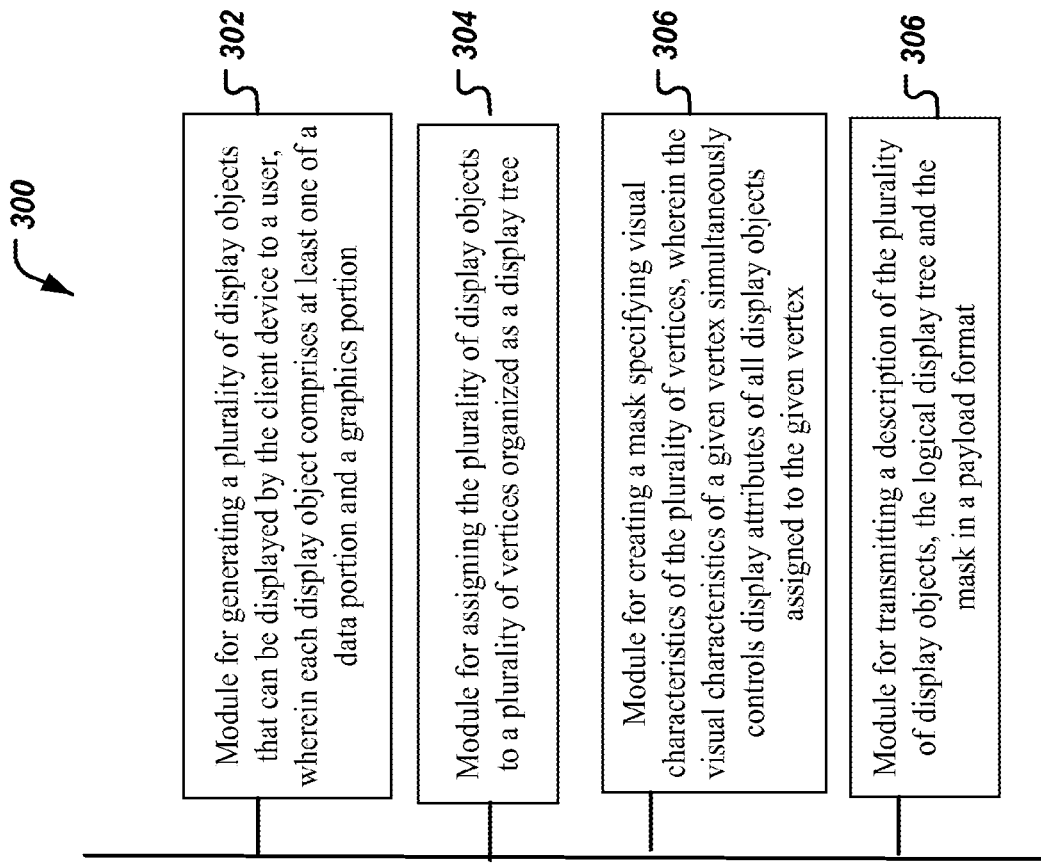


FIG. 31

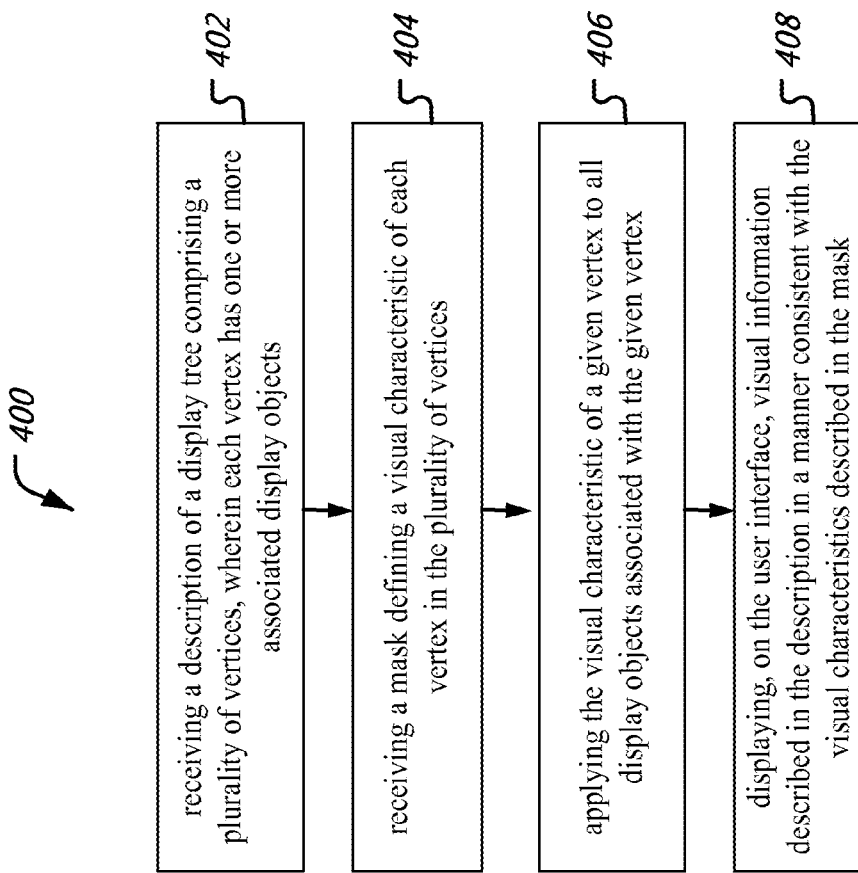


FIG. 32

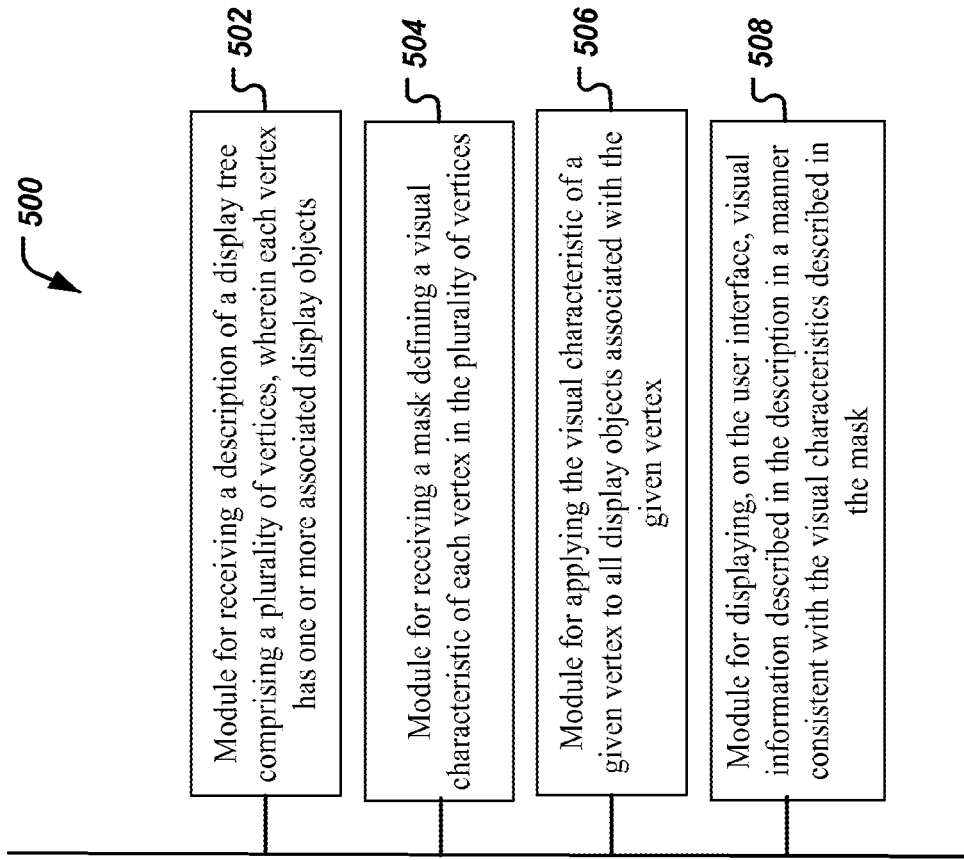


FIG. 33

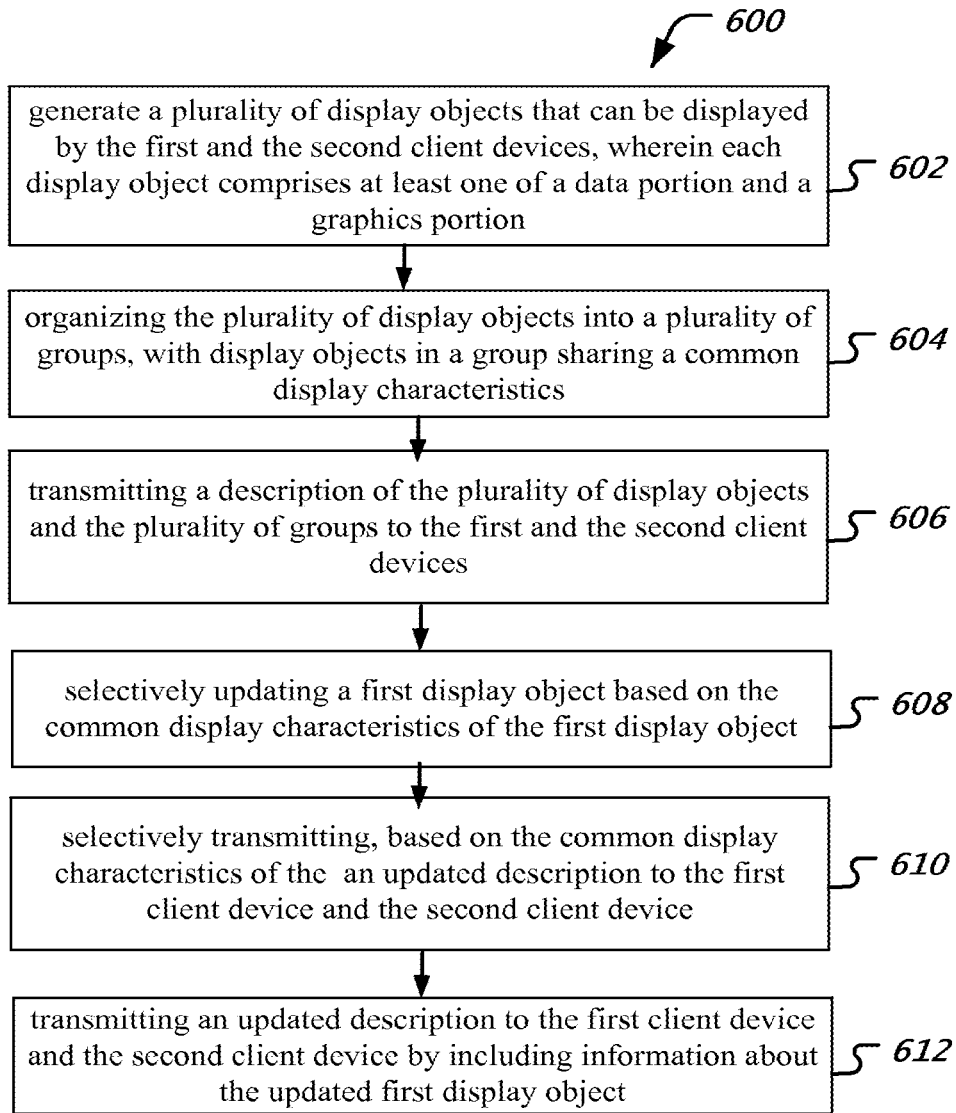


FIG. 34

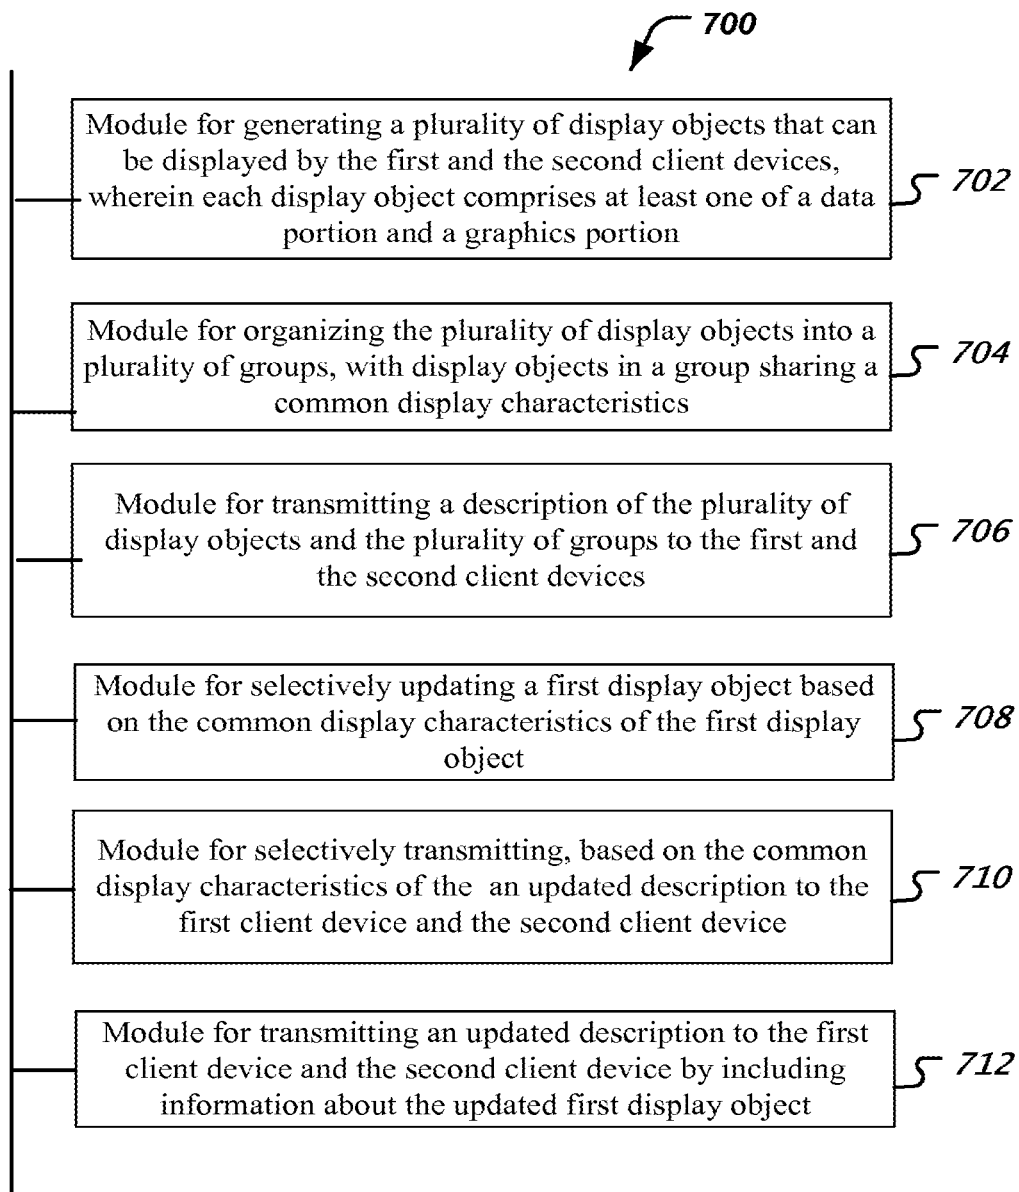


FIG. 35

DATA STRUCTURES FOR VISUALIZATION OF HIERARCHICAL DATA

CROSS REFERENCE TO RELATED APPLICATION

[0001] This patent document is a continuation of and claims the benefit of priority to U.S. patent application Ser. No. 14/769,041 entitled "DATA STRUCTURES FOR VISUALIZATION OF HIERARCHICAL DATA," filed on Aug. 19, 2015, which is a 35 U.S.C. § 371 National Stage application of PCT Application No. PCT/US2014/017209 entitled "DATA VISUALIZATION TECHNIQUES," filed on Feb. 19, 2014, which further claims the benefit of priority of U.S. Provisional Patent Application No. 61/766,531 entitled "DATA VISUALIZATION TECHNIQUES," filed on Feb. 19, 2013. The entire content of the aforementioned patent applications are incorporated by reference as part of the disclosure of this application.

