

This document includes some recent decisions of the EPO in 2014 with regards to software related inventions and shows relevant extracts from the respective decisions.

T 0426/09 (Alert notification/SAP) of 9.9.2014

Alert Notification Management

**Inventive step - (main request Inventive step - no)
Remittal for additional search (auxiliary request)**

Application number: 04101335.0
IPC class: G06F 17/60
Applicant name: SAP SE

Board: 3.5.01

<http://www.epo.org/law-practice/case-law-appeals/pdf/t090426eu1.pdf>

European patent application number 04 101 335.0 (publication number 1 475 734) relates to a method of informing parties of alert situations in business or industrial applications, alert situations like the termination of a contract, a delivery problem for an order, a decline in revenue, or some technical problems in a system.

T 2348/09 (Anreichern und Verdichten/SAP) of 4.12.2014

Bestimmen einer Kennfunktion aus einer Matrix nach vorbestimmtem Schema

Erfinderische Tätigkeit - (nein) (Naheliegende Computerimplementierung eines Verfahrens - Technische Aufgabe nicht glaubhaft gelöst)

Anmeldenummer: 01130883.0
IPC-Klasse: G06F 17/60, G06F 17/16
Name des Anmelders: SAP SE

Kammer: 3.5.01

Relevante Rechtsnormen:

Agreement on Trade-Related Aspects of Intellectual Property Rights Art 27(1)
Angeführte Entscheidungen: T 0641/00

<http://www.epo.org/law-practice/case-law-appeals/recent/t092348du1.html>

Gegenstand des vorliegenden Anspruchs 1 ist ein "(c)omputer-implementiertes Verfahren () zum Bestimmen einer Kennfunktion aus einer Matrix mit einer vorgegebenen Zahl von Datensätzen mit klassifizierbaren Merkmalen".

Anspruch 1 in der als Hauptantrag bezeichneten Fassung lautete wie folgt:

"Computer-implementiertes Verfahren (402) zur Datenverarbeitung mit programmgesteuertem Prozessor (100/910) zum Bestimmen einer Kennfunktion (Y) aus einer Matrix (X0) mit einer vorgegebenen Zahl (ZO) von Datensätzen (DS) mit klassifizierbaren Merkmalen (M) durch Bearbeiten (430) der Matrix und Berechnen (440) der Kennfunktion (Y) unter Berücksichtigung der Mehrzahl der Datensätze der bearbeiteten Matrix (XH), wobei die zu verarbeitenden Daten Matrizen mit Zeilen und Spalten sind, wobei eine Zeile einem Datensatz (DS) entspricht, und wobei das Verfahren (402) die folgenden Schritte umfasst

Bereitstellen eines Prozessor-Steuerprogramms (100), von Regeln (300) und Matrizen (X) auf einem Computer (900),

Bearbeiten (430) durch Anreichern (431) der Matrix mit Klassen (K) für die Merkmale (M) unter Beibehaltung der Zahl der Datensätze (DS), und durch

Verdichten (432) entsprechend den Klassen (K) unter Verringerung der Zahl der Datensätze (DS), wobei Bearbeiten (430) und Berechnen (440) an Matrizen erfolgt, die im Hauptspeicher implementiert sind und durch

Interpretieren (450) eines vorbestimmten Schemas (200) durch das Prozessor-Steuerprogramm (100), in dem die Reihenfolge der einzelnen Ausführungen von Anreichern (431) und Verdichten (432) sowie die Regeln für die jeweils einzelnen Ausführungen festgelegt sind,

wobei der Verfahrensschritt Berechnen (440) der Kennfunktion (Y) mit der letztmalig durch Verdichten (432) bearbeiteten Matrix (XH) ausgeführt wird."

Vom nächsten Stand der Technik unterscheidet sich das beanspruchte Verfahren in folgenden Merkmalen (Nummerierung hinzugefügt):

- (1) Der Berechnung der Kennfunktion liegt eine Matrix mit einer vorgegebenen Zahl von Datensätzen mit klassifizierbaren Merkmalen zu Grunde.
- (2) Der Berechnung liegt ferner ein vorbestimmtes Schema zu Grunde, dass von dem Verarbeitungsprogramm interpretiert wird und in dem die Reihenfolge der einzelnen Ausführungsschritte wie Bearbeiten von Matrizen und Berechnen der Kennfunktion sowie die Regeln für die jeweils einzelnen Ausführungsschritte festgelegt sind.
- (3) Das Bearbeiten erfolgt durch Anreichern einer Matrix mit Klassen für die Merkmale unter Beibehaltung der Zahl der Datensätze und durch Verdichten entsprechend den Klassen unter Verringerung der Zahl der Datensätze.
- (4) Das Berechnen der Kennfunktion wird mit der letztmalig durch Verdichten bearbeiteten Matrix ausgeführt.

Diese **Merkmale haben keinen Bezug zu einer spezifischen Anwendung**. Die einzelnen Ausführungsschritte sind bildhaft als Anreichern und Verdichten skizziert und **stehen, über die einfache Computerimplementierung hinaus, in keinem Zusammenhang mit einem technischen Prozess**. Zieht man die Beschreibung der Erfindung zusätzlich in Betracht, so lassen sich diese **Merkmale bestenfalls als die Implementierung eines betriebswirtschaftlichen Berechnungsverfahrens** verstehen.

Die **Implementierung eines abstrakten oder bestenfalls betriebswirtschaftlichen Berechnungsverfahrens mittels bestimmungsgemäßer Verwendung der Komponenten einer herkömmlichen EDV-Anlage begründet zwar die Technizität des computerimplementierten Verfahrens, leistet aber keinen erfinderischen Beitrag zum Stand der Technik**.

