

(19) World Intellectual Property Organization
International Bureau



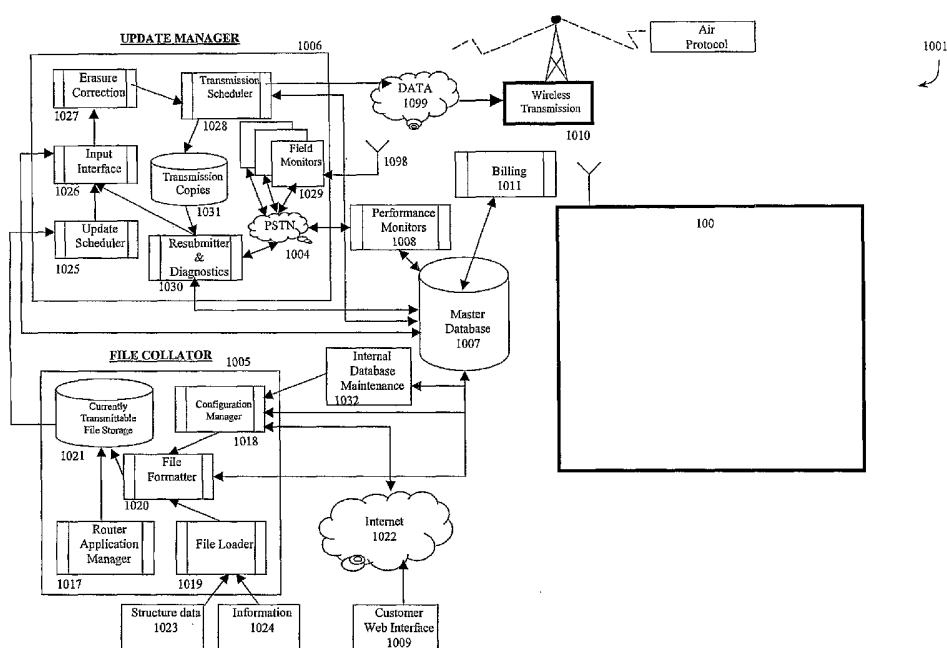
(43) International Publication Date
5 January 2006 (05.01.2006)

PCT

(10) International Publication Number
WO 2006/000038 A1

- (51) International Patent Classification: ⁷ **G06F 9/445**, H04Q 7/00
- (21) International Application Number: PCT/AU2005/000922
- (22) International Filing Date: 24 June 2005 (24.06.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
2004903457 24 June 2004 (24.06.2004) AU
60/608,379 10 September 2004 (10.09.2004) US
2004905262 13 September 2004 (13.09.2004) AU
- (71) Applicant and
(72) Inventor: **DONAGHEY, Andrew, Paul** [AU/AU]; 95-99 Terrara Road, VERMONT SOUTH, Victoria 3133 (AU).
- (72) Inventors; and
(75) Inventors/Applicants (for US only): **MCDONALD, Ian, Kenneth, Francis** [AU/AU]; 57 Beauford Street, HUNTINGDALE, Victoria 3166 (AU). **BEARD, David, Lewis** [AU/AU]; 37 Bemboka Road, WARRANWOOD, Victoria 3134 (AU).
- (74) Agents: **WEBBER, David, Brian** et al.; DAVIES COLLISON CAVE, 1 Nicholson Street, MELBOURNE, Victoria 3000 (AU).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: CLIENT PROCESSOR DEVICE



(57) Abstract: A client processor device including a communications receiver for receiving file fragments for the device, a micro-controller system including a microcontroller and a programmable non-volatile memory system, for building and storing application and data files from the fragments, and executing an application of the device by processing at least one application file and associated data identified by configuration instructions included in at least one of the fragments.

- 1 -

CLIENT PROCESSOR DEVICE

FIELD

- 5 The present invention relates to a client processor device that is able to remotely receive application, configuration or data files to establish, adjust or control the application executed by the device.

BACKGROUND

10

Consumer electronic devices are now commonly used for a wide variety of applications. Many of the devices now include an embedded processor architecture, which normally comprises a microprocessor and associated memory. Instructional code and data for the device is stored in the architecture, and defines and controls the application for which the device has been specifically built. For example, a compact disc (CD) player and a cordless telephone each include a basic processor architecture, but the similarities then normally cease as the CD player is specifically configured to perform a number of tasks associated with processing the audio data on a CD, whereas a telephone is specifically configured to perform a number of call handling tasks. Similar considerations apply to other electronic devices, such as video game consoles, mobile telephones, electronic toys, cameras, etc. The device is specifically configured for its chosen application, and that application cannot be changed, particularly by the receipt of any remote instructions. Even more general processing devices, such as a wireless personal digital assistant (PDA) or a personal computer connected to a computer network, have defined operating systems with a defined set of APIs that restrict the applications that can be performed on the devices, and restrict the extent to which the devices are remotely configurable. Restrictions are also imposed by the security constraints associated with the network that can be used to connect to the device, and any configuration is performed on a one to one basis, ie only one client device can be configured remotely at a given time.

30

It is desired to address the above or at least provide a useful alternative.

- 2 -

SUMMARY

In accordance with the present invention there is provided a client processor device including:

- 5 a communications receiver for receiving file fragments for said device;
- a microcontroller system including a microcontroller and a programmable non-volatile memory system, for building and storing application and data files from said fragments, and executing an application of said device by processing at least one application file and associated data identified by configuration instructions included in at
- 10 least one of said fragments.

The present invention also provides a client processor device, including:

- a wireless communications receiver for receiving files for said device;
- a microcontroller system for processing and storing said files, and automatically
- 15 performing an application defined by said files without any user intervention or interaction.

The present invention also provides a client processor device, including:

- a microprocessor;
- a programmable non-volatile memory system addressable by the microprocessor;
- 20 a communications receiver for receiving file fragments for said device; and
- an operating system (OS) for controlling the device to identify fragments for said device, build files from said fragments, store said files in said memory system, and process the files as including application, configuration or data code.

25 The present invention also provides a client processor device, including:

- a communications receiver for receiving communications packets including transmitted application data sent to said device; and
- a microcontroller system for processing a first application based on stored application data of said device and subsequently processing a second application based on
- 30 said transmitted application data.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the present invention are hereinafter described, by way of example only, with reference to the accompanying drawings, wherein:

- 5 Figure 1 is a block diagram of a preferred embodiment of a client processor device;
 Figure 2 is a block diagram of compilation components for firmware of the client processor device;
 Figure 3 is a block diagram of compilation components for an application of the client processor device;
- 10 Figure 4 is a block diagram of firmware management components for the client processor device;
 Figure 5 is a block diagram of the logic architecture of the client processor device;
 Figure 6 is a block diagram of the software architecture of the client processor device;
- 15 Figure 7 is a block diagram of a flash memory system of the client processor device;
 Figure 8 is a block diagram of the flash file system architecture of the client processor device;
 Figure 9 is a block diagram of an application management system of the client processor device;
- 20 Figure 10 is a block diagram of a file transmission system of the client processor device;
 Figure 11 is a block diagram of an erasure correction fragment generation process of the client processor device;
- 25 Figure 12 is a block diagram of a Lagrange polynomial interpolation process of the client processor device;
 Figure 13 is a flow diagram of an initialisation procedure for the flash memory system;
- 30 Figure 14 is a flow diagram of a directory entry initialisation procedure for the initialisation procedure of the flash memory system;
 Figure 15 is a flow diagram of a receive process of the flash memory system;

- 4 -

Figure 16 is a flow diagram of an update quanta process of the receive process of the flash memory system;

Figure 17 is a block diagram of the client device receiving files from the server transmission system via a variety of network interfaces;

5 Figure 18 is a flow diagram of a file house keeping process of the flash memory system;

Figure 19 is a flow diagram of an incomplete file removal process of the flash memory system;

Figure 20 is a flow diagram of a file deletion process of the flash memory system;

10 Figure 21 is a flow diagram of a building missing fragments process of the flash memory system;

Figure 22 is a block diagram of a remote relay unit for communicating with the client processor device;

15 Figure 23 is a block diagram of a server transmission system for communicating with the client processor device;

Figure 24 is a block diagram of the client processor device for a call management application;

Figure 25 is a block diagram of a logical architecture of the device for the call management application;

20 Figure 26 is a block diagram of a carrier selection main process loop of the system for the call management application;

Figure 27 is a block diagram of an initialisation process of the system;

Figure 28 is a flow diagram of an exchange code resolution process of the system for the call management application;

25 Figure 29 is a flow diagram of a file loader process of the system for the call management application;

Figure 30 is a flow diagram of a best national carrier process of the system for the call management application;

30 Figure 31 is a flow diagram of a redirection option process of the system for the call management application; and

- 5 -

Figure 32 is a flow diagram of an interface option process of the system for the call management application.

5 DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS OF THE INVENTION

1. The Client Processor Device

A client processor device 100, as shown in Figure 1, has a housing 101 that contains a printed circuit board on which three semiconductor chips 110, 120 and 150 are mounted to provide a receiver 110, a non-volatile programmable memory 120, in the form of data-flash memory, and a microcontroller 150. Other support components for the chips are included on the board within the housing 101, such as a battery power system, and any additional input/output (I/O) components 130 that are required, such as a LED or LCD display and control input buttons or keys. The microcontroller 150 is connected to the receiver 110 and the memory 120 via a serial interface 117 and to an address and data bus 114 for connection to the I/O components 130, and has an asynchronous serial interface 115 for external peripherals. The housing 101 is relatively small, and the footprint of the device 100 occupies less than 20 cm². The electronic components of the device 100 also consume less than 100 mW.

The microcontroller 150 is an 8-bit microcontroller with flash read only memory (ROM) 106 (about 128 kbytes), electronically erasable and programmable read only memory (EEPROM) 109 and random access memory (RAM) 108 (less than 4 kbytes), such as the PIC 18 Series Microcontrollers produced by Microchip Technology Inc (<http://www.microchip.com>). The processor core 102 of the microcontroller 150 has a hardware return stack of 32 levels, and no formal data stack, making it primarily suitable to handling only one processor task. Offset addressing is not supported and only two hardware interrupts are provided, one low and one high priority interrupt, without any support for software interrupts. The return stack stores a return address only.

The memory 120 provides data storage on a data flash device, such as the AT45 Series of

- 6 -

data-flash memories produced by Atmel Inc (<http://www.atmel.com>), using a hierarchical structure of sectors, blocks and pages as discussed in detail below. A complete page, comprising 264 or more bytes depending on device type, is written at once, whereas at least one block is erased, with a block comprising eight pages of data. The data-flash
5 memory 120 is managed by a flash file system, described in detail below.

The receiver 110 is an RF transceiver, such as the NT29xx series produced by Numa Technologies, Inc (<http://www.numatechnologies.com>), and is controlled by the microcontroller 150 (via the serial interface 117) to provide an analog output 118, for
10 sampling by the processor core 102, of the radio frequency (RF) signal received by the receiver 110. The receiver provides an air interface, and is frequency agile being phase locked loop synthesised. The receiver uses diversity antennas 111 and 112 connected to the transceiver chip 110, under control of the processor 150 using a control signal 113 to switch the signal from one or the other antenna, to improve performance under marginal
15 signal conditions. A switching arrangement is used, but it is also possible to use dual receivers and receive paths, which may be tuned to different frequencies.

The firmware production process is shown in Figure 2. The firmware 2003 stored in the ROM 106 of the device 100 is primarily written as source code 2001 in the C
20 programming language and compiled using the PICC-18 ANSI C compiler 2002 provided by HI-TECH Software LLC. The compiler uses a compiled stack to deal with the omitted data stack from the processor core 102. The data stack would normally hold auto-allocated (local) variables and function parameters. Instead, the complete calling hierarchy of the code is evaluated at compile time by the compiler to produce a call-graph tree, as described
25 below. The auto-allocated variables and function parameters are allocated to a fixed memory bank in a section of the RAM 108, being designated parameter RAM by the compiler, and the allocations are shared where possible to minimise the overall parameter RAM requirements. For example, the maximum possible parameter RAM size may be 256 bytes, a compromise between the bank switching limitations of the processor and the speed
30 and size of the code. This maximum allocation size is normally more than sufficient for a single main task and routines called using the low and high priority interrupt handlers. The

- 7 -

compiled stack is only designed to handle a single task without supporting re-entrant code. The compiler defines three separate roots in a call graph, one for main execution and one for each of the interrupts.

- 5 The application production process is shown in Figure 3. An application 2012 sent to and stored in the flash memory 106 of the device 100 is primarily written as source code 2010 in a programming language, such as C, and compiled into byte code for a virtual machine (VM) using a compiler 2011 for the device 100.

10 *Operating System (OS)*

The operating system (OS) 210 can be considered to reside on and be supported by the hardware 202 of the device 100, as shown in Figure 5. The code for the OS and its basic tasks 220 are stored in the flash ROM 106. The OS on boot, sets the interrupts, initialises
15 semaphores (described below) and scheduling, and supports and runs a start-up task 227 which in turn starts the remaining basic tasks 220, including the virtual machine task 221 on which the application code 230 relies. The number of tasks is limited to eight due to the word length of one byte of the processor 102 and because an increase would substantially increase overhead. The OS provides a cooperative multi-tasking environment to allow
20 asynchronous events to be handled independently, with much reduced code complexity, enabling the device 100 to be more responsive and reliable. The asynchronous events are normally those either caused externally to, or not under control of, the device, such as button presses, characters received on a serial port interface of the I/O 130 or a hardware peripheral connected to the I/O 130 completing an operation. Most of the events require a
25 response in a timely manner.

The operating system is able to share the time of the processor 102 between the tasks by rapidly context switching between them on a cooperative basis. The tasks are:

- (i) A clock task 225. This maintains time of day and a calendar.
- 30 (ii) A virtual machine (VM) task 221. This processes the byte code of the application files as discussed below.

- 8 -

- (iii) A command task 222. This establishes a set of commands available via the local asynchronous serial interface 115.
- (iv) A flash file system task 223. This handles building and management of all files within the flash memory 120.
- 5 (v) An air task 224. This establishes communications and control of the transceiver 110 and supports an air protocol of the device 100, as discussed below.
- (vi) Additional tasks 226. This represents any other tasks as required, for example, to process data received from a global positioning system (GPS) receiver.
- (vii) The start-up task 227. This establishes the remaining tasks in a controlled fashion
10 as discussed below.

The tasks are each assigned a unique priority depending upon the required processing rate and latency, when started. When a context switch is performed, the OS moves the processor stack contents to a memory area allocated for the task being swapped out, and
15 then restores it with the stored stack for the new task to be run. By allocating stack memory levels required for each task, memory requirements are contained. The storage level can be allocated from 2 to 32 depending on the task, when it is started. A context switch may involve calling a function (e.g. the OS_Yield() function) in the OS, which causes the running task to yield control to the OS and therefore allowing another task to
20 run. Preferably, context switching is achieved only by calling the OS_Yield() function. A common scratch-pad area of the RAM 108 is used for temporary processing and is only relevant within a function call, as the contents are ignored if another routine is called, which may overwrite the contents. Accordingly, when the OS accomplishes a task switch, the scratch pad area does not need to be saved thereby providing a significant saving of
25 RAM and processor overhead.

Many software modules and hardware resources (e.g. the data-flash memory 120, the EEPROM 109, and the bus 117) need to be shared between tasks, and each shared resource is allocated a binary semaphore (or token). If a resource is used and another task attempts
30 to use it, then this is prohibited until the first task finishes with the resource. Locking of the shared resources protects them from re-entrant calls. In the device 100, normally 20 to

- 9 -

30 semaphores are allocated to lock the resources. Binary semaphores are controlled as a bit field in bit addressable memory of the RAM 108, and counting semaphores are used to control linear or multi-level resources (such as buffers). The semaphores are allocated when the firmware 2003 is compiled by the compiler 2002.

5

To protect the shared code (and resource) via locking, the following procedure is used:

- (a) Request access to the module/resource via the semaphore.
- (b) Lock out other tasks using the semaphore associated with the module/resource once access is obtained;
- 10 (c) Call functions for the module/resource; and
- (d) Allow other tasks access by freeing the semaphore associated with the module/resource.

These steps (a) and (d) may be implemented as functions in the corresponding software
15 modules to save code space (e.g. as *Lock* and *Unlock* functions).

Task re-entrant code can be written to allow multiple tasks to call the same routine. A particular task may only call the one routine once. To do this the routine's parameters are saved if they must be accessed after the context switch. This can be achieved using an
20 array to provide context sensitive storage for each task as follows:

```

char storage[8];                // Parameter storage for each of 8 tasks
...
void shared(char var) {
25  storage[<current_task>] = var;    // save the parameter
...

OS_Yield();                    // context switch
var = storage[<current_task>] ;    // Restore parameter...
30  }
```

- 10 -

The OS also includes a scheduler that is responsible for determining the processor time allocated to each task. The highest priority task (e.g. task "1") is allocated up to 50% of the processor time, given the opportunity to run every second context switch, whilst the next high office priority task (e.g. task "2") is allowed to run 25% percent of the context switches and so on. For example, the task may be switched based on a priority level as follows: 1 2 1 3 1 2 1 4 1 2 1 3 1 2 1 5, etc. If a task cannot be run because it is blocked, then the next lowest priority task is run and so on. This allows resource intensive or critical tasks to be run more often or with lower latency whilst ensuring slower tasks are not able to stall. The scheduler in addition is able to prevent a potential priority inversion hang without requiring a complex priority inheritance procedure. While lower priority tasks are guaranteed to run, a higher priority task waiting on a resource in use by a lower level task will not be indefinitely blocked.

The low and high priority interrupts of the processor 100 are treated as being slow and fast, respectively. The fast interrupt is reserved for handling time critical tasks, such as analog to digital conversion of a sampled analog signal. The fast interrupts generally do not relate to the use of shared routines or those of the OS, and the slow interrupts are used to handle asynchronous events, such as timer interrupts, serial port events and lower level packet decoding that directly relate to OS controlled tasks. The low interrupt handlers use OS functions, such as semaphores to communicate with tasks.

A timer module 316 of the OS provides, shown in Figure 6, accurate timing of intervals and delays, and allows tasks to be synchronised to a period such as one second for a clock of the device 100. The timer module allows a variety of different high and low resolution durations to be measured, and a number of timers, can be allocated when the firmware is compiled.

Software Architecture

The software architecture employed by the firmware stored in the ROM 106 distinguishes between software modules 300 and the tasks 220, as shown in Figure 6. The software

- 11 -

modules 300 are available for use by the tasks 220, the arrows indicating dependencies. An object-oriented approach is used in the module decomposition process, and some modules are given responsibility at a given time for a particular shared resource (e.g. the serial interface, EEPROM, etc.) as required.

5

There are four classes of software modules 300. The first type 310 includes the OS core 312 that manages scheduling and task switching, and the semaphore 314 and timing 316 modules. These modules are typically used by all tasks and written to guarantee that they can be called at any time (i.e. not require locking). The second type are shared modules
10 that must be locked to ensure conflicts do not occur. These include an SPI module 323 that manages data transfer on the serial peripheral interface and an EE module 320 that allows access to the processor's EEPROM 109. These modules typically contain a number of available routines, which are locked as a group (i.e. these routines are locked/unlocked using the same semaphore).

15

The third type are modules run autonomously from a task. These include the virtual machine 321, command processor 322, air interface 324 and start-up 327 modules. These are each dedicated to a particular task, not shared and therefore do not require locking.

20

The fourth type is a combination of the latter two. These have components that are run only from a single task, for example in the flash file system (FFS) module 330, the module 331 is run from a FFS task 223 whereas a set of shared functions, 332 and 333 are available for use by external modules, including the internal component 331. The shared functions are locked for use. The clock module 340 is another example, where a dedicated
25 module 341 is run from the clock task 225, and a set of shared functions 342, 343 are also available for use by other modules.

25

The main() routine 328 is responsible for initialising all modules as required and calling an OS 312 routine to run the start-up task 227.

30

Call Graph Management

Manipulation of the call graph supports a multi-tasking environment. Call-graph allocation is performed using assembler directives unique to the firmware compiler 2002.

- 5 These include directives to define a root, and to allocate each routine to a particular root. Any of the modules that have contained routines that run to completion and do not require a separate call-graph root are defined to use a default main root to save parameter space.

- 10 A flattening call-graph approach could be used but typically results in the very limited parameter space allocation being easily exceeded. Rather, the call graph is broken into a number of separate trees, whereby the main function of each task forms the base of a tree or root. This allows the overlaying of auto variables between functions and parameter passing. For routines that are not run to completion, assembly directives in the task code are used. These are used internally by the compiler to ensure interrupt code starts at a
- 15 separate base and interrupts can run at any time. Routines or functions that are shared across tasks are always assigned to a root. Shared functions which always run to completion between context switches may all share the same root as they will not run concurrently, but they may not share a root with tasks that do not always run to completion. Any shared functions that may incur a context switch before they complete
- 20 are assigned a unique root. Groups of such functions or routines that are guaranteed never to be called concurrently by more than one task can be assigned the same root to conserve parameter space.

- Any functions solely called either directly or indirectly by the main() routine 328 need not
- 25 be formally placed in the call-graph as they will be placed in the main root of the call-graph by the compiler. These would typically include once-off initialisation and configuration functions.

Firmware Management

- 13 -

The device 100 can be structured so that only a small boot routine 160 is permanently held in the device ROM 106 as depicted in Figure 4. The boot routine 160 can program the remaining firmware components 170 into the ROM with data obtained via the local interface 115, or transfer execution to the remaining firmware components once they are
5 programmed.

Once the remaining firmware components are in place and operating, the flash file system of the client processor device 100 receives and stores files in the flash memory 120 for processing by the virtual machine 321 of the OS. The files normally relate to an
10 application or a version of the application, but the files may also include updates to the remaining firmware components. The boot routine 160 may read file information from the flash memory 120 but without the need to implement the flash file system to the full extent. The boot routine may compare the versions of the firmware components programmed in the device ROM with files identified to be firmware (e.g. 121, 122) in the
15 flash memory 120 and automatically update the firmware components in the device ROM, if it determines they are newer.

The device 100 is entirely remotely configurable, without a user having to interact directly with the device, e.g. by selecting a button.

20 *Flash File System*

The flash file system of the client processor device 100 receives and stores files in the flash memory 120 for processing by the virtual machine 321 of the OS. The flash file system 400, as shown in Figure 7, is a linear file system that allows up to 64 files of up to 64 kbyte
25 to be stored in the flash memory 120. The file system allows storage of files (e.g. files 411, 412), in a write-protected area 410 and optionally in a read-write area 420 (e.g. file 421, 422, 423, 424), as shown in Figure 7. The write-protected area is read-only if the device is not in an initial configuration state, otherwise write access is allowed. The flash file system supports 512kbyte and 1Mbyte data flash devices 120 and this does not preclude other
30 device sizes or numbers of devices, if required.