BACKGROUND

[0002] This application relates to data management and display.

[0003] The amount of data that can be processed and stored by one or more computers has grown multi-fold over the last few years. The explosive growth in the data managed and processed by computers can be witnessed in application areas such as web servers, e-commerce servers, financial databases, multimedia content servers, and so on.

SUMMARY

[0004] The present document describes techniques for organizing display of large, complex data to a user. A user is able to navigate through data by applying various data classification criteria. After the application of a category, the data display can be updated in vertical or horizontal directions to display an attribute of the data that meets the applied criterion.

[0005] In one aspect, techniques are provided for providing visualization data to a client device. A server is controlled to generate a plurality of display objects for selectively displaying at the client device to a user, wherein each display object comprises at least one of a data portion and a graphics portion. Using the server, the plurality of display objects are assigned to a plurality of vertices organized as a logical display tree. The server is controlled to create a mask specifying visual characteristics of the plurality of vertices, wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex. The server is controlled to transmit a description of the plurality of display objects, the logical display tree and the mask in a payload format.

[0006] In another aspect, techniques are provided for displaying data on a user interface. A description of a display tree comprising a plurality of vertices, wherein each vertex has one or more associated display objects is received. A mask defining a visual characteristic of each vertex in the plurality of vertices is received. The visual characteristic of a given vertex is applied to all display objects associated with the given vertex. Visual information described in the description is displayed on a user interface screen in a manner consistent with the visual characteristics described in the mask.

[0007] In yet another aspect, a method of providing visualization data to a first client device and a second client device includes generating a plurality of display objects that can be displayed by the first and the second client devices, wherein each display object comprises at least one of a data portion and a graphics portion, organizing the plurality of display objects into a plurality of groups, with display objects in a group sharing a common display characteristics, transmitting a description of the plurality of display objects and the plurality of groups to the first and the second client devices, selectively updating a first display object based on the common display characteristics of the first display object, selectively transmitting, based on the common display characteristics of the an updated description to the first client device and the second client device, and transmitting an updated description to the first client device and the second client device by including information about the updated first display object.

[0008] The details of above aspects and their implementations are set forth in the accompanying drawings, the description and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1 is an example screen shot of a data browsing graphical user interface (GUI).

[0010] FIG. 2 is an example screen shot of a menu for data browsing.

[0011] FIG. 3 is another example of a data browsing GUI.

[0012] FIG. 4 is another example of a data browsing GUI.

[0013] FIG. 5 is another example of a data browsing GUI.

[0014] FIG. 6 is another example of a data browsing GUI.

[0015] FIG. 7 is another example of a data browsing GUI.

[0016] FIG. 8 is another example of a data browsing GUI.

[0017] FIG. 9 is another example of a data browsing GUI.

[0018] FIG. 10 is another example of a data browsing GUI.

[0019] FIG. 11 is another example of a data browsing GUI.

[0020] FIG. 12 is another example of a data browsing GUI.

[0021] FIG. 13 depicts a GUI control widget.

[0022] FIG. 14 depicts a GUI control widget.

[0023] FIG. 15 depicts a GUI control widget.

[0024] FIG. 16 depicts a GUI that provides visual menu indication to a user.

[0025] FIG. 17 is another example of a data browsing GUI.

[0026] FIG. 18 is another example of a data browsing GUI.

[0027] FIG. 19 is another example of a data browsing GUI.

[0028] FIG. 20 is another example of a data browsing GUI expanded along multiple directions.

[0029] FIG. 21 is another example of a data browsing GUI expanded along multiple directions.

[0030] FIG. 22 is another example of a data browsing GUI expanded along multiple directions.

[0031] FIG. 23 is another example of a data browsing GUI.

[0032] FIG. 24 is another example of a data browsing GUI expanded along multiple directions.

[0033] FIG. 25 is another example of a data browsing GUI expanded along multiple directions.

[0034] FIG. 26 is another example of a data browsing GUI.

[0035] FIG. 27 is another example of a data browsing GUI.

[0036] FIGS. 28A, 28B and 28C depict example relationships among data attributes.

[0037] FIG. 29 shows a system for data organization and presentation.

[0038] FIG. 30 is a flow chart representation of a method of providing visualization data to a client device.

[0039] FIG. 31 is a block diagram of an apparatus for providing visualization data to a client device.

[0040] FIG. 32 is a flow chart representation of a method of displaying data on a user interface.

[0041] FIG. 33 is a block diagram representation of an apparatus for displaying data on a user interface.

[0042] FIG. 34 is a flow chart representation of a method of providing visualization data to a first client device and a second client device.

[0043] FIG. 35 is a block diagram of an apparatus for providing visualization data to a first client device and a second client device.

DETAILED DESCRIPTION

[0044] Traditional data visualization programs such as Microsoft Excel provide a functionality-rich software application for users to store, manipulate and visualize data. However, these techniques are typically designed for use by a single user and do not generally allow simultaneous read/write access by multiple users. New document and data sharing techniques such as GoogleDocs, on the other hand, provide a solution for data viewing/editing by multiple users but often are provided limited functionality.

[0045] Furthermore, shared data viewing techniques such as GoogleDocs do not provide for mechanisms that provide multiple views of the same data to multiple users. For example, when users A and B are editing a GoogleDocs document, another user C can either participate or observe the changes, but cannot independently be working on the same data. In addition, the data sharing experience is not typically adjusted to be specific to a user's hardware platform.

[0046] The techniques disclosed in this document address the above-discussed limitations and others. In one exemplary aspect, a data transfer format called ROST (reference/optimistic/screen/typed value) is disclosed. Example advantages offered by the ROST format include reduction of data bandwidth when transferring a data view from a server to a client, the ability to perform fast refreshes at the server-side, supporting data view sharing by multiple users, resolving conflicts among data edits when multiple users attempt to concurrently change value of a given data field at the server, and so on.

[0047] The techniques disclosed in this document further facilitate viewing of data by multiple users based on synchronized views. The synchronized views may be used by one set of users while another set of users may be independently able to work on the same data.

[0048] In yet another example, the disclosed techniques allow for synchronized viewing by multiple users using client devices having different display capabilities. The ROST format, e.g., allows for client device-specific viewing of data to each user. These, and other, aspects are further disclosed below.

[0049] The techniques described in the present document can be implemented on a user's client devices using a Javascript or AJAX or similar platform. For example, the disclosed functionalities may be downloadable to a user device as a web browser plug-in or may be installed at a user device as an application (e.g., a Python application). In some implementations, no special installation or download at the user device is performed, with the functionalities simply relying on user device's HTML5 or Javascript compatibility.

[0050] Examples and implementations of techniques are provided for organizing data that includes multiple related records, and presenting the data according to a user's viewing choices are described. In some disclosed embodiments, data is classified using two or more hierarchical classifications. A user is presented with data based on a first level of break down, or data classification, specified by the user. The data thus presented can be further broken down into a second level of break down. In some disclosed implementations, the second level of break down can be accomplished either in a horizontal direction on a graphical user interface (GUI) or in a vertical direction on the GUI. The software tool that enables the above functions for browsing large sets of data based on multiple criteria is referred to as an Organizer.

[0051] Techniques are disclosed to simplify the task of browsing through large sets of data items. The disclosed graphical tool combines the visual and interactive advantages of tree and table displays. Visual cues are provided to a user to help with browsing of data on a computer platform that includes a processor and a display.

[0052] Large quantities of data, including text and numbers, can be presented as expandable lists (sometimes called trees) in which a list can be condensed to a top level headers (e.g., for a Microsoft Word file in the "document map" format). The top level headers can be individually expanded to reveal additional data, which is further expandable to additional levels. While the expandable lists offer a convenient way to manage data complexity by allowing a user to selectively zoom into data, such data expansion is limited to opening/closing additional data underneath a level in a single direction (typically vertical). Furthermore, there is no provision to apply multiple different classifications to achieve expansion of a list. When a user action desires expanding/collapsing data below a certain level, the additional data includes data that fits one category only. It is not possible in various existing, conventional or other software systems to achieve the expansion/collapsing based on multiple data attributes.

[0053] Complex data can also be presented to a user using a technique that is sometimes called "a pivot table." A pivot table can automatically sort, count, total or give the average of the data stored in one table or spreadsheet and display the results in the form of a second table called a pivot table. However, the depth of data sorting, or the level to which a user can dive into data, is only one layer deep in the pivot table technology.

[0054] With the emergence of complex computer systems and storage devices that can store large amounts of data, there is an ever-growing need for data classification, sorting and presentation to a user. The traditional tree/pivot table techniques fall short of addressing today's needs to be able to classify and present large quantities of data efficiently.

[0055] In the present document, examples of implementations of the described techniques are disclosed for orga-

nizing, categorizing and displaying data to a user. The disclosed techniques can be implemented on a user's computer, at a web server, or any other user device such as a tablet, a smartphone, a notebook computer, etc. In some implementations, the techniques may be used to transfer data between a server and client device and used to provide views of the data to a user.

[0056] As a non-limiting example, in some implementations, the disclosed technology could be used at a e-commerce server that offers thousands of catalog items for sale. Using the disclosed techniques, e.g., a user accessing the e-commerce server over the Internet, may be able to quickly sort through available merchandise using applying multiple classifications.

[0057] As another non-limiting example, in some implementations, an individual investor, or an investment professional, may be able to sort through financial data using various search criteria (e.g., price to earnings ratio, closing stock price, profit growth, market cap value etc.), by applying these criteria to a database of financial data.

[0058] Visually browsing large sets of data items is commonly performed. Just a few examples are: files on a computer, articles for sale on an online retailer's web-site, a person's credit card transactions for the past year, a factory's inventory. While database systems are a common solution for storing and querying data, presenting information in human readable form is a different task.

[0059] Two common ways to visualize data items are "tree+filter" and table. An example of "tree+filter" is browsing items on an online store: products are grouped by nested categories, and the view can be restricted by using various criteria. For a table display, each line and column represents a category, and the table shows some information regarding the items that match each line and column combination. Both ways can be unsatisfactory, particularly for large, complex sets of items.

[0060] Many limitations of current displays stem from giving some criteria preferential treatment, which often leads to relying primarily on one visual format. Computer files e.g. are easier to see by location than by other characteristics, and are not commonly shown as a table: it is not known in advance which attributes should be used for rows and columns.

[0061] Specifically, a graphical tool is provided for browsing data sets subject to any number of criteria. In one aspect, the graphical tool combines the visual and interactive advantages of tree and table displays, and allows the user to break down data by prioritizing criteria himself, with none privileged a priori. Aggregate measures for sets of items matching desired combinations of characteristics are easy to obtain. The disclosed examples for implementing the Organizer pertain to visually browsing data items and include, among others, the following aspects or features, (1) the techniques for data the Organizer works with; (2) user interaction including features of what a display user interface look like, and commands that can be given, and effects of such commands; (3) how a state of the display user interface is represented, either for volatile (memory) or permanent storage (disk or equivalent); and (4) the effect of commands on the display representation.

[0062] In the following sections, section numbers are used for the ease of understanding and cross-referencing.

[0063] 1. Examples of Browsing Data Using Multiple Criteria

[0064] The approach relies on combining a number of insights:

[0065] Many classification criteria can be manipulated in a uniform way by regarding them as directed graphs, a generalization of tree graphs.

[0066] A display composed of nested tables can be kept intelligible with proper formatting (colors, borders, alignment, etc.).

[0067] The interaction between data items and visualizations can be simplified by using aggregate measures for subsets of data.

[0068] Nested arrays are suitable for representing combinations of directed graphs.

[0069] 1.1 Criteria and Categories

[0070] For understanding the disclosed techniques, one start by analyzing characteristics independently of particular data sets. As it is not practical to enumerate all potential uses, it makes sense to look at a characteristic's intrinsic properties, with the understanding that data may or may not have it. For example, zip codes are applicable primarily to addresses, but things like parks or average daily temperatures possess locations as well; it would be incorrect to assume that zip codes are a characteristic of addresses only. If technically needed, an N/A category can be added to handle items for which the characteristic does not apply.

[0071] Zip codes exhibit a common problem: at approximately 43,000 of them, it is hard to make use of one without some looking up. However, zip codes are not random numbers; grouping of them into regions, states, counties, cities etc. can be used to make sure that at every level there are sufficiently few categories. This is helped by the fact that someone who is looking at the sub-categories of Chicago is likely to know what River North is.

[0072] Another example of a human-usable characteristic is time. With 2000 years being approximately 63 billion seconds, referring to event times by second alone is feasible for computers, but not humans. The solution is calendars and clocks, based on groupings into centuries, years, months etc. With no group too big, everyone can easily read and understand dates and times.

[0073] This grouping approach can be formalized by using directed graphs, a well understood mathematical notion with many applications to computing. A directed graph is a set of vertices, some of them connected by arrows.

[0074] For example FIG. 28A, represents a 2012 calendar as a directed graph. The vertices of the graph are referred to as categories and arrows join sub-categories to each parent category. If y can be reached by arrows (perhaps 0) from x, x is said to be narrower than y, and that y is said to be broader than x (by convention x is narrower/broader than itself). For example, Jan 1st is narrower than 2012, without being its sub-category. Finally, the depth of a category is the minimum number of arrows needed to reach the top category.

[0075] Throughout this description a criterion will refer to a classification characteristic that has been organized as a directed graph. A directed graph has a top category conventionally named all, and that it contains no cycles. Graphs like (FIG. 28B) are disallowed, as they do not model progressively narrower categories useful for classification.

[0076] The reader already knows many examples, as tree graphs fit the above requirements, and have the additional property that every category belongs to a unique parent category. Files organized by directory, dates and times, and

zip codes as above, are all trees. Not every criterion is a tree. Consider the classification of users of a website illustrated in FIG. 28C: with Joe both an administrator and a reader, this criterion is not a tree.