Die Beschwerdeführerin hat vorgetragen, die Erfindung dürfe nicht auf die einfache Computerimplementierung eines x-beliebigen Verfahrens reduziert werden. Die Erfindung löse mit technischen Mitteln, nämlich einem iterativen Prozess aus Anreichern und Verdichten von Matrizen und der abschließenden Berechnung der Kennfunktion aus der zuletzt verdichteten Matrix, ein ganz wesentliches technisches Problem der begrenzten Prozess- und Speicherressourcen von Computersystemen. Die Verarbeitung sehr großer Matrizen mit Millionen von Datensätzen stoße an die Grenzen der Leistungsfähigkeit auch von modernen Computersystemen. Die Erfindung erlaube eine entscheidende Reduzierung von Rechenzeiten und Systemressourcen, denn es sei mittels der Erfindung nicht mehr nötig, eine große Matrix mit Millionen von Datensätzen in komplizierter Weise zu bearbeiten. Die Kennfunktion werde nun in effizienter Weise aus einer in wenigen Iterationsschritten berechneten kleinen Matrix bestimmt.

Dieser **behauptete Erfolg der beanspruchten Erfindung hält jedoch einer Überprüfung nicht stand**. Wie aus der Beschreibung der Ausführungsbeispiele ohne weiteres erkennbar ist, steht das **Schema von Anreichern und Verdichten in unmittelbarem Zusammenhang mit der betriebswirtschaftlichen Aufgabenstellung**. Beispielsweise werden die Datensätze mit Postleitzahlen "angereichert" und die Datensätze für solche Postleitzahl-Regionen zusammengefasst um wirtschaftliche Kennzahlen für solche Regionen zu berechnen (siehe beispielsweise Figur 12, Spalte S und Figur 14).

Sinn und Zweck der Verfahrensschritte Anreichern und Verdichten ist ausschließlich die betriebswirtschaftliche Rechnung; eine positive Auswirkung auf Berechnungszeiten und Speicherbedarf ist nicht erkennbar. Die betriebswirtschaftliche Aufgabenstellung legt fest, ob das beanspruchte Berechnungsverfahren aus Anreichern und Verdichten überhaupt sinnvoll anwendbar ist oder im algorithmischen Nirgendwo endet. Die iterative Erzeugung von Zwischenmatrizen reduziert auch per se nicht den Speicherbedarf oder die Anzahl der erforderlichen Rechenschritte, wenn man die gesamte Iterationskette mit der Erzeugung immer neuer Matrizen in Betracht zieht. Wenn daher schon bei den in der Anmeldung angegebenen Beispielen Zweifel berechtigt sind, ob das anvisierte Ziel einer Reduzierung des Zeit- und Ressourcenbedarfs überhaupt erreicht wird, so gelten diese Zweifel umso mehr für den sehr breiten Schutzbereich, der effektiv ohne Einschränkungen oder Vorgaben für die Berechnung von Kernfunktionen beansprucht ist und auch die Natur der Daten und der Kennfunktion völlig offen lässt. Die im Anspruch geforderte Klassifizierbarkeit von Merkmalen ist keine Einschränkung, da selbst das Nichts klassifizierbar wäre. **Spätestens also in der Breite des Anspruchs kann das von der Beschwerdeführerin behauptete**

technische Problem nicht als gelöst gelten und muss daher bei der Prüfung der erfinderischen Tätigkeit unberücksichtigt bleiben.

Die Beschwerdeführerin hat ferner mit Verweis auf Artikel 27 (1) TRIPS vorgetragen, der Technikbegriff werde von der Kammer zu eng ausgelegt und der Beitrag nichttechnischer Merkmale bei der Prüfung auf erfinderische Tätigkeit nicht ausreichend berücksichtigt.

Diese mit dem Vortrag der Beschwerdeführerin aufgeworfenen grundsätzlichen Fragen sind jedoch für das vorliegende Urteil über die erfinderische Tätigkeit unerheblich, da sich dieses Urteil nicht auf eine fehlende Technizität des beanspruchten Verfahrens stützt, sondern darauf, dass das von der Beschwerdeführerin formulierte Ziel tatsächlich nicht erreicht wird. Aus diesem Grund erübrigt sich eine weitergehende Diskussion des Begriffs der Technizität.

T 2185/10 (Image file management/CANON) of 21.10.2014

File management system of image data

Inventive step - after amendment (yes)

Application number: 97117880.1
IPC class: G06F 17/30, H04N 1/21
Applicant name: Canon Kabushiki Kaisha

Board: 3.5.07

Cited decisions: T 0962/98, T 1644/11

<http://www.epo.org/law-practice/case-law-appeals/pdf/t102185eu1.pdf>

The invention

The application relates to a method for managing the image files in a memory card of an image pickup apparatus, for example an electronic camera, so that each file name is unique for the lifetime of the camera, even when the power is switched off or when the memory card is changed (original description, page 1, lines 5 to 12; page 10, lines 15 to 26).

Unique file names are automatically created for the generated image files. Each file name includes the current numerical value of a counter stored in a non-volatile memory of the camera (page 1, lines 5 to 12; page 10, lines 15 to 26). Each time a new image file is generated, it is stored with a name including the current counter value and the counter is increased by one (page 11, line 2, to page 12, line 8).

The files are stored in different directories, each directory name also including a numerical value (page 19, line 21, to page 20, line 12; figure 5). The numerical value for the current directory name is the result of the division of the highest used file number by a predetermined

number of files per directory, rounded down to a whole number (page 19, line 24, to page 20, line 2; figure 8).

When the memory card of the camera is changed or the power source is turned on, a set-up process is performed for initialisation. This process calculates the current file number and current directory number to be used to generate and store the next file. The current file and directory numbers are calculated based on the non-volatile counter and taking into account the contents of the recording medium (page 14, line 12, to page 19, line 12, figures 4A and 4B).

Each file name includes a limited number of digits (e.g. five). When the maximum value is achieved (e.g. 99999), no new filenames can be generated. The invention foresees exception processes to deal with that problem (page 20, line 13, to page 21, line 14).