- 14 -

The data-flash memory 120 is divided up into 264 byte pages, where 256 bytes are used for actual data, and the 8 spare bytes are used for indicating what portions of the page contain valid data. A structure is used, as shown in Figure 8, where 508 represents RAM components and 504 the components stored in the data-flash memory 120. One or two
5 flash blocks are treated as an allocation unit, depending on the device capacity, so that a total of 256 allocation units are available. An allocation unit may be a directory unit 502, or a data unit 505. The directory units 502 each include directory entries 510 having a header 512 and a table 509. The header 512 includes information about the file, such as its local creation date 541, CRC 542, release date 543, length 544, data quanta count 545,
10 erasure quanta count, 546, access mode 547, name 548, type 549, quanta size 550, quanta count 551, erasure correction 552. Table 514 includes pointers 509 to the data allocation units 505 that actually hold the data for the file. Each data allocation unit holds one or two blocks, being eight or sixteen pages 530 depending on the flash device capacity. A page 530, containing valid data, consists of 256 bytes of data 531, 4 bytes of quanta 532, and a
15 32-bit CRC. The quanta 532 represents up to 32 indicator flags, which indicate which bytes of the data 531, are valid. The RAM 108 maintains a set of tables 520, which each have pointers 507 to predefined number of directory units 502 which each include 16 or 32 directory entries 510 per unit, depending on flash device capacity.

20 Allocation units in write-protected flash and those in read-write flash are allocated from pools 521 and 522, respectively, representing free allocation units, maintained in RAM 108.

A table 523 is maintained for the state of all directory entries 510 in the directory
25 allocation units 502 in the write-protected and read-write areas of the flash memory 120. The state defines whether the directory entry 510 contains an allocated file header 512, and if so, the file's name and flags indicating, whether it is complete, and any request that it be deleted. For a complete file additional state flags indicate whether it is released, for use by the virtual machine. The state also indicates whether a released the file is locked and/or
30 open.

- 15 -

A count of the total number of writes to each allocation unit is maintained in a table in EEPROM 109. The count is used to ensure the next allocation unit with the lowest count is used thereby achieving even wear of the flash. If this count exceeds a threshold, then the file system will not reallocate the allocation unit, and it can be treated as allocated in order to avoid further access to it. A free allocation unit with the least number of writes to it is allocated first. The file system updates the pointers accordingly if allocation units are moved due to a worn unit.

The flash file system task 223 on initialisation executes a start up procedure 7000, as shown in Figure 13. The start up procedure 7000 begins at step 7002 by searching the write-protected area 410 of the flash memory 120 for the expected number of directory allocation units 502. Step 7004 determines whether such directory allocation units are found. If all units are found, step 7004 proceeds to step 7006 to search the read-write area 420 of the flash memory 120 for directory allocation units. Then, if step 7008 determines that directory allocation units are found in the read-write area, step 7008 proceeds to step 7080, where all directory entries are processed.

However, if step 7004 determines that the expected number of directory allocation units are not found in the write-protected area, step 7004 proceeds to step 7024, to determine whether the unit is in an initial configuration state. Initial configuration occurs at manufacture being indicated by a dedicated hardware signal. If step 7024 determines that the unit is not being configured, step 7026 records the error and, at step 7028, a signal is generated to indicate that the flash file system has failed to start.

If step 7024 determines that the unit is being configured, step 7024 proceeds to step 7030 to format the flash memory 120 (e.g. by deleting data from both the write-protected and read-write areas of the memory 120). At step 7032, header units in the write-protected area are created, and at step 7036, header units in the read-write area are similarly created. Step 7036 proceeds to step 7040 to process all directory entries, as shown in Figure 14.

30

If step 7008 determines that directory allocation units 502 are not found in the read-write

- 16 -

area, step 7008 proceeds to step 7034 to format only the read-write area. Step 7034 proceeds to step 7036 to create header units in the read-write area. Step 7036 proceeds to step 7040 to process all directory entries, as shown in Figure 14.

5 Once all the directory entries have been processed, step 7040 proceeds to step 7010 to determine whether any of the directory allocation units 502 in the read-write area of the memory 120 are missing. If none are missing, step 7010 proceeds to step 7022 to start the file system. Otherwise, step 7010 proceeds to step 7012 to search for a free allocation unit in table 52 (e.g. in the read-write area). Step 7014 determines whether a free allocation
10 unit has been found. If found, step 7014 proceeds to step 7016 to assign the free allocation unit as a directory entry. Step 7016 then proceeds to step 7010 to continue processing other directory entries. If step 7014 determines that a free allocation unit cannot be found, step 7014 proceeds to step 7018 to record the error, and at step 7020, generates a signal indicating that the file system has failed to start.

15

The process shown in Figure 14 processes all directory entries in order to ascertain the files resident in the flash memory, the state and integrity of each file (i.e. whether complete, incomplete or corrupted) and the allocation units used. No assumptions are made about the state or number of files, or data allocation units prior to this process.

20

Step 7042 is the first step for individually processing each directory entry 510 in the directory allocation units 502 in the write-protected area, and also the directory allocation units in the read-write area of memory 120. Each directory entry 510 potentially contains data related to a file (e.g. 411, 421). Step 7042 selects a directory entry for processing. At
25 step 7044, the file name 548 represented by the selected directory entry (contained in header of the selected directory entry), and the file location (e.g. write-protected or read-write), and attributes such as completeness are entered into a RAM directory table 523. Step 7052 then determines whether the selected directory entry contains file data. If step 7052 determines that the file contains no data, step 7052 proceeds to step 7046 which
30 determines whether all the directory entries identified by the directory allocation units (eg. 502) in memory 120 have been processed. If so, step 7048 proceeds to step 7050

- 17 -

indicating all directory entries have been processed. Otherwise, step 7046 proceeds to step 7048, which selects another directory entry from the directory allocation units in memory and proceeds to step 7044.

- 5 If step 7052 determines that the file contains data, step 7052 proceeds to step 7054 to select a data allocation unit 505 corresponding to the file (e.g. by retrieving the pointer to the first data allocation unit (e.g. 505) from the table (e.g. 514) for the selected directory entry). Step 7054 then proceeds to step 7052 to process the selected data allocation unit by determining whether that unit has been allocated in memory 120 by checking table 522. If
- 10 so, step 7056 proceeds to step 7058 to first determine whether the file can entry can be deleted from the corresponding directory allocation unit (as the allocation unit is incorrectly in use by another file) by determining whether the file is in write-protected space. Otherwise, step 7056 proceeds to step 7064 to check the contents of the selected data allocation unit.
- 15
- If step 7058 determines the file is in write-protected space, step 7058 proceeds to step 7060 which determines if the device is in an initial configuration state. Otherwise, step 7058 proceeds to step 7062 to clear the file entry. If step 7060 determines the device is not in an initial configuration state, step 7060 proceeds to step 7094 to indicate an unrecoverable
- 20 error. Otherwise, step 7060 proceeds to step 7062 to clear the file entry by de-allocating all data allocation units referred to in the by the file's header table 514, except the data allocation unit already in use by another file, then clearing the file's directory entry 510. Step 7062, then proceeds to step 7046 which determines whether all the directory entries have been processed. If so, step 7046 proceeds to step 7050 indicating that all directory
- 25 entries have been processed. Otherwise, step 7046 proceeds to step 7048 to select a new directory entry and proceeding to step 7044..

Step 7064 proceeds to check the contents of the selected allocation unit 505 by selecting the unit's first page and proceeding to step 7066. Step 7066 determines whether the

30 selected page 530 contains any data by checking whether the data quanta flags 532 indicate any data bytes are present. If so, step 7066 proceeds to step 7068 to check the CRC of the

- 18 -

selected page by evaluating the CRC of the data 531 and quanta flags (e.g. 532) Otherwise, step 7066 proceeds to step 7078 to check whether all the pages in the allocation unit have been processed.

- 5 Step 7068 proceeds to step 7070 to check whether the evaluated CRC in step 7068, matches the stored CRC 533. If not, step 7070 proceeds to step 7072 to record the error and, at then proceeds to step 7058 to delete the file entry from the corresponding directory allocation unit, if allowed. Otherwise, step 7070 proceeds to step 7074 to set a flag in table 522 indicating that the data allocation unit 305 has been allocated.

10

Step 7074 proceeds to step 7076 to record a running total of the amount of data quanta contained in the data pages, and proceeds to step 7078 to check whether all the pages in the allocation unit have been processed.. If not, step 7078 proceeds to stop 7080 to select another page in the data allocation unit for processing, and proceeds to step 7056.

- 15 Otherwise, step 7078 proceeds to step 7082 to determine whether all the data allocation units for the file (as defined in the directory entry) have been processed. If not, step 7082 proceeds to step 7094 to select another data allocation unit for the same file (e.g. another pointer to a data allocation unit for that file), and the new selected data allocation unit is processed at step 7056.

20

If step 7082 determines that all the data allocation units for the file have been processed, step 7082 proceeds to step 7084. Although all the data allocation units for a file (as defined in the directory entry) have been processed, the file may still be incomplete. Step 7084 determines whether the data for the file is complete by whether the quanta accumulated in step 7076 matches the expected quanta count (e.g. 545) in the file header 540. If not, step 7084 proceeds to step 7046 to select another directory entry. Otherwise, step 7084 proceeds to step 7086 to check the CRC data for the file as a whole. Step 7088 determines whether the CRC check at step 7086 indicates that all the data contained in the file is consistent with the file CRC543 contained in the header 540. If not, step 7088

25

proceeds to step 7058 to delete the file entry from the corresponding directory allocation unit, if allowed. Otherwise, step 7088 proceeds to step 7090.

30

- 19 -

Step 7090 determines whether the file presently being processed is a duplicate of another file in the flash file system by comparing file header (e.g, 512) parameters. If step 7090 determines that the file is not a duplicate, step 7090 proceeds to step 7046 to select a new
5 directory entry. Otherwise, step 7090 proceeds to step 7092 to determine whether the date/time stamp on the present file 543 represents a time that is earlier than the date/time of another file of which the present file is a duplicate. If so, step 7092 proceeds to step 7058 to check if the file entry can be deleted from the corresponding directory allocation unit. Otherwise, step 7092 proceeds to step 7046 to 7046 to check whether all the directory
10 entries have been processed.

The failures identified in flash memory initialisation processes 7000 and 7040 may occur as a result of corruption or loss of data in one or more flash memory pages. Possible causes of such corruption may be power failure or system re-boot whilst the data is being written
15 to the device 100, or random data bit-errors due to wear and/or cross-talk between pages. It is beneficial to guard against these modes of failure and reducing their severity if they do occur. This can be achieved by modifying the file system to:

- (a) wear the flash memory as evenly as possible;
- (b) ensure loss of a single data page never affects more than one file; and
- 20 (c) ensure corruption, or loss, of a single page can be recovered.

The first two are satisfied by employing distributed file directory headers throughout the flash memory (e.g. within the first data allocation unit of each file), rather than aggregating them in a few allocation units. The last can be addressed, at least in part, by using duplicate directory headers for each file, each on separate flash memory pages, to allow a corrupted
25 header to be recovered from the intact header. To recover a lost page from the data area of a file, a simple erasure correction can be implemented by storing an additional page (e.g. along with the duplicated directory headers), to provide a page-wise parity of all data pages in the file. A corrupt header or data page can be identified by the failure of the CRC associated with each page written to flash.

30

- 20 -

The flash file system processes file fragments received by the receiver 110 to perform a receive process 7100, as shown in Figure 15. Space in the flash 120 is only allocated for the actual fragments received, to reduce overhead, and each file fragment is time stamped with its local creation date 541, in addition to a release date 543 determined by the transmitting server. Fragments can be received in different orders and multiple files received at once, which allows for interleaved transmission and prompt transmission of high priority files.

The receive process 7100 as shown in Figure 15 begins at step 7102 on receiving a file fragment. Step 7102 searches the RAM directory table 523 for an unreleased file (i.e. an incomplete file or completed file not in the released state) with the same name 548 as the file fragment received. Step 7102 then proceeds to step 7104 to determine whether such an incomplete file exists in the RAM directory table 523. If found, step 7104 proceeds to step 7106, which compares the header parameters of the fragment with those of the existing file. If an unreleased file is not found, step 7104 proceeds to step 7134 to search the RAM directory table 523 for a complete file in the released state with the same name as the file fragment received.

Step 7106 then proceeds to step 7108, to determine whether the header parameters (i.e. 542 to 552) of the fragment received and the existing file match. If so, step 7108 proceeds to step 7110, which checks whether the file is completed. Otherwise, step 7108 proceeds to step 7116, which determines whether the file is obsolete by checking whether the release date 543 of the fragment received is newer than that of the existing incomplete file.

If the file is determined as complete in step 7110, step 7110 proceeds to step 7148, to end process 7100. Otherwise, step 7110 proceeds to step 7112 which determines whether the existing unreleased file is in the write-protected section of flash memory 120. If step 7112 determined that the file is write-protected, step 7112 proceeds to step 7114, which determines whether the device is in the initial configuration state which allows writes to the write-protected section of flash memory. Otherwise, step 7112 proceeds to step 7150; a process which updates the incomplete file in flash memory with the data received.

- 21 -

If step 7114 determines that the write-protected section is writable, step 7114 proceeds to step 7150, to update the incomplete file in write-protected flash memory. Otherwise, step 7114 proceeds to step 7148, to end process 7100, with the file fragment received being
5 discarded, as the write-protected section 410 of the flash memory 120 is read-only.

If step 7116 determines that the fragment is newer, step 7116 proceeds to step 7118, which checks whether either the file fragment has a read-only access flag (e.g. 547), indicating that it should be stored in the write-protected area, or that the existing file is in the read-
10 only section of flash memory 120. Step 7118 ensures a file of the same name may not exist in both write-protected and read-write areas of the flash memory.

If step 7118 indicates neither the fragment nor the existing file is read-only access, step 7118 proceeds to step 7122, which deletes the existing incomplete file as it has been made
15 obsolete by the newer file fragment. Otherwise, step 7118 proceeds to step 7120 which checks whether write access to the write-protected area of the flash memory system is available. If so, step 7120 proceeds to step 7122 to delete the existing incomplete file by setting the delete flag in the corresponding entry in the directory table 523. Otherwise, step 7120 proceeds to step 7148, to end process 7100. Step 7122 proceeds to step 7124, which
20 attempts to search for a free directory header in the appropriate write-protected 410 or read-write 420 areas of the flash memory depending on the read-only access flag in the received file fragment. The result of the search in step 7124 is checked in step 7126. If the search located a free header step 7126 proceeds to 7128 to initialize the fields (e.g. 542 to 552) in a new directory header 512 with the header information from the received file
25 fragment, and the local time field 541 with the current time maintained by the clock task 225. Otherwise, step 7128 proceeds to step 7130 to record a failure to locate a free header, which proceeds to step 7132 to indicate a file system failure.

Step 7128 also flags the status of the new file as incomplete in the RAM directory table
30 523, and proceeds to step 7150 to commence storing the fragment as a new file.

- 22 -

Based on the search performed at step 7134, step 7136 determines whether a released file with the same name as the file fragment received exists in memory 120. If no such file exists, step 7136 proceeds to step 7144. However, if a released file is found, step 7136 proceeds to step 7138 to compare the release dates of the file and the file fragment. If the
5 file fragment is newer, step 7138, proceeds to step 7140 to check whether the device is in an initial configuration state. If so, step 7140 proceeds to step 7124 commence storing the file fragment as a new file. Otherwise step 7140 proceeds to step 7142, which checks whether either the file fragment has a read-only access flag (e.g. 547), indicating that it should be stored in the write-protected area, or that the existing file is in the write-
10 protected section of flash memory 120. If step 7142 indicates neither the fragment nor the existing file is read-only access, step 7142 proceeds to step 7124 to commence storing the file fragment as a new file. Otherwise, step 7142 proceeds to step 7148 to end process 7100.

15 Step 7144 determines whether the data in the file fragment received belongs to a file intended for storage in the write-protected area of memory 120 (e.g. based on access mode data stored in the file fragment). If so, step 7144 proceeds to step 7146 to determine whether the unit is in an initial configuration state, which allows write-access to the write-protected section 410 of the flash memory. Otherwise, step 7144 proceeds to step 7124 to
20 commence storing the file fragment as a new file. If step 7146 determines the unit is in initial configuration, step 7146 proceeds to step 7124 to commence storing the fragment as a new file. Otherwise, step 7142 proceeds to step 7148 to end process 7100.

The process 7150 to update an incomplete file with the data quanta received in a file
25 fragment is shown in Figure 16. The update quanta process begins at step 7152 by retrieving the pointer to the page 530 in the data allocation unit where the data contained in the file fragment is destined to be written. The data allocation unit is identified by indexing the table of data allocation unit pointers (e.g. 514) in the file's directory entry. The index is calculated by using the starting quanta number (e.g. 741) in the file fragment received, the
30 quanta size (e.g. 550) and the number of quanta held by each data allocation unit.

- 23 -

It is assumed quanta in the file fragment do not overlap pages 530 in data allocation units, however it is straightforward to enhance the procedure to cater for file fragments that contain data that may overlap pages.

5 Step 7152 proceeds to step 7154, which determines whether a data allocation unit identified by step 7152 has been allocated in the data allocation unit table (e.g, 514) in the file's directory entry. If so, step 7154 proceeds to step 7156, which determines if the page identified in step 7152 contains all data quanta present in the file fragment received. Otherwise, step 7154 proceeds to step 7164, which searches the relevant pool of allocation
10 units (e.g. 521 or 522) for a free data allocation unit depending on the read-only access flag of the file (e.g. 547). If step 7156 determines all data in the fragment are present in the file, step 7158, following step 7156, proceeds to step 7162 to end the update quanta process 7150. Otherwise, step 7158 proceeds to step 7160 to write the data in the fragment to the page and end the update quanta process 7150 at step 7162. Step 7164 proceeds to
15 step 7166, which determines whether the search for a free allocation unit was successful. If so, step 7166 proceeds to step 7168, which adds a pointer (e.g 509) to the free allocation unit to the index in the file's header table (e.g. 514), and proceeds to step 7160 to write the fragment to the allocation unit. Otherwise, step 7166 proceeds to step 7170, which records an error, and proceeds to 7172 to end the update quanta process 7150 with a failure.

20

Figure 17 shows the client device 100 receiving files from the server transmission system as file fragments 2501, 2502 and 2503. The file fragments 2501, 2502 and 2503 may be sent to the client device 100 via a different communications interface 2506, 2508, 2510 (e.g. interfaces for a paging network, a wireless network, the Internet, or a direct
25 connection with the client device 100). The file fragments 2501, 2502 and 2503 do not have to be sent or received in any particular sequence. When the client device 100 receives all the file fragments 2501, 2502 and 2503, the device 100 reorganises all the fragment into their proper sequence for processing (e.g. based on the fragment identifier for each file fragment 2501, 2502 and 2503 indicating their sequence).

30

- 24 -

Similarly, commands and configuration data can be sent to the client device 100 via a different communications interfaces.

The file system also implements erasure correction recovery of missing fragments. All
5 files have an associated 32-bit cyclic redundancy check (CRC) and a CRC calculated locally and stored for every page to check the integrity of the data stored in the flash memory page. These erasure correction pages follow the end of the file. A file missing
core fragments also has provision for Erasure Error Correction (EEC) data at the end, and when the necessary file data is received or able to be reconstructed using the EEC data, any
10 EEC data is freed and file marked as "complete".

Once the file is complete, it cannot be altered or patched by the VM task 213, but can only be deleted by the flash file system.

15 The flash file system supports a locking mechanism to prevent them from being deleted during a garbage collection process performed by the file system task, particularly if a new version of the file is available. A locked file gives an application guaranteed access to that file. The flash file system always locks existing files on initialisation and newly completed files. Files can only be unlocked by the application.

20

The file system housekeeping process 7200 is shown in Figure 18. Once the flash file system is started at step 7202, process 7200 executes a number of procedures in turn to remove incomplete files (at step 7204), remove deleted files (at step 7206) and to check completed files for deletion (at step 7208). The function of steps 7204, 7206 and 7208 are
25 described in further detail with reference to processes 7300, 7400 and 7500 respectively, as depicted in Figures 19, 20 and 21. The actions of determining which file to delete and actually deleting them are made separate by the use of a deletion flag to indicate to the deletion process that the file is a delete candidate. This allows the relatively slower deletion process shown in Figure 20 to be handled by the housekeeping task 223 thereby
30 speeding up any process which wishes to delete a file.

- 25 -

Figure 19 is a flow diagram of the process 7300 for removing incomplete files if they are not completed within a specified time. If the number and type of free allocation units (i.e. write-protected or read-write), related to the access mode of the file, fall below a preset threshold less time is allowed for incomplete files to become complete to free up more allocation units. Process 7300 begins at step 7304 by determining whether the Real Time Clock (RTC) of the microprocessor 102 is returning a valid time value. If the RTC output is not a valid time value, process 7300 ends and execution returns to process 7200. Otherwise, if the RTC output is valid, step 7304 proceeds to step 7306 to locate a first directory entry in the flash memory 120 for processing. Step 7306 proceeds to step 7308 to retrieve data stored in the directory entry and determine whether the directory entry corresponds to a file. If the directory entry does not correspond to a file, step 7308 proceeds to step 7330. Otherwise, step 7308 proceeds to step 7310 to determine whether the file has been marked for deletion. If the file has been marked for deletion step 7310 proceeds to step 7330. Otherwise, step 7310 proceeds to step 7312 to determine if the file is complete. If so, step 7312 proceeds to step 7330. Otherwise, step 7312 proceeds to step 7314 to determine if the file is stored in the write-protected section of the flash memory 120 (i.e. read-only memory if not in initial configuration).

If the file is stored in write-protected memory, step 7314 proceeds to step 7324 to determine if there are sufficient free allocation units in the read-only memory for storing more data. If the write-protected memory has sufficient capacity, step 7324 proceeds to step 7330. Otherwise, step 7326 generates the time difference between the local creation date of the file (i.e. the time when the first portion of the file was stored in memory 120) and the time at which the most recent portion of the file was received. Then, step 7328 determines whether the time difference generated at step 7326 exceeds the predetermined timeout value for the write-protected memory. If the generated time difference exceeds the timeout value for the write-protected memory, then step 7328 proceeds to step 7330. Otherwise, step 7328 proceeds to step 7322 to mark the file for deletion.

If the file is stored in the read-write section of the flash memory 120 (i.e. read-write memory), then step 7314 proceeds to step 7316 to determine if there are sufficient free

- 26 -

allocation units in the read-write memory for storing more data. If the read-write memory has sufficient capacity, step 7316 proceeds to step 7330. Otherwise, step 7318 generates the time difference between the local creation date of the file (i.e. the time when the first portion of the file was stored in memory 120) and the time at which the most recent portion of the file was received. Then, step 7320 determines whether the time difference generated at step 7318 exceeds the predetermined timeout value for the read-write memory. If the generated time difference exceeds the timeout value for the read-write memory, then step 7320 proceeds to step 7330. Otherwise, step 7328 proceeds to step 7322 to mark the file for deletion.

10

Step 7322 marks the file for deletion (e.g. by setting the deletion flag for the file) and proceeds to step 7330 to determine whether all directory entries in the memory 120 have been processed. If so, process 7300 ends and execution returns to process 1200. Otherwise, step 7330 proceeds to step 7308, where a different directory entry is selected and processed.