[0077] Turning a characteristic into a good criterion may benefit from an understanding of its intended use. A representative example is the calendar, which is not just an arbitrary breakdown of seconds, but also relates months to seasons, days to daylight etc. In many cases criteria incorporate domain knowledge and significant work; the Organizer derives value from allowing that work to be easily shared.

[0078] 1.2 Multi-Descriptions

[0079] Some presently disclosed techniques use criteria in combination. A choice of one category from each criterion under consideration will be called a multi-description, and be denoted $\{\text{crit}_1: \text{cat}_1, \text{crit}_2: \text{cat}_2, \dots\}$.

[0080] Multi-descriptions delimit subsets of data. In FIG. 11, the 59 files that are directories, sit under GUIkit and were last changed in 2012, share the multi-description $\{\text{dir: GUIkit, type: /, user: all, last: 2012}\}$. Another use of multi-descriptions is specifying changes to the display. The transition FIG. 10→FIG. 11 breaks down the multi-description $\{\text{dir: GUIkit, type: all, user: all, last: all}\}$ by the last criterion.

[0081] A multi-description x will be said to fit another multi-description y if every criterion's category appearing in x is narrower than the same criterion's category appearing in y . Equivalently, x is narrower than y and y is broader than x . The depth of a multi-description will be the sum of the depths of the categories comprising it.

[0082] 1.3 Nested-Table Displays

[0083] A nested-table display can be structured to include rectangular arrays and each array includes a total and a breakdown. The breakdown can be 1-dimensional (vertical or horizontal) or 2-dimensional. Numbers and labels can be arranged and formatted in a way that makes obvious what is a total and what is a breakdown, even when there are quite a few arrays.

[0084] One possible formatting scheme is used in the figures. Labels, rather than sub-aggregates, are aligned with the total, and different colors insure labels and values cannot be confused. When a child array is created, rows and columns are enlarged appropriately to maintain the alignment of parent arrays.

[0085] To make breakdowns visually distinct, they are enclosed in a border and given a different background color, with the exception of a breakdown that has no data, in which case a dash is displayed with no border, using the parent's background color. The breakdown background color loops through 4 shades of gray, which keeps adjacent areas distinct, without using too many colors. Labels, values and the hide/include category colors (see Section 3.2) seen clearly on these 4 shades of gray; light blue, white, orange and black work well.

[0086] Besides the rectangular array, the nested-table display generalizes the tree display, which can be simulated using 1-dimensional vertical breakdowns alone. That can be seen in FIG. 27, albeit with different formatting.

[0087] 1.4 Data Items and Aggregate Measures

[0088] The purpose of working with criteria is to apply them to data sets. The two are linked via a computer-performed calculation that, for each criterion and data item, provides the category the data item belongs to. For unifor-

mity, if a data item is not related to a criterion, a simple function that returns all or N/A can be used.

[0089] Given a data item x , a multi-description m_x can be obtained by performing this calculation for every criterion under consideration. m_x can be used to decide whether the item belongs to a set defined by a multi-description d , by checking if m_x fits d . This can be done via known algorithms for directed graphs.

[0090] In some cases a data item belongs to more than one category, e.g. a rain jacket sold under both casual and mountaineering wear. In that case, what is obtained is not one multi-description m_x , but a set of multi-descriptions, one for each combination of the several categories the item belongs to for each criterion. In this case, the item fits a multi-description d if at least one of its associated multi-descriptions fits d .

[0091] When working with large sets of data, a user tends to be interested in an overall characteristic of the set rather than individual items. Such overall characteristics are referred to as aggregate measures. Related to e.g. inventories, some familiar aggregate measures are total cost, total weight, and average time since purchased. Given a multi-description d , a data item set and an aggregate measure, an aggregate can be computed for the items whose multi-description fits d . All display values are aggregates.

[0092] Some aggregate measures are applicable to all data sets. The number of items in the set is an important one, often of interest. Another is whether the set is empty or not. While more abstract, it helps better reproduce tree displays with nested-table displays. Tree displays often show a

☐ when there are items underneath, but not otherwise. This can be replicated using the empty-or-not aggregate measure, and displaying ☐ or ☐ appropriately.

[0093] 1.5 Putting it Together

[0094] Nested-table displays are suited to working with multiple criteria. This is easiest shown on a concrete case, e.g., the transition FIG. 9→FIG. 10. For that transition, the type criterion was dropped onto the 426 files under the GUIkit and devel categories, i.e. the files that fit the multi-description $d=\{\text{dir: GUIkit, type: devel, user: all, last: all}\}$. As they are all devel files, breaking down again by type shows the sub-categories of devel: C/C++, python and hg (which can be seen in FIG. 27). For each sub-category, a new multi-description can be obtained from d by replacing devel, and an aggregate measure computed. The display is updated by adding a 1-dimensional, horizontal breakdown.

[0095] Generalizing, a state of the display is represented by a tree with one vertex for each aggregate shown. For example, the tree that generates FIG. 2 has a root vertex with 33 children, one for each line in the breakdown. Each vertex can contain:

[0096] The aggregate's multi-description.

[0097] Criteria assigned to the horizontal and vertical directions, if any. There are 3 possibilities:

[0098] The vertex has no criteria associated with it. Such a vertex will have no further breakdowns, and will be a leaf of the tree (no descendants).

[0099] The vertex has exactly one criterion associated with it. Such a vertex will have a 1-dimensional array of child vertices indexed by the sub-categories of the category of the assigned criterion appearing in the multi-description of the vertex. Geometrically, the child aggregates will be arranged along the direction the criterion was assigned to.

Child multi-descriptions are obtained by replacing the appropriate sub-category in the multi-description of the vertex.

[0100] The vertex has both a horizontal (C_H) and a vertical criterion (C_V) assigned. Denote by d the multi-description of the vertex. The child vertices will form a 2-dimensional array indexed by pairs of subcategories (s_h, s_v), where s_h spans the sub-categories of the category c_h of C_H appearing in d , and s_v spans the sub-categories of the category c_v of C_V appearing in d . Geometrically the new aggregates are displayed in a 2-dimensional array, with column labels provided by s_h and row labels provided by s_v . Child multi-descriptions are obtained from d by replacing c_h, c_v with s_h and s_v .

[0101] Whenever a criterion is assigned or unassigned, the vertex is recomputed and all child vertices are discarded. Any breakdowns in child vertices will be lost. If one criterion of a 2-dimensional breakdown is unassigned, the breakdown becomes 1-dimensional.

[0102] To interact with the display, commands to assign a criterion to a direction of an aggregate, and to undo a prior assignment, are needed. For some sample interface examples, described herein, this functionality consists of the $D+$ and $D-$ commands, with the limitation that a drop only works if the desired breakdown direction is free. To avoid confusion regarding figures, the sample interface may incorporate the features described with respect to “resolving drop conflicts.”

[0103] 2. Additional Example Refinements

[0104] While the discussion above describes some implementation examples, a number of variations are possible, each requiring some degree of work and insight. The sample interface described herein incorporates some, but not all of the features below.

[0105] 2.1 Displaying Item Lists

[0106] Besides displaying aggregate measures, it may be useful to display the items comprising an aggregate. An example interface for this are the $L+$ and $L-$ commands, with effects occurring in FIG. 15, FIG. 16, FIG. 17 and FIG. 18. A way to store this information is given by “aggregates currently displayed list” feature disclosed below.

[0107] 2.2 Simplifying the Display

[0108] By design, the number of sub-categories of a category in a criterion should be kept small for the display to be easily read.

[0109] Even so, not all the data may be of interest. A solution to simplify the display is to designate categories as included or hidden. A hidden category will not be shown, and its items will be excluded from aggregates and item lists.

[0110] To be precise, for aggregates below the breakdown where the hiding takes place, the display should behave as if the arrow in the directed graph, going from the hidden category to its parent category appearing in the breakdown’s multi-description, has been severed. While for a tree, items fitting the hidden category will no longer be tallied in the aggregate, for a general directed graph that is not true; narrower categories can be reached more than one way.

[0111] A possible user interface for managing included/hidden categories consists of the editing commands in Section 3.2. For examples, refer to FIG. 3, FIG. 4, FIG. 5, FIG. 6, FIG. 7, FIG. 8 and FIG. 9.

[0112] Another way to reduce the complexity of the display is provided by restrictions, which allow the display to behave as if the set of items was reduced to those matching

a multi-description. Any visible aggregate can be used to restrict; the relevant Section 3.2 commands are $R+$ and $R-$, put at work in FIG. 12, FIG. 13, FIG. 14, FIG. 15, FIG. 16, FIG. 17, FIG. 18 and FIG. 19.

[0113] To implement these two features, the editing and restriction information can be stored as part of the display state. For one convenient way see Section 2.5.

[0114] 2.3 Resolving Drop Conflicts

[0115] It can be desirable to perform more than one breakdown with a single command. As an example, look at the $D+$ command used to reach FIG. 11. Its intent is to break down all the aggregates on the GUIkit line using the last criterion. This requires 3 break-downs, resulting in a 1-dimensional array for the first two columns, and in converting a horizontal 1-dimensional array into a 2-dimensional one for the third.

[0116] If one command can result in multiple breakdowns, a feature is to be able to revert. For $D-$ to accomplish the transition FIG. 24→FIG. 25, it becomes desirable to represent how breakdowns are related.

[0117] Complex cases can occur. If the breakdown into $C/C++$, python and hg in FIG. 11 had been vertical, the additional last breakdown could not have been done directly, as the vertical direction is already occupied. Passing it down to children works, but creates further complications. Should the $C/C++$, python and hg breakdown be removed, the conflict disappears and child breakdowns would have to be consolidated. It is not clear how that should be detected.

[0118] One way of avoiding representing such complex relationships is to keep track only of desired drops, which are breakdown commands by the user, and translate them into applied drops, corresponding to displayed breakdowns. A desired drop can result in several applied drops, as seen above (FIG. 11). No applied drops is also possible, when the desired drop was based on the results of another desired drop that was later removed (e.g. FIG. 25→FIG. 26).

[0119] A desired drop can be characterized by the following:

[0120] target: The multi-description of the aggregate or label onto which the criterion was dropped.

[0121] direction: Horizontal or vertical.

[0122] criterion: The criterion that was dropped.

[0123] index: A unique integer number for every drop; earlier drops have lower numbers.

[0124] The relationship between desired drops and applied drops can be defined by 3 aspects:

[0125] No desired drop is applied above its target.

[0126] Every desired drop is applied once to data that fits its target. More precisely, every aggregate that is not broken down, and whose multi-description fits the target of a desired drop, have one applied drop coming from that desired drop, above it in the tree.

[0127] A desired drop is passed on to children only if it conflicts with a higher priority drop. Two desired drops conflict if they have the same direction or the same criterion, as they cannot be displayed correctly by a rectangular table of values with labels on the outside. Preference is given to drops with low target depths; for equal depths, a lower index will be applied first, i.e., the earlier drop.

[0128] Based on the desired drops, a tree in which every vertex corresponds to a displayed aggregate can be constructed iteratively. A vertex will have its own breakdown and child vertices if it has any applied drops. Visually the result looks the same as for Section 1.5; the difference is, in

essence, that in Section 1.5 applied drops are managed by the user interface directly, while here they are computed from desired drops.

[0129] Each vertex in the tree is characterized by:

[0130] content: A multi-description. Data items that fit the content are tallied in its aggregate.

[0131] applied drops: 0, 1 or 2 drops. They determine if the vertex is a single aggregate, a 1-dimensional breakdown, or a 2-dimensional breakdown. If there are two drops, they cannot conflict (see requirement (B) above).

[0132] For each vertex v in the tree, the following sets of drops are computed:

[0133] $I(v)$: The set of inherited drops that children of v tries to accommodate.

[0134] $P(v)$: The set of potential drops that might be needed below v . These are desired drops, for which v 's content does not fit the target. As such, by requirement (A) these drops cannot be applied yet.

[0135] $F(v)$: The set of found drops, potential drops of the parent which, absent conflicts, could be applied at v . Found drops are a subset of the required drops.

[0136] $R(v)$: The set of desirable drops that are to be applied to v or all its children.

[0137] $U(v)$: The set of redundant drops, which have not been applied yet but have become redundant at v or below.

[0138] $A(v)$: The set of applied drops, which will be used for the break-down at v .

[0139] The calculation starts with a single vertex root, whose content consists of criteria's all categories. For each vertex v , the above sets are computed from the sets $P(p)$ and $I(p)$ belonging to its parent p ; in the case of root, the set of desired drops and the empty set are used instead of $P(p)$ and $I(p)$. The iteration rules are:

[0140] $F(v)$ will consist of the drops in $P(p)$ with target broader than the content of v .

[0141] $R(v)$ will be $I(p)$ together with $F(v)$.

[0142] $U(v)$ will consist of the drops in $R(v)$ for which the category in the content of v corresponding to the drop criterion has no sub-categories.

[0143] $A(v)$ will consist of however many drops from $R(v)$, not in $U(v)$, can be accommodated without conflicts, taken in the priority order (see requirements (B) and (C) above).

[0144] $I(v)$ will consist of the drops in $R(v)$ that are not in $A(v)$ or $U(v)$.

[0145] $P(v)$ will consist of the drops in $P(p)$ not in $F(v)$.

[0146] After computing the above sets for a vertex v , its children are computed based on $A(v)$. With no conflicts among them, there can be at most two drops. The cases mimic those of Section 1.5:

[0147] $A(v)$ is empty. The vertex has no children.

[0148] $A(v)$ consists of exactly one element. The children will form a vector indexed by the sub-categories of the category in v 's content corresponding to the drop criterion.

[0149] $A(v)$ consists exactly of two elements. The children will form a 2-dimensional array.

[0150] This construction meets requirements (A), (B) and (C) above, and it can be shown that any other calculation satisfying (A), (B) and (C) will produce the same results.

[0151] With this approach, reverting a drop is accomplished by simply deleting it from the desired drop set, and the consequences are figured out by the calculation above.

[0152] 2.4 Mixing Aggregate Measures

[0153] It is desirable to be able to mix more than one aggregation measure in the same display. For example, after isolating a number of files by looking at file counts, the user could see how much the disk space they occupy by switching the aggregate to show disk space.

[0154] Aggregate measure drops can be represented by a triplet (target, aggregation measure, priority), similar to the description of desired criterion drops. The calculation of Section 2.3 can be adapted as follows:

[0155] Measure drops do not conflict with each other, or with criteria drops.

[0156] Measure drops do not cause the creation of child vertices.

[0157] Applied measure drops are left in the set of inherited drops, as they need to be applied to all child vertices.

[0158] Of the measure drops applied to a vertex, the display will use the one with least priority. This means the measure with the deepest target, or in the case of equal depths, the most recent one. Switching an aggregation measure will change child vertices, unless they have aggregate measure drops of their own.

[0159] The user interface could have a list of aggregation measures at the top of the display, and allow aggregate measures to be dropped wherever a criterion can be.