Claim 1 of the sole request reads as follows:

"A file management method for managing files of image data photographed with an image pickup apparatus (7), the files being classified into directories every predetermined number of files in accordance with numerical values included in their file names, the method comprising

- a) a counting step of updating a current numerical value every image data recording,
- b) a generating step of generating the file name including as a part thereof said current numerical value when data of one of the photographed images is recorded as a file on a recording medium (33) and storing the data in a file with the generated file name, and
- c) a storing step of storing the current numerical value updated in said counting step, into a nonvolatile memory (37) of said image pickup apparatus (7),
- d) adding 1 to the current numerical value in said counting step every image data recording, and wherein the method further comprises,

when the power source of the image pickup apparatus is turned on or the recording medium is changed:

- e) obtaining (S6) a maximum directory number of directories stored on the recording medium for storing a recorded image;
- f) comparing (S12) a current directory number based on the current numerical value with the maximum directory number; and

g) if the current directory number is not larger than the maximum directory number:

obtaining (S13) the maximum value of the numerical values included in the file names already stored in the directory having the maximum directory number,

comparing (S14) the obtained maximum value with the current numerical value, and

changing (S15) the current numerical value to be the obtained maximum value plus 1 in case the current numerical value is not larger than the obtained maximum value of the file names."

Inventive step

Document D1 relates to storing digitised image signals as image files in devices such as electronic still video cameras (column 1, lines 25 to 34; column 3, lines 44 to 49). An embodiment described in columns 3 and 4 is directed to generating unique file names in order to prevent the processing of the wrong image file in the camera or external devices (column 3, lines 23 to 31; column 4, lines 5 to 17). According to this embodiment of document D1, the solution is to keep track of the number of recorded pictures in a counter and add it "as a discrimination number" to the image file of an obtained picture when it is stored (column 3, lines 23 to 43). Additionally, in order to avoid a reset of the counter when a battery of the device is removed (column 3, lines 50 to 56), the counter is stored in a non-volatile memory (column 4, lines 18 to 22).

Document D1 also discloses using directories in some embodiments, e.g. in an embodiment described as the combination of a "main operating routine" including steps for storing image data (figures 2 and 3, column 15, line 53, to column 18), and the routines of figures 4 and 8. However, these features are not described for the embodiment of columns 3 and 4.

The Board therefore agrees with the Examining Division that document D1 discloses the combination of features a) to d).

The claimed subject-matter differs from the embodiment of columns 3 and 4 of document D1 in that the files are "classified into directories every predetermined number of files in accordance with numerical values included in their file names" (see first lines of the claim) and in the performance of steps e) to g) when the power source is turned on or the recording medium is changed.

Steps e) to g) are performed for finding out the maximum directory number of existing directories and, if necessary, recalculating the current file number.

The **distinguishing features solve the technical problem of avoiding overwriting image files in a recording medium previously used in another camera.**

The use of directories in the claimed manner has the **further advantage of decreasing the average file access time by the file system, since it is faster to search a file in a directory with fewer files** (see also page 26, lines 6 to 9 of the original description).

In the opinion of the Board, the **skilled person would not arrive at the claimed solution without inventive skill**. None of the available prior-art documents discloses the distinguishing features. The Board does **not** consider the solution **obvious in the light of document D1** alone either.

Considering the above technical problem, the **skilled person could choose among several possibilities to solve it**. Examples of the most immediate solutions would be checking whether the file name of the file to be created already existed and, if so, incrementing the counter, or taking the lowest non-used number for the next file name. Compared to these

solutions, the **claimed invention has the advantage of fast storage of obtained pictures, since no checks have to be performed each time data is stored.**

Regarding the advantage of efficient search, other embodiments of document D1 also disclose the use of directories and it is obvious for the skilled person that file search is faster in directories having fewer files. However, the skilled person would still have to conceive a way of combining the distribution of files among directories with the file naming scheme of document D1.

Furthermore, the Board is of the opinion that **the advantage of efficient file access cannot be seen as a second independent problem in the context of the present application. The distinguishing features are functionally interdependent.** The directory number is a function of the numerical value of the file names and this relationship influences the other features, in particular the way the current numerical value for the next file is updated.

The subject-matter of independent claim 1 therefore involves an inventive step.

T 1329/10 (Managing multiple keys/MICROSOFT) of 8.12.2014 **Systems and methods for managing multiple keys for file encryption and decryption**

Inventive step - (yes)

Application number: 06100077.4
IPC class: G06F 21/00, H04L 9/08
Applicant name: Microsoft Corporation

Board: 3.5.06

<http://www.epo.org/law-practice/case-law-appeals/pdf/t101329eu1.pdf>

The context of the invention

The application relates to managing data files, for instance in an encrypted file system (EFS). The data files may be stored on a single device (see figure 1) or on different devices linked by a communications network - see figure 2 (not claimed). Each file has been encrypted using a key provided, for instance, on a smartcard. Smartcards may be replaced over time, meaning that the key they contain changes. The description also mentions the user changing the current key; see paragraph [0049], second sentence.

The issue arises of how to provide long-term access to files as the current encryption key changes, since each file can only be decrypted using the key with which it was originally encrypted; see figure 5. Long-term file access is achieved by storing a list of the respective previous cryptographic keys, the list itself being encrypted using the current cryptographic key.

Claim 1 reads as follows:

"A computer-implemented method for managing multiple keys for file encryption and decryption, the method comprising: decrypting (1000) a list (610, 700) of previously used keys using a prior current key (59B, 701) with which the list of previously used keys has been previously encrypted; adding (1001) the prior current key to the list of previously used keys; and

encrypting (1002) the list of previously used keys containing the added prior current key using a new current key."

Claim 1 of the main request is directed to the method, shown in figure 10, of adding a new current key to the key list. This involves decrypting the list using the prior current key, adding the new current key to the list and then encrypting the list using the new current key; see paragraphs [0069] to [0072].

The prior art

Document D1

D1 relates to storing data in a database managed inter alia by a security administrator. If a user or the security administrator designates a database column as encrypted (see figure 5 and page 8, line 12, to page 9, line 5) then data is automatically encrypted - in a manner which is transparent to the user - before being stored in that column; see figure 6. If the designated column is already encrypted, then it is decrypted using the "previous key" and re-encrypted using a "new key"; see page 3, lines 17 to 21, and figure 5, steps 510 to 518.