15

Figure 20 is a flow diagram of the process 7400 for removing files marked for deletion from the memory 120. Process 7400 begins at step 7404 by locating a first directory entry in the flash memory 120 for processing. Step 7404 proceeds to step 7406 to retrieve data stored in the directory entry and determine whether the directory entry corresponds to a file. If the directory entry does not correspond to a file, step 7406 proceeds to step 7416. Otherwise, step 7406 proceeds to step 7408 to determine whether the file has been marked for deletion by examining the file's status in the directory table 523. If the file has not been marked for deletion step 7408 proceeds to step 7416. Otherwise, step 7408 proceeds to step 7410 to determine if the file is open (e.g. currently being opened or used). If so, step 7410 proceeds to step 7416. Otherwise, step 7410 proceeds to step 7412 to determine if the file is locked (e.g. locked so it will not be deleted). If so, step 7412 proceeds to step 7416. Otherwise, step 7412 proceeds to step 7414 to delete the file. Step 7414 proceeds to step 7416 to determine whether all directory entries in the memory 120 have been processed. If so, process 7400 ends and execution returns to process 7200. Otherwise, step 7416 proceeds to step 7406, where a different directory entry is selected and

20

25

30

- 27 -

processed. The main file system housekeeping loop is also responsible for checking incomplete files and building any missing fragments using received erasure correction fragments where possible, as depicted in Figure 21. A CRC is performed on any newly completed files to check their integrity. The process of assembling files is decoupled from
5 the receive fragment process shown in Figure 15, to simplify both.

Figure 21 is a flow diagram of the process 7500 for marking incomplete files for deletion from the memory 120. Process 7500 begins at step 7504 by locating a first directory entry in the flash memory 120 for processing. Step 7504 proceeds to step 7506 to retrieve data
10 stored in the directory entry and determine whether the directory entry corresponds to a file. If the directory entry does not correspond to a file, step 7506 proceeds to step 7530. Step 7530 determines whether all the directory entries in the memory 120 have processed. If so, process 7500 ends and execution returns to process 7200. Otherwise, step 7530 selects another directory entry from memory 120 and proceeds to step 7506 to process the
15 new directory entry.

If step 7506 determines that the directory entry corresponds to a file, step 7506 proceeds to step 7508 to determine whether the file has been marked for deletion. If so, step 7508 proceeds to step 7530. Otherwise, step 7508 proceeds to step 7510 to determine whether
20 the file is complete. If the file is complete, step 7510 proceeds to step 7530. Otherwise, step 7510 proceeds to step 7512 to generate the core quanta present in the file (i.e. the amount of data for the file, excluding the erasure correction fragments) and the core quanta still required (i.e. the amount of data still required for the file to be complete). The core quanta required may be determined, for example, by subtracting the core quanta present in
25 the file from the file length. Step 7512 proceeds to step 7514 to verify whether the quanta present for the file (determined at step 7512) is actually stored in the file system. If so, step 7514 proceeds to step 7518. Otherwise, step 7514 proceeds to step 7516 to attempt to record core quanta of the file using the error correction fragments. Step 7516 then proceeds to step 7518.

30

- 28 -

Step 7518 determines whether the core quanta for the file has been stored in memory 120 (e.g. including the fragments recovered at step 7516). If not, step 7518 proceeds to step 7530. However, if all the core quanta is present, step 7518 proceeds to step 7520 to determine whether the file has now been marked for deletion. If so, step 7520 proceeds to step 7530. Otherwise, step 7520 proceeds to step 7522 to generate CRC data for the file. Then, at step 7524, the CRC data generated at step 7522 is compared with the CRC data stored in the header of the file. If the CRC data match, step 7524 proceeds to step 7526 to mark the file as completed, and truncates the erasure correction fragments for the file. Otherwise, step 7524 proceeds to step 7528 to mark the file for deletion. Both steps 7526 and 7528 proceed to step 7530.

Virtual Machine

The virtual machine (VM) 321 is a software module executed by the processor core 102 as a task, VM task 221, which runs applications stored in the flash file system. The VM is implemented as a byte-code interpreter that allows a single application to be run at any time, as the interpreter is single threaded. The VM may include, for example, two separate program counters and two separate stacks in order to run multiple applications. The use of byte-code allows the execution of applications from the data-flash memory that otherwise would not be able to be run by the processor 102.

A VM application is primarily event driven. It polls and acts on an event queue which accepts asynchronous events from user I/O (e.g. buttons), scheduled timer events, etc.

A particular application normally resides in a single identified file, as discussed below, but the interpreter supports call and return of functions in other files, which allows the application code to be maintained in a modular fashion, and thereby reducing file size downloads.

At any time there may be many application and data files resident in the flash file system. The following process determines what application to run and also manages the removal of

- 29 -

redundant files.

The VM 321 when started in the VM task 221 first attempts to run a file with name "0" (i.e. the "init" application 611, as shown in Figure 9) typically stored in the write-protected area 410 of the flash file system 400. The init application 611 reads a configuration "init-tab" table file 621 which determines what file or files correspond to applications to be executed, the data files required for each application and the conditions under which the applications should each be executed (e.g. by using the pointers 631 to those files stored in the init-tab file 621). The init-tab file 621 is stored in the read-write area 420 so it can be regularly updated as the application and data files 622, 623, 624 stored in the flash memory 120 are updated. The init application 611 is also responsible for identifying and flagging files 625 that are candidates for deletion to help ensure the flash file system always has space for new files. Any files not identified in init-tab file 621 are delete candidates.

Event	Application	Data Files
A	12	2,3,7,8,9,13
	⋮	
A	6	2,3,5,7,8,9,10,11
A, B	1	2,3,4,5

Table 1

Table 1 provides an example of the data in an init-tab file 621 which refers to three applications, respectively stored in files 12, 6 and 1. The first column identifies the events from the event queue that may be used to trigger the running of the application file in the second column. The last column determines data files required by each application.

The init-tab file 621, as shown in Table 1, states that an initial version of the application governing the operation of the device 100 was placed in file "1", and used data files 2, 3, 4

- 30 -

and 5. After some time, a new version of the application was produced with expanded functionality, and transmitted. This new version of the application used all files that original version used, with the addition of new data files 7, 8, 9, 10 and 11. The new version of the application was placed in file 6. Finally, the latest version of the application
5 was created and transmitted to the device 100. The latest version of the application eliminated several of the old data files (4, 5, 10 and 11) and added new data file 13. The latest application file has been placed in file 12. The init-tab file 621 endeavours to run application file 12 if possible, otherwise application file 6 is run, or otherwise application file 1. An application file is only run if the init-tab file 621 is able to verify that all
10 required data files for execution of a particular application file are available.

The init-tab file may also be structured to identify additional files for each event that need not be present to run the corresponding application, but must not be identified as delete candidates. This allows optional files to be present to be used by the application, if desired,
15 without them being removed by the init application 611. If the space for new files or the number of free directory entries is below a certain threshold, file system resource may be recovered by removing one or more sets of files in init-tab corresponding to the lower priority events depending on the amount of resource to be recovered.

20 New data files 622, 623, 624 for an application may be received using receive process 1100 in Figure 15, and once received, the files are locked by the application (e.g. by setting a read-lock flag corresponding to each file). Such locking prevents that application from using other data files which are not part of the current version being used by the application (e.g. files for a previous version). The application may use data files for a
25 previous version by first unlocking the files for the current version being used, and then locking the files for the version selected for use.

The byte-code instructions of an application can be read into a small cache maintained in the RAM 108 to reduce the flash access overhead. The cache size being empirically set to
30 balance cache load overhead with cache misses. The VM implements a data stack to store the local temporary variables and parameters. The design of the byte-code instructions

- 31 -

used by the application files and interpreted by the VM are orthogonal to ease compiler design whilst minimising the number of instructions. They are grouped into 8-bit, 16-bit, 32-bit, floating-point groups and by number of operands for execution efficiency. A common set of operations is supported in all arithmetic groups (such as arithmetic, memory operations etc.), and additional operations may be supported for each group (e.g. bit tests for integers or for floating-point numbers etc.). Other operations, which are not part of the VM (such as trigonometric operations) are implemented, as desired, by function calls to specific files. Hardware or software resources are accessed using system calls, such as flash read/write. The function and system calls are executed using VM instructions involving unique function and system operational codes, respectively. A single operand is accepted to determine the operation to be performed, and any required parameters are passed via the data stack. The system calls are handled externally of the VM and this allows extensions that do not affect the core of the virtual machine. Calls to the flash file system, for example, allow directory information and file data information to be read.

15

An application can make a system call to the VM to reboot the device 100, if this is necessary for the application to now be run or take advantage of any new OS components sent to the device 100.

20 *Air Protocol*

The receiver 110 and the processor 102 utilise an air protocol to receive information including commands, time of day and file fragments. Information is transmitted using the air protocol by a server transmission system that has a server communications application to implement the air protocol and which can use an existing pager communications network infrastructure for transmission. The air task 214 controls communications between the processor 102 and the receiver 110. The air protocol is a communications protocol that, in accordance with the OSI layer model, has different attributes for a physical layer, a data link layer, a network layer and a transport layer

30

The device 100 can use an existing communications network infrastructure to receive

- 32 -

information, for example a paging network. Accordingly the link layer of the air protocol can be tailored to support transmission of text messages as defined by the Post Office Code Standardisation Advisory Group (POCSAG) standard.

- 5 The physical layer is simply the bit stream that is broadcast and received by the receiver 110. All the receivers 110 of different client devices 100 are normally tuned to the same carrier frequency for an application. This allows all of the client devices 100 to be updated with application files, as required. A unique ID or address for each device 100 (such as the serial number for the device 100), is available to filter the packets for an individual device
- 10 100. The device 100 is therefore identified by a paging address, and transmissions can be made for a number of paging addresses. The link layer supports a sub-set of the printable 7-bit characters, as discussed below.

- The network layer supports the transmission of multicast packet data. The packets include
- 15 address information to send the data to a single device address, a group of devices or all devices. A type parameter is used to identify the data carried, and the data types include commands, time and file segments, ie fragments. The network layer also enables data authentication and validation.

- 20 The transport layer supports the carriage of an arbitrary number of data octets of any length up to a maximum length supported by the link layer, which for the device 100 is 170 octets. The transport layer also provides error detection and optional correction so that corrupted packets can be identified and corrected.

- 25 For the transport layer, the data is pre-pended with its length in octets. The length is represented by an octet representing an unsigned 8-bit number. A CRC-32 is calculated from this and appended in little endian format as a quad-octet. The resulting message is then broken into quad-octets, padding with zeroes as necessary. Each quad-octet is converted to group of 5 ASCII characters using a base 85 encoding scheme, similar to
- 30 RFC1924, as discussed below.

- 33 -

The quad-octets are each treated as a little endian 32-bit unsigned number. To convert repeatedly they are divided by 85, taking the remainders at each step. The 5 remainders are placed in little endian order. They are each mapped to a printable character by using a mapping scheme that has monotonically increasing values, which allows for simpler decoding (ie without using a look-up table).

The resulting ASCII message is then padded out to a multiple of 20 characters by padding with '}'. A '}' is used because a group of 5 is an illegal BASE-85 encoding. This is to ensure that the message occupies a complete number of POCSAG codewords (each contains 20-bits, 20 characters = 7 codewords). The group is encoded also using the base 85 encoding scheme.

For the network layer, an air protocol packet has the following structure:

<packet> ::= <Protocol Id><payload>
 <Protocol Id> ::= <octet>
 <payload> depends on <protocol Id> which designates the protocol type.

Protocol Id has a value that identifies the protocol as follows.

<Protocol Id> = 1 allows sending of a data of a specified type. This can be sent to specified processors by the parameter <Processors> if the device 100 has more than one processor 102. More than one processor may be included in the device 100 in situations where a GPS receiver of the device 100 has a processor that is considered to be a separate processor for which files may need to be downloaded. Alternatively, an additional processor 102 may be included to be dedicated to a particular OS task to enhance the cycle time for that task and in turn enhance the speed of the device 100. Then

<payload> ::= <Address><Processors><Type><Data>

<Protocol Id> = 2 allows transmission of a specified data type to specified processors in a group of devices with a timestamp.

<payload> ::= <Address><Processors><Type><Time stamp><Data>

- 34 -

<Protocol Id> = 3 allows transmission of authenticated data to specified processors with specified type to a group of devices with a timestamp.

<payload> ::= <Address><Processors><Type><Time stamp><Authentication><Data>

5

<Protocol Id> = 4 allows transmission of a specified data type to specified processors in a group of devices with a timestamp and a local time offset.

<payload> ::= <Address><Processors><type><Time stamp><Local time offset><Data>

10 <Protocol Id> = 5 allows transmission of specified data type to a set of units in a specified product and platform range with timestamp and local time offset.

<payload> ::= <Product Id><Platform_range><Address><Type><Data>

<Protocol Id> = 6 allows transmission of specified data type to a set of units in a specified product and platform range with timestamp.

<payload> ::= <Product Id><Platform_range><Address><Type><Time

15 stamp><Data>

<Protocol Id> = 7 allows transmission of specified data type to a set of units in a specified product and platform range with timestamp and authentication.

<payload> ::= <Product Id><Platform_range><Address><Type><Time stamp><Authentication><Data>

20 <Protocol Id> = 8 allows transmission of specified data type to a set of units in a specified product and platform range with timestamp and local time offset.

<payload> ::= <Product Id><Platform_range><Address><type><Time stamp><Local time offset><Data>

25 For product Id:

The product Id allows the particular pager address (cap code) to support multiple products by identifying which product should receive the packet:

<Product Id> = <octet>

30 For platform range:

- 35 -

The platform range determines what releases of the platform should receive the packet by means of a range of acceptable platform build numbers:

<Platform_range> ::= <platform_min><platform_max>

<platform_min>,<platform_max> = <platform> = <octet_LSB><octet_MSB>

5

For addressing:

<address> ::= <broadcast>|<group>|<serial> = <octet><octet><octet><octet>

<broadcast> ::= 0xFF, 0xFF, 0xFF, 0xFF

<group> ::= 0xFF,0xFF, <group Id>

10 <group Id> ::= <octet_LSB><octet_MSB> = 0x0000 .. 0xFFFFE

<serial> ::= <octet_LSB><octet><octet><octet_MSB>

Specifies all devices, a particular device or a group of devices.

<Processors> ::= <octet> = <processor1:lsb>..<processor8:msb>

15 Specifies a set of processors from a maximum of 8.

Message types specifies how the system should handle the data being transmitted:

<type> ::= <octet> = <file fragment> | <set time> | <set groups> | <user data> | <user alarm>

20 < file fragment> ::= 1

<set time> ::= 2

<set groups> ::= 3

<user data> ::= 4

<user alarm> ::= 5

25

The actual message data transmitted:

<data> ::= {<octet>} | <file fragment> | <set time> | <set groups> | <user data> | <user alarm>

30 For a file fragment sent to a specified memory:

<file fragment> ::= <header>{<byte>}

- 36 -

<header> ::= <CRC32><date><length><data><erasure><number><access><name>
<size><count><correction>

<CRC32> ::= <LSB:octet><octet><octet><MSB:octet> = CRC32 of core file
contents

5

<date> ::= <time> = Release date

<length> ::= <LSB:octet><octet><octet><MSB:octet> = file length in bytes

<data> ::= <LSB:octet><MSB:octet> = length of all data quanta and erasure
correction quanta in bytes

10

<erasure> ::= <LSB:octet><MSB:octet> = length of all erasure correction quanta
only

<number> ::= <LSB:octet><MSB:octet> = first quanta number in this fragment

<access> ::= <section:LSN><storage:MSN>

15

<name> ::= <byte> = filename

<size> ::= <byte> = quanta size

<count> ::= <byte> = quanta count

<correction> ::= <nibble:E><nibble:c> = erasure correction, E = k and c = n-k

<section> ::= <FLASH_RO> | <FLASH_RW> = flash section type

20

<FLASH_RW> ::= 0 = Read/write section

<FLASH_RO> ::= 1 = Read only section

<storage> ::= 0 .. 15 = memory id to store fragment

The parameter set groups allow up to eight group addresses to be set for a device 100
25 identified by its 10, ie its serial number.

<set groups> ::= <group1:group> .. <group8:group>

<group> ::= <LSB:octet><MSB:octet>

The parameter set time sets time using a previous time reference to reduce error.

30

<set time> ::= <time reference><time correction><Local time offset>

<time reference> := <time>, previous timestamp in payload

- 37 -

<time correction> ::= <octet> = 0 .. 240 (seconds = 4 minutes)

The format of user data

<user data> ::= {<offset><length><data>}

- 5 <offset> ::= 16 bit unsigned
 <length> ::= 16 bit unsigned
 <data> ::= <length> octets.

For user alarms to be set:

- 10 <user alarm> ::= {<set ualarm> | <cancel ualarm>}
 <set ualarm> ::= 'S'<ualarm number><ualarm record>
 <cancel ualarm> ::= 'C' <ualarm number>
 <ualarm number> ::= <octet> = 0 .. 3 (device may use less than this)

A <ualarm record> is large enough for 4 <time>s however only 2 are defined:

- 15 And the rest may or may not be used for <time>s in the future, they may be just data.
 <ualarm record> ::= <ualarm time><ualarm useby><ualarm 0><ualarm 0>
 <ualarm time> ::= <time> (the alarm is to go off)
 <ualarm useby> ::= <time> (the time the alarm is not to be attempted after)
 <ualarm 0> ::= a zero of number of octets in <time>.

- 20 The time format, which defines days, minutes, seconds since 01-JAN-2000, GMT is:

<time> ::= <octet_LSB><octet><octet><octet_MSB>
 = <ddddddddddddddmmmmmmmmmmsssss>

days most significant 15 bits, minutes next 11 bits, seconds least significant 6 bits

- 25 <Time stamp> ::= <time>
 <Local time offset> ::= <octet>

specifies signed difference in half-hour increments between local and UTC time.

E.g. 20 = +10 hours

- 30 To supply an authentication keyword:

<authentication> ::= <octet_LSB><octet><octet><octet_MSB>

Command Processor

The command processor task 222 and associated module 322 implements a command line interface for a computer attached via the asynchronous serial interface 115.

For example a “LS” command would produce, as shown in the example Listing 1 below, first a header output and then a listing of the files of the device 100 written out one file per line. The fields in the lines correspond to the fields in the header, and what each field means is described in Table 2 below:

Field	Description
Na	Directory name, the name in the directory written out as a <byte2> since the file may be incomplete and so that is not printable.
N	Name of the file written as a <graphicchar>. If the name is not a graphic char then the “~” is written instead so that the columns stay correctly spaced. Not “~” is 0xxx so it may actually be the name of the file.
DN	Directory number of the file written out in <byte8>.
TY	File type written out in <byte8>.
Length	This is the length of the file in bytes and the value is written out in <byte8>.
CRC32	This is the CRC32 (32 bit CRC) of the file and its value is written out as <byte8>.
R Date	This is the Release Date of the file written out in <time>.
C Date	This is the Create Date of the file written out in <time>.
DQ	The number of Data Quanta for the file in <byte4>.
EQ	The number of Erasure Quanta for the file in <byte4>.
QS	The Quanta Size for the file in <byte2>.
QC	The Quanta Count for the file in <byte2>.
EC	The Erasure Correction byte for the file in <byte2>.
A	The Access for the file written as a <char>. A “W” means it is read and writeable on the media and “R” means that it is read only on the media.
FG	The flag byte which flags various information about the file. The value is written out as <byte2> and the value is the sum of various flags. Their values are: Deleted: 0x01 Incomplete: 0x02 Complete: 0x04 Released: 0x08

- 39 -

Field	Description
	Locked: 0x10 Open: 0x20

Table 2

An example file listing is shown below in Listing 1:

```

5  >ls
   Na N DN TY Length CRC32 R_Date C_Date DQ EQ QS QC EC A FG
   51 Q 00 00 000004E6 12B0DD16 0F922003 0F926EC0 000A 0000 80 01 10 W 38
   4D M 01 00 00002A4C D2E930D3 0F921FDC 0F926E95 0055 0000 80 01 10 W 38
   00 ~ 02 00 000004F6 FD16F9E8 0F92689E 0F926E91 000A 0000 80 01 10 W 08
10  31 1 03 00 00006950 76BED359 0F92689E 0F926E28 00D3 0000 80 01 10 W 38
   32 2 04 00 00006D08 CD49541C 0F92689E 0F926DC0 00DB 0000 80 01 10 W 08
   33 3 05 00 0000736A 23BCF15C 0F92689E 0F926D4E 00E7 0000 80 01 10 W 08
   3A : 06 00 00000B4B 44658709 0F92689E 0F926D43 0017 0000 80 01 10 W 38
   40 @ 07 00 0000001E 7A09C6EA 0F92689E 0F926D43 0001 0000 80 01 10 W 18
15  41 A 08 00 00002A0C 58EC9718 0F92689E 0F926D18 0055 0000 80 01 10 W 08
   5F _ 09 00 00000001 D202EF8D 0F92689E 0F926D17 0001 0000 80 01 10 W 18
   4B K 0A 00 0000EA60 944B18FD 0F9261E0 0F9468AE 01D5 0000 80 01 10 W 02
   >>ls Q
   Na N DN TY Length CRC32 R_Date C_Date DQ EQ QS QC EC A FG
20  51 Q 00 00 000004E6 12B0DD16 0F922003 0F926EC0 000A 0000 80 01 10 W 38

```

Listing 1

2. Transmission System Infrastructure

2.1 File Transmission

The files to be transmitted to the device 100 are each first broken up into a size that can be encapsulated by the air protocol and sent using a server transmission system, described below and shown in Figure 23. The file fragments produced in accordance with the air protocol can also be transmitted to the device 100 using a remote relay unit (RRU) 1900, as described below and shown in Figure 22.

The file to be transmitted 701, as shown in Figure 10, is treated as an arbitrary sequence of bytes, 710, up to the limit supported by the flash file system described above. Each byte, 711 is an 8-bit binary value. Consecutive sequences of bytes 721 are assembled into equal length groups of quanta 720. The last quanta 722 to be transmitted is padded with zeros, as

- 40 -

required, to ensure that the last quanta 722 has the same size as the other quanta (e.g. 720) being transmitted. Quanta are all equal length for a particular file and are numbered consecutively starting at zero.

- 5 A file fragment 730 suitable for transmission consists of a header 742 followed by data 743. The data is a sequence of quanta 720. The data in the last fragment is padded with zeros as required.

The header 742 contains information related to the data 743 in a file fragment 730, and
10 consists of a number of parameters (described as follows) transmitted in little endian order for compatibility with the representation used by the firmware compiler 2002. Each header 742 starts with a 32-bit cyclic redundancy check 731 calculated for the data 743 contained in the file fragment 730. This is followed by a 32-bit date/time stamp 732 for the file 701 in the format described in the air protocol section above, and a 32-bit length 733
15 representing the total length of the file 701 (e.g. in bytes), both of which may be derived from the file's 701 directory information. This is followed with a 16-bit count of the total number of data quanta 734 required to send the whole file 701 (i.e the core data quanta, or the data corresponding to the actual data in the file 701), and the total number of erasure correction quanta 735 generated for the file 701, as described below. A 16-bit quanta
20 number 736 identifies the number assigned to the data quanta that first appears in the current file fragment 730 being transmitted. If the quanta number 736 in the current file fragment 730 is greater than the number of data quanta required to send the whole file 701 (i.e. as represented by 734), then the data 743 in the fragment 730 is erasure correction data. An 8-bit access field 736 determines if the file 701 is to be placed into read/write 420
25 or write-protected 410 section of the file flash system 400, based on the value represented by the access field 737 (e.g. as "0" for read/write access and "1" for write-protected access). Files (e.g. 701) are given a single 8-bit name 737 (e.g. an 8-bit binary number) of which typically the last 255 values represented by the 8-bits are available for use. The number of bytes per quanta size and the number of quanta per fragment are represented by
30 8-bit values 738 and 739, respectively. The correction information 741 represents an 8-bit value and identifies the level of erasure correction used, the nibbles at the upper and/or

- 41 -

lower nibble locations, representing as k and $n-k$ respectively (where k is an integer with a predetermined value and n is the total number of fragments in the group). For a particular file all parameters in the header are held constant with exception of the quanta number 736 which corresponds to the file location of the data being transmitted.