[0160] 2.5 Adapting the Display to Changing Data

[0161] In most realistic situations data is dynamic. Consider a tool to browse the articles available on an online store. When a display is saved for later use, if available articles or their categorizations change, the display will be lost unless it can be adjusted. The purpose of this section is to present a storage format that, in addition to supporting the features discussed so far, can be adjusted after changes in the data set, among the categories within criteria, and of the criteria list itself.

[0162] An instance of the storage format, called a display state, is comprised of the following:

[0163] Desired drops: For each drop, the following are stored:

[0164] target: A multi-description, see Section 2.3.

[0165] criterion: The breakdown criterion.

[0166] direction: The breakdown direction, one of horizontal or vertical.

[0167] index: A unique integer. Earlier criterion drops have lower numbers.

[0168] hiding data: Information about included/hidden categories. All occurrences of a category produced by a desired drop are hidden together. An example would be hiding one of the three 2012's in FIG. 11, making all 3 disappear.

[0169] Aggregate measure drops: Each consists of:

[0170] target: A multi-description, see Section 2.4.

[0171] measure: The aggregate measure.

[0172] index: A unique integer. Earlier measure drops have lower numbers.

[0173] Aggregates currently displaying item lists: Identified by the aggregate's multi-description.

[0174] Sets of categories currently under edit: Identified by the break-down's multi-description, and the set's geometric direction.

[0175] Restrictions: Represented similarly to aggregate measure drops, each consists of:

[0176] target: A multi-description.

[0177] index: A unique integer. Earlier restrictions have lower numbers. The above restriction data can be processed

as the tree of aggregates is constructed, and a sequence of multi-descriptions, each narrower than its predecessor, computed. The display can be adjusted based on this sequence.

[0178] As the display state does not reference data items, it can be updated for changes in data by simply recomputing it. For criteria and category changes, the adjustment is as follows:

[0179] Add a criterion: All multi-descriptions used throughout the display state representations can be extended with the all category of the new criterion. This creates a display state compatible with the new criterion list.

[0180] Remove a criterion: Any criterion drop, measure drop, restriction, edit, or item list expansion, that contains a multi-description referencing any category of the deleted criterion other than all, is deleted.

[0181] Add categories: A new category can not appear in existing multi-descriptions. It does however need a choice of hidden or included, and either is valid.

[0182] Remove categories: Again, display state components containing multi-descriptions that reference a category that no longer exists are deleted.

[0183] Importantly, the above adjustments produce a correct display state because there are no constraints among components. This is a key feature of the display state format, and of the calculations that render the display based on it.

[0184] 2.6 Minimizing Clutter

[0185] With many nested tables, the display can get complicated, and it helps to keep buttons and other non-essential features to a minimum. A system of drag-and-drops and context sensitive cursors makes the interface intuitive, visually simple and less error prone. The same drag-and-drop works with both horizontal and vertical breakdowns, and the user knows which will occur before a drop is performed (see Section 3.2). Likewise, distinct cursors for restriction and undoing drops make things easier.

[0186] 2.7 Filtering

[0187] Range filtering is useful when browsing data. For example, one could be interested in files changed from March 2009 through February 2012. While this can be accomplished by hiding sub-categories, it is tedious.

[0188] While in general a criterion does not have a natural order, trees in which every category's sub-categories are ordered can be used to define ranges. This is the case for numerical criteria, such as time. The order is not entirely free of convention, as for example one has to decide whether a year comes before its January sub-category, after its December sub-category, or somewhere in between. While before is quite reasonable, the user will have to remember that 2009 date 2012 means 1/1/2009, . . . , 12/31/2011. Ordering trees is a well understood topic in the prior art.

[0189] Filters can be stored similar to included/hidden category choices.

[0190] 3. Example User Interfaces

[0191] Example user interfaces are described through a file browsing example.

[0192] 3.1 Display Elements

[0193] Visually the display is composed of one or more of the following elements:

[0194] criterion: Available criteria are shown at the top.

[0195] category: Labels for rows and columns. Shown in light-blue if not being edited, and orange/black in edit mode.

[0196] aggregate: The aggregate value of items that match a combination of categories. Shown in white. For the file example, it is the number of files (see Section 1.4).

[0197] total: The aggregate for all data items under consideration. It is the very top-left value.

[0198] item list: A list of the data items comprising an aggregate.

[0199] 3.2 Summary of Commands

[0200] Here are the commands supported by the sample interface, with a mnemonic in parenthesis:

[0201] (D+) Drop: Drag and drop a criterion on a aggregate or category. When hovering over the drop area, the cursor will change to 4 or showing the direction that will be used for the breakdown.

[0202] (D-) Undo Drop: Drag and drop a category to any aggregate. The cursor will be x.

[0203] (L+)/(L-) See/Hide item list: Click on aggregate.

[0204] (R+) Restrict total: Clicking in the upper-left area of an aggregate causes the aggregate to become the new total. The hovering cursor is \square .

[0205] (R-) Undo restrict: Press Escape key. Removes the most recent restrict.

[0206] Editing categories: Each category can be excluded from aggregates and hidden from view. Included categories are shown in orange, hidden ones in black. Items from excluded categories are filtered out of any item lists.

[0207] (E) Edit: Click on a category. Aligned categories will be edited as well.

[0208] (Q) End edit: Press Enter. Saves include/Hide selections for the entire display.

[0209] (i)/(h) Include/Hide one category: Click on category.

[0210] (I)/(H) Include/Hide all categories: Press Equal/Minus key. Affects all categories under edit.

[0211] 3.3 Screen-Shot Sequence

[0212] The sequence will demonstrate browsing files organized by four criteria: location (dir), type of file (type), owner of file (user), and date when the file was last accessed (last). The aggregation measure in this case is the number of files matching a combination of categories. At the start, the display shows the criteria and the total number of files:

[0213] To see the breakdown of files by directory, the dir criterion can be dragged and dropped over the total (D+ command). If the drop is towards the left/lower corner of the total, the cursor will show 4, and the breakdown will be displayed vertically, as seen below:

[0214] To add a breakdown by file type, the type criterion is dragged and dropped towards the right/upper corner of the total (D+ command). The cursor changes to \Rightarrow . The display becomes a table with directories as lines, and file type categories as columns:

[0215] The user might not be interested in all sub directories. Sub directories (categories in general) can be hidden. The first step is to click on a category, e.g., econBoost (E command). This will put econBoost and its sister categories in edit mode; categories will be shown in orange or black, depending on whether they are included or hidden. In this case all categories are included:

[0216] If only a few sub-categories are of interest, the user can hide all sub-categories first by pressing the minus key (H command). Excluded categories show up in black:

[0217] To add individual categories back, the user can click on them (i command) in turn. The total will update after every change:

[0218] The Enter key exits the edit mode (Q command). The display now shows only the included categories:

[0219] Columns also can be hidden. After clicking on any column label, e.g. exe (E command), clicking on any included category hides it (h command). In this example N/A, docs and media were hidden:

[0220] The Enter key again exits the edit mode (Q command).

Drops (D+ command) can be applied to any aggregate, not just the total. Applying type horizontally to the files that match both the GUIkit location and the devel file type (the number 426) will expose the sub-categories of the devel.

[0221] Drops (D+ command) can also be applied to a label, in which case a breakdown will be performed for all aggregates under that label. Here the last criterion was applied to GUIkit vertically:

[0222] It is possible to restrict the total to an aggregate. In this case the 103 python files last changed in 2012 were chosen, by clicking in the upper left corner of the aggregate (R+ command). While hovering around the upper-left corner, the cursor becomes \. The display simply shows that aggregate as a total, maintaining the coloring:

[0223] Drops can be performed after restricting. Here the user criterion was applied (D+ command):

[0224] Successive restrictions are possible. Here the total was again restricted to the 103 files belonging to the user andrei (R+ command):

[0225] Drops can still be performed after successive restrictions. Here the last criterion is applied again to the 103 files (D+ command), showing months as the sub-categories of 2012:

[0226] The list of items comprising an aggregate can be displayed at any time by clicking on the aggregate (L+ command). The list is shown with a scroll bar, and the surrounding display is expanded to make room for it. In this example the 19 files for November are shown:

[0227] Restrictions can be undone, starting with the most recent one (R-command), by pressing the Escape key. The display maintains all other properties, e.g., the item lists displayed.

[0228] The list of items for an aggregate can be hidden by clicking again on the aggregate (L-command).

[0229] Undoing restrictions can continue as long as there are any remaining. Here Escape removes the first restriction (R-command), showing the full display:

[0230] The visual layout constructed for a category is remembered if a category is hidden.

[0231] The first step to demonstrate that is putting GUIkit and its sister categories in edit mode, by clicking on GUIkit (E command). Both hidden and included categories are shown:

[0232] Clicking on GUIkit again hides it, and the color turns black (h command):

[0233] Hitting the Enter key (Q command) exits the edit mode, and the line corresponding to GUIkit disappears:

[0234] GUIkit can be restored by clicking on any of its sister categories, e.g., econBoost (E command) and then on GUIkit (i command):

[0235] Enter again exits the edit mode (Q command), and hidden categories disappear:

[0236] As it can be seen, the visual details of the GUIkit line have been preserved.

[0237] Drops can be undone (D-command). This is done by dragging any category that resulted from that drop onto

any aggregate. For example, dragging the 2012 under GUIkit and/onto the 59 to its right removes the drop of last over GUIkit that led to FIG. 11. That drop caused breakdowns in 3 places (see FIG. 11), and all those break-downs are now removed together. This is important, as it would be tedious to revert all break-downs caused by a single drop one by one. The breakdown of 103 by user also disappears, as that relied on the category 2012 which became visible as a result of the drop now being reversed. The same applies to the breakdown by months resulting from applying last to the 103 files belonging to andrei.

[0238] The breakdown of GUIkit and devel by type survives, as it did not rely on the breakdown by last which was just reverted. A further D-command by dragging utils onto the aggregate 9 to its right leaves the display with a breakdown by type (some categories still hidden):

[0239] Another D-command returns the display to the initial state of FIG. 1.

[0240] An Example Remote Visualization Framework

[0241] The present document also describes a remote visualization framework that allows for visualization of hierarchically complex data. Only for the ease of explanation, the framework is referred to as “QE View.” However, it is to be understood that different embodiments that mix and match the disclosed features are possible.

[0242] In some embodiments, the visualization framework may focus on multi-user interaction and provide a highly responsive user interface. Web pages do not naturally fit this task, and any good solution faces data consistency and hardware differentiation challenges.

[0243] QE View separates visualization data into the hardware-specific and hardware-independent portions. The latter permits asynchronous updates, enabling real-time interaction among users, as well as handling many client events optimistically. The result is a responsive, real-time user interface, even with communication delays.

[0244] While the concepts behind QE View are not web-specific, they can be used in conjunction with web technologies such as AJAX and HTML5. This affords familiar advantages, such as sharing views as links and eliminating software downloads.

[0245] For some data, complexity stems from nested components, with each component relatively simple. Examples abound: investment portfolio, engineering designs, store and plant inventory, many scientific models. By contrast, other types of data, e.g., video content, have no identifiable components. Yet another possibility is exhibited by Wikipedia: while there are components, there is no natural hierarchy.

[0246] Not surprisingly, different types of information require different tools. Web pages are not naturally suited to hierarchical data, but fit Wikipedia extremely well. QE View is similar: while it is intended for hierarchical information, it can work with other data, albeit providing with fewer benefits.

[0247] A revealing example is Google Earth. Determining the content of a map fits under the mask concept (Section 7.1), and QE View could be used to provide multiple users with real-time interaction. However, whatever the demand for such a feature, it represents only a small fraction of the overall usefulness of the product.

[0248] Hierarchical data has specific visualization needs. Even more so than complex information in general, the screen layout is important, and should not be discarded when

an application quits. Many applications (e.g., Bloomberg Launchpad and the Eclipse IDE) go beyond one step further, and provide different layouts for various tasks.

[0249] Handling visual information becomes more difficult in a multi-user setting, as asynchronous changes can cause inconsistencies. Supporting different hardware adds to the challenge; window positions acquired from one computer are not useful on a smaller monitor or mobile device. As a result, retaining visual information is uncommon for web applications, despite their having acquired many of the advantages of stand-alone applications via technologies such as AJAX.

[0250] QE View handles these difficulties. Its features are, among others, to:

[0251] (1) Provide a highly responsive user interface.

[0252] (2) Allow saving display layouts.

[0253] (3) Enable real-time, multi-user interaction.

[0254] (4) Work with different hardware.

[0255] (5) Be efficient with respect to network traffic, memory requirements, and compute time (there are tradeoffs).

[0256] The benefit for using QE View in an application is proportional to its hierarchical complexity and degree of interactivity. While these requirements are more common in the commercial sector, many retail concepts fit to considerable degree. On-line store's inventory and product information, airline reservations, managing personal information, credit card statements are a few examples.

[0257] Concretely, QE View is a development framework. It exposes a set of classes that can be customized, and provides a client-server application for conforming data. The intent is that most complex aspects are isolated by the framework, resulting in a straightforward development process. Terminology-wise, the methods of a provided class that are to be implemented by any derived class, or likely to be customized, will be called the interface of the class.

[0258] 4.2 Example Ideas Behind QE View

[0259] QE View's basic approach to hierarchical complexity is expanding and collapsing objects. This is both simple and not new; what is original is realizing that by using a suitable storage format and adequate event handling, it can be made scalable, efficient and highly interactive, as well as compatible with small-screen devices, thus becoming an effective alternative to web-page navigation.

[0260] A fundamental concept introduced by QE View is that of mask/link tree (ML-tree). It is a view storage format that allows associating any view to any data, preempting inconsistencies (see Sections 6.3 and 6.2). ML-trees provide a common abstraction for expanding/collapsing sub-objects, scrolling large arrays, and conditional displaying of sub-components, such as showing features based on a map's resolution.

[0261] QE View's responsiveness relies on handling certain events optimistically, using update managers (Sections 7.2 and 7.3). By maintaining a cache of information on the client, and delaying updates from the server, it is possible for commands to take effect without waiting for a server response, with differences subsequently reconciled. The ROST update manager (Section 7.5) provides an original solution to editing scalars (single values) asynchronously without locking the interface, even with lengthy server delays.

[0262] Hardware differentiation can be done by isolating exactly which part of a view's description is hardware-

dependent. QE View's notion of detachment data, stored locally, (Sections 6.1 and 8), addresses window position concerns, and provides a navigation method for mobile devices.

[0263] 4.3 A Brief Overview

[0264] Section 5 describes, among other things, features for data QE View works with. The mask and payload concepts are introduced, and the DisplayObject class detailed.

[0265] Section 3 follows the steps to render a client display:

[0266] (1) Constructing a tree of display items: Section 6.1 describes the view storage format, Section 6.2 the information ML-trees transport, while 6.2.1 details some technical aspects of shared views.

[0267] (2) Sending client updates: Section 6.3 describes ML-trees, the key concept behind the client-server communication used.

[0268] (3) Updating widgets: Section 6.4 describes the Widget class and the various steps involved in patching widgets.