The security administrator stores the encryption keys in a keyfile and may also select the mode of encryption, create the keyfile (see figure 3) and establish how many keys are to be stored in the keyfile; see page 3, lines 9 to 11. The keyfile (figure 1; 120) can be stored as an encrypted file in the database system or at a location separate from it; see page 3, lines 13 to 14. The security administrator also moves an obfuscated (i.e. difficult to read) copy (figure 1; 116) of the keyfile to a volatile memory within a server associated with the database system; see page 3, lines 15 to 16. The security administrator can also cause a specified database column to be decrypted and then re-encrypted using a new key; see page 3, lines 17 to 21.

The encryption function in the database server uses keys from the obfuscated keyfile to encrypt data received from the client for storage by a storing function in a row of the database; see figure 6 and page 9, lines 7 to 18. Correspondingly a retrieving function in the database server retrieves data from a row in the database which, if the request is from an authorized user, is then decrypted using keys from the obfuscated keyfile and passed to the client; see page 6, lines 11 to 15 and 20 to 21, page 9, line 20, to page 10, line 9, and figure 7. Metadata stored in the database (see figure 2; metadata 222) records which columns of the database are encrypted and, if so, the key identifier for the key in the obfuscated keyfile that is used to encrypt data in that column and the encryption mode, such as DES; see page 6, lines 22 to 27.

According to the reasons for the appealed decision, the keyfile (120) stored as an encrypted file in the database system can be regarded as the claimed list of previously used keys, the

subject-matter of claim 1 of the main request only differing from the disclosure of D1 in the step in which the key, with which the list of keys had been previously encrypted, called the "prior current key", was also stored in the list of keys.

3.1.5 The appellant has disputed whether D1 discloses a prior current key within the meaning of the claims and pointed out that, although D1 mentions the keyfile being stored as an encrypted file, no details are given of the key used to encrypt the keyfile.

The board finds that D1 discloses the following features set out in claim 1 of the main request: a computer-implemented method for managing multiple keys for encryption and decryption, the method comprising: decrypting a list of previously used keys using a prior current key with which the list of previously used keys has been previously encrypted and adding a key to the list.

Hence the subject-matter of claim 1 differs from the disclosure of D1 in that:

a. said multiple keys are used to encrypt and decrypt files;

b. said added key is said prior current key and

c. encrypting the list of previously used keys containing the added prior current key using a new current key.

The board notes that difference "b" corresponds to the difference feature identified in the reasons for the decision. In view of the additional difference features "a" and "c", the disclosure of D1 is less relevant than stated in the decision for the assessment of inventive step (see below).

Document D4

According to point 6 in section "IV Further Remarks" of the decision, the claimed subject-matter is known from D4, in particular claim 1 in conjunction with paragraphs 14, 18, 25 and 26. The board does not accept this assessment.

The cited passages of D4 relate to the encryption and decryption of a key to provide a backup in case the user forgets it or is unavailable. The encryption/decryption occurs in two stages. The key is first encrypted using a key derived from the hash of private information, for example the mother's maiden name, and a symmetric encryption algorithm, such as DES. The result is then encrypted using the public key of a trusted party, for example a certificate authority, and an asymmetric encryption algorithm. The original key can be recovered from the resulting "key recovery file" by two corresponding decryption steps, i.e. asymmetric decryption using the private key of the trusted party, followed by symmetric decryption using the key derived from private information.

Since D4 does not disclose an encrypted list of previously used keys, set out in claim 1 according to the main request, it follows that, contrary to the statement in the decision, the subject-matter of that claim is new, Article 54(1,2) EPC 1973, in view of the disclosure of D4.

Document D3

According to point 5 in section "IV Further Remarks" of the decision, the claimed subject-matter is not inventive in view of the background art acknowledged in D3. The board is not convinced by this reasoning (see below).

The invention in D3 relates to preventing a previously authorized user from accessing an encrypted file in a shared file system, also termed user revocation, without having to re-key (i.e. decrypt and then re-encrypt using a new key) the file. The problem is solved by giving an authorized user with the old password a "private share" with which he can generate a new cryptographic key based on an old password, the "private share" and a "rotation catalyst" published on a shared bulletin board.

The "background art" section of D3 mentions in paragraph [0005] "lazy revocation" where files are only re-keyed (making them inaccessible to revoked users) when they are updated. Keys are stored in an encrypted file called a "lockbox". In the event of user revocation all of the lockboxes accessed by the revoked user are marked as dirty and any subsequent update to a dirty file causes the file to be re-keyed. To prevent revoked users from accessing unchanged files, paragraph [0006], right-column, lines 3 to 9, teaches re-encrypting all the lockboxes, this necessitating additional key storage for the new lockbox keys.

Hence D3 discloses an encrypted list of keys. There is however no suggestion that the old lockbox encryption key is added to the contents of the lockbox.

Inventive step, Article 56 EPC 1973

Starting from D1

As set out above, the subject-matter of claim 1 differs from the disclosure of D1 in that:

a. said multiple keys are used to encrypt and decrypt files;

b. said added key is said prior current key and

c. encrypting the list of previously used keys containing the added prior current key using a new current key.

Regarding **difference "a"**, the board can see **no obvious problem or solution which would lead the skilled person starting from D1 to use the keys used to encrypt/decrypt specified columns of the database to also encrypt/decrypt files.** The board regards the **encryption of files as a technical activity which can contribute to inventive step.** In view of difference feature "a", the subject-matter of claim 1 involves an inventive step in view of D1.

Difference features "b" and "c" solve the problem of ensuring that users in possession of the new current key can access files encrypted using previous current keys. For the purposes of this decision there is no need to consider whether these features contribute to inventive step.

Inventive step starting from D3

In the light of the above analysis, the background art acknowledged in D3 discloses the following features of claim 1 of the main request: a computer-implemented method for managing multiple keys for file encryption and decryption, the method comprising decrypting a list (lockbox) of previously used keys using a prior current key with which the list of previously used keys has been previously encrypted and subsequently encrypting the list of previously used keys using a new current key.

Hence the subject-matter of claim 1 differs from the disclosure of D3 in that, before the list is re-encrypted, the prior current key is added to the list of previously used keys.