5

File fragments may be transmitted in any order and there is a benefit to be gained by interleaving fragments to improve system performance if a number of consecutive fragments are lost by the transmission system. The system also allows a file fragment to be transmitted repetitively to increase the probability the complete file is received.

10

If desired, erasure correction fragments (ECFs) can be additionally generated from the fragments derived from the file data. ECFs allows lost file fragments to be recovered more efficiently than relying on multiple transmissions of the sample fragment. This redundancy greatly improves reception performance under adverse signal conditions. The level of redundancy is set on a file by file basis depending on the reliability desired based on transmission loss, number of fragments, etc.

15

Figure 11 is a schematic diagram showing how erasure correction fragments are generated. A set of k consecutive fragments (e.g. 810, 820), are used to generate one or more (ie $n-k$) erasure correction fragments 840 using an ECF process performed by an ECF module 830. In addition to the data in the sets of fragments 810, 820 received by the module 830, the module 830 accepts further input for determining the number of k groups of fragments 831 in each set and the desired number of erasure correction fragments 832 to be generated based on the groups of fragments in each set. The size of the erasure correction fragment 840 is the same size as the original fragments as received by the ECF module 830.

20

25

The ECF module 830 generates the desired erasure correction fragment using finite field arithmetic on a nibble-by-nibble basis. For example, lower nibbles from the same locations from each of the original fragments in each set (e.g. nibbles 811 to 821) are each used to generate a single corresponding erasure correction nibble (e.g. 841). Similarly, the upper nibbles at locations (e.g. 812 to 822) are used to generate a corresponding upper erasure

30

- 42 -

correction nibble (e.g, 842). A field size of 16 is therefore sufficient and means an efficient implementation of the finite field arithmetic using 8-bit arithmetic natively supported by the processor 150.

- 5 The ECF process may use a Lagrange polynomial function to generate erasure correction data. For example, each integer value of x (along the x -axis, as shown in Figure 12) corresponds to a nibble in their order of appearance in each set (e.g. 811 to 821). The value on the y -axis of Figure 12 at each x -location corresponds to a numeric value represented by each nibble respectively. The polynomial coefficients of the Lagrange
- 10 function are determined such that the function uses the data values of the nibbles at their equivalent x -locations 901, 902, 903, 904, as depicted in Figure 12 (i.e. the coefficients define an approximated function that passes through the y -axis value corresponding to each x -location 901, 902, 903, 904, etc.). The erasure correction data may also be generated by evaluating the polynomial at other points 905, 906, 907 as required. The
- 15 finite field arithmetic limits the sum of the number of erasure correction fragments and original fragments to less or equal to 16.

- In the client processor device 100, the Lagrange polynomial interpolation may be used in a similar fashion to retrieve any missing fragments provided that at least a unique k (as
- 20 defined by the correction information 741) of the original number of n transmitted file fragments 730 have been received.

2.2 *Web Server System*

- 25 A user is able to control one or more client devices 100, and these can be allocated to an identified group using a group Id, as discussed above for the air protocol.

A web server system 1005, as discussed below, is provided for users to access and enable them to:

- 30 (i) Create a group;
(ii) Change a group's details;

- 43 -

- (iii) Change devices in a group;
- (iv) Remove a group;
- (v) Configure a group of devices 100;
- (vi) Display data received from one or more devices 100.

5

2.3 *Transmission System*

The infrastructure includes a server transmission system, as shown in Figure 23, having a file collator 1005, an update manager 1006, a master database 1007, performance monitor
10 1008 and a wireless transmission apparatus 1010. The server transmission system also provides a customer web interface 1009 via a web server provided by the file collator 1005.

The file collator 1005 (or central transmission server) loads and maintains information that
15 controls the operation of the device 100, and to create and store instruction and data files used by that device. The data is stored in a relational database 1007, and data file suitable for transmission to the device are created by a file formatter 1020 and kept on disk in a currently transmittable file storage database 1021 from where they can be retrieved by an update scheduler 1025 as required. Instruction files used by the device 100 are also stored
20 in the currently transmittable file storage database 1021. The file collator 1005 includes a router application manager 1017, a configuration manager 1018, a file loader 1019 and a file formatter 1020, in addition to the currently transmittable files storage database 1021.

The master database 1007 interfaces with a billing module 1011, which enables the data
25 files for transmission to the devices to be updated with a range of billing data. For example, the master database 1007 may retrieve call tariff rates for different carriers, carrier rate structure (e.g. peak and off peak rates), carrier billing data and customer billing data via the billing module 1011 for storage in the master database 1007.

30 The router application manager 1017 generates files for transmission to the device 100. Modifications to the manner in which a task is performed by the device 100 are achieved

- 44 -

by sending a new version of the relevant file to the device 100. The file is written directly to the current transmittable file storage database 1021 from the file formatter 1020. From this database 1021, the update scheduler 1025 forming part of the update manager 1006, retrieves the file for transmission at an appropriate time.

5

The configuration manager 1018 enables manual changes to be made to the information maintained in the master database 1007. The configuration manager 1018 includes a number of screen layouts that can be used to review and change a different section of the database 1007. Each user of the system is issued a user name and password that
10 determines which screen the user can access. Screens that are used to maintain a particular client's information are accessible only to that client, using the Internet 1022, and a suitable web browser constituting the customer web interface 1009, and authorised company personnel. The configuration manager 1018 also interfaces with the internal database maintenance module 1032, which maintains the consistency of the data in the
15 master database 1007 and may provide the configuration manager 1018 with an updated view of the database 1007.

The configuration manager 1018 includes two sub modules (not shown in Figure 23), namely a transmission manager and a client web interface. The transmission manager
20 enables user control of the schedule for transmission of files maintained in the currently transmittable file storage database 1021 to the device 100. Files can be transmitted regularly either at a set interval or at a specific day and/or time, or a file can be transmitted on a one off basis. Files can be assigned a priority when transmitted. The client web interface enables users to maintain personal information and device configuration
25 information, as discussed previously.

The file loader 1019 provides a consistent process by which data contained in externally sourced files is loaded into the master database 1007 (via file formatter 1020), and is tested and released to production in a controlled manner. Given the differing nature of the
30 various files that need to be loaded, the file loader 1019 is modular in nature so that each type of file can be handled individually when necessary, whilst assuring the same overall

- 45 -

process is followed. Externally sourced files, such as structure data files 1023, and external entity information files 1024, are shown in Figure 23, and may be provided by an infrastructure provider or service provider. The information files may include costing information and other data.

5

The file loader 1019 ensures that the following steps are undertaken during the loading of a file:

- (1) Parsing. Errors in file formats are detected and handled.
- (2) Preliminary load. The data is first loaded into a “staging” area within the database so that it can be verified.
- 10 (3) Redundancy checks. New data is compared against data maintained in the master database 1007, if there is no difference, the new data is discarded. If differences are found however, the nature of the differences are gauged so that appropriate levels of testing can be performed.
- 15 (4) Testing. The new data is tested in an appropriate fashion.
- (5) Approval. Authorised personnel check the results of the previous stages and approve, or not, the new data as suitable for production.
- (6) Release. The new data is released to production and the file formatter 1020 is instructed to regenerate the affected files in the current transmittable file storage database 1021. This typically involves scheduling one or more files for
- 20 transmission to devices 100, depending upon the importance of the new data.

The file formatter 1020 provides a consistent framework for the generation of files that are required for operation of the device 100. Similar to the file loader 1019, the file formatter

25 1020 provides a consistent structure for the files whilst still allowing for the individual requirements of each type of file. The file formatter 1020 places the generated files in the current transmittable file storage database 1021 where the files are accessible by the process responsible for initiating the transmission of the files (e.g. the update scheduler 1025). The current transmittable file storage database 1021 forms part of the directory

30 hierarchy on a main server's disk which is accessible to all processes involved in the generation and transmission of files, namely the router application manager 1017, the file

- 46 -

formatter 1020, and the customer web interface 1009.

The update manager 1006 selects the appropriate transmissions medium, controls the packaging of data and schedule delivery of the data (via the selected transmission medium) to the wireless transmission apparatus 1010 and/or the Remote Relay Unit (RRU) 1900, for subsequent delivery to the device 100. The update manager includes an update scheduler 1025, an input interface 1026, and an erasure correction routine 1027, a transmission scheduler 1028, in-field monitors 1029, and a re-submitter and diagnostics system 1030, in addition to a transmission copies database 1031. The update scheduler 1025 is a software process that runs continuously reading the master database 1007 to retrieve the schedule for file transmission and sending requests to the input interface 1026 as required. The input interface 1026 is an executable program that extracts prioritised and formatted data from the current transmittable file storage database 1021. This data is passed directly through to the erasure correction routing 1027. As data is passed through this routine, it is prefixed with the appropriate device header information.

The erasure correction routine 1027 transforms the input data into a number of redundant blocks (or error correction fragments) that improve the reliability of reconstruction of the original data, provided a given number of the redundant blocks are received unchanged by the device 100. The redundant blocks may be generated using the process described above in relation to Figures 11 and 12. The redundant blocks are packaged and sent to the transmission scheduler 1028 individually. The transmission scheduler 1028 selects the appropriate medium for transmission (e.g. a wireless network used by the wireless transmission apparatus 1010). The transmission scheduler 1028 receives and stores packets passed from the erasure correction routine 1027, and periodically forwards the highest priority data packets to, for example, the wireless transmission apparatus 1010 for transmission via the selected transmission medium (e.g. a wireless network). A copy of the packets sent to the wireless transmission apparatus 1010 is written into the transmission copies database 1031 and removed from storage. The in-field monitors 1029 are modified versions of the client device 100, with modems attached to a dedicated link to, for example, the PSTN network. These monitors receive all data transmissions to the devices

- 47 -

100 in their locality. This data is forwarded to the resubmitted and diagnostics system 1030 via the PSTN network.

5 The resubmitter and diagnostics system 1030 includes multiple input lines from the PSTN network 1004, and executes a comparison process per input line and a single resubmit process. The comparison process for each input line compares data received by the corresponding in-field monitor 1029 with data written to the transmission copies database by the transmission scheduler 1028. The resubmit process performed by the resubmitter and diagnostics system 1030 periodically reviews the comparison results and resubmits to 10 the input interface 1026 any packets or files as required. The in-field monitors 1029 also transmit information for the purpose of assessing the performance of the selected transmission medium/network (e.g. by transmitting data via antennas 1098). All information of this nature is stored in the master database 1007 for analysis. The information is used on an ongoing basis to ensure that the selected transmission 15 medium/network performs at a required specification.

The file loader 1019 and configuration manager 1018 maintain the information stored in the master database 1007. The file loader 1019 performs routine loading of data contained in files with known formats. The configuration manager 1018 is used for manual entry of 20 data that is not available in files, and for manual adjustment of other data as required via an internal database maintenance module 1032.

The performance monitor 1008 is a group of processors that monitor the operation of the support systems on a day to day basis. The performance monitor 1008 provides a means of 25 visual inspection of resource usage levels, as well as standard reports at periodic intervals. It also issues warnings when the system becomes overloaded in any way. Components that are monitored in this way include:

- (1) Files transmission. The number of errors and hence retransmissions are monitored to determine transmission quality in the various areas covered by the network.
- 30 (2) Database host. System resources on the database host machine(s) such as CPU, disk and memory are monitored to ensure performance targets are met.

- 48 -

- (3) Database server. Inbuilt capabilities of the database service software are used to ensure that the database is running efficiently.
- (4) Web host. System resources on the web-server host machine(s) such as CPU, disk and memory are monitored to ensure performance targets are met.
- 5 (5) Web server. The inbuilt capabilities of web server software are used to ensure that the database is running efficiently.

The wireless transmission apparatus 1010 uses the air protocol and enables information to be efficiently distributed from a central location to a large number of devices 100 using,
10 for example, the existing pager communication network.

2.4 *The Remote Relay unit (RRU)*

The remote relay unit 1900, as shown in Figure 22, is a handheld device that can be used
15 by service personnel or a user of the device 100 to transmit file fragments to the device 100. The RRU 1900 is particularly useful in situations where the wireless communications network, eg the pager network, is unable to communicate with the device 100, for example if coverage of the network does not extend to the location of the device 100. The RRU 1900 is able to obtain the file fragments from the update manager 1006 using online
20 connection methods, or by direct connection, to the transmission system 1910 (e.g. via the data network 1099 as shown in Figure 23). Once the RRU has received the file fragments, the fragments can be sent to the device 100 using a short range wireless communications protocol, such as Bluetooth or ZigBee (<http://www.zigbee.org>), or an infrared protocol, such as one that complies with the standards of the Infrared Data Association (IrDA)
25 (<http://www.irda.org>). One RRU 1900 may be used to communicate and configure a number of devices 100. The RRU 1900 may also be used to program and/or control other items in a user's household, such a VCR, DVD, TV, lights etc, either autonomously or by user intervention. The control of these items, together with the devices 100, can be based on data dynamically available and updated from the central transmission system 1910.
30 The data and file fragments can be delivered to the RRU 1900 based on user preferences

- 49 -

stored as part of the transmission system, or based on direct instructions from the user, eg by interrogating the transmission system over the Internet.

Th RRU 1900 is also configured to enable information to be extracted from local units, such as the household items and the devices 100, so as to obtain diagnostic status, statistic and event log data. A wireless interface, such as an IrDA or ZigBee interface, can be used for retrieval of the information, but again may be under the direct control of the user or performed autonomously when a connection can be established. Information retrieved from the local unit is stored on the RRU 1900 and then made available to be sent to the central transmission system 1910 or for upload to a user's computer device, e.g. personal computer system or PDA. The computer device is able to access the RRU 1900 using an IP connection.

The RRU 1900 includes a communications module 1902 that enables the RRU 1900 to connect to the transmission system using at least one of the following: GPRS; CMDA-1X; Ethernet; WLAN; or IP Dial-up. The connection normally would be always online. If not, the connection is either established by either the central server 1005 or the RRU 1900 on demand. The central server 1005 would typically use the IP based protocols for connection, but a simplex download connection to the RRU 1900 would suffice.

20

The RRU 1900 includes a flash file system 1904 that allows information to be received locally in addition to being downloaded from the central server transmission system 1005. The file system 1904 is scaled appropriately to cater for the number and size of files required. The size of this file system 1904 may directly relate to the number of local units 100 associated with it.

25

The RRU 1900 has a second communications module 1903 to allow information to be distributed to its associated local units, ie the devices 100, and support retrieval of information from the local units. The second module 1903 provides the wireless communications interfaces, such as ZigBee, IrDA, and/or pager (POCSAG) interface.

30

- 50 -

The RRU 1900 has a microcontroller system 1905 that maintains a real-time clock for synchronising the time in the associated local units when required. The clock is set using a network time protocol from the server transmission system 1005. The system 1905 includes a user display that can be used to show status information and to facilitate the user interface. The system 1905 also has buttons and/or a keypad to provide user interaction/control.

A RRU 1900 can be associated with a number of user local units that include third-party consumer devices (e.g. VCRs, TVs) in addition to the client processor devices 100. The user registers the make and model number of each of the consumer devices on the central server transmission system 1005 via the web interface 1009. Depending on support for the third party devices, configuration information is generated on the central server 1005 and sent to the RRU 1900 to allow the devices to be controlled. The serial number of each of the client processor devices 100 is similarly registered on the web interface 1009. The central server 1005 controls the number and type of associated devices 100 depending on the resources (power, memory and communications bandwidth) associated with the remote relay unit 1900. A particular device 100 is normally associated with a single remote relay unit 1900.

In order that a relay unit 1900 can communicate with a particular device 100 it may, depending on the transmission method, need to be in the vicinity or close line-of-sight of the device. This means the RRU 1900 is normally a portable device similar to a mobile phone. It is returned to a cradle 1906 to recharge and/or make a wired connection to server transmission system 1005, if required.

25

In order that the client processor devices 100 can communicate with the relay unit 1900, they are provided with the I/O module 130. This provides a corresponding communications interface 1908, such as ZigBee or IrDA. These interfaces support relatively high-speed bidirectional communication and this ensures that the data is transferred in a reliable and timely manner. The device 100 monitors the alternate interface 1908 and communications over this interface 1908 would normally override that

30

- 51 -

received on the pager receiver 110.

After the RRU 1900 and associated local units are registered, the file system 1904 is then able to accept updates from the central server 1005. The user may chose to make changes (i.e. configuration or otherwise) using the web interface 1009 or new data may be available from the server 1005 depending on the service(s) the user is subscribed to. Subscription data indicating the services subscribed by a user may be maintained, for example, on the central server 1005. If new data is available the central transmission server 1005 checks the online status of user's the RRU 1900 and indicates this status to the user during the time the user is using the web interface. Once the changes are sufficiently completed the server 1005 attempts to download the data to the remote unit 1900. Once the data is fully downloaded, the RRU 1900 indicates this as fresh data available (e.g. as a flashing light or textual display) and then proceeds to automatically transfer the data to the associated local units. If the RRU 1900 cannot be contacted from the server 1005, the user needs to take action such as manually starting the connection by returning the RRU 1900 to the cradle 1906. If the RRU 1900 has data for an associated client processor device 100 but cannot connect to it, the user is instructed to take appropriate action (i.e. to bring the remote relay unit 1900 to the client processor device 100). The client processor device 100 may similarly indicate that it has received fresh data (e.g. by flashing a light).

20

3. Telecommunications Signal Carrier Selection

The client processor device 100, when used as part of a call management system 1001 acts as a customer interface device 1002, as shown in Figure 24, for connection between a telephone or other local telecommunications device 1003 and a telecommunications network 1004, such as a Public Switched Telephony Network (PSTN). The call management system includes the server transmission system having a file collator 1005, an update manager 1006, a master database 1007, performance monitor 1008 and a wireless transmission apparatus 1010. The server transmission system also provides the customer web interface 1009.

30

- 52 -

The client device 1002 is the physical and logical component of the call management system that enables a user of local telecommunications device 1003 to gain real time access to the most cost effective telecommunications signal carrier available at any given time. The device 1002 has a line interface 1016 that operates as an I/O unit 130 and is
5 connected between the telecommunications device 1003 and the network 1004, as shown in Figure 24.

As shown in Figure 24, the line interface 1016 of the client processor device 1002 communicates with the processor core 102 of the microcontroller 150 on a data and
10 address bus 114 and a control bus 2221. The line interface 1016 receives control signals generated by the processor core 102 from the control bus 2221, which triggers the line interface 1016 to send or receive data via the data and address bus 114. The client processor device 1002 also includes an alpha numeric display 1033 for the user to view current information available, and the display 1033 communicates with processor core 102
15 via the data and address bus 114.

The virtual machine VM of the client device 1002 executes an application (e.g. 2330 as shown in Figure 25) transmitted from the central server 1005 to perform the selection of an optimal telecommunications signal carrier by controlling the line interface 1016 at the time
20 when a call is desired to be placed from the local telecommunications device 1003. For example, the VM restarts the application whenever the telecommunications device is detected as going "off hook".

The line interface 1016 comprises electronic circuitry interposed between the local
25 telecommunications device 1003 (such as a telephone, facsimile machine, modem etc) and a telephony exchange, PABX system or other network mode connecting the local telecommunications device 1003 to the telecommunications network 1004. The line interface 1016 detects the "of hook" or line looping state, either of the local telecommunications device 1003 or of other equipment if the phone line is in use
30 elsewhere. The line interface 1016 detects and decodes DTMF digits dialled by a user from the local telecommunications device 1003. The line interface device 1016

- 53 -

additionally isolates the audio path on the telephone line from the local telecommunications device 1003 when instructed to by the virtual machine, so as to be able to intercept and prevent dialled digits from reaching the telephone exchange, PABX or other network node that would ordinarily receive the call. The line interface device
5 1016 additionally dials DTMF digits when instructed to do so by the virtual machine executing in the processor core 102, in order to make calls autonomously or whilst the local telecommunications device 1003 is off-hook and isolated from the telephony exchange or PABX.

10 The master database 1007 is used to record all structure information used to control operation of the call management system.

Typically, the following information is stored in the database 1007.

- (1) Call rates for each carrier.
- 15 (2) Carrier details.
- (3) Exchange codes.
- (4) Customer details and preferences.
- (5) Carrier rate structure (for example, peak and off peak rates)
- (6) Field monitor locations and connection details.
- 20 (7) File transmission states
- (8) File transmission schedules.
- (9) Customer billing data.
- (10) Carrier billing data.

25 While the device 1002 shown in Figure 23 enables the interconnection of the local telecommunications device 1003 with a PSTN network 1004, the client processor device 1002 may further enable connection of the local telecommunications device 1003 to a remote telecommunications device via a selected telecommunications signal carrier within a same telecommunications network (such as the PSTN network 1004), or across two or
30 more different telecommunications networks (such as a mobile, ISDN, frame, Internet, satellite, voice-over-ip, or internal network). Accordingly, the carrier identifiers, variable

- 54 -

costing information and other data maintained in the memory 120 of the client processor device 1002 (as updated from the central server 1005) may relate to a single telecommunications network, or to different telecommunications signal carriers for different types of telecommunications networks (i.e. telecommunications networks that conform to differing protocols). Both types of client processor devices have the same hardware and logical architectures as shown in Figure 24 and 25, respectively. Additional tasks are included for sue by the application files 2330, such as a line task 2321 for controlling the line interface 1016 to connect calls over one or more different telecommunications networks.

10

In operation, the device 1002 performs the following basic steps whenever a call is attempted to be established from the local telecommunications device 1003 to a remote telecommunications device:

- 15 (1) Determine any parameters supplied by the user for the call, such as the maximum/minimum length of call or call cost. These parameters come from the user at the time the call is made, for example, by entering a special key sequence.
- (2) Determine the customer's general parameters such as call blocking, average call time or preselected carrier. These parameters are stored in the information storage system.
- 20 (3) Read any key pressed by the user until the nature of the call, for example local, neighbourhood, national, international, etc, can be determined.
- (4) Determine the destination of the call – this process may differ depending upon the nature of the call cost for each carrier available to the user.
- (5) Determine the cheapest carrier for use in placing the call.
- 25 (6) Send an override code which will select the cheapest carrier to the line interface 1016, then send the units dialled so far and place the device 1002 in a stage which will send all further digits straight through to the line interface 1016 so the number can be completed (if it has not already).

30 The above described operation relies upon detecting the digits or other identifiers entered by a user at the local telecommunications device 1002, and the comparison of these digits

- 55 -

(as they are dialled in real time) with data maintained in the information storage system 120. The data maintained in the information storage system 120 indicates the identity of various telecommunications signal carriers, the variable costing information for each carrier, and the location of telephone exchanges or other billing calculation nodes within one or more telecommunications networks.