[0269] Events are covered in Section 7. QE View's messaging system is discussed first, followed by optimistic events and update managers. A specific example of optimistic mask update shows the interaction between events and ML-trees.

[0270] Section 8 discusses hardware differentiation using detachment data, for both desktop and mobile devices. Shared views are covered in Section 8.2, and the need for the DI-server of Section 6.2.1 explained.

[0271] Finally, the development framework is outlined describing the general approach and specific packages provided, with their additional requirements for display objects.

[0272] 5 Example Features for Data

[0273] QE View works with items of data called display objects, each corresponding to a client-side widget; widgets are organized as a tree. Their interaction with QE View is mediated by masks and payloads:

[0274] 5.1 Masks

[0275] Masks

[0276] Masks are a way to limit which descendants of a vertex are visible. A mask is a boolean function $mask(link)$, returning True if, given mask is in effect for the parent, a descendant characterized by link is visible. A few examples:

[0277] (1) Expansion mask: True/False for an expanded/collapsed object. For an expanded object, all descendants are visible, for a collapsed one, none.

[0278] (2) Range mask: Scrolling a window of size s over a vector is represented by masks $m_n(i)=n \leq i < n+s$, for $n \in \mathbb{N}$; the link of an entry is its position.

[0279] (3) Map mask: (long, lat, width, height, scale), with links $(x, y, maxScale)$; it evaluates to True if (x,y) belongs to the viewed rectangle and $scale \leq maxScale$. It is a 3-dimensional range mask.

[0280] Masks and links can be serializable.

[0281] 5.2 The Payload Class

[0282] A payload carries the information useful to build client side widgets. To minimize network traffic, only differences are sent. The approach is similar to using the diff/patch utilities for UNIX files, relying on the guarantee that $patch(x, diff(y, x))=y$.

[0283] The Payload interface consists of:

[0284] (1) pl_diff (other): Raising an error causes the corresponding widget to be rebuilt.

- [0285] (2) `pl_patch(diff)`: Should not raise an exception; if the patch cannot be done, `pl_diff` fails on the server.
- [0286] Payload classes is serializable.
- [0287] 5.3 The DisplayObject Class
- [0288] Display objects are derived from DisplayObject:
- [0289] (1) `id`: Unique identifier for each object, used to recover top objects in a view.
- [0290] (2) `stamp`: An integer describing when the object was last modified, analogous to database timestamps. If the stamp hasn't changed, it is assumed all methods below would produce the same results. (`descendants(mask)` will be run only if the mask has changed).
- [0291] (3) `oneValue()`: A scalar and/or icon, to visually reference the object, e.g. when it is detached or in window titles.
- [0292] (4) `Widget()`: The type of the widget to be constructed on the client; see Section 6.4.
- [0293] (5) `payload()`: Data useful to build the object's widget, exclusive of descendants.
- [0294] (6) `descendants(mask)`: The visible descendants given mask, indexed by link. For each descendant the following information is provided:
- [0295] a. `object`: The descendant display object.
- [0296] b. `placing`: Describes how to place the child widget in its parent, e.g., grid coordinates and labels.
- [0297] c. `mask`: Initial mask of descendant; use `None` to revert to the descendant's `defaultMask`.
- [0298] d. `format`: Initial format of child widget.
- [0299] e. `detach`: Detachment data for child widget; see Section 8.
- [0300] (7) Top object initialization: When an object is opened (added as a top object of the view), there is no parent to provide the mask, format and detachment data. `defaultMask()`, `defaultDetach()` and `defaultFormat()` are used instead.
- [0301] (8) Static methods:
- [0302] a. `calls()`: Returns a dictionary of callable objects providing access to application logic. The arguments are provided by widgets when the client generates events.
- [0303] b. `getObject(objectId)`: Recovers the display object with `id=objectId`.
- [0304] 6. Rendering a Client Image
- [0305] In broad terms, the steps are as follows:
- [0306] (1) The display item tree (containing the last-used visualization) and the display object tree (current data to be displayed) are combined into a new tree, with display items created or deleted as needed.
- [0307] (2) The result `T` is compared against a synchronized copy `S` of widget information, kept both on the client and on the server. The difference `T-S` is sent to the client.
- [0308] (3) `S` is patched with `T-S`, resulting in an `S'`, both on the client and on the server. It is not true that `S'=T`; `S'` is a superset of `T`, with `S'-T` consisting entirely of invisible vertices. `S` acts as a cache, by reducing network traffic from the server to the client, and allowing certain events to be anticipated on the client.
- [0309] (4) Widgets are updated on the client based on the new `S`.
- [0310] 6.1 The View Storage Format
- [0311] A view consists of a set of top objects, which appear in a window of their own. Each top object is associated a unique display item, which describes which sub-objects are to be displayed, and how. A display item consists of the following:
- [0312] (1) `id`: Unique integer across all display items.
- [0313] (2) `link`: Link to be used with the parent's mask.
- [0314] (3) `parent`: The parent display item, or `None` for top objects.
- [0315] (4) `mask`: The current mask to be applied to descendants.
- [0316] (5) `format`: Formatting information, interpreted by the widget.
- [0317] (6) `detach`: Default detachment data, per device class. See Section 8.
- [0318] Given the parent field, display items represent a tree. When rendering a client image, the display object and display item trees are reconciled, using existing display items as much as possible. Changes in object's descendants dictionaries result in adding or deleting display items.
- [0319] Display items and their association to top objects are persistent; links, masks, formats and detachment data are serializable.
- [0320] The application can open objects by the `openObject(objectId)` function provided by the framework. The effect will be to recover the object using `DisplayObject.getObject`, create a new display item with `link=parent=None` and `mask`, `format` and `detach` data provided by the object's `defaultMask`, `defaultFormat` and `defaultDetach` methods, and insert the new pair in the view's list of top objects.
- [0321] 6.2 Generating a Payload Tree
- [0322] The payload tree consists of a vertex for each client-side widget (instance of `Widget`) to be displayed; each vertex also corresponds to a display object. A vertex appearing in the payload tree is visible according to the masks.
- [0323] To generate it, a queuing algorithm may be used. The queue consists triples (`displayObject`, `displayItem`, `placing`), and is initialized with the top objects of the view, using `placing=None`. Processing a queue entry (`obj`, `di`, `plc`) is done according to the following steps:
- [0324] (1) Add a vertex to the payload tree: Each payload tree vertex contains an instance of the class `VertexState`. Following fields and values may be used:
- [0325] a. `id`: Set to `di.id`.
- [0326] b. `mask`: Set to `di.mask`.
- [0327] c. `link`: Set to `di.link`.
- [0328] d. `oneValue`: Set to `obj.oneValue()`, or from cache retrieve from cache if `obj.stamp` is not changed (see below).
- [0329] e. `payload`: Set to `obj.payload()`, or from cache.
- [0330] f. `Widget`: Set to `obj.Widget()`, or from cache.
- [0331] g. `format`: Set to `di.format`.
- [0332] h. `detach`: Set to `di.detach`.
- [0333] i. `placing`: Set to `plc`.
- [0334] j. `message`: Get messages addressed to `di.id` from the mailbox. See Section 7.1.
- [0335] (2) Generate descendants: The following steps may be followed:
- [0336] a. The descendants dictionary `descObject` is obtained by calling `obj.descendants(di.mask)`, or retrieved from cache.
- [0337] b. A dictionary `descDI` of display items with `parent=di.id` is constructed based on the permanent storage.
- [0338] c. For every link in `descObject` such that `di.mask(link)=True`, do:
- [0339] i. If link appears in `descDI`, let `childDI:=descDI[link]`.

[0340] ii. Otherwise, let childDI be a new display item with link:=link, parent:=di.id, mask:=descObject[link].mask (replace with obj.defaultMask() if None), format:=descObject[link].format, detach:=descObject[link].detach. Add childDI to permanent storage (see Section 3.2.1).

[0341] iii. Add (descObject [link].object, childDI, descObject[link].placing) to the queue.

[0342] d. For every link in descDI which is not in descObject, such that di.mask (link)=True, mark the display item as obsolete (see Section 6.2.1).

[0343] (3) Remove triple from queue.

[0344] The payload tree is complete when the queue is empty.

[0345] To speed up calculations, a cache stores a reference to the object, the mask last used, the stamp of the object, and the results of calling payload, descendants etc. If the stamp and the mask have not changed, cached results are used.

[0346] 6.2.1 The Display Item Server

[0347] Ideally, generating a view were a read-only operation, especially to deliver multi-user views. That is however not the case, as display items are added to permanent storage when they are first used, and deleted when they are no longer needed.

[0348] This can be solved by routing these operations through a display item server (DIS). Whenever a display item is needed, an in-memory representation is created to be used for that one tree construction, and an event is sent to the DIS. The DIS will create (or restore an obsolete one) or mark obsolete a display item only once. Importantly, it is not necessary to change a display item during payload tree generation.

[0349] Some cases require care. Suppose an object has d_0, d_1, d_2 as descendant dictionaries for stamps 0, 1, 2 (in the sense of the DisplayObject field stamp, Section 5.3), and there is a link l such that $l \in d_0, d_2, l \notin d_1$. Two refreshes R_1, R_2 are started, based on stamps 1 and 2, but R_1 finishes after R_2 . R_1 finds l obsolete and requests the DIS to mark it so. The DIS knows the request should be ignored.

[0350] This may be accomplished by including the stamp of the object in the request to the DIS. The DIS can compare the received stamp with the current stamp of the object, and ignore the request if the two are different. In this case $1 \neq 2$.

[0351] Only requests to mark as obsolete may be refused. Creating or restoring should be done all the time, as it has minimal costs. In the case of fast-changing objects, this avoids the potential problem of rejecting too many create requests.

[0352] 6.3 Diff/Patch for Mask/Link Trees

[0353] Sending a full payload tree to the client every time any change has occurred is impractical for large displays. Updates may be limited to visible changes that have occurred since the last client refresh, using the mask/link mechanism.

[0354] 6.3.1 ML-Trees

[0355] It is clearer to first describe the approach in an abstract setting. Let $0 \in \mathbb{C} \subseteq \mathbb{N}$ be a set of positive integers, thought of the set of all display item IDs used by the application. Let L an arbitrary set of links, and $l: I \rightarrow L$ the link of each display item (links are immutable). A mask will be any subset $M \subseteq L$, i.e. $M \in \mathcal{P}(L)$, where $\mathcal{P}(X)$ denotes the set of subsets of a given set X .

[0356] The set of possible items may be denoted by V (in practice $V = \text{VertexState}$). If V has diff/patch operations of its own, it may be possible to reduce network traffic by computing differences of elements of V . The term $v-v$ may represent the result of V -diff.

[0357] An ML-tree will be a set $0 \in T \subseteq I$, together with functions $p: T \rightarrow T$, $M: T \rightarrow \mathcal{P}(L)$ and $v: T \rightarrow V$. $p(x)$ is the parent of the vertex x ; the root is always 0, and $p(0)=0$. $v(x)$ is the useful value of the vertex, and $M(x)$ its mask.

[0358] A tree T' will be a subtree of T , and this may be written as $T' \leq T$, if $T' \subseteq T$ and $p=p'$, $M=M'$, $v=v'$ on T' . For a set $0 \in X \subseteq I$ and functions $p: X \rightarrow I$, and m, v as above (which is not a tree as $p(x) \in X$ is not guaranteed), the set of vertices reachable from 0 forms a tree $T[X]$. Thus, $T[X] = \{x \in X \mid p^n(x) \in X, \forall n \in \mathbb{N}\}$.

[0359] For any ML-tree T its visible part T^* may be defined to be equal to $T[\{x \in T \mid l(x) \in M(p(x))\}]$; it consists of vertices for which none its parents are masked. In this case, T^* is an ML-tree and $T^* \leq T$.

[0360] While only vertices in T^* are displayed on the client, it is useful to keep track of more than visible vertices. Consider scrolling a large array of size N . Let $I = \{0, 1, \dots, N\}$ and $l(i) = i$, $p(i) = 0$ for $i = 1, \dots, N$. Define ML-trees T, T' using $M = \{1, \dots, 100\}$, $M' = \{51, \dots, 150\}$ as masks at $0 \in I$. The transition T to T' corresponds to scrolling down 50 entries. After the scroll, the client has to wait for entries 101 through 150 to be received from the server. However, it is useful to be able to scroll back and display the old values for entries 1 through 50 right away, and receive updates only for changed values.

[0361] Consider construction of operations, \ominus, \oplus (diff and patch), such that for T_0, T_1 ML-trees with $T^*_1 = T_1$ give:

$$T_1 = (T_0 \oplus (T_1 \ominus T_0))^* \quad (1)$$

[0362] A sequence of ML-trees $T_n = T_n^*$ can then be displayed by iteratively constructing $S_0 = T_0$ and $S_n = S_{n-1} \oplus (T_n \ominus S_{n-1})$, which satisfy $S_n^* = T_n$. Synchronized copies of S_n are kept both on the server and the client. The construction also makes sure updates changes are only sent to visible vertices.

[0363] The approach can be thought of as a cache strategy. The server cache reduces network traffic, while the client cache allows some events to be handled optimistically. The identity $S_n^* = T_n$ allows the client to rely on the mask to determine what vertices should be displayed.

[0364] The \ominus, \oplus operations below are one possible cache strategy. The S_n 's can be limited in size by retaining vertices most recently visible. The client cache can be different, as long as it includes all vertices in the server cache; if the server assumes the client already has a vertex, it possibly still exists on the client. As memory and CPU constraints are likely more stringent on the server, a larger client cache can be beneficial.

[0365] To define \ominus, \oplus let T_0, T_1 be ML-trees with $T_1^* = T_1$. Define $\Delta_+ = T_1 - T_0$, Δ_- to consist of $x \in T_0 \cap T_1$ for which at least one of $p_0(x) \neq p_1(x)$, $v_0(x) \neq v_1(x)$, or $M_0(x) \neq M_1(x)$ is true, and $\Delta_- = \{x \in T_0 - T_1 \mid p_1(x) \in T_1, l(x) \in M(p_1(x))\}$. Δ_+ consists of new vertices, and Δ_- of changed vertices. The vertices of Δ_- do not appear in the new tree, yet their parent does, and the link is visible according to the new mask, and will be deleted.

[0366] Define $T_1 \ominus T_0 = ((p|_{\Delta_+}, v|_{\Delta_+}, M|_{\Delta_+}), (p|_{\Delta_-}, v_1 - v_0, M|_{\Delta_-}), \Delta_-)$. For patching, let $X_1 = (T_0 - \Delta_-) \cup \Delta_+ \cup \Delta_-$ by patch-

ing p_0 , v_0 , m_0 with the functions in $T_1 \ominus T_0$, giving preference to the latter whenever there is an overlap. Define $T_0 \oplus (T_1 \ominus T_0) = T[X_1]$.

[0367] It may be seen that $T[X_1]$ is an ML-tree, $T_1 \leq T[X_1]$ and $T_1 = T^*_{1 \leq T[X_1]}$. To see that $T[X_1]^* = T_1$, notice that if that weren't the case, any element of $T[X_1]^* - T_1$ of minimal distance from 0 would be in Δ_- , contradicting the definition of X_1 .