The board can see **no obvious problem or solution which would cause the skilled person starting from D3, before the list is re-encrypted, to add the prior current key to the list of previously used keys**, in particular because neither the lockbox nor any of the encrypted files is still encrypted with the prior current key. **This difference feature solves the technical problem of allowing file decryption and thus can contribute to inventive step.**

Hence the subject-matter of claim 1 of the main request involves an inventive step.

T 1338/10 (Server affinity/ORACLE) of 13.11.2014 **SYSTEM AND METHOD FOR SERVER LOAD BALANCING AND SERVER AFFINITY**

Inventive step - (yes)

Application number: 04714154.4
IPC class: G06F 9/00, G06F 9/46, G06F 15/16
Applicant name: Oracle International Corporation

Board: 3.5.06

<http://www.epo.org/law-practice/case-law-appeals/pdf/t101338eu1.pdf>

Overview of the invention

The application relates to a method of distributing method calls from external client computers via client-side stubs to server instances (i.e. server programs running on a server computer). The server instances may be RMI (remote method invocation) objects like JMS (Java message service) or EJB (enterprise Java bean) interfaces (original description paragraph [29], second sentence). For method calls from clients to services not configured for server affinity, a conventional load balancing algorithm (LBA) like round-robin is used for both internal and external connections ([22], paragraph 2). For method calls from external clients to services configured for server affinity, no LBA is used, but server affinity which means that the client-side stub attempts to choose a server instance to which it is already connected ([21], second and third sentence; [22], paragraph 5). Furthermore, if a server instance becomes unavailable, then the client-side stub fails over to another already connected

server instance ([21], last sentence; [32]). And an LBA which an administrator configures for a service overrides the default LBA for the cluster ([23], fourth sentence).

Claim 1 of the main request reads as follows:

"1. A system for server load balancing that includes server affinity, comprising: a cluster that includes a plurality of server instances providing services, and wherein each of the services provides a plurality of method calls;

a load balancing and affinity processor that assigns server instances from said cluster to service client requests in the form of method calls from external clients;

a client-side stub on an external client obtained for a service, wherein the client-side stub attempts to choose a server instance to which the external client is already connected, and the client-side stub continues to use the same server instance and the same connection for method calls of that service;

wherein if the server instance becomes unavailable, the stub fails over to a server instance to which the client is already connected and which provides said service; and

wherein the cluster is adapted to use a load balancing algorithm that includes server affinity to govern connections between external clients and server instances and wherein the load balancing algorithm is overridden by a user-configured load balancing algorithm for the service maintained in the client-side stub."

Independent method claim 10 of the main request reads as follows (additions or modifications with respect to claim 10 of the refused main request are marked in italics; deletions are [deleted: struck through]):

"10. A method for server load balancing that includes server affinity, comprising the steps of:

providing a plurality of server instances as a cluster providing services, and wherein each of the services provides a plurality of method calls;

assigning server[deleted: s] instances from said cluster to service client requests in the form of method calls from external clients;

wherein said step of assigning includes using a client-side stub on an external client obtained for a service, wherein the client-side stub attempts to choose a server instance to which the external client is already connected, and the client-side stub continues to use the same server instance and same connection for method calls for that service;

wherein if the server instance becomes unavailable, the stub[deleted: s] fails over [deleted: if possible] to a server instance to which the external client is already connected and which provides said service; and

wherein the cluster uses a load balancing algorithm that includes server affinity to govern connections between external clients and server instances and wherein the load balancing

algorithm is overridden by a user-configured load balancing algorithm for the service maintained in the client-side stub."

Inventiveness of claim 10 of the main request

The board considers D1 to be the closest prior art.

During oral proceedings the appellant explained that the application and D1 differentiate between external and internal clients: External (or "normal" or "classical") clients are computers at the user's side, i.e. computers which are no servers (e.g. the client in figures 1, 2 and in paragraphs [31], [32] of the application; or client computer 250 in figure 4 of D1), whereas internal clients are themselves servers (e.g. managed server MS4 in figure 3 and in paragraph [33] of the application; or web server client 250 in figure 4 of D1). The claim only related to external clients as shown in figure 2 and not to internal clients as in figure 3. The appellant then filed amended claims for the main request which explicitly specify "external clients" as requesting services and having client-side stubs performing the assignments of server instances and the failover handling.

The board agrees that this differentiation is important. Therefore, one cannot combine the embodiments of an internal web server plug-in performing the load balancing (D1, page 12, line 1 to page 13, line 12) with the passage in D1 about "sticky" load balancing by an external client (page 15, line 27 to page 16, line 26). Since the latter passage in D1 (about sticky load balancing) relates to server-side load balancing (lines 35-38: "depending on the outcome of the load balancing decisions described above"; the preceding section is about server-side load balancing), it does not disclose a (client-side) stub on an external client performing load balancing.

Furthermore, there is no disclosure about a stub obtained for a service in the above mentioned passage of D1. There might be something similar to a stub on page 16 (lines 14-15: "... the client computer(s) may instead be operable to maintain information regarding sticky requests so that requests are sent directly to the correct application server."), but not for a specific service and not doing load balancing when no connection to a server instance exists.

Accordingly, D1 neither discloses a stub on an external client performing load balancing, nor a stub obtained for a service.

Current claim 10 differs from D1 in that for the initial request of an external client for a service and during failover, the client-side stub chooses, if possible, a server instance to which it is already connected, whereas D1 chooses a server determined by a (server-side) load balancing for the initial request (page 15, lines 35-38). For failover no procedure how to choose a server instance is disclosed in D1

As to the similarities between the claim and D1, in both of them further requests from the same client for the same service stick to the chosen server instance.

The objective technical problem as formulated in the decision (2.1.4), i.e. how to minimise the number of connections, does not seem to be appropriate for the current claim, since it implies already the solution in it and it does not cover all the differences.

The board formulates the technical problem as how to reduce the time to serve an initial request from a client for a service and to react in a failure situation.