Figure 26 is a flow diagram providing a high level overview of the process 7600 performed by the device 1002 for selecting an optimal telecommunications signal carrier. Process 7600 begins at step 1060, where a connection is established between the local telecommunications device 1003 and the PSTN 1004 or other telecommunications network to which the local device 1003 is to be connected, and the use has not yet entered a remote telecommunications device identifier (such as dialled numerical string of digits). At step 1061, a remote telecommunications device identifier is entered from the local telecommunications device 1003 by a user. As each digit is entered, an attempt is made to resolve the sequence of digits thus far entered until a corresponding telephone exchange or other network billing node is able to be identified at step 1062. If it is determined at step 1063 that resolution of the entered string of digits was possible and a billing calculation node is able to be identified, then target data for that node is retrieved from the memory system 120 in step 1064. The telecommunications signal carrier providing the cheapest call cost is then determined in step 1065. At step 1066, a determination is made as to whether the device 1002 is adapted to route the call from the local telecommunications device 1003 via multiple types of networks (Figure 32). If this is the case, each of these interface options are examined at step 1067 to determine whether a cheaper call may be obtained via these interface options.

25

If the interface options have not been installed, a carrier identifier string corresponding to the selected optimal telecommunications signal carrier is then transmitted to a telephone exchange associated with that carrier, together with the string of digits already composed by the user at the local telecommunications device, at step 1068. At step 1069, further digits required to identify the remote telecommunications device with which a connection is to be established (e.g. an area code or country code) are transmitted directly to the

30

- 56 -

telephone exchange associated with the selected carrier. The process 7600 is terminated at step 1070.

5 If it was determined at step 1063 that resolution of the string of digits entered at the local telecommunications device by a user to identify a telephone exchange or other billing calculation node is not possible, the digits received are transmitted at step 1064a the telephone exchange of the default telecommunications carrier, and then step 1069 is carried out.

10 Figure 20 shows a flow diagram of the process 7700 for examining the application files stored in the memory 120 of the client processor device (e.g. 1002) for validity prior to execution. Process 7700 begins at step 1071. Once the device 1002 is initiated, step 1071 proceeds to step 1071a, where the next application file in the init-tab file maintained in the memory system 120 is read and a determination is made as to whether that application file
15 is valid. If the application file is not valid, process 7700 finishes at step 1074. If the application files are valid, then at step 1072 the VM determines using the init-tab file whether all data files required for execution of that latest application file are valid and complete. If this is the case, then step 1072 proceeds to step 1073 to execute the current application, and then process 7700 finished at step 1074. However, if step 1072
20 determines that the set of data files is not complete or valid for a particular current application file, then step 1072 proceeds to step 1071a, where the next most recent application file is examined and a determination made as to whether the data files required for execution of the next version of the application file are current and valid.

25 In the case of a PSTN network, call costing is determined on the basis of the geographical location of the telephony exchange to which the local telecommunications device 1003 is directly connected. In Australia, customer telephone numbers are grouped into exchange service areas, which are in turn grouped into zones. Zones are grouped to form charging districts and call costs are calculated based on the distance between the zone and/or district
30 charging point of the A-party number and the B-party number. The chargeable distance

- 57 -

for a call is based on the distance between the zones or districts the parties are in, rather than the actual distance between the callers. That is, each zone or district has just one centrally located geographically point for measurement of the distance to other zones or districts. Accordingly, step 1062 in Figure 26 requires the resolution of the digits entered
5 by the user of the local telecommunications device 1003 in order to enable identification of the telephony exchange to which the local telecommunications device 1003 is directly connected. A determination can then be made of the district or zone within which that telephony exchange is located in order to derive a call distance upon which a call charge may be calculated. As each digit is dialled by the user, an attempt is made to resolve the
10 string of digits already entered so that the resolution to identify the telephone exchange in question, and hence determine the corresponding billing calculation location, without requiring entry of the entire string of digits identifying the remote telecommunications device.

15 In Australia, there are some 4,500 telephony exchanges, grouped into various precincts, zones and districts. There are only a small number of precincts (currently 118) but over 2,000 zones in which the telephony exchanges may be located. Each telephony exchange is identified by an exchange code. The exchange code forms part of the telephone number identifying the telecommunications devices connected to the network. For example, in
20 Australia the telephone number +61-03-9851-0000 includes an international access code (represented here by the symbol "+"), the country code 61 for Australia, the state (or area) code of (0)3 for Victoria, the exchange code 9851 and a terminal code 0000 to identify the communications device connected to each exchange. An exchange code file is maintained in the information storage system 120. The exchange code file contains a series of
25 encoded exchange codes along with the identification of the exchange for each code. The exchange code resolution process determines whether a given string of digits identifies a specific exchange and if so, finds the identifying details for that exchange.

The exchange code file is not a simple list of exchange codes, but includes several space
30 optimisations. The exchange code file takes advantage of the fact that assorted series of

- 58 -

codes can be more efficiently expressed as a series of differences to the preceding code. Each code entry in the exchange code file therefore consists of a number of characters to retain from the previous code, and the new characters to add to that code. Ranges of codes are compressed into a single entry since all codes consist of decimal digits only, pairs of
 5 digits can always be expressed as a single byte using four bits for each digit.

For example, the series of codes 0234, 0235, 02351, 02353, 02354, 02355, 02366 and 02367 if stored in an uncompressed form, would require $4+4+5+5+5+5+5+5 = 38$ bytes. If differences only are considered, the following stored codes can be used:

10

0, 0234	Keep none, add 4 = 0234
3,5	Keep 3 (023), add 5 = 0235
4,1	Keep 4 (0235), add 1 = 02351
4,3	Keep 4 (0235), add 3 = 02353
15 4,4	Keep 4 (0235), add 4 = 02354
4,5	Keep 4 (0235), add 5 = 02355
3,66	Keep 3 (023), add 66 = 02366
4,7	Keep 4 (0236), add 7 = 02367

20 This requires only $5 + (6 * 2) + 3$ bytes = 20.

Considering ranges, where a range is a series of codes with successive values that belong to the same exchange and have the same length, gives:

25 0,0234-5	Keep none, add 0234, run the last digit to 5 = 0235
4,1-1	Keep 4 (0235), add 1 = 02351, no range
4,3-5	Keep 4 (0235), add 3 = 02353, run the last digit to 5 = 02354, 02355
3,66-7	Keep 3 (023), add 66 = 02366, run the last digit to 7 = 02367

30 This requires only 16 bytes.

- 59 -

The space can be halved again by considering the above sequence in 2 character pairs:

00 23 45 41 14 35 36 67

- 5 These pairs can be represented as hexadecimal byte values, eg 00 = 0, 23 = 35, 45 = 69, etc, and one byte written for each pair.

The data section of the exchange file is essentially made up of a sequence of blocks with the same internal structure. That structure is as follows:

- 10 Code Block: <Keep><New Digits><Range Digit><Exchange Type><Exchange ID>, where

<Keep>: Number of digits to keep from the previous entry, first entry in file must have Keep = 0.

<New Digits>: Sequence of digits to add to those kept from the previous entry.

- 15 <Range Digit>: If non-zero, signifies that entry should match all the codes between the code derived from Kept + New and Kept + Range_End, where Range_End is New with the last digit replaced by the Range Digit.

<Exchange Type>: Indicates whether this exchange is a Zone, Precinct or Other (eg ISD, emergency, mobile, etc).

- 20 <Exchange ID>: The index of the exchange within its type; if the Type is Zone, the Exchange ID will be a Zone Index.

An example of the exchange code file is shown in Table 2:

Exchange Code	Type	ID
000	Other	1
0011	Other	2
03980	Zone	1
039802	Zone	2
03985	Precinct	1

- 60 -

03986	Precinct	1
03987	Precinct	1
0399	Zone	3
03992	Zone	4
03993	Zone	4
0419	Other	3
131	Other	4

Table 2

In order to simplify the example, the final list of codes is not converted to hexadecimal byte values. The Exchange Type values used is 1 for Other, 2 for Zone and 3 for Precinct.

- 5 The data set shown in Table 2 is used to generate the following data string:

0000011211012139800215201245731390234232414190130131014

Splitting this string into code blocks:

0000011 211012 13980021 52012 45731 39023 42324 1419013 0131014

Tagging each part of each block with K for Keep, N for New digits, R for Range

- 10 Digit, T for Exchange Type and I for Exchange ID:

0000011 211012 13980021 52012 45731 39023 42324 1419013 0131014

KNNNRTI KNNRTI KNNNNRTI KNRTI KNRTI KNRTI KNRTI KNNNRTI
KNNNRTI

- 15 The data is processed from the beginning from a block with a Keep value of 0, since all other blocks depend on the preceding blocks back to the most recent Keep=0 block.

- 20 The device 1002 uses the exchange code file to determine the exchange code applicable to a string of digits (the "Target") entered by the user. Each code block is read in turn, the complete exchange code derived for that entry by combining the previous code with the Keep and New Digits values for the new block. The target string is then compared to the file exchange code, possibly including additional file codes as dictated by the Range Digit. There are four possible outcomes from this comparison:

- 61 -

- (1) Continue; the point in the file has not been reached where the target string should occur, and the next block.
- (2) Incomplete; read past the point where the target string should have occurred, but adding more digits to the target string and re-starting the search may result in an exact match.
- (3) Valid; found an exact match and should use the exchange data for the current block
- (4) Invalid; read past the point where the target string should have occurred, and the target string cannot be made to match by adding to digits to it.

10

In practice, the client processor device 1002 attempts to resolve the exchange after each DTMF digit is received. Each attempt will produce one of outcomes 2,3 and 4 above; code resolution will continue to take place until one of outcomes 3 and 4 arise.

- 15 For an example, consider the phone number (03) 9851 0000. As this number is dialled, resolution will be attempted after each new digit, using all the current digits each time. The results from these attempts are shown in Table 3 below:

New digit	Target string	Outcome	Last block read	Last exchange code
0	0	Incomplete	000011	000
3	03	Incomplete	13980021	03980
9	039	Incomplete	13980021	03980
8	0398	Incomplete	13980021	03980
5	03985	Valid	45731	03985-7

Table 3

20

- The pattern of a successful resolution is shown above. Repeated "Incomplete" outcomes, where the target string is a subset of the exchange code of the last block read. The outcome of the comparison between the target and all the codes preceding the final one would have been "Continue", hence the next block was read. An exit is made after these final blocks,
- 25 since the file exchange code is greater than the target code, but the potential for a match is

- 62 -

always there (since the target string and the last exchange code do match over the length of the target), and does eventually occur (outcome “Valid”).

For a second example, consider the outcomes for phone number 0398810000.

5

New digit	Target string	Outcome	Last block read	Last exchange code
0	0	Incomplete	000011	000
3	03	Incomplete	13980021	03980
9	039	Incomplete	13980021	03980
8	0398	Incomplete	13980021	03980
8	03988	Invalid	39023	0399

Table 4

As shown in Table 4, the outcomes for the first 4 digits are the same, as they would be. After the fifth digit is entered, the final block comparison returns “Invalid”, as there is no way a match for a target starting with “03988” will occur if one has not been found before file code “0399”, regardless of what subsequent digits may be entered.

10

The application automatically adds the local state code (eg 03 for Victoria) to numbers that do not already begin with 0 or 1 before attempting exchange code resolution.

15

Figure 28 shows a flow diagram of the process 7800 for resolving the exchange code as described above. Process 7800 processes a series of digits as currently entered by the user (defined as the “target” string). Each time a user enters a new digit, the target string is updated to include the new digit, and process 7800 processes the updated target string again starting from step 1080. Process 7800 begins at step 1080 by accessing the target string of digits entered by the user. Step 1080 proceeds to step 1081 to attempt to retrieve an entry from the exchange code file, the entry corresponding to a block of code (as described above) that represents an exchange code. If an entry cannot be found in the exchange code file at step 1081, process 7800 ends by determining that the target string is

20

- 63 -

invalid. Otherwise, if an entry is found, step 1081 proceeds to step 1082.

Step 1082 generates an exchange code based on the code for the current entry retrieved at step 1081, and determines whether the exchange code for the current entry is past the point where the target string should have occurred. As shown in Table 2, the entries in the exchange code file are arranged in descending order based on the exchange code value for each entry. Step 1082 determines that the exchange code for the current entry is past the point where the target string should have been if the exchange code for the current entry includes other digits in addition to all the digits in the target string. For example, where a target string includes only the digits "03", a current entry having a corresponding exchange code of "03980" is past the point where the target string should have occurred in the descending order of exchange code values in the exchange code file. If the exchange code for the current entry has not passed the point where the target string should have occurred, step 1802 proceeds to step 1803 to determine whether the digits in the target string exactly matches the exchange code for the current entry. If so, the data characterising the exchange identified as corresponding to the current entry is retrieved at step 1084 (e.g. by using process 7900 as shown in Figure 28), and process 7800 finishes by determining that the target string are valid.

However, if there is no match at step 1083, step 1083 proceeds to step 1081 to attempt to retrieve the next entry in the exchange code file as the new current entry, which corresponds to another exchange code. If the next code is unable to be read at step 1081, process 7800 ends and the series of digits entered by the user is determined to be invalid. If it is determined at step 1082 that the exchange code corresponding to the current entry has past the point where the target string should have been in the exchange code file, determination is then made at step 1085 as to whether additional digits in the target string would assist in resolving the code. For example, step 1085 assesses whether the exchange code for the current entry includes additional digits which do not present exist in the target string. If so, step 1085 determines that additional digits in the target string might be helpful to resolve the exchange code, and process 7800 ends by determining that the digits entered by the user is incomplete, and process 7800 starts again on receiving an updated

- 64 -

target string that includes a new digit. However, if step 1085 determines that this is not the case, process 7800 ends and the digits entered by the user are determined to be invalid.

Once the exchange code corresponding to the exchange to which the remote
5 telecommunications device is connected has been thus identified (e.g. at step 1064 in
process 7600 and step 1084 in process 7800), data corresponding to that exchange is
retrieved by the virtual machine from the memory system 120. An exemplary file for an
exchange generated from the retrieved data may be as follows:

Zones:

10 Aldersyde

Exchange codes: 089600(0-2), 089642(6-7)

Adjoining zones: Brookton, Bullaring, Pingelly, Pingelly East, South Quairading,
Yealering

District: York

15 Charge Point: 32<B0>25'30"S 117<B0>18'30"E

Community access: Narrogin

Alectown

Exchange codes: 0268005, 0268613, 026865

20 Adjoining zones: Baldry, Bindogundra, Bogan Gate, Bruie Plains, Parkes, Peak
Hill

District: Parkes

Charge Point: 32<B0>56'30"S 148<B0>13'0 "E

25 From this file, the charge point from which the chargeable distance for the call will be
calculated is identified.

The exchange code to which the local telecommunications device 1003 is directly
connected is identified in an identical manner to the process described above, and the
charging point of the zone, district, or precinct within which that telephony exchange is
30 located similarly determined. Comparison of the charging points in the zones, districts
within which the telephony exchanges to which the local and remote telecommunications

- 65 -

devices are directly connected are then used to determine the chargeable distance for the call.

The manner in which the file loader 1019 extracts data from the exchange structure data
5 1023 in order to derive the one or more files, each categorising a different telephony
exchange, is illustrated by process 7900 as shown in Figure 28. The file loader 1019 is
responsible for extracting data from one or more files (e.g. flat text files or HTML files)
provided by telecommunications infrastructure providers, and inserting the extracted data
into the database 1007. Examples of information contained in these files includes
10 exchange data, carrier call rates and lists of current international dialling codes. Carrier
information 1024 is also extracted in this same manner.

Process 7900 begins at step 1090, where the file loader process 1019 is initiated by
receiving a file from a telecommunications infrastructure provider. At step 1091, a first
15 line of data is extracted from the relevant file from the telecommunications provider. If it
is determined at step 1092 that the data being examined is district data, then that data is
formatted at step 1093 into district data. Similarly, if it is determined at step 1094 that the
data characterises the zones adjoining the zone or other region under consideration, that
data is extracted and formatted at step 1095 as zone data. A similar process is followed at
20 steps 1096 to 1101 for charge point data, community access data and any other types of
data in the file. Steps 1093, 1095, 1097, 1099 and 1101 all proceed to step 1091 to select
the next line of data from the file. When step 1102 determines that the end of file has been
reached, the formatted data is loaded into the database at step 1103, and the file loader
process 7900 finishes at step 1104. Otherwise, if step 1102 determines that the end of the
25 file has not been reached, step 1102 proceeds to step 1091 to select the next line of data
from the file.

Once the chargeable distance for the call in question has been determined, the carrier
providing the cheapest call cost for that chargeable distance is then selected. Figure 30 is a
30 flow diagram of the carrier selection process 8000 for achieving this. Process 8000 begins
at step 1121 by determining the current date and time (e.g. by querying the clock task 225

- 66 -

of the client processor 1002 for the current time). Step 1121 proceeds to step 1122 to retrieve data for a particular telecommunications carrier (e.g. the data generated by process 7900 that is stored in the memory 120 of the client device 1002 by way of update data transmissions from the update manager 1006). For example, step 1122 retrieves the call
5 tariffs at different time bands (e.g. peak and off peak), the distance segment, and the rate for the time band and distance.

At step 1123, the current time determined at step 1121 is used to search the data retrieved at step 1122 for a particular carrier to determine the call rates applicable for the current
10 time band. Then, at step 1124, the distance segment of the call for the carrier is determined (e.g. based on the charge points as described above). Step 1125 then determines the carrier's rate based on the time band and distance determined at steps 1123 and 1124 respectively. The rate determined at step 1125 is used to generate the call cost at step 1126. Step 1127 stores the call cost generated at step 1126 for each carrier, and
15 compares the call cost generated for the carrier currently being analysed with the stored call cost data. Based on this comparison, step 1127 determines whether the call cost generated for the carrier currently under analysis is the best (i.e. the lowest) cost determined so far. If this is the best cost yet, as determined at step 1127, the carrier is recorded at step 1128 as being the best or optimal carrier. If this is not the case, step 1127
20 proceeds to step 1122 to retrieve data for another carrier, and process 8000 repeats again until the data for all the carriers (as stored in the database 1007 and memory 120) have been processed. Once no more carriers are available for processing, any digits subsequently dialled by the user are passed directly through to the telephony exchange associated with that optimal carrier at step 1129. The carrier selection process is
25 terminated at step 1130.

Figure 30 illustrates the steps of the process 8100 performed by the client processor device 1002 when an incoming call is received, and redirected to a remote telecommunications device 1050 (as shown in Figure 23) associated with the owner of the local
30 telecommunications device 1003. In this example, the device 1050 enables connection of the local telecommunications device 1003 to the remote telecommunications device 1050

- 67 -

via multiple telecommunications networks, such as a mobile telecommunications network or a Voice-Over-IP (VoIP) network. Process 8100 starts at step 1140, where an incoming call is received at the local telecommunications device 1003. The client processor device 1002 then acts to redirect that incoming call to a remote telecommunications device 1050
5 associated with the owner of the local telecommunications device 1003. Initially, at step 1141, the customer interface device 1050 determines the called party (or B-party) number to which the incoming call has to be redirected. At step 1142, the digits forming part of the called party (or B-party) number are analysed, in the manner described in relation to Figure 28, in order to resolve the telecommunications exchange or other telephony node to
10 which the call is to be routed.

At step 1143, a determination is firstly made as to whether the device 1002 includes interfaces to multiple telecommunications networks. A determination is also made as to whether the B-party number is supported by any of the telecommunications networks to
15 which the customer interface device is able to be connected. The device 1002 also determines (in a manner similar to that described in relation to Figure 30) whether the various alternate interfaces providing access to telecommunications networks other than the conventional PSTN network 1004 will provide a cheaper call establishment cost.

20 If the device is determined in step 1143 as not supporting multiple telecommunications networks, a carrier identifier string corresponding to the optimal selected telecommunications signal carrier within the PSTN network 1004 is then transmitted to telephone exchange associated with that carrier, together a string of digits from the calling party (A-party) number, at step 1144.

25

At step 1145, further digits required to identify the remote telecommunications device to which the call is to be redirected are transmitted directly to the telephone exchange associated with the selected carrier. This process is terminated at step 1146.

- 68 -

- However, if it is determined at step 1143 that multiple interfaces are available, then a determination is made at step 1147 as to whether the selected optimal interface is associated with a mobile network. If this is the case, a connection is established at step 1148 to the mobile interface, then, at step 1149, the B-party number is transmitted to the mobile network for call completion. At step 1150, further digits required to identify the remote telecommunications device to which the call is to be routed are then transmitted directly to an exchange within the mobile network. This process 8100 is terminated at step 1150.
- 10 If step 1147 determines that the interface associated with the telecommunications network selected at step 1143 was not associated with a mobile network, step 1147 proceeds to step 1152 to determine whether the interface associated with the telecommunications network selected at step 1143 was associated with another type of interface (e.g. a GPRS network). If so, step 1152 proceeds to step 1153, where the appropriate sequence of steps are performed by the device 1002, depending upon the telecommunications network selected. For example, this may involve performing steps similar to steps 1148 and 1149 by connecting to the other chosen interface (e.g. a GPRS interface) and sending identification information to that interface for the purpose of establishing a call. However, if step 1152 determines that a Voice-over-IP (VoIP) interface has been selected at step 1143, then step 1152 immediately proceeds to step 1154.

At step 1154 it has been determined the call is to be routed via a VoIP network. The A-party number is resolved into a network address at step 1155. This network address is used at step 1156 to establish a connection with a computing device or other network entity within the VoIP network, and the incoming call then is redirected to that computing device.

Figure 31 is a flow diagram of the process 8200 for handling an alternative scenario to that shown in Figure 30, in which a telephone call is initiated from the local

- 69 -

telecommunications device 1003 and may be routed via one of multiple telecommunications networks interfaced with the client processor device 1002, rather than an incoming call being redirected by the device 1002 as was previously described in relation to Figure 30. In Figure 31, at step 1170, a remote telecommunications device
5 identifier string (such as a telephone number associated with a remote telephone 1050 or other telecommunications device) is composed at the local telecommunications device 1003 by a user. At step 1171, a determination is made as to whether an interface to an alternate network other than the PSTN network 1004 is available. If no alternate telecommunications network is available, then at step 1172, a carrier identifier string
10 corresponding to the selected optimal telecommunications signal carrier within the PSTN network 1004 is then transmitted to a telephone exchange associated with that carrier, together with the string of digits already composed by the user at the local telecommunications device 1003. At step 1173, further digits required to identify the remote telecommunications device with which a connection is to be established are
15 transmitted directly to the telephone exchange associated with the selected carrier. The process 8200 is terminated at step 1174.

At step 1171, a determination is also made as to whether the alternate telecommunications networks are able to offer a cheaper call connection cost, and whether the B-party number
20 is supported by any of the telecommunications networks to which the customer interface device is able to be connected. At step 1175, if it is determined that an acceptable alternate telecommunications network through which the call is to be routed is a mobile network, then at step 1176 a connection is established to the mobile interface, and, at step 1177, the digits dialled by the user at the local telecommunications device 1003 are transmitted
25 through the mobile interface. At step 1178, further digits that may be entered by a user at the remote telecommunications device (for example in an auto attendance application) are then transmitted directly to the telephone exchange or other routing devices associated with the selected carrier in the alternate telecommunications network. The process 8200 is terminated at step 1179.

- 70 -

At step 1180, a determination is made as to whether the selected alternate telecommunications network is neither a mobile network or VoIP network (e.g. a GPRS network). If this is the case, then, at step 1181, an appropriate sequence of operations are performed by the device 1050 in order to direct the call through that alternate network.

- 5 Alternatively, if the alternate telecommunications network is a VoIP network, as determined by step 1182, then the number dialled at the local telecommunications device 1003 is resolved into a network address at step 1183, and a connection then established from the local telecommunications device 1003 to a computing device or other network entity (for example a VOIP telephony device) at that network address at step 1184.