[0368] 6.3.2 Using ML-Tree diff/patch for Client/Server Interaction

[0369] The above operations translate into the following algorithm for updating a client display:

[0370] (1) Generate a payload tree, as in Section 6.2. The result T_n is an ML-tree with $V = \text{Vertex State}$.

[0371] (2) Compute $\delta n = T_n - T_{n-1}$. This is done by parsing T_n , and checking the conditions for Δ_+ , $\Delta \neq$ and Δ_- . Relying on the Payload interface, VertexState implements methods diff and patch, thus $v_n - v_{n-1}$ can be used.

[0372] (3) Send δn to the client.

[0373] (4) Compute $S_n = S_{n-1} \delta n$ on both the server and the client. This includes deleting vertices no longer reachable from 0.

[0374] (5) (optional on the server) Check that $T_n = S$. This should always be true if, have been implemented correctly, but it does not hurt to verify during testing.

[0375] Notice the client response is not delayed by steps 4 and 5 on the server.

[0376] 6.3.3 Features of an Underlying Transport Layer

[0377] QE View uses objects sent between the server and the client. The patch mechanism above relies on a sequence of the updates, therefore data sent by the server arrives in order. This is not required for client-generated events.

[0378] In some embodiments, the above can be implemented by establishing an ssh encrypted tunnel over TCP/IP, and using Python's pickle as a serialization tool.

[0379] For a web-based solution, AJAX/Comet provide the tools to establish asynchronous communication. Long events can be used to accomplish the ordered queue going from the server to the client. POST http requests can be used for sending events from the client to the server.

[0380] 6.4 Updating the Client Display

[0381] The first step is patching the client ML-tree, as per above. Vertices that appear in the diff tree are marked as "touched".

[0382] Once the ML-tree is correct, widgets are updated or build as needed. The process relies on the Widget class, exposing the following:

[0383] (1) Constructor: Takes two arguments:

[0384] a. parent: The parent vertex in the tree.

[0385] b. vertexState: The vertex information in the ML-tree, which includes the Payload instance associated with the vertex. See Section 6.2.

The constructor of a widget is run before those of its children. For operations that is to be done after the children have been constructed, $w_afterChildren$ can be used.

[0386] (2) $w_patch(\text{diff})$: Corrects the widget based on the result of p_l_diff . It can raise an exception, causing the widget to be rebuilt. The um_update methods of update managers used by the widget are called here.

[0387] (3) Geometric placement methods:

[0388] a. $w_place(\text{placing})$: Place a widget.

[0389] b. w_remove : Remove a widget, forgetting its position; reverts w_place .

[0390] c. $w_withdraw$: Make a widget invisible, but remember its placement.

[0391] d. $w_restore$: Restore a withdrawn widget.

[0392] (4) Optional methods (provided but can be overridden):

[0393] a. $w_afterChildren$: A function to perform tasks that use information from children. This is run whenever the vertex or its children change, or the widget is constructed.

[0394] b. $w_showAlert$: A function to update event alerts.

[0395] c. $w_setMask$: Withdraws/restores child widgets according to the mask. It relies on $w_withdraw$ and $w_restore$, and is exposed so that additional actions can be taken.

[0396] d. $w_updateMask$: Processes a change of mask in the client ML-tree. As optimistic mask events are used, it calls the mask update manager's um_update method. The call to $w_setMask$ may be delayed. See Sections 7.2 and 7.3.

[0397] e. $w_setPlacing(\text{placing})$: Places the widget according to placing. It relies on w_place and w_remove , and is exposed so that additional actions can be taken.

[0398] f. $w_updatePlacing$: Processes a changed placement by calling the widget's placing update manager, as widgets allow placing to be changed optimistically.

[0399] g. $w_destroy$: Destructor.

[0400] For each vertex that was touched in the patch process, or whose widget is to be built, the following steps are followed (recursively):

1. Patch the widget, unless:

[0401] (a) It is a new vertex.

[0402] (b) A parent widget was destroyed, causing the vertex' widget to be destroyed.

[0403] (c) p_l_diff failed on the server.

[0404] (d) The widget type has changed.

[0405] (e) w_patch fails. An existing widget that cannot be patched is destroyed along with its children.

2. If desired, build the widget.

3. Update the mask by calling the widget's $w_updateMask$.

4. Update the placement by calling the widget's $w_updatePlacing$.

5. Delete vertices that have disappeared from the tree.

6. Build widgets for new children.

7. Update child vertices that were touched by ML-tree patch.

8. Call event messages. Doing this step after w_patch and child updates insures um_update calls have already run, as required by update managers (see Sections 7.1 and 7.3).

9. Run $w_afterChildren$.

10. Run $w_showAlert$.

[0406] The update is complete once every new or touched vertex has been processed.

[0407] 7. Event Handling

[0408] 7.1 Sending and Tracking Events

[0409] While the server distinguishes between events that affect display items only, such as mask and format changes, from running callable objects provided by the application, on the client the mechanism is the same.

[0410] The Widget class provides the $putEvent(\text{event}, \text{handler})$ method for sending events. A unique ID is returned, and the event is registered with the vertex. The server provides an outcome message, and when received the client calls the handler with the message as argument. Which vertex an event is registered to affects when it is delivered, as only visible vertices receive messages; this insures visual

alerts are seen as they occur, e.g. flashing value updates. To avoid delays, an event can be registered to the root or a parent (but not child, see Step 8 in Section 6.4).

[0411] Messages resulting from processing events are conveyed to the payload tree generation thread via a message box server, which stores messages by display item id and event id.

[0412] 7.2 Optimistic Events

[0413] A display can be made more responsive by performing some changes before events are sent. While in general it is not possible to know what the effect of an event will be, in many cases there is a likely outcome that is worth displaying.

[0414] This makes sense for events altering the view, for example collapsing an object. Even in the unlikely case the event fails, the effect is easily reverted. When re-expanding a collapsed object, if child vertices are cached, the old information can be initially displayed. It should be possible for sequences of expand/collapse to proceed with no server delay.

[0415] QE View accomplishes this by distinguishing information received from the server, called reference data, from screen data that is actually displayed. The two are reconciled using event messages. A number of schemes are possible; the corresponding QE View concept is the update manager.

[0416] 7.3 Update Managers

[0417] Update managers encapsulate the logic of what to display and how to handle server updates. The SetAndSuspend manager used for mask events (expand/collapse, scrolls) blocks updates when a vertex has pending events (Section 7.4). The more complex case of inputting values uses the ROST manager (Section 7.5). Finally, non-optimistic events can use JustAlert to perform no visual changes but keep track of pending events.

[0418] Update managers share a simple interface:

1. A constructor that takes the following arguments:

[0419] (a) scrGet. A function to read the current display value.

[0420] (b) scrSet. A function to set the current screen value.

[0421] (c) val. The initial value.

2. um_update(val) Changes the reference value. Should be called from w_patch. The messaging system guarantees it is called before handlers, as long as events are registered with the vertex or one of its parents.

3. um_pending() Returns True whenever an alert should be displayed, typically when pending events are present.

[0422] Update managers are quite self-contained. Aside from expecting handlers to always run after um_update, the only feature is that any method that involves communication with the server, e.g. saving a value, is passed a function putEvent(event,handler) that returns a unique event ID.

[0423] From an allocation standpoint, update managers should be members of the ML-tree vertex rather than the widget, as to not be lost should the widget be destroyed (see Section patch-client).

[0424] 7.4 The SetAndSuspend Update Manger

[0425] Used by mask and format events, this manager is suitable for values that can be changed instantaneously (e.g., expand/collapse), but not inputting values. It is essentially a counter of events, and changing the reference value is ignored as long as the count is positive. The screen value can

be changed at any time by calling set(value,putEvent), whether events are in progress or not.

[0426] To clarify how update managers interact with ML-trees, here are the steps of an optimistic expansion of an object:

[0427] (1) A bound action (e.g. a mouse click) calls SetAndSuspend.set(True,widget.putEvent). This results in an immediate visual update, showing the children of the previously collapsed vertex. The screen mask is changed, but the reference mask is not.

[0428] (2) Visual mask updates are suspended until the event returns. If another user toggles the object, the effect is not shown, and similarly for a quick sequence of expand/collapse by the same user. The latter eliminates the object seeming to toggle “by itself”.

[0429] (3) The event is processed on the server, and the mask is updated.

[0430] (4) The new mask is picked up by ML-tree diff, sent to the client, and the client tree is patched. um_update is called, but as the event’s handler has not been called yet, the counter is still positive, and there is no visual effect. The order of calling w_patch and handlers is important here (Step 8 in Section 3.4 and Section 4.1).

[0431] (5) The event handler is called. Assuming only one mask event, this will end update suspension, and the screen will show the reference value. If the event succeeded, the screen is already correct and no update is performed.

[0432] 7.5 The ROST Update Manager

[0433] The ROST update manager (an acronym of the reference/optimistic/screen/typed values used) is designed for editing a single scalar value. An important application of is the formEditor package (Section 9.2).

[0434] Inputting values raises problems not presented by mask events. Updates while typing is in progress, attempts to edit while saving, and error values are handled. ROST has the following properties:

[0435] 1. Editing can always be started instantaneously.

[0436] 2. Unfinished edits are never accidentally lost.

[0437] 3. The display can always display the most up-to-date server value, as well as recent optimistic values.

[0438] 4. Values resulting an error are displayed in a visually distinct way, and can be re-edited.

[0439] One property of the messaging system may be that a message is generated after the event has been run. It is not assumed that processing events finishes in the order they were sent; the server can handle long events asynchronously.

[0440] ROST is a generalization of SetAndSuspend, whose set method can be replicated by an “edit+typing+save” sequence.

[0441] 7.5.1 Display State Representation

[0442] The display state consists of the following:

R Reference Value: The last scalar value received from the server. At any time, there is exactly only one reference value.

H History: A stack of triples, each consisting of:

[0443] 1. value: Either a reference value, or a value for which an attempt to save has been made (successful or not). Typed but unsaved values do not appear in the stack.

[0444] 2. id: For optimistic values, the id of the save event sent; None otherwise.

[0445] 3. message: The message received from the server. <OK> signifies success. A stack entry with a non-None id, and None message is a pending save, i.e. a response has

not been received from the server yet. If a save fails, the history retains the error message.

[0446] The stack obeys the constraint:

(SC) If the top entry is not pending, $H[\text{TOP}].\text{value}=\text{R}$, and $H[\text{top}].\text{message}$ is either None or $\langle\text{OK}\rangle$.

The message restriction insures that when a save event fails, but the saved value is equal to R (can happen if a value was resent, but the server rejects redundant updates), the client is not stuck with an error message.

[0447] T Typed value: A value that was typed in and not saved, but is currently not on the screen.

[0448] C Cursor: A token describing what the screen shows:

[0449] REF The reference information.

[0450] TYPING The the display shows an unsaved value. $\text{T}=\text{None}$ in this case.

[0451] TOP The top of the history stack, which is either the last save sent (pending or not), or the reference value.

[0452] int An integer pointing to a stack entry that is not top.

[0453] S The screen value. Whenever the cursor changes, S is updated accordingly.

[0454] 7.5.2 Events

[0455] The user interaction is provided by events. Following is their description and effect on the display state: (E) Edit. Accepted whenever $\text{C}\neq\text{TYPING}$. The state changes to $\text{C}=\text{TYPING}$, $\text{T}=\text{None}$, with R, H, S unchanged.

[0456] (D)Discard typing. Allowed only when the $\text{C}=\text{TYPING}$. The effect is $\text{C}:=\text{TOP}$, $\text{S}:=\text{H}[\text{TOP}].\text{value}$.

[0457] (S) Save. Receives `putEvent`. Allowed only when $\text{C}=\text{TYPING}$. The value to save is S, and the following steps are followed:

[0458] 1. An instance `event=ROST.SaveEvent(S)` is created.

[0459] 2. An `eventId` is obtained by calling `putEvent(event,handler)`. Calling handler triggers the (M) event below.

[0460] 3. The triple (`id=eventId`, `value=S`, `message=None`) is added to H.

[0461] 4. $\text{C}:=\text{TOP}$.

[0462] (U) Update. Receives `val`. Can occur for any cursor. The reference value is updated $\text{R}:=\text{val}$. (SC) may no longer hold. If that is the case, a triple (`value=R`, `id=None`, `message=None`) is added at the top of H, and an int cursor is replaced with TOP.

[0463] (M) Message return. The handler receives `eventId` and `message`. Can occur for any cursor. The (M) event has an effect on both H and C:

[0464] The stack is changed in 3 steps:

[0465] 1. The unique integer *i* such that $\text{H}[i].\text{id}=\text{eventId}$ is found.

[0466] 2. $\text{H}[i].\text{message}:=\text{message}$.

[0467] 3. If *i* points to the top of the stack, i.e., $\text{H}[i]=\text{H}[\text{TOP}]$, (SC) may no longer hold. If that is the case a triple (`value=R`, `id=None`, `message=None`) is added at the top of H.

[0468] The cursor is updated according to the following rules:

[0469] 1. If $\text{C}=\text{TOP}$, $\text{H}[i]=\text{H}[\text{TOP}]$ and $\text{message}\neq\langle\text{OK}\rangle$, $\text{C}:=i$. This has the effect of showing the error value (perhaps in a different format); *i* is no longer the top of H.

[0470] 2. If a new entry has been added to the stack to satisfy (SC), and C is int, then $\text{C}:=\text{TOP}$. This discards looking at old values in favor of showing an updated reference value.

[0471] (R) Switch to reference. Accepted when $\text{C}\neq\text{REF}$. The effect is $\text{C}:=\text{REF}$.

[0472] (B) Go backwards. Always accepted. The new cursor depends on the initial cursor as follows:

[0473] TY TOP.

[0474] PING:

[0475] RE TOP.

[0476] F:

[0477] TO Index of the stack item right bellow top.

[0478] P:

[0479] n: $n-1$, or no change if at the bottom of H.

[0480] (F) Go forward. Acceptable when the cursor is not TYPING. The new cursor depends on the initial cursor as follows:

[0481] RE TYPING if T None.

[0482] F:

[0483] TO TYPING if T None

[0484] P:

[0485] n: $n+1$ or TOP.

[0486] (C) Clear history. If $\text{T}=\text{None}$, $\text{C}:=\text{TYPING}$ and $\text{C}:=\text{TOP}$ otherwise.

[0487] If desirable for memory limitations, the stack can be trimmed at the bottom, limiting the history, with appropriate changes in event implementations.

[0488] 7.5.3 Visual Alerts

[0489] The display can be made more user friendly using visual alerts. The following system works well:

[0490] 1. When the concerned value is on screen, show a hard flash (strong colors):

[0491] (a) A hard/normal flash, to show a value being updated, e.g. R changes and there are no events pending; the new value is visible immediately.

[0492] (b) A hard/error flash, to show an error has occurred when $\text{C}=\text{TOP}$ and $\text{H}[\text{TOP}]$ is pending; the error-causing value remains visible.