The solution is to reuse an existing connection of a client to a server instance for any initial request for a service from a client, or for requests when the previously used server instance failed. To achieve this **in D1, the method disclosed** on page 15, line 27 to page 16, line 26 **would have to be completely rebuilt: replace the server-side load balancing by a client-side "connection-first" strategy (called "server affinity" in the application) for the initial request; do this by a stub obtained for that service; for failover also select this connection-first strategy.**

The concept of a component similar to a stub disclosed in another embodiment in D1 (namely the web server plug-in 242 in figure 4) **is not directly applicable for solving the problem,** since this is situated in a very different context of a two-level architecture (figure 2A) with a web server 104 placed between the (external) client computer 100 and the application servers 108A and 108B. The claim explicitly specifies a one-level architecture with a direct connection between the external client and the server instances (i.e. the application servers). Furthermore, also this web server plug-in does not disclose a connection-first strategy, but mere load balancing. The only hint to a connection-first strategy in D1 would be the continued usage of the server instance previously selected by the server-side load balancing in the sticky load balancing example. However, as argued by the appellant during oral proceedings, the reason for that continued usage seems to be to avoid the migration of the data structure "ShopCart" from one server instance to another, and not to minimise open connections (i.e. sockets) or to avoid the time to open a connection.

Therefore, claim 10 of the main request is inventive in the sense of Article 56 EPC 1973.

T 2035/11 (Navigation system/BEACON NAVIGATION) of
25.7.2014

Navigation system with user definable cost values

Inventive step - main request (no)

Remittal to the department of first instance - (yes)

Application number: 97942552.7

IPC class: G06F 17/00, G08G 1/137, G01C 21/20

Applicant name: Beacon Navigation GmbH

Cited decisions: T 0154/04, T 1784/06, T 1670/07

Board: 3.5.07

<http://www.epo.org/law-practice/case-law-appeals/pdf/t112035eu1.pdf>

The invention

The application mainly relates to navigation systems that can be tailored to a user's particular wishes. The focus of the application is on the route-planning functionality of a navigation system, and the application in fact discloses that the navigation system need not be installed in a vehicle, but can be located in a personal computer (see page 1, lines 11 and 12, and page 4, line 29, to page 5, line 4).

The navigation system comprises a database of road segments, a system for selecting a beginning point, a user input device for selecting a desired destination, and a system for determining a route between said beginning point and said desired destination. Examples of user input devices for selecting the desired destination include a mouse and a keyboard (page 4, lines 19 and 20). The system for selecting a beginning point may be a GPS system, but may also be a user input device (page 2, lines 27 to 29). The database of road segments includes, for each road segment, an "estimated time of travel" and an "estimated length of travel". In addition, each road segment may be associated with a "congestion level" and a "road type". Based on this data and on routing criteria set by the user, the "system for determining a route" evaluates the cost of relevant road segments and determines (and recommends) the route having the lowest total cost.

The congestion level of a road segment determines a factor by which the estimated time and length of travel of the road segment are multiplied for the purpose of evaluating the cost of that road segment. The application explains that the user may increase the congestion level of an area or road that he wishes to avoid "for purely personal reasons" (page 6, lines 20 to 30). The application further discloses that congestion levels may be updated in response to received congestion information, either manually by the user listening to radio reports on traffic conditions or automatically by a data receiving device. The routing criteria described in the application are a "time/distance factor", a "highway preference/avoidance factor" and a "toll roads preference/avoidance factor". These may be used to influence the evaluation of the cost of a road segment as follows:

- the time/distance factor is used to weight estimated time of travel against estimated length of travel;
- the highway preference/avoidance factor is used to adjust the cost of road segments with road type "highway";
- the toll roads preference/avoidance factor is used to adjust the cost of road segments with road type "toll road".

Claim 1 of the main request reads as follows:

"A navigation system comprising:

a database of road segments to be travelled by a user, said database including a cost associated with each said road segment, wherein an estimated time of travel and an estimated length of travel are [sic] associated as costs with each said road segment;

a system for selecting a beginning point relative to said database of roads;

a user input device for selecting a desired destination relative to said database of roads;

a system which is configured to allow the user to adjust said cost of at least one selected said road segment proportionally to its estimated time of travel or estimated length of travel utilizing said user-input device; and

a system for determining a route from said database of roads between said beginning point and said desired destination based upon said adjusted cost of said road segments, wherein the navigation system evaluates the cost of the road segments to be traveled in several potential routes and recommends the potential route having the lowest total cost."

Claim 1 of (new) auxiliary request 1 reads as follows:

"A navigation system comprising:

a database of road segments to be travelled by a user, said database including a cost associated with each said road segment, wherein an estimated time of travel, an estimated length of travel and a congestion level are associated as costs with each said road segment;

a system for selecting a beginning point relative to said database of roads;

a user input device for selecting a desired destination relative to said database of roads;

a data receiving device which is configured to receive congestion information regarding congested road segments and automatically update congestion levels associated with appropriate road segments;

a system which is configured to allow the user to adjust said congestion level of at least one selected said road segment utilizing said user-input device; and

a system for determining a route from said database of roads between said beginning point and said desired destination based upon said adjusted congestion level of said road segments, and which multiplies the cost of each road segment by the congestion level for that road segment,

wherein the navigation system evaluates the cost of the road segments to be traveled in several potential routes and recommends the potential route having the lowest total cost."

Claim 1 encompassed a personal computer programmed to execute a particular optimisation algorithm. This algorithm takes the following input:

- a graph of which the edges represent road segments;
- cost values ("estimated length of travel", "estimated time of travel") associated with the edges of the graph;
- multiplication factors ("congestion levels") associated with at least some of the edges;
- a first graph node ("beginning point") and a second graph node ("destination").

The aim of the optimisation algorithm is to determine a path (consisting of a series of graph edges) connecting the first graph node to the second graph node and **having the lowest total cost**.

The claim leaves it undefined how this total cost is based on the "estimated length of travel" cost values and the "estimated time of travel" cost values, but in the light of the description it is clear that the total cost of a path may be based for example solely on the "estimated length of travel" costs (adjusted by the "congestion level" multiplication factors), or solely on the (adjusted) "estimated time of travel" costs, or on a weighted average of both.