10

4. Other Applications

Many other applications are available for the client processor device 100 and the server transmission system. For example, photocopier software can be upgraded remotely.

- 15 Automotive electronic system firmware can be upgraded without need to recall the vehicle. new features can be ordered for home appliances over the Internet and delivered when the customer wants them. Electronic billboards can be provided with new display information. Supermarket price tags can be updated across the store or across the country. RFID readers can have new standards installed in the field. Building and home systems can be
20 switched off and on, such as watering systems, alarm systems or lighting, using the devices 100.

- Many modifications will be apparent to those skilled in the art without departing from the scope of the present invention as herein described with reference to the accompanying
25 drawings.

CLAIMS

1. A client processor device including:
a communications receiver for receiving file fragments for said device;
5 a microcontroller system including a microcontroller and a programmable non-volatile memory system, for building and storing application and data files from said fragments, and executing an application of said device by processing at least one application file and associated data identified by configuration instructions included in at least one of said fragments.
10
2. A client processor device as claimed in claim 1, wherein said configuration instructions instruct said device to quit execution of a current application and commence execution of a new application immediately.
- 15 3. A client processor device as claimed in claim 1, wherein said fragments include new operating system components for said device, and said configuration instructions cause said device to reboot.
4. A client processor device as claimed in claim 1, wherein the device stores
20 application and data files for different versions of an application, and said configuration instructions select said versions for execution by said device.
5. A client processor device as claimed in claim 1, wherein the communications receiver uses a wireless communications protocol, such as POCSAG, Bluetooth, or
25 GSM/SMS, to receive said file fragments, configuration data and commands.
6. A client processor device as claimed in claim 1, wherein said file fragments are broadcast by a transmission system.

- 72 -

7. A client processor device as claimed in claim 1, wherein said communications receiver receives commands and configuration data for said device, in addition to said file fragments, for direct processing by said microcontroller.
- 5 8. A client processor device as claimed in claim 1, wherein said device is remotely configured by a user accessing a server, such as a web server, and defining the application, configuration and/or data files to be sent to said device in said fragments.
9. A client processor device as claimed in claim 1, wherein said microcontroller
10 system provides a file system structure, having directory and allocation units, for managing, storing, deletion and use of said files based on said configuration instructions.
10. A client processor device as claimed in claim 1, wherein said microcontroller
15 includes a boot routine adapted to install firmware components received by said communications receiver.
11. A client processor device as claimed in claim 9, wherein said boot routine substitutes firmware components.
- 20 12. A client processor device as claimed in claim 1, having a footprint less than 20 cm².
13. A client processor device as claimed in claim 1, wherein said microcontroller comprises an eight bit processor.
- 25 14. A client processor device as claimed in claim 1, wherein the device is connected to customer premises equipment, and said application includes processing at least part of the destination address of a telecommunications terminal selected on customer premises equipment to select a communications service to use for connecting to the destination terminal.

- 73 -

15. A client processor device as claimed in claim 14, including an interface to said telecommunications terminal for receiving call setup signals from said terminal and processing said signals to access said destination address and insert an override code in said signals to select said communications service.
- 5
16. A client processor device as claimed in claim 15, wherein said files include charge rate data for communications services, and said communications service is selected based on said charge rate data.
- 10 17. A client processor device, including:
a wireless communications receiver for receiving files for said device;
a microcontroller system for processing and storing said files, and automatically performing an application defined by said files without any user intervention or interaction.
- 15 18. A client processor device as claimed in claim 17, wherein a new application, as defined by said files, is performed by the device without any user intervention or interaction.
19. A client processor device as claimed in claim 17, wherein the device and the device
20 is reconfigured, as defined by said files, without any user intervention or interaction.
20. A client processor device, including:
a microprocessor;
a programmable non-volatile memory system addressable by the microprocessor;
25 a communications receiver for receiving file fragments for said device; and
an operating system (OS) for controlling the device to identify fragments for said device, build files from said fragments, store said files in said memory system, and process the files as including application, configuration or data code.

- 74 -

21. A client processor device as claimed in claim 20, wherein the received fragments for the device include an initial file defining the remaining files as being of a type of application, configuration or data, and defining reading of the files.
- 5 22. A client processor device as claimed in claim 21, wherein the operating system includes a number of tasks to control basic processes performed by the microprocessor and the tasks include at least one of:
- (i) A clock task for implementing a real time clock and calendar.
 - (ii) A virtual machine task for initiating a virtual machine, and causing reading of said
10 initial file; and executing of a file.
 - (iii) A command task for establishing device commands for production testing and initial configuration.
 - (iv) A file system task for processing said fragments and said files and removing unwanted and redundant files.
 - 15 (v) A communications interface task for controlling said receiver and processing data received by the receiver.
 - (vi) Tasks for processing of input and output signals for I/O peripherals, such as input buttons, or a GPS receiver of said client processor device.
- 20 23. A client processor device as claimed in claim 22, wherein the operating system includes a scheduler to allocate each task to be processed by the microprocessor, and allocates semaphores to resources of said device to allow the microprocessor to identify resources allocated to a running task.
- 25 24. A client processor device, including:
- a communications receiver for receiving communications packets including transmitted application data sent to said device; and
 - a microcontroller system for processing a first application based on stored application data of said device and subsequently processing a second application based on
30 said transmitted application data.

- 75 -

25. A server transmission system for generating and transmitting said file fragments for a device as claimed in claim 1 or 20.

26. A server transmission system as claimed in claim 25, including a web server for providing a user interface allowing a user to submit instructions used to determine the application, configuration, and data files, to be sent to said device in said fragments.

27. A server transmission system as claimed in claim 25, wherein said fragments are broadcast to a plurality of said device.

10

28. A handheld communications device for transmitting said file fragments to a device as claimed in claim 1 or 20.

29. A handheld communications device as claimed in claim 28, wherein said fragments are received by said handheld device from a server transmission system as claimed in claim 25, 26 or 27 prior to transmission to said client processor device.

15

30. A client processor device for processing at least part of the destination address of a telecommunications terminal selected on customer premises equipment to select a communications service to use for connecting to the destination terminal.




20

31. A client processor device as claimed in claim 30, including an interface to said telecommunications terminal for receiving call setup signals from said terminal and processing said signals to access said destination address and insert an override code in said signals to select said communications service.

25

32. A client processor device as claimed in claim 31, including files with charge rate data for communications services, and said communications service is selected based on said charge rate data.