[0493] 2. When $\text{C}=\text{TYPING}$, and an event concerns a value not currently on screen, show a soft flash (pale colors) without leaving edit mode:

[0494] (c) A soft/normal flash, to alert the user the $\text{H}[\text{TOP}].\text{value}$ value has changed.

[0495] (d) A soft/error flash, to alert the user the pending save at $\text{H}[\text{TOP}]$ failed.

[0496] 8. Hardware Differentiation

[0497] QE View is intended as a portable framework. The ability to support different hardware has become increasingly important, driven by the desire to access information from work, home and mobile devices. A lot of industry effort is devoted into this issue, and resulting technologies like HTML5 and Javascript can be used by QE widgets.

[0498] Handling different monitor sizes (or multiple monitors), however, is not solved by the above technologies. This issue is particularly relevant for hierarchical data: expanding an object “in-place” is useful on a large monitor, but awkward on a smaller device as it requires too much scrolling. Web design runs into this problem as well, and web-sites dedicated to mobile are common.

[0499] QE View addresses differences in display size by separating detachment data from view storage. Every display item can be toggled between two modes:

[0500] 1. Embedded: The parent display item is enlarged to accommodate it.

[0501] 2. Detached: Shown as a separate window.

[0502] Detachment data consists of the mode and the window coordinates, and is stored on the client machine. QE View has a concept of installation, each tied to one device only, and containing a device class characterization. Detachment data is stored keyed by display item and installation, thus different installations on the same machine do not interact.

[0503] Each display item is constructed with default detachment data, one per device class. An expansion can be in-place for a desktop or tablet, but detached on a phone. When a view is first opened by an installation, the default detachment data is downloaded; subsequent changes are stored locally.

[0504] Local detachment data can be uploaded to the server to replace the device class defaults, or conversely, the defaults reloaded. A local setup can be replicated on a device in the same class by combining the two, without affecting other device classes.

[0505] 8.1 Small Devices

[0506] The meaning of detachment coordinates is for each device class to decide; any string, and by extension any serializable object can be used. For a small enough device, window coordinates make no sense, as each window takes up the whole screen. While a phone will never have the same capability as a large monitor to show complex information, detachment data can provide a way to organize and access objects via a navigation map and shortcuts.

[0507] The navigation map consists of an icon for each object, viewed similarly to apps on a smart phone. The grid position of the object's icon is stored as detached data, along with shortcuts and priority in the object stack (determines which object is on top). Detachment data translates into an icon configuration and shortcuts; customizing and sharing configurations works similarly to window positions for the desktop class.

[0508] Bear in mind a view's masks and formats are shared among installations. The navigation map visually differentiates collapsed and expanded objects, and touching a collapsed object expands it. A synchronized desktop will see the toggle as well.

[0509] Finally, the notion of activation provided by the selection package (Section 9.2), which on a desktop has the effect of changing the color of the object's top bar, is interpreted on a small device as raising the object to the top of the stack; users can show objects to one another.

[0510] 8.2 Shared Views

[0511] Shared views rely on a number of design features presented so far. Display items are stored in a database, allowing concurrent updates. The DIS (Section 3.2.1) reconciles create/delete display item requests from multiple sources, while changes are distributed automatically via ML-trees. Finally, storing detachment data locally allows masks and formats to be shared among users, while adapting the views to specific hardware.

[0512] Additional features can be added. Especially for a large number of users, it is desirable to restrict user's ability to change a view. Three modes are provided:

[0513] 1. Free: Any user can act at any time.

[0514] 2. Grab: Only one user has focus, but anyone can grab it at any time.

[0515] 3. Conference: Only one user has focus, and he can cede it at someone else's request.

[0516] The user interface provides a share bar, showing information about users currently on. It consists of:

[0517] 1. A left button, used to request control. Shows own status. Grey=Idle, White=Focus, Flashing White=Requesting.

[0518] 2. Middle buttons, for other user's status. When another user requests, the button splits in two. Top part flashes white, bottom is solid red. Clicking the top grants control, bottom denies.

[0519] 3. Right button. Deny all.

[0520] Every view has an owner who can set the mode, and grab control for himself

[0521] 9. Development Framework

[0522] The framework is currently for Python. Using a Django/jQuery/JavaScript stack to deploy as web technology, the client side can be ported to JavaScript/HTML5, to run within the browser environment.

[0523] 9.1 Example Principles when Developing with QE View

[0524] Data is separated into categories, reflecting how QE View handles it:

1. Core data: Application data that makes sense absent displays. Core data complies with the DisplayObject interface, requiring Payload and Widget classes.

2. Display data: This is the part of the visual information shared among users, and is represented in terms of masks and formats.

3. Local adjustments: Can be either volatile or permanent, e.g., using cookies or HTML5's localStorage.

[0525] Secondly, optimistic events should be identified, and an update manager decided upon. As with any concurrent problem, users can attempt conflicting things. A few principles help:

[0526] 1. Make the information granular. This reduces the number of collisions among events.

[0527] 2. Have events be incremental. If a format consists of ordering a number of items, events should be transpositions rather than the new permutation; this prevents stale data from being reinstated.

[0528] 3. Events should commute whenever possible. For example, expand/collapse is implemented as a toggle rather than setting the mask to value, as toggles commute, while sequences of set do not.

[0529] While for many applications the classes provided eliminate the desire to write client-side scripts, QE View is designed as a general purpose tool. In many cases it will be useful to customize existing classes, or implement the interfaces from scratch.

[0530] New masks/links can be created, subject to being serializable. The JustAlert, SetAndSuspend and ROST managers cover a large number of situations, but new update managers can be implemented based on the messaging system; the principles described in Section 7.3 should be followed.

[0531] 9.2 Packages

[0532] The framework provides a number of packages that simplify common tasks. As expected, they impose requirements in addition to those of DisplayObject.

[0533] 1. selection: Implements a copy/paste mechanism, and a concept of active widget. Objects implement:

[0534] (a) pl_Type: A field within the payload class. Used by the client to perform some type checking as

part of the copy/paste functionality. For example, an integer should not be confused with an object id.

[0535] (b) set (link,value): For the paste event.

[0536] (c) choicePath() Returns a list of links. This allows visually highlighting a sub-display item, e.g. a choice in a drop-down list. Using a list of links as the return, rather than just one link, allows handling cases where the selection is not directly visible. An example is making a selection within a tree, where the item selected can belong to an unexpanded vertex; in that case the object that is to be expanded to find the selected item is highlighted instead.

[0537] 2. execution: Executing methods for visual objects. Relies on selection. Objects implement a methods (mask) method, returning a dictionary of methods with bindings as keys, in a format understood by the corresponding widget. Entries should expect the following arguments:

[0538] (d) mask: The object's mask in the context where the method is called. There is no way for the object to recover this information independently.

[0539] (e) path: The succession of links to reach the selection from the object whose method is being executed, if within the object. One use is container operations.

[0540] (f) selection: Selection on the client. Allows e.g., inserting a value into a list.

[0541] (g) pl_Type: The type of the selection.

[0542] The return can be either None or an object id. The view will open the id as a top object, if it is not a top object already, and bring its window to the foreground.

[0543] 3. formEditor: Implements optimistic form editing using the SetAndSuspend and ROST update managers. Relies on selection.

[0544] A form can be composed of scalars or sub-objects. The Form class provides functionality to request editing rights from the server; such requests are handled optimistically by a SetAndSuspend manager. Each individual scalar has its own ROST manager. The interaction is linked: attempting to edit a scalar causes a request to put the form in edit mode, and a save command causes all sibling scalars to be saved as well.

[0545] The combination of SetAndSuspend and ROST for scalars provide a robust solution to a problem that otherwise requires complex state diagrams. A fairly large number of cases are covered, e.g., when a scalar edit resulted in an error, but the user loses the right to edit the object. The formEditor provides editing without delays in all situations, difficult to achieve with state diagrams. Indeed, the ideas of ROST were arrived at after repeated attempts at the problem using state diagrams.

[0546] Implementation-wise, the package contains the class ScalarWidget, which wraps a ROST manager, providing hooks the form can use to take actions when e.g., a scalar is edited, allowing the above linkage between forms and scalars.

[0547] 4. array: 1 and 2-dimensional arrays with optimistic scrolling, using range masks.

[0548] FIG. 29 depicts an example system 2900 comprising client computers 2902, 2903 communicatively coupled over a communication network 2904 with a server 2906 and a database 2908. The various techniques described in this document can be implemented at the client computer 2902, the server 2906 or partly between the client computer 2902 and the server 2906.

[0549] The communication network 2904 can be a suitable network, such as a wired or wireless network, e.g., the Internet, Ethernet, wireless cellular network such as 3G, Long Term Evolution (LTE) network, WiMax, etc. In various embodiments, the client computer 2902 may be a computer, a smartphone, a tablet device, using a suitable operating system.

[0550] The server 2906 may include a processor, instruction memory and a suitable operating system. The database 2908 may be implemented using storage such as hard drive, flash memory, etc. In various embodiments, the server 2906 and the database 2908 may be implemented on the same hardware platform (e.g., sharing the same power source) or may comprises two different platforms connected over a local connection such as an SCSI interface.

[0551] Many systems provide database viewing access to users using client devices via a server, while the underlying data itself may also be changing. The changes may be caused by other applications running on the server or by some users of the system. One challenge is to provide the most current or accurate data to a user, while the data may be changed. Several techniques are possible. For example, in some implementations, e.g., when a user is using a web browser to view data, the user may be able to refresh the web page loaded in the web browser (either manually or automatically) to receive and view the latest data from a server. This solution may be tedious to users and may be unacceptable because of the constant user interaction required. In some implementations, the server may "push" live data, e.g., stock quotes, to the client device. However, scaling such a solutions to millions of users may be run into operational issues such as computational power and bandwidth. Furthermore, streaming view of data may not allow a user to modify the data and propagate the modifications to other users.

[0552] As a non-limiting example, consider the case of viewing airlines seating chart while making flight reservations. Some airlines web sites allow users to click on a menu item that opens up an airlines seating chart showing available/taken seats. Often, multiple users may be looking at the same seating information from the airlines' server database and some users may be attempting to make bookings of their seats. In some implementations, a server that provides a view of the airlines seating chart to the client may need to provide a way for the client to be able to view any updates to the data in real time (e.g., within 1 to 10 seconds). It may further be advantageous for the server to provide the data in a format such that only data that is in current active view of the user is refreshed with current values at the server. The ROST data format and the display tree structure based conveyance of viewable data provides a format that meets these objectives, and others. While the above-discussed airlines seating chart example is one example of an application that may benefit from the disclosed viewable data format, the applications of this format are easily scalable and useful in many other situations, including but not limited to, web browsing on an e-commerce web site, accessing financial data, etc.

[0553] While viewing data, certain data may not remain in a user's active view. For example, the user may scroll past the data so that the data is outside the active display area that is being viewed by the user. Another example may be that the user may collapse the viewable data by clicking on a "hide" control widget (e.g., a "+" mark in a listing). The server may track a current state of the viewing session and

may either refresh slowly or not refresh at all the data values that are not currently in the user's active view, compared to when the values are in the user's active view. This hide/refresh technique may be applied to all data values that can be collectively navigated in or out of the user's active view and may therefore be treated similarly (e.g., all data objects belonging to a vertex of a display tree).

[0554] FIG. 30 is a flowchart representation of a process 200 of providing visualization data to a client device. The process 200 may be implemented, e.g., on the server 206.

[0555] At 202, a server is controlled to generate a plurality of display objects for selectively displaying at the client device to a user, wherein each display object comprises at least one of a data portion and a graphics portion. For example, in some embodiments, the display object may include alpha-numerals such as stock values, or data entries. In some embodiments, the display objects include icons, images, an airlines seating chart, and such.

[0556] At 204, using the server, the plurality of display objects are assigned to a plurality of vertices organized as a logical display tree. As explained before, and in, e.g., sections 6 and 7, the logical display tree, organized as a plurality of vertices, provides an advantageous way of selectively refreshing or changing a subset of all display objects.

[0557] At 206, the server is controlled to create a mask specifying visual characteristics of the plurality of vertices, wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex. The visual characteristics may specify whether or not each of the plurality of vertices are visually occluded. In some implementation, the visual characteristics may specify a transparency level, a font type, a size of display of the vertices. In some implementations, the visual characteristics of a vertex are automatically imposed upon the data objects that are assigned to the vertex. For example, when a vertex is hidden from a user's view, all data objects assigned to the vertex are also hidden from the user's view.

[0558] At 208, the server is further controlled to transmit a description of the plurality of display objects, the logical display tree and the mask in a payload format. For example, in some implementations, the reference/optimistic/screen/typed value (ROST) format disclosed in the present document is utilized. In some implementations, the description is transmitted as a difference from the previously transmitted description. In one advantageous aspect, this transmission of differences ("cliffs") can reduce the amount of data transmitted from the server 2904 to the client device 2902.

[0559] FIG. 31 is a block diagram representation of an apparatus 300 for providing visualization data to a client device. The module 302 (e.g., a generator module) generates a plurality of display objects for selectively displaying at the client device to a user, wherein each display object comprises at least one of a data portion and a graphics portion. The module 304 (e.g., an assignment module) assigns the plurality of display objects to a plurality of vertices organized as a logical display tree. The module 306 (e.g., a mask creation module) creates a mask specifying visual characteristics of the plurality of vertices, wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex. The module 308 (e.g., a data transmission module)

transmits a description of the plurality of display objects, the logical display tree and the mask in a payload format.

[0560] As previously discussed, the ROST format provides a novel way for a server to provide views and updated of views to data objects stored on the server to a client device. From a client device's perspective, the ROST format provides the advantages, among others, that the received data is in a compact format so that relatively lighter processing of data is performed at the client device. Furthermore, in some implementations, changes to data objects since the last refresh are organized as vertices that provide a collective state of multiple data objects, thereby simplifying rendering task. A mask may, e.g., provide information about which vertices of a display tree are to be exposed to a user's view and which are not.

[0561] FIG. 32 is a flowchart representation of a method 400 of displaying data on a user interface. The method 400 may be implemented on a client device 2902.

[0562] At 402, the method 400 receives a description of a display tree comprising a plurality of vertices, wherein each vertex has one or more associated display objects. The description may be received using the ROST format. The description may be received at the client via a message received over the network 2904.

[0563] At 404, a mask defining a visual characteristic of each vertex in the plurality of vertices is received. Various implementations of masks are disclosed in Sections 6, 7, 8 and 9.

[0564] At 406, the visual characteristic of a given vertex are applied to all display objects associated with the given vertex.

[0565] At 408, visual information described in the description is displayed on a user interface in a manner consistent with the visual characteristics described in the mask.

[0566] FIG. 33 is a block diagram representation of an apparatus 500 for displaying data on a user interface. The module 502 is for receiving a description of a display tree comprising a plurality of vertices, wherein each vertex has one or more associated display objects. The module 504 is for receiving a mask defining a visual characteristic of each vertex in the plurality of vertices. The module 506 is for applying the visual characteristic of a given vertex to all display objects associated with the given vertex. The module 508 is for displaying, on the user interface, visual information described in the description in a manner consistent with the visual characteristics described in the mask.