According to claim 1, the graph is provided by means of a "database of road segments". This database also stores "estimated length of travel" and "estimated time of travel" cost values for each road segment. The claim does not define the database in terms of technical features and it may therefore be assumed to be a conventional storage means.

The "beginning point" graph node, "destination" graph node and "congestion level" multiplication factors may be input by a user using a conventional "user input device". The determined optimal path is "recommended", which for example could take the form of displaying the ordered list of graph edges on a conventional display device.

Optimisation algorithms are mathematical methods, excluded "as such" from patentability under Article 52(2) and (3) EPC. Mathematical algorithms may contribute to the technical character of an invention only in so far as they serve a technical purpose (see e.g. decision T 1784/06 of 21 September 2012, reasons 3.1.1).

In the present case **the purpose of the algorithm is the mere display of an optimal path to the user for cognitive processing.** The **user may act on the information, but does not need to.** As stated in decision T 1670/07 of 11 July 2013, reasons 13, a technical effect may arise from either the provision of data about a technical process, regardless of the presence of the user or its subsequent use, or from the provision of data (including data that on its own is excluded, e.g. produced by means of an algorithm) that is applied directly in a technical process. In the present case the **data is produced by means of an algorithm and is not applied directly in a technical process,** so that neither possibility applies.

The Board hence considers that the optimisation algorithm of claim 1 does not serve a technical purpose and therefore does not make a technical contribution.

The present invention is not concerned with the specifics of the constitution of road maps and corresponding databases. The database of road segments of claim 1 is essentially a graph comprising nodes and edges with cost values assigned to the edges. Calculating a route having the lowest total cost does not require knowledge of environmental conditions or traffic routing and planning. The database of road segments serving as input to the algorithm may reflect real-world facts, but **the mere fact that the input to an algorithm has a particular meaning in the real-world is insufficient for the algorithm to make a technical contribution** (see e.g. decision T 154/04, OJ EPO 2008, 46, reasons 20).

The appellant argued that it made no difference whether the route planning algorithm is implemented in a general-purpose computer or in a navigation system installed in a vehicle or

used on a bicycle or while walking. It was also not relevant whether the beginning point or destination were received by GPS or entered manually. **The view that the (claimed) context in which an algorithm is used is of no relevance for the question whether the algorithm contributes to the technical character of the claim is simply incorrect.** Features that as such are excluded from patentability may provide a technical contribution to the extent that they interact with the technical subject-matter of the claim for solving a technical problem (see decision T 154/04, reasons 5, under (f), and reasons 13). **If an invention resides in the application in a technical process of data produced by an algorithm** (the example given in decision T 1670/07, reasons 13, see point 5.1.35.1.3 above), **the application of the produced data in the technical process should be properly reflected in the claim.**

he technical content of claim 1 consists in a computing system comprising conventional storage means and conventional input means. Since, as the Examining Division correctly considered, such computing systems were notorious knowledge at the filing date of the application, the subject-matter of claim 1 lacks an inventive step within the meaning of Articles 52(1) and 56 EPC.

For the sake of completeness, the Board notes that the same conclusion of lack of inventive step is reached if the term "navigation system" is given the narrower interpretation of a system comprising route-planning functionality as well as a position-determining device and being configured to provide route-guidance information in dependence on the actual real-world position of the system. The reasoning is however different. **Providing real-time route-guidance information to a user in dependence on the user's real-world position is a technical task.** It involves an interaction between the user and the navigation system, wherein the navigation system continuously measures the user's position using technical means and, on the basis of these measurements, provides the user with information aimed at enabling the user to manage the technical task of moving a vehicle to a desired destination. Although the **completion of this technical task depends on the user acting upon the provided route-guidance information and hence on an intervention by the user,** it does not rely on subjective considerations by the user or on psychological effects. The user may still decide to ignore the route-guidance information, but that **does not detract from the technical character of the navigation system as a technical tool to be used interactively in a technical process and not merely in a preparatory phase as a substitution of what could also be done using pencil and paper.**

It follows that **a mathematical route-planning algorithm, when used in a navigation system, contributes to the technical character of the system at least to the extent that it produces information that enables the route-guidance functionality.** A route-planning algorithm does this by producing a route in the form of an ordered list of road segments based on real-world map data.

In this context, **features resulting in a different choice of route may contribute to the technical character of the invention, but only in so far as they are based on technical considerations which influence the provided route-guidance information.** In other words, while moving a vehicle along a different route always has physical consequences, for the purpose of determining the technical contribution of the route-planning algorithm such consequences are only to be taken into account as technical effects to the extent that they result from the choice of a route based on technical considerations. In the case of claim 1, **basing the route calculation on estimated time of travel and/or estimated length of**

travel, both of which arguably reflect technical characteristics of a real-world road network, would seem to be an example of such a technical aspect.

The navigation system of claim 1 differs from the acknowledged prior art in that (in accordance with points 4.34.3 and 3.53.5 above) it further comprises a system configured to allow the user to adjust the "congestion level" of at least one selected road segment utilising the user-input device, thereby determining a factor by which the estimated time and length of travel of the road segment are multiplied for the purpose of evaluating the cost of that at least one selected road segment. By increasing or decreasing the congestion level associated with a particular road segment, the user may make it less or more likely for the road segment to be included in the route calculated by the route-planning algorithm. In this way, the user may express his subjective preference for the road segment. The description of the application on page 6, lines 23 to 25, confirms that the motivation for adjusting the congestion level of a route may reside in "purely personal reasons". At the oral proceedings, the appellant also stressed that the invention gave the user more options to influence the route calculation in accordance with the user's personal needs. **The Board considers that the idea of allowing the user to set congestion levels for selected road segments which serve as multiplication factors in the route-calculation algorithm is not technical and can therefore not contribute to an inventive step.** This is because accommodating the user's personal preferences is as such not technical and the Board does not see any other (technical) purpose that the congestion levels serve in the context of claim 1. Although the user of the claimed navigation system could certainly use this option for a technical purpose (for example in order to avoid a particular road for which he knows that his car is technically less suitable), the claim is not limited to such uses. The implementation of this non-technical idea essentially consists in enabling the user to utilise the conventional user input device to select at least one road segment and to adjust the congestion level for the selected at least one road segment. The Board considers this implementation to be obvious, and the appellant has not argued otherwise.