30

**Figure 1****Figure 2****Figure 3**




Figure 4




Figure 5





Figure 6

**Figure 7**




Figure 8




Figure 9




Figure 10




Figure 11




Figure 12

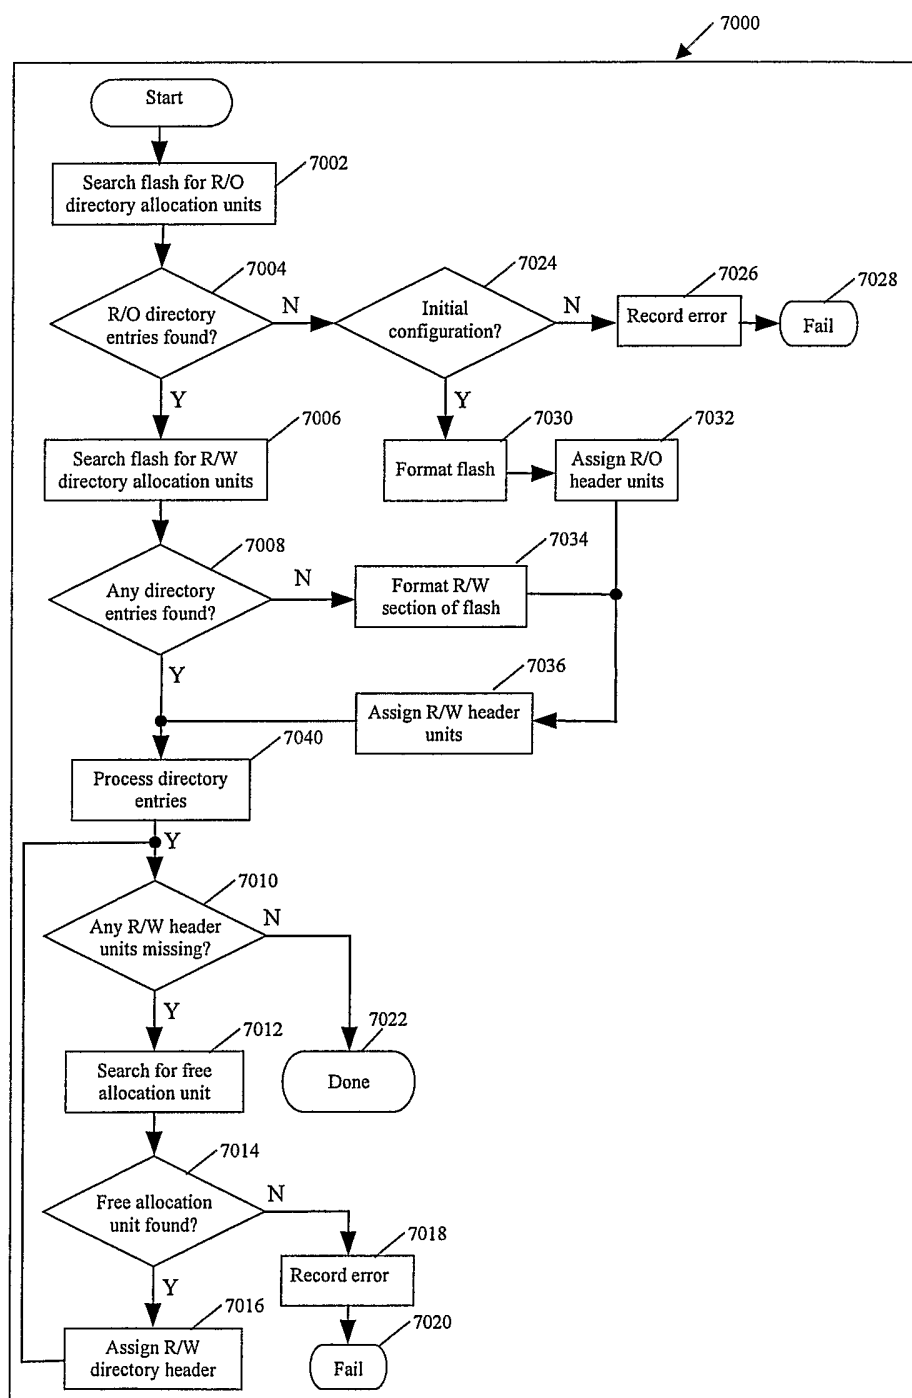


Figure 13

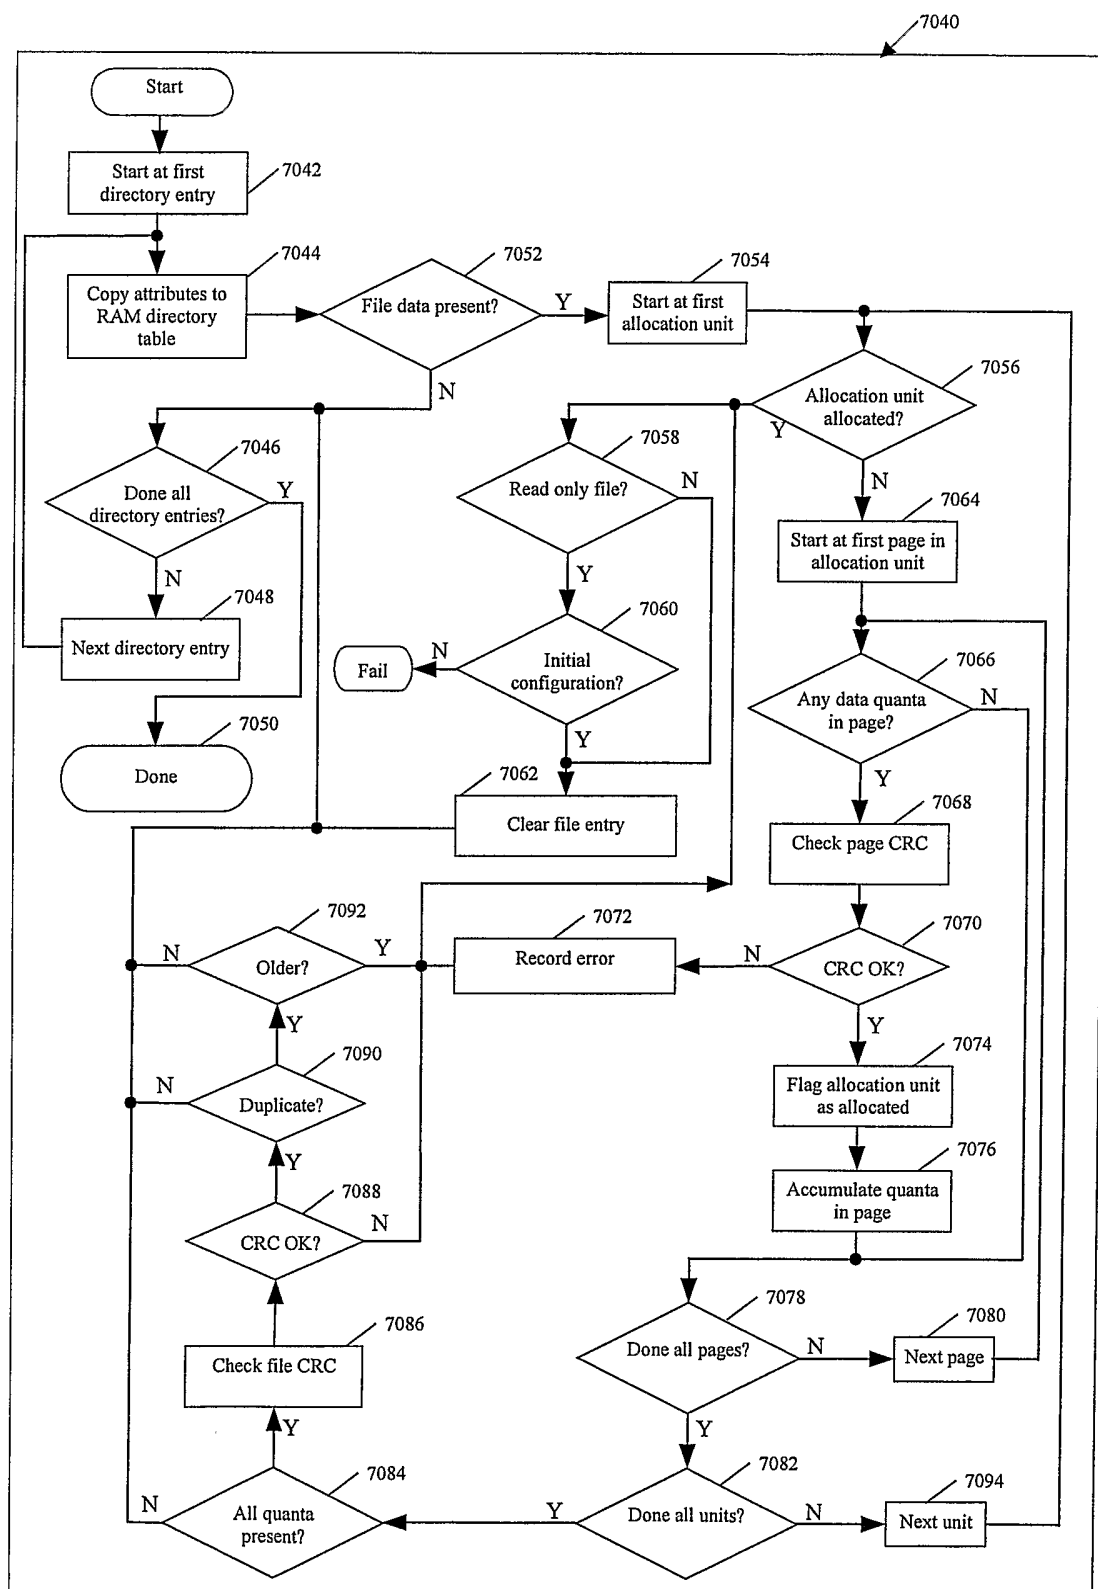


Figure 14

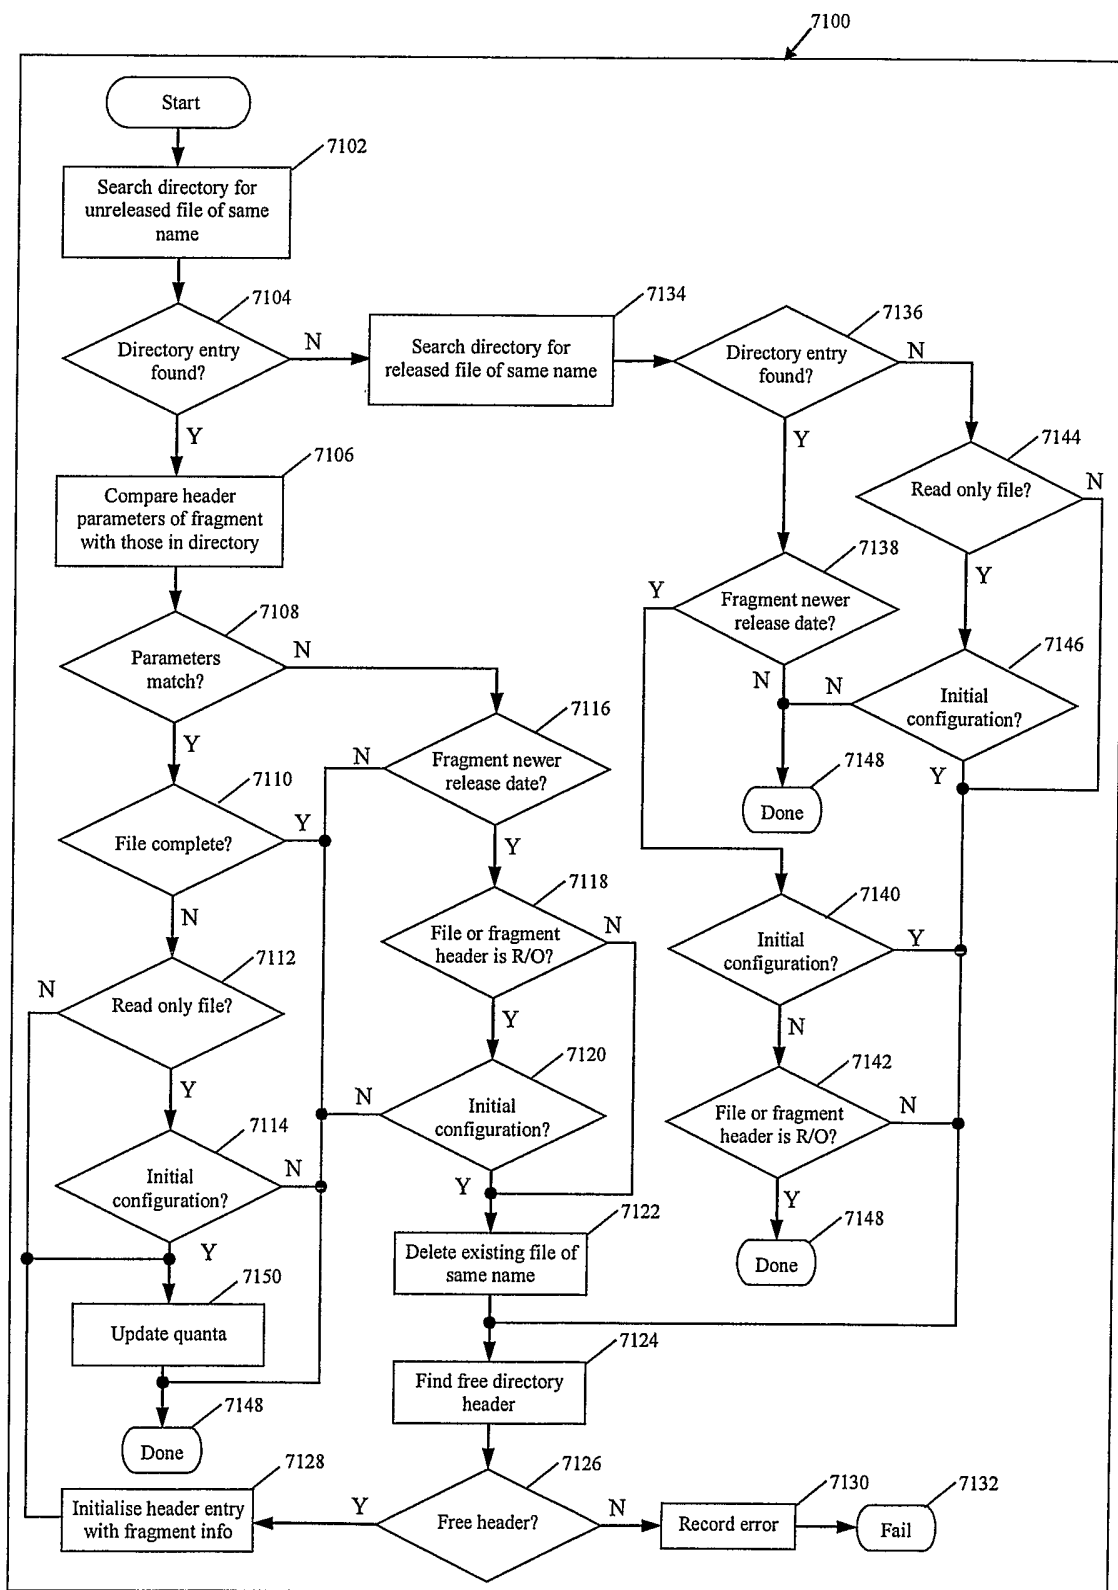


Figure 15

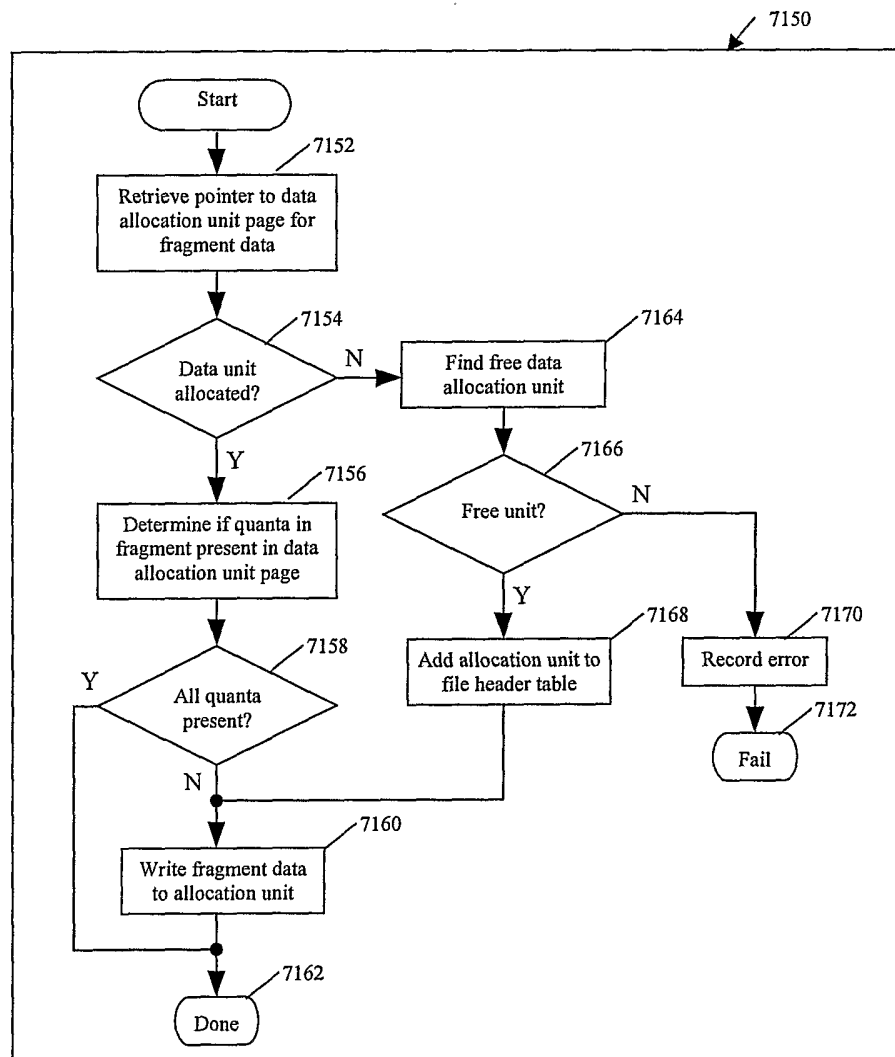
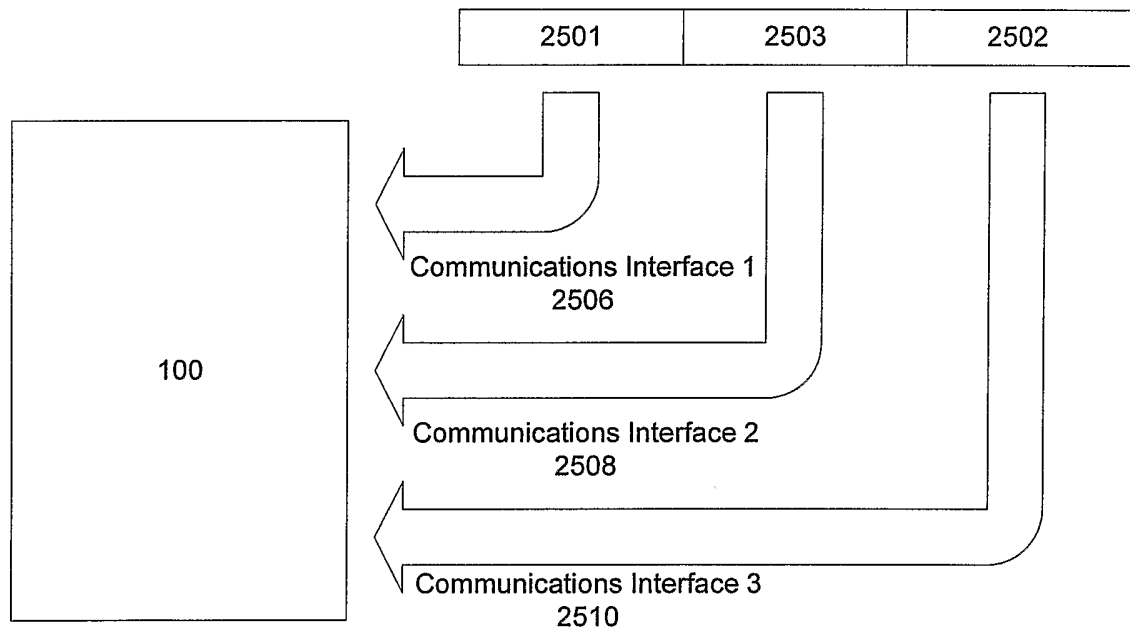