[0567] The above-described use of the display tree structure and the ROST format may advantageously be used in situations where multiple users share or view the same data from a server. One example scenario is when user A in San Diego may want a user B in Phoenix to be able to jointly browse through a web site that sells toys so that user A can choose a toy as a gift to user B's daughter. Another example scenario may be when a stockbroker may want a client to be able to view data and graphs as the stock broker walks through the client's portfolio and provides explanation or choices for making trades. Yet another example scenario may include a joint collaborative research effort between two researchers that are geographically separated and looking at or charting certain data in a database, while a third researcher, in an unrelated activity, may be either modifying the data or copying the data for a research paper he is editing.

[0568] In some implementations, the above-described display tree and ROST formats may be useful in providing a shared view to data to multiple users. For example, in some implementations, each display tree, at its root (e.g., the starting point), may be identified by a unique id. This id may be sent by the server to the first user when the first user initiates a viewing session. The first user may then be able to communicate the unique id to the second user. The second user can then point his viewing application (e.g., web browser) to the server using the unique id. At this time, the server may become aware that two users are now sharing view of the same data (e.g., session) and may therefore provide views and view refreshes to the two users in a contemporaneous manner. The previously discussed display masks and vertexes can be advantageously used to reduce the amount of refresh data to be sent to each viewer and also provide a way in which views to multiple users are synchronously managed (i.e., the state of a vertex—e.g., a one byte information—may provide viewing status for hundreds of data items in a compact and consistent manner to all users).

[0569] FIG. 34 is a flowchart representation of a process 600 of providing visualization data to a first client device and a second client device.

[0570] At 602, a server is controlled to generate a plurality of display objects that can be displayed by the first and the second client devices, wherein each display object comprises at least one of a data portion and a graphics portion.

[0571] At 604, the server is controlled to organize the plurality of display objects into a plurality of groups, with display objects in a group sharing a common display characteristics. For example, one display characteristic may include whether or not the group is in active view of a user. Another display characteristic may include whether or not a third user is able to edit or view the data. Yet another display characteristic may include display aspects such as display pixel dimensions, color, etc.

[0572] At 606, the server is controlled to transmit a description of the plurality of display objects and the plurality of groups to the first and the second client devices. In some implementations the server is controlled to provide autonomous periodic refreshes. For example, a one second to 5 second periodic refresh rate may be satisfactory for some applications. In some implementations, the server may refresh based on a user activity (e.g., increase refresh rate when a user is actively interacting with the data).

[0573] At 608, the server is controlled to selectively update a first display object based on the common display characteristics of the first display object. For example, the update may include providing new or changed values of the display object (e.g., in the above-described airlines reservation scenario, when a seat is purchased by someone else).

[0574] At 610, the server is controlled to selectively transmit, based on the common display characteristics of the an updated description to the first client device and the second client device. In some implementations, each of the multiple users sharing the same data view may receive the updated description based on the display size of their devices. For example, a user on a large screen device may receive data updates more frequently (because there is a greater amount of data visible on her screen) than another user sharing the same view using a smaller screen sized device.

[0575] At 612, the server is controlled to transmit an updated description to the first client device and the second client device by including information about the updated first display object. The server thus may provide synchronized viewing updates to the multiple client devices.

[0576] FIG. 35 is a block diagram representation of an apparatus for providing visualization data to a first client device and a second client device. The module 702 is for generating a plurality of display objects that can be displayed by the first and the second client devices, wherein each display object comprises at least one of a data portion and a graphics portion. The module 704 is for organizing the plurality of display objects into a plurality of groups, with display objects in a group sharing a common display characteristics. The module 706 is for transmitting a description of the plurality of display objects and the plurality of groups to the first and the second client devices. The module 708 is for selectively updating a first display object based on the common display characteristics of the first display object. The module 710 is for selectively transmitting, based on the common display characteristics of the an updated description to the first client device and the second client device. The module 712 is for transmitting an updated description to the first client device and the second client device by including information about the updated first display object.

[0577] In some embodiments, a technique for facilitating viewing of data by a user operating a client device includes, opening a viewing session with the client device, wherein the viewing session is identified with a unique identification. The technique further includes transmitting an initial description of data for viewing at the client device, wherein the initial description comprises data organized as an information tree having a plurality of vertices, with each vertex comprising corresponding visual data that whose visual characteristics are simultaneously manipulatable by a single user action (e.g., simultaneously expand or collapse all objects from the view using GUI control widget). The technique further includes monitoring a visual state of the viewing session, the visual state indicating which vertices of the information tree are in active view and which vertices are not in active view at the client device. The technique further includes selectively updating a given visual data only when visual state indicates that the given visual data is in active view. The technique further includes communicating the updated given visual data to the client device as a difference from a previously sent description of data, wherein visual data includes at least one of an alphanumeric character, a graphical object or a graphical instruction.

[0578] In the above-discussed scenarios when a user is viewing and possibly modifying data from a database at a server, other users may also be simultaneously changing value of certain data. One example is a scenario such as GoogleDocs where multiple users can edit the same document. However, unlike GoogleDocs, the document or data may be being edited simultaneously. In these situations, several actions may be happening to a data field at the same time. For example, the server may have one value for a given data field. A first user may be in the process of editing the data field by editing a second value and may therefore be viewing the data field as it is being edited by the first user. Similarly, a second user may be in the process of editing the same data field with a completely different third value and may be seeing an entirely different data field during his editing process. Without rules to streamline this process, a

race condition may arise in which both users or one of the users may either overwrite each other or lose their edits without becoming aware of the failure.

[0579] To overcome such situations, and other problems, the present provides for techniques in which a behavior is defined for local caches of a user's edits to shared data. In some implementations, the local memory includes a record for all edits attempted by the user and the status of whether these edits were successfully accepted by the server or not. In some implementations, visual cues such as green/red colors may be used to indicate to the user whether his edits were successful or not. In some implementations, the user can access the locally stored editing history stack to check whether or not his edits were successful and, if not, why they were not successful. For unsuccessful edits (e.g., when another user beats the user to changing that data field), the user may be able to recall his edit and attempt to re-edit the field that has now been changed by another user.

[0580] In some embodiments, a method implemented at a client computer for visual editing of a data item stored on a server includes receiving a first value of the data item from the server (e.g., a value of the data item that is stored in the server's record), displaying a second value of the data item on a display (e.g., the value that is being edited by the user—e.g., this value may have started as the first value but the user may be mid-edit), receiving a third value input for the data item from a user (i.e., when the user has finished editing the value and indicates that he wishes to commit the change), checking with the server whether or not the server can commit to replacing the first value at the server with the third value, receiving a response from the server indicating success/failure in replacing the first value with the third value, and updating a local memory stack based on the response received from the server.

[0581] The local memory stack may be updated with an error message, e.g., when the response indicates that the replacing was unsuccessful. In some cases, while the user is editing the field (i.e., while the second value is displayed in the field), the user may receive a view refresh from the server and a third value may overwrite the user's mid-edit value. In some implementations, the user device may, upon this overwriting, provide a visual alert (e.g., flash the screen or flash red color in the data field) to the user indicating that the user's edit was lost (i.e., the user was not able to successfully replace the first value with his edited second value).

[0582] Many real life examples of data sharing include multiple users, having a mixed level of read/write authorization to a database, sharing the same database. Some example applications include traders sharing the same stock or money data, users in a business organization sharing the same files, data (e.g., payroll), medical practitioners sharing patient data, e-commerce sellers and buyers using the same inventory database, and so on. The client device **2904** used by each user may have different display capabilities (e.g., screen size, graphics capabilities, etc.). It may be desirable to provide a visual access to the shared data using techniques that allow adaptation of database access based on client device capabilities. For example, in some cases, new windows or data viewing panes may be opened when a user desires to zoom in or zoom out on data, or manipulate data to draw graphs etc. This processed data, called detachment data, could be displayed to the user on a full screen window, or in a second window (e.g., a pop-up window) and so on.

No techniques presently provide for a streamlined control of shared view to the same data by multiple users operating client devices with different capabilities.

[0583] In some embodiments, a technique for providing access to a database includes receiving a display profile from a client device, receiving a data view request from a user of the client device, starting a data viewing session based on the received data view request, associating the data viewing session with the user, specifying to the client device a first display property of data being viewed based on the display profile of the client device, receiving another data view request for the same session from the user using another client device, and specifying to the another client device a second display property of data being viewed based on a second display profile of the another client device.

[0584] One of skill in the art will also appreciate that several techniques for viewing of data are provided. In some examples, a user using a client device receives views of a database from a server in a compact format organized as a display tree comprising vertices. At least some visual characteristics of all data belonging to a given vertex are controlled by simply specifying the visual characteristic at the vertex level. In one advantageous aspect, the vertex level specification is compact, because it makes it unnecessary to individually list visual characteristics for possibly hundreds of data objects under the vertex.

[0585] One of skill in the art will further appreciate that techniques are provided for multiple users to share views of the same data. A server provides synchronous updates to the data to all the users sharing a same session, e.g., the same id of a display tree.

[0586] It will further be appreciated that techniques are provided for local storing of a user's data edits so that no data is lost when a user is unsuccessful to edit a shared data field on the server.

[0587] It will further be appreciated that techniques are disclosed for a server to provide data views to client devices having different hardware/software capabilities.

[0588] The disclosed and other embodiments, modules and the functional operations described in this document can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this document and their structural equivalents, or in combinations of one or more of them. The disclosed and other embodiments can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a computer readable medium for execution by, or to control the operation of, data processing apparatus. The computer readable medium can be a machine-readable storage device, a machine-readable storage substrate, a memory device, a composition of matter effecting a machine-readable propagated signal, or a combination of one or more of them. The term "data processing apparatus" encompasses all apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them. A propagated signal is an artificially generated signal, e.g., a machine-generated electrical, opti-

cal, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus.

[0589] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0590] The processes and logic flows described in this document can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0591] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for performing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0592] While this patent document contains many specifics, these should not be construed as limitations on the scope of an invention that is claimed or of what may be claimed, but rather as descriptions of features specific to particular embodiments. Certain features that are described in this document in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the

claimed combination may be directed to a sub-combination or a variation of a sub-combination. Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results.

[0593] Only a few examples and implementations are disclosed. Variations, modifications, and enhancements to the described examples and implementations and other implementations can be made based on what is disclosed.

1. A method of providing visualization data to a client device, comprising:

receiving, by a client application on the client device, a user input to change a visual characteristic of a display object tree;

transmitting, to a server comprising a server cache and in response to the user input, an event message comprising an indication of the user input;

updating, using a client cache and based on the user input, the display object tree on the client device;

receiving, from the server, a description of a difference tree, wherein the server is configured to generate the difference tree based on the user input and a display tree stored in the server cache;

patching, based on the description of the difference tree, the display tree in the client cache; and

selectively updating, subsequent to the patching and based on the display tree, the display object tree on the client device.

2. The method of claim 1, wherein updating the display object tree using the client cache incurs no server delay.

3. The method of claim 1, wherein selectively updating the display object tree on the client device comprises refraining from updating the display object tree upon a determination that the user input is a valid user input.

4. The method of claim 1, wherein a plurality of display objects is assigned to each of a plurality of vertices of the display object tree, wherein the user input comprises a mask, wherein the mask specifies visual characteristics of the plurality of vertices, and wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex.

5. The method of claim 4, wherein the mask comprises (i) a screen mask corresponding to the visual characteristics of the display object tree displayed on the client application, and (ii) a reference mask corresponding to the visual characteristics of the display tree stored in a server cache.

6. The method of claim 5, further comprising:

updating, based on the user input, the screen mask; and refraining, prior to receiving the description of the difference tree, from updating the reference mask.

7. The method of claim 5, further comprising:

updating, based on the user input prior to receiving the description of the difference tree, the screen mask; and updating, based on the description of the difference tree, the reference mask.

8. The method of claim 5, wherein the screen mask and the reference mask are Boolean functions.

9. The method of claim 1, wherein the visual characteristics include information about whether or not each of the plurality of vertices are visually occluded.

10. The method of claim **1**, wherein the description of the difference tree is transmitted using a reference/optimistic/screen/typed value (ROST) data transfer format.

11. An apparatus for providing visualization data to a client device, comprising:

a client application; and

a client cache coupled to the client application,

wherein the client application is configured to:

receive a user input to change a visual characteristic of a display object tree,

transmit, to a server comprising a server cache and in response to the user input, an event message comprising an indication of the user input,

update, using the client cache and based on the user input, the display object tree on the client device,

receive, from the server, a description of a difference tree, wherein the server is configured to generate the difference tree based on the user input and a display tree stored in the server cache,

patch, based on the description of the difference tree, the display tree in the client cache, and

selectively update, subsequent to the patching and based on the display tree, the display object tree on the client device.

12. The apparatus of claim **11**, wherein updating the display object tree using the client cache incurs no server delay.

13. The apparatus of claim **11**, wherein selectively updating the display object tree on the client device comprises refraining from updating the display object tree upon a determination that the user input is a valid user input.

14. The apparatus of claim **11**, wherein a plurality of display objects is assigned to each of a plurality of vertices of the display object tree, wherein the user input comprises a mask, wherein the mask specifies visual characteristics of the plurality of vertices, and wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex.

15. The apparatus of claim **14**, wherein the mask comprises (i) a screen mask corresponding to the visual characteristics of the display object tree displayed on the client

application, and (ii) a reference mask corresponding to the visual characteristics of the display tree stored in a server cache.

16. The apparatus of claim **15**, wherein the screen mask and the reference mask are Boolean functions.

17. A computer program product comprising a computer readable storage medium having code stored thereon, the code, when executed by a processor, causing the processor to implement a method of displaying data to a user, the data comprising a plurality of data items, the method comprising:

receiving, by a client application, a user input to change a visual characteristic of a display object tree;

transmitting, to a server comprising a server cache and in response to the user input, an event message comprising an indication of the user input;

updating, using a client cache and based on the user input, the display object tree on the client device;

receiving, from the server, a description of a difference tree, wherein the server is configured to generate the difference tree based on the user input and a display tree stored in the server cache;

patching, based on the description of the difference tree, the display tree in the client cache; and

selectively updating, subsequent to the patching and based on the display tree, the display object tree on the client device.

18. The computer program product of claim **17**, wherein updating the display object tree using the client cache incurs no server delay.

19. The computer program product of claim **17**, wherein selectively updating the display object tree on the client device comprises refraining from updating the display object tree upon a determination that the user input is a valid user input.

20. The computer program product of claim **17**, wherein a plurality of display objects is assigned to each of a plurality of vertices of the display object tree, wherein the user input comprises a mask, wherein the mask specifies visual characteristics of the plurality of vertices, and wherein the visual characteristics of a given vertex simultaneously controls display attributes of all display objects assigned to the given vertex.

* * * * *