The Board's reasoning does not rely on a particular suggestion in the prior art, but on the finding that **the proposed algorithmic change has no technical motivation and that its implementation is, uncontestedly, trivial.** In so far as it is the appellant's view that technicality is irrelevant for the assessment of inventive step, this view is contrary to the established case law. It follows that the subject-matter of claim 1 lacks an inventive step within the meaning of Articles 52(1) and 56 EPC even if the term "navigation system" is given a narrow interpretation.

The Board notes that the German Federal Supreme Court has come to a similar conclusion in respect of a navigation system offering the user the possibility to exclude road segments based on a user-selectable property such as the road segment being a toll road (see BGH, 18 December 2012, X ZR 3/12, GRUR 2013, 275 -Routenplanung).

Independent claim 1 of auxiliary request 1 differs from claim 1 of the main request essentially in that the claimed navigation system further comprises a data receiving device configured to receive congestion information regarding congested road segments and automatically update congestion levels associated with appropriate road segments. The Board considers that the inclusion in claim 1 of a data receiving device configured to receive congestion information excludes the broad interpretation of the term "navigation system" covering a personal computer executing route-planning software and that the claim is now restricted to a

navigation system operating in a vehicle. Document D1 does not disclose a navigation system configured to automatically receive congestion information. The Board is of the view that the received congestion information provides an indication of real-time real-world road conditions and that basing the calculation of a route on such conditions in the context of actual navigation is technical. There is therefore the potential of an inventive step. Auxiliary request 1 hence overcomes the reasons for the refusal.

T 0330/10 (Resetting passwords/AVAYA) of 1.10.2014

Method and system for resetting passwords

Patentable invention - independent method claim (no)

Application number: 07251625.5
IPC class: G06F 21/20
Applicant name: Avaya Inc.

Board: 3.5.06

<http://www.epo.org/law-practice/case-law-appeals/pdf/t100330eu1.pdf>

The invention

The application relates to the situation that a user has lost or forgotten a password and needs it to be reset. It discusses known ways for resetting a password and their advantages and disadvantages and proposes a new way meant to avoid the drawbacks of the prior art: Specifically, the invention is meant to increase secrecy of the new password by avoiding the risk that the new password is intercepted during transmission and by eliminating the need to involve a trusted third party (see original application, pars. 18-19). The central idea of the invention is that a newly created password is transmitted to the requesting user in (at least) two portions via two possibly different channels. The transmission of the "first portion" is essentially unrestricted (see claim 1), the "second portion" (or any further ones) should be transmitted to a "location having restricted access" but provides access at least to the requesting user. The user is thus able to retrieve both portions and recreate the complete new password from them. The way in which a password is divided into portions is disclosed as "completely arbitrary" (par. 48).

Independent method claim 1 reads as follows:

"A method for resetting passwords comprising:

receiving a request from a purported user to reset a previously set password;

authenticating (106) said purported user as an actual user;

establishing (108) a new password and resetting the previously set password as the new password if the purported user is authenticated as the actual user;

dividing (112) said new password into at least first and second portions;

transmitting (114) said first portion of said new password to said actual user; and

transmitting (116) said second portion of said new password to a location having restricted access;

whereby said actual user has access to said location and is enabled to retrieve both said first and second portions of said new password, and thereby may recreate said new password securely, whereby said actual user has access to said location before said request is received and said new password is established."

The invention as claimed requires that a request to reset a password is received and that the requesting user is authenticated. The description discloses that a user can request resetting the password by calling a responsible "central office" over the phone (see para. 36) and that the authentication can be done "in any known fashion", for instance "by interrogating the [user] with one or more security questions" which may be answered orally (see para. 37). The first portion of the new password may be given immediately, e.g. orally over the telephone, while the second portion may be sent to the user's mobile telephone, voice mailbox or email account (paras. 39 and 43). It is disclosed that, alternatively, the second portion may also be sent to the user's supervisor, e.g. via telephone or email, for personal delivery (para. 44).

Technical character, Article 52 EPC

The description discloses that the user's request to reset a password may be handled by a central office which the user calls over the telephone. The request and the user authentication can then be performed orally (paras. 36 and 37), and also the first portion of the new password can be communicated orally (para. 39). The independent claims do not exclude the possibility that the user, instead of calling a central office, walks up to a service desk and thus does not even use a telephone. Eventually, the user may recreate the new password from the obtained portions in his mind.

The password itself may be just a numerical secret, a PIN, and **the steps of setting and resetting a password may be satisfied by memorising a new password instead of a previously valid password.**

The board is of the opinion that **the concept of a "restricted access" location by itself does not imply any physical means such as a letter box with a lock or an access-controlled email or voice mail account.** Rather, the board considers that a person revealing a secret only to a number of authorised persons may also be considered a "location having restricted access" in the sense of the claims. In the board's judgment this is consistent with the application, which discloses that, if users do not have access to a secure location themselves, the second portion of the new password may be sent to a supervisor instead. This interpretation was presented in the annex to the summons to oral proceedings and was not challenged by the appellant.

Based on this interpretation the board comes to the conclusion that **the method of claim 1 as drafted subsumes a method of communicating a secret between people which does not require any technical means:** the claimed steps of setting a password and establishing a new

password, of dividing the new password into two portions, and of recreating the new password from the two portions relate to activities which users can carry out in their minds and which thus define a mere mental activity (cf. Article 52(2)(c) EPC). The remaining steps of receiving a request, of authenticating a user, of transmitting the two portions to the user and to "a location having restricted access", respectively, and of retrieving the portions can be carried out by persons interacting with each other without using any technical means.

Due to the lack of any technical feature, method claim 1 thus does not have technical character and is therefore not to be regarded as an invention in the sense of Article 52(1) EPC.