Figure 16

**Figure 17**

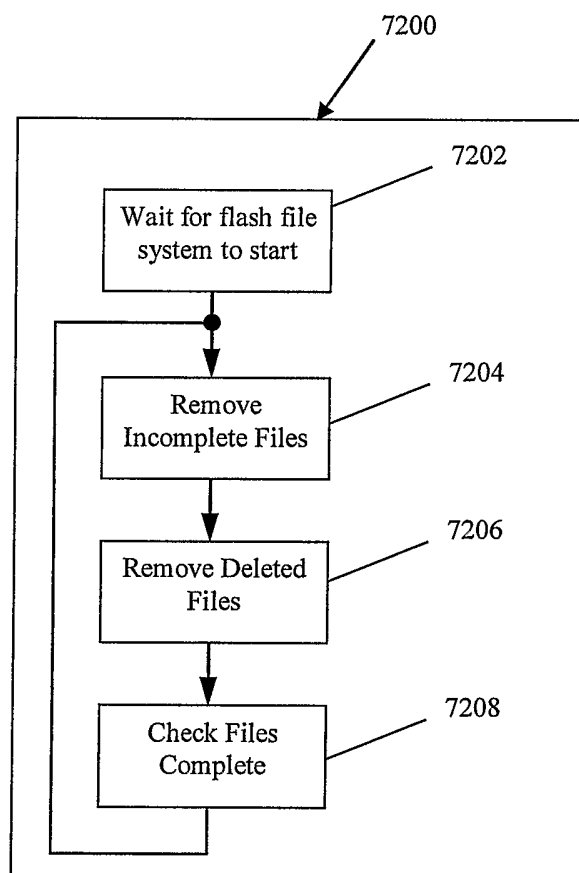


FIGURE 18

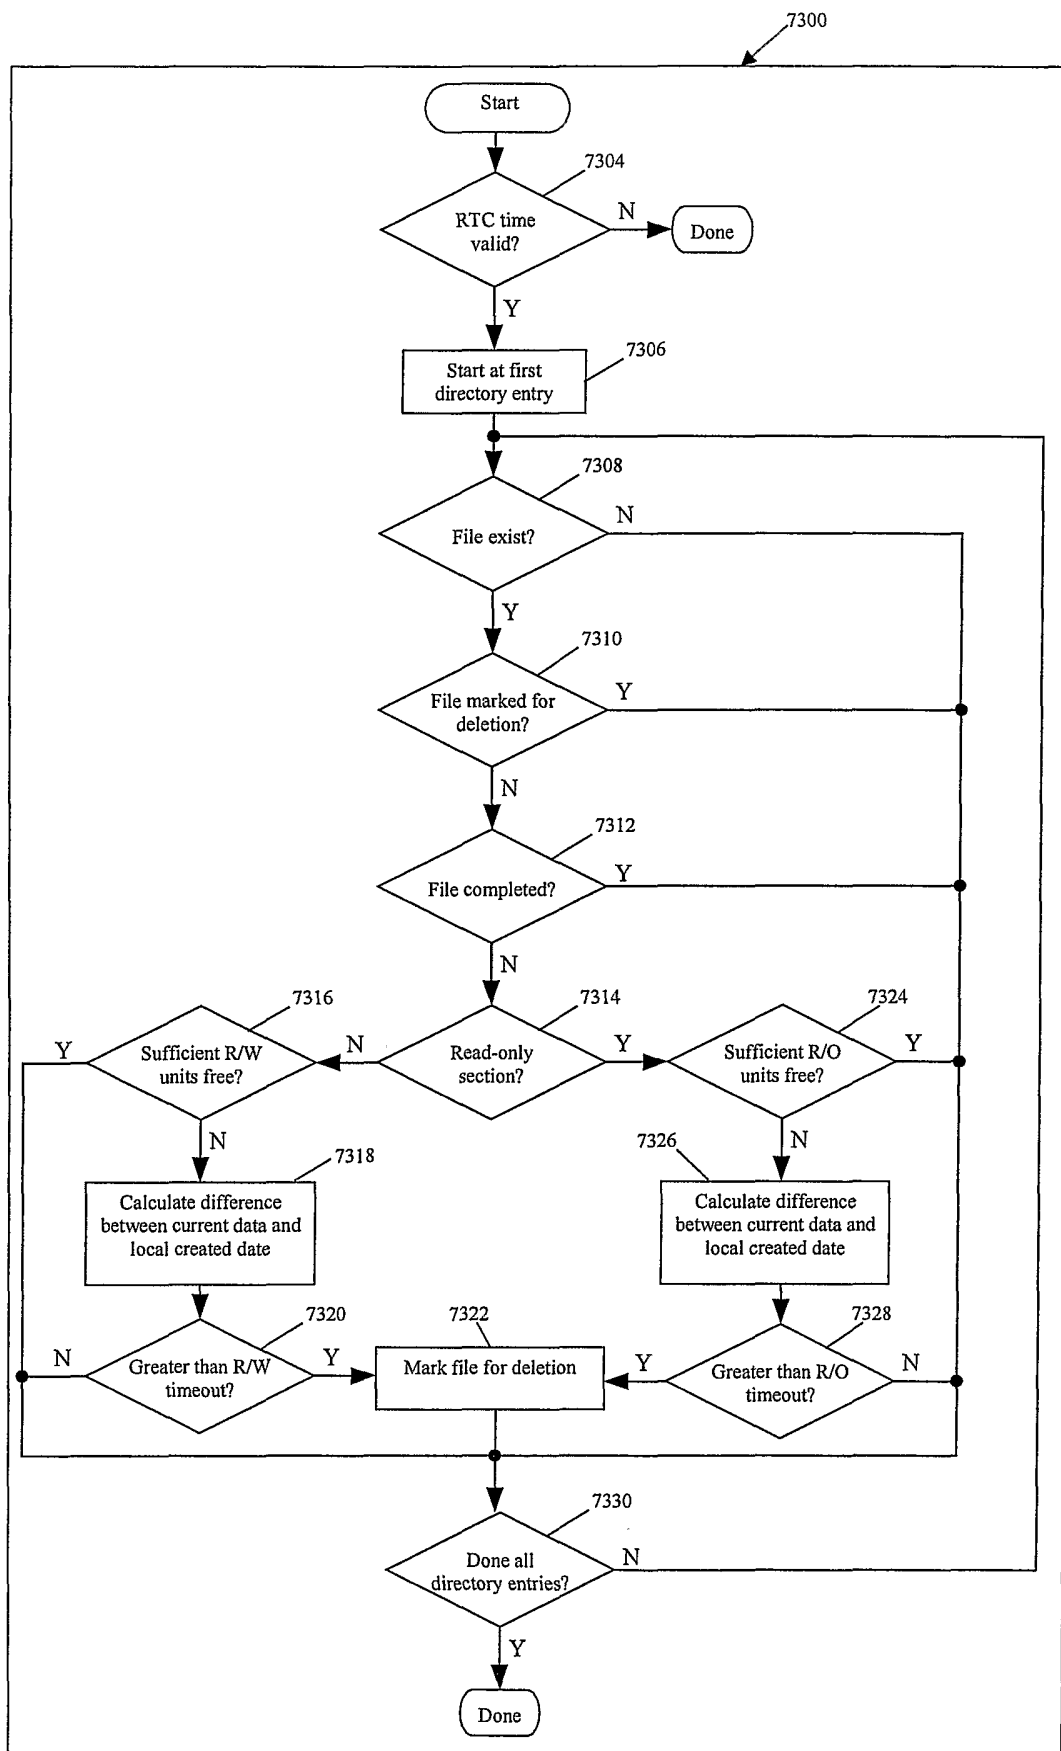


Figure 19

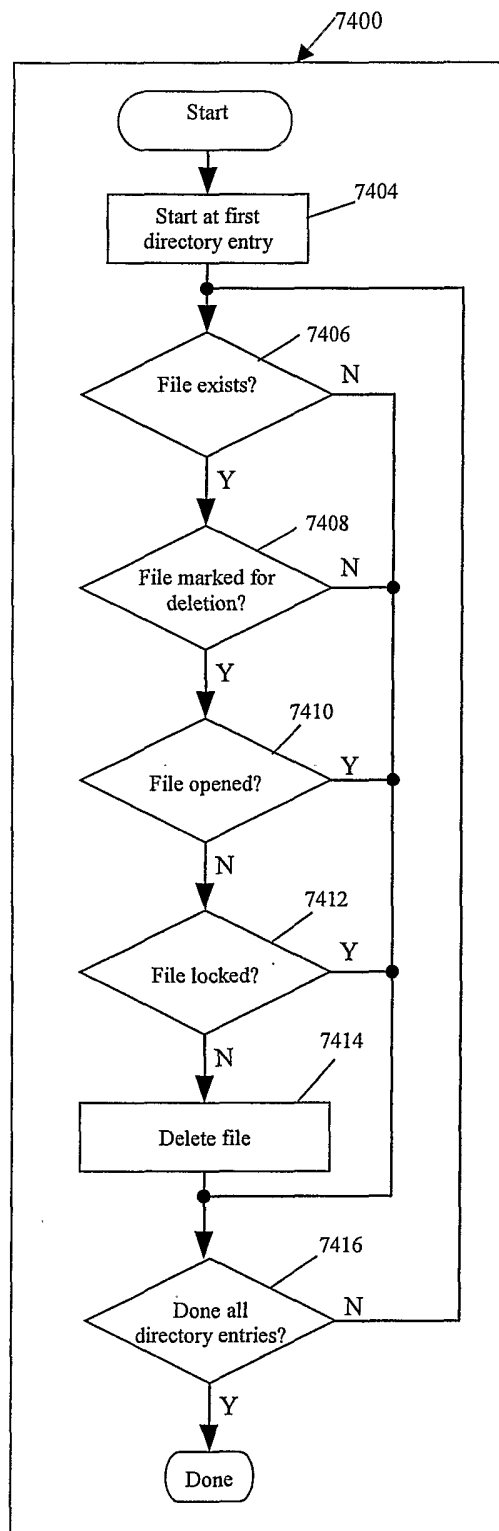


Figure 20

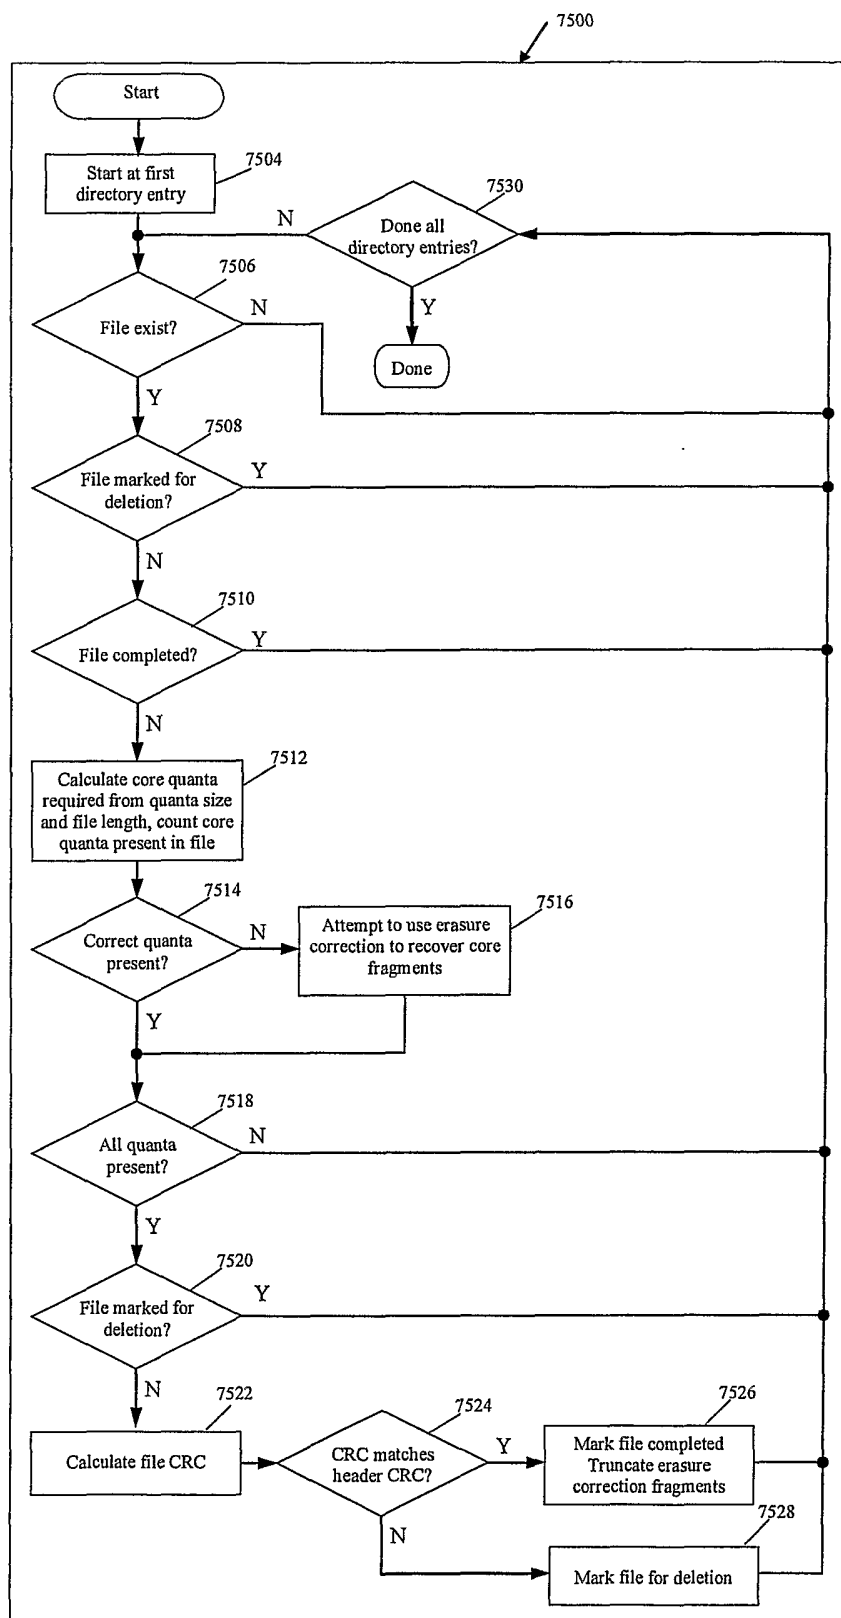


Figure 21




Figure 22




Figure 23




Figure 24




Figure 25

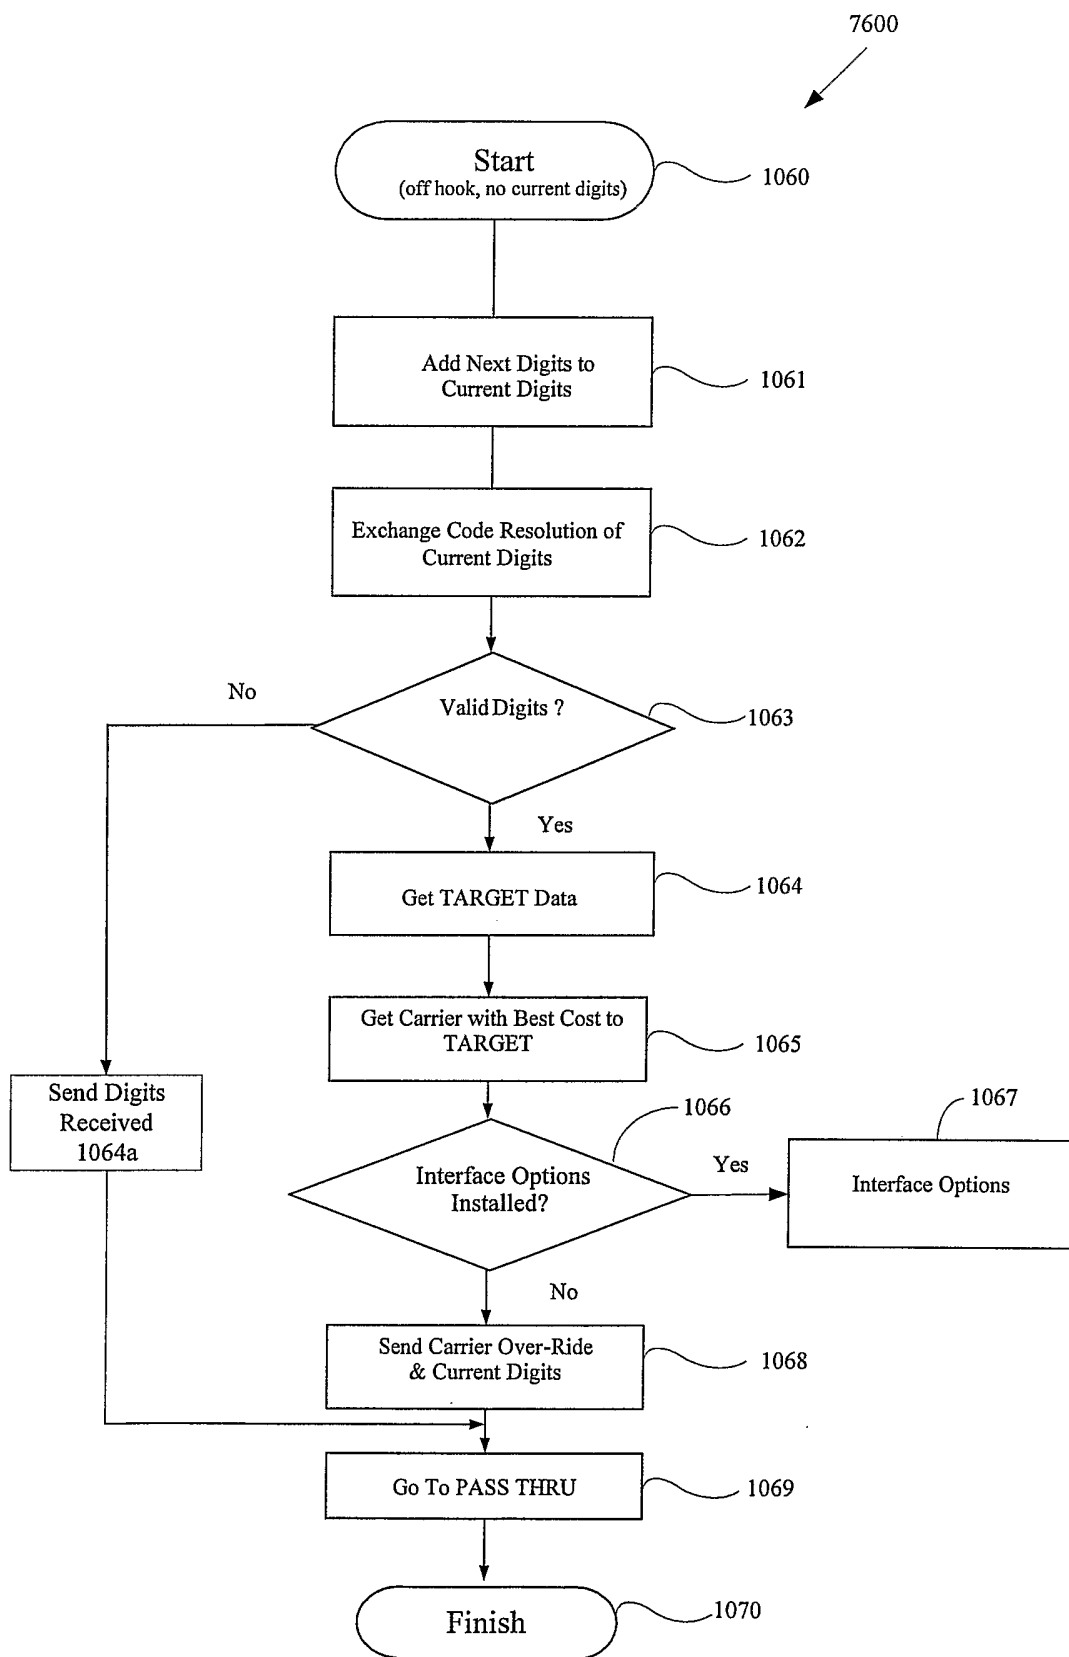


Figure 26

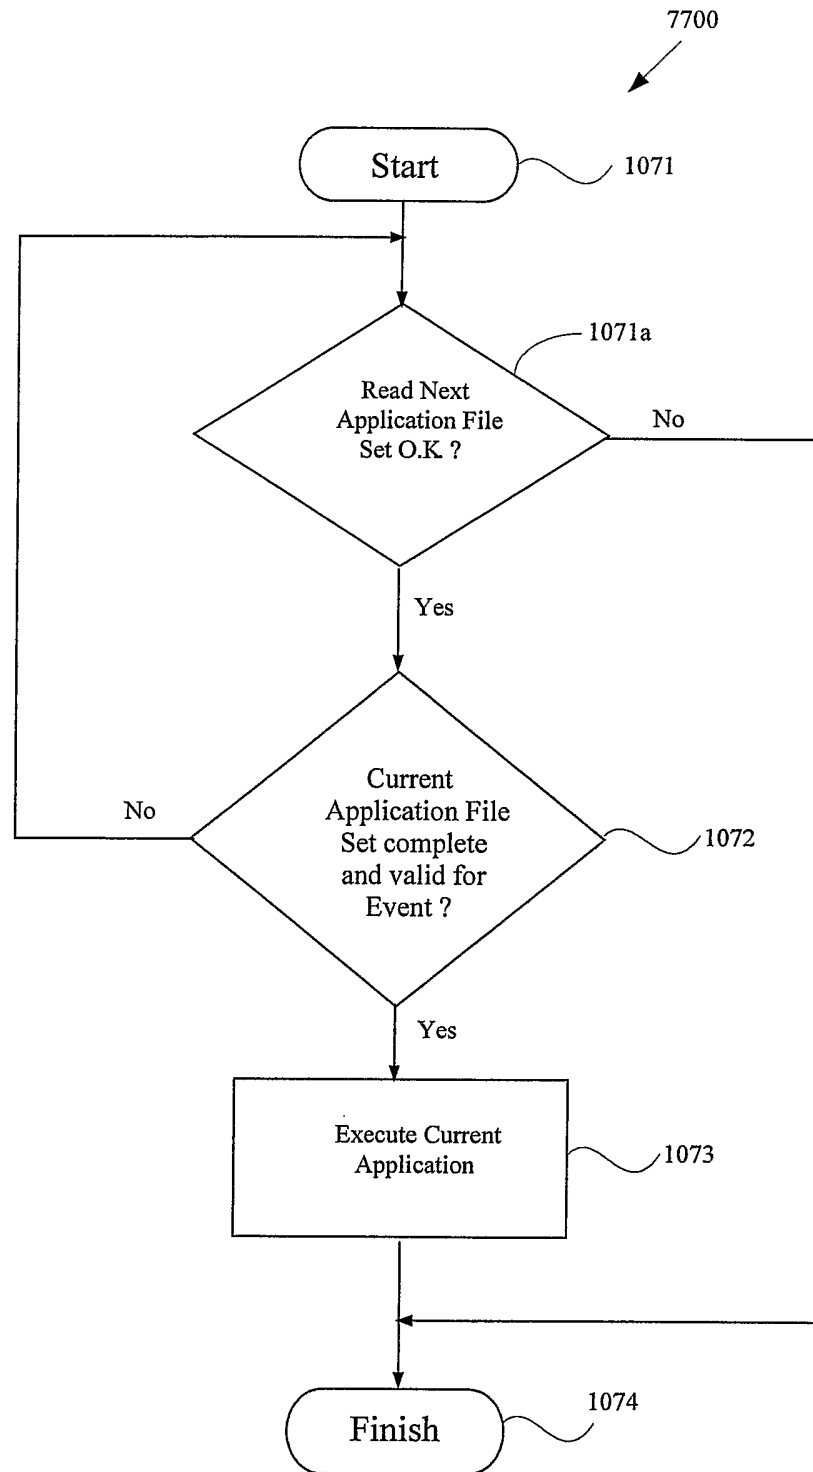


Figure 27

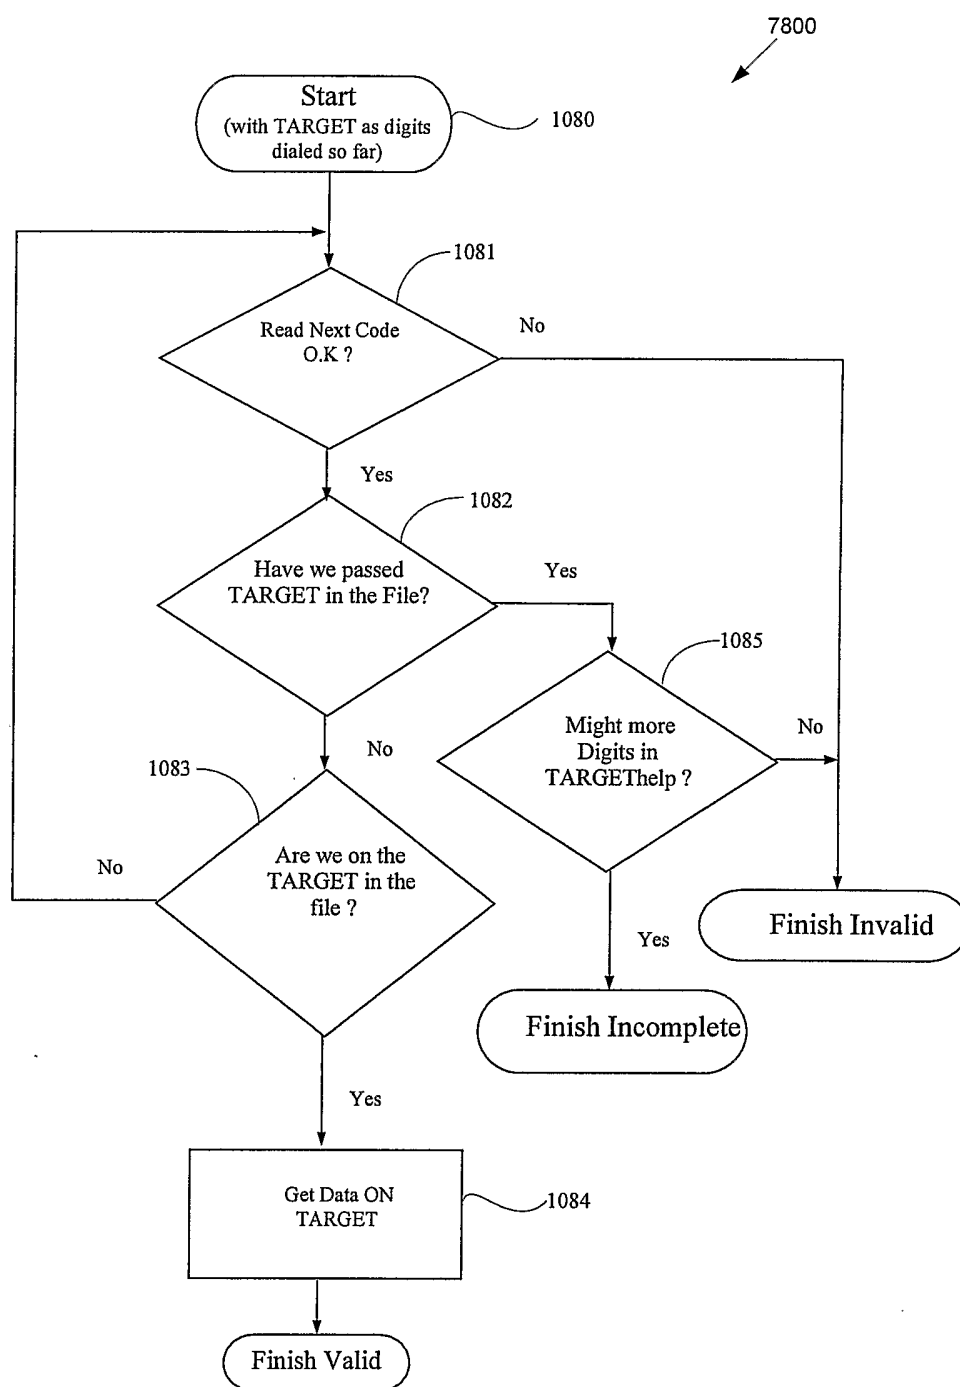


Figure 28

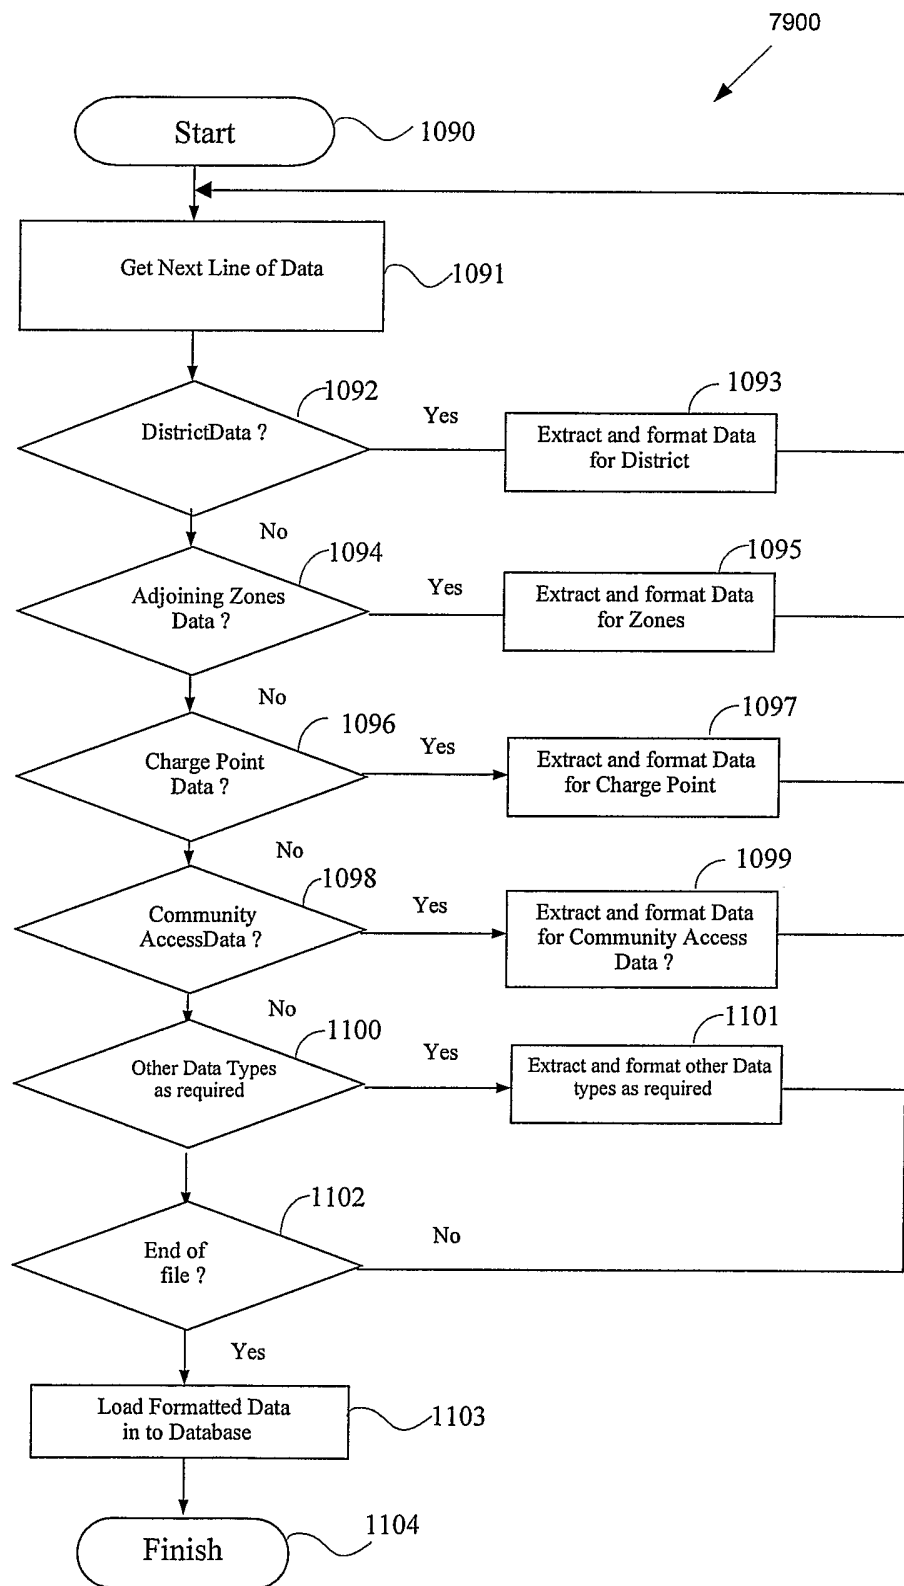


Figure 29

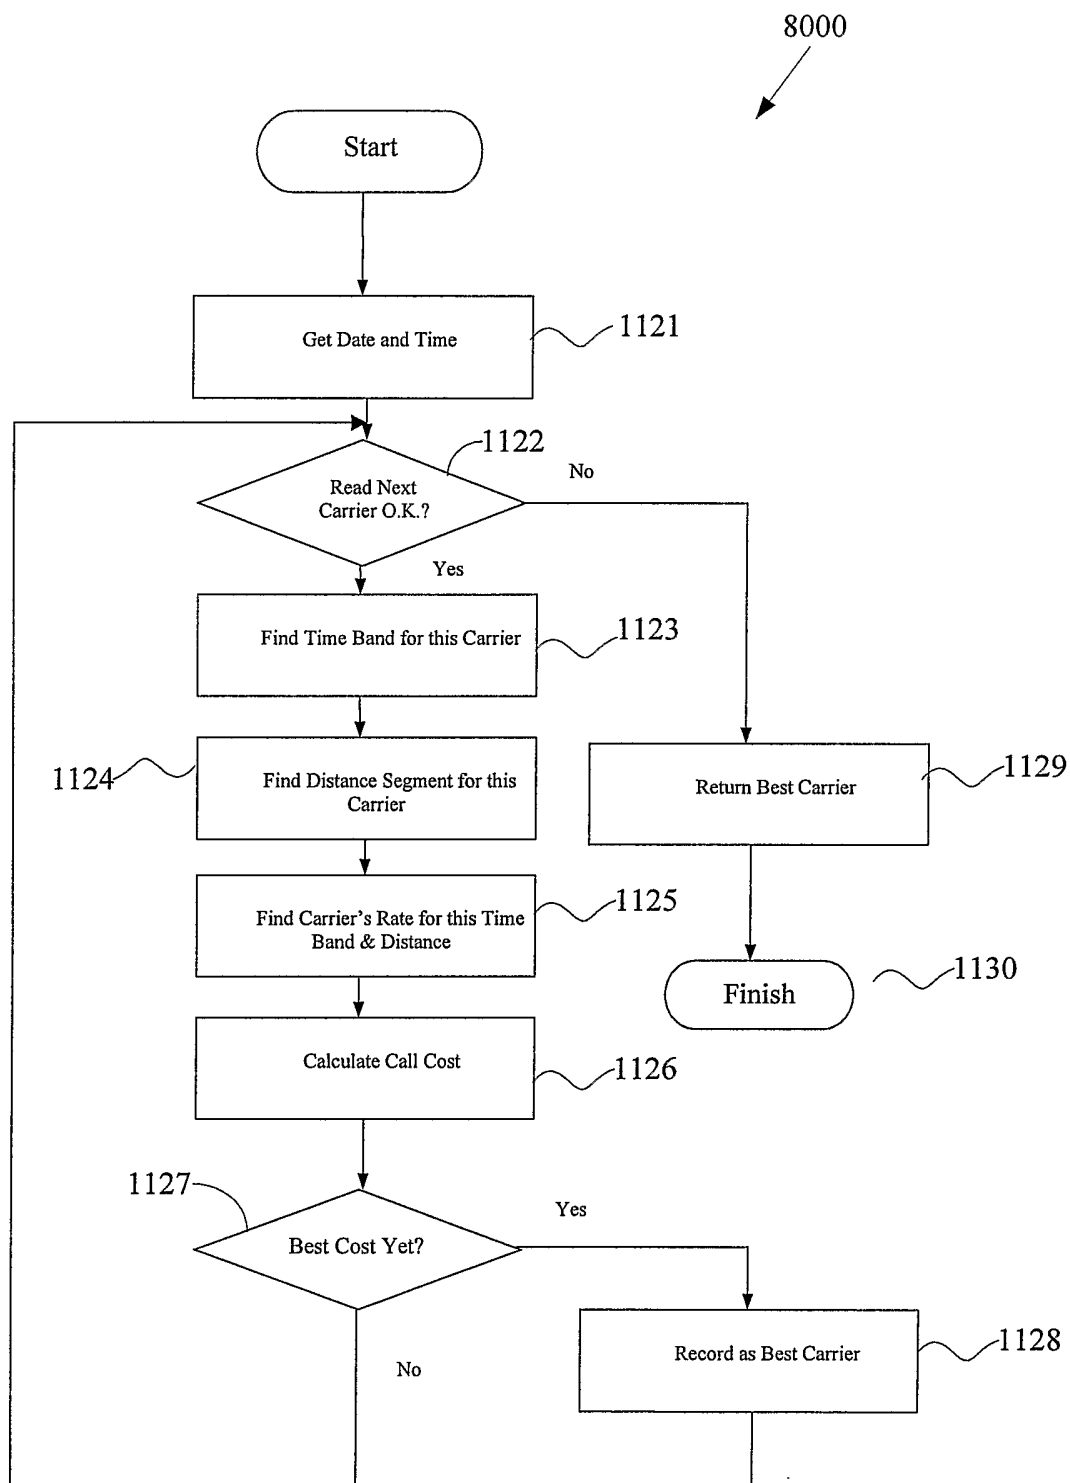


Figure 30

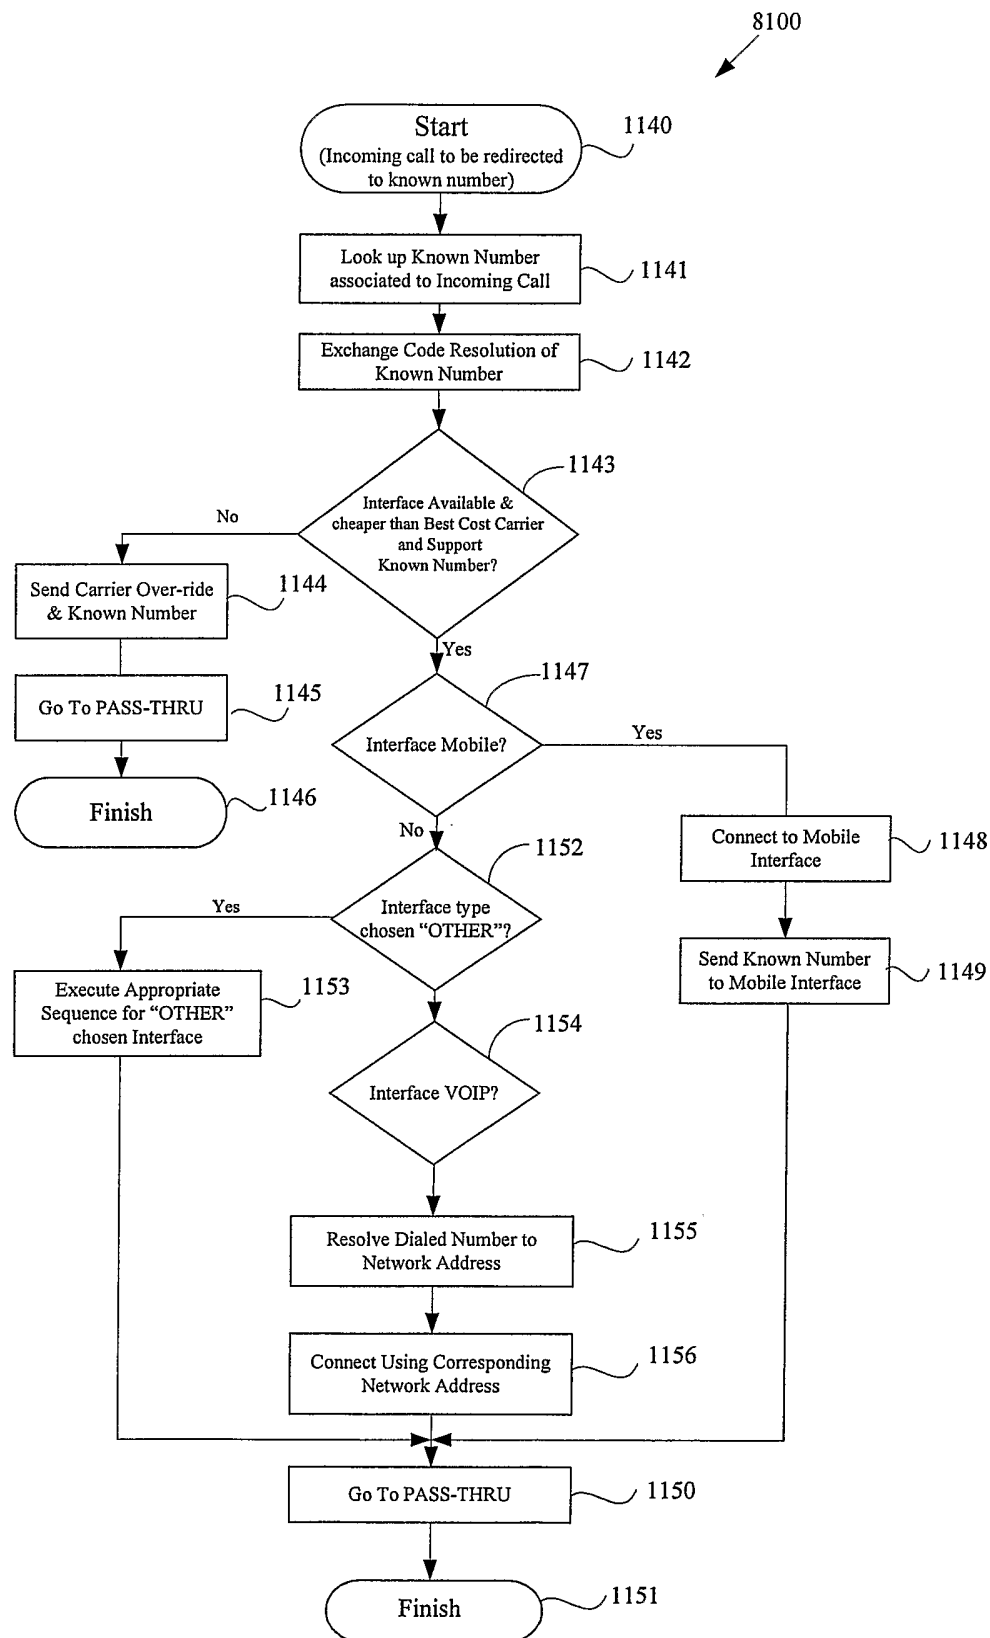


Figure 31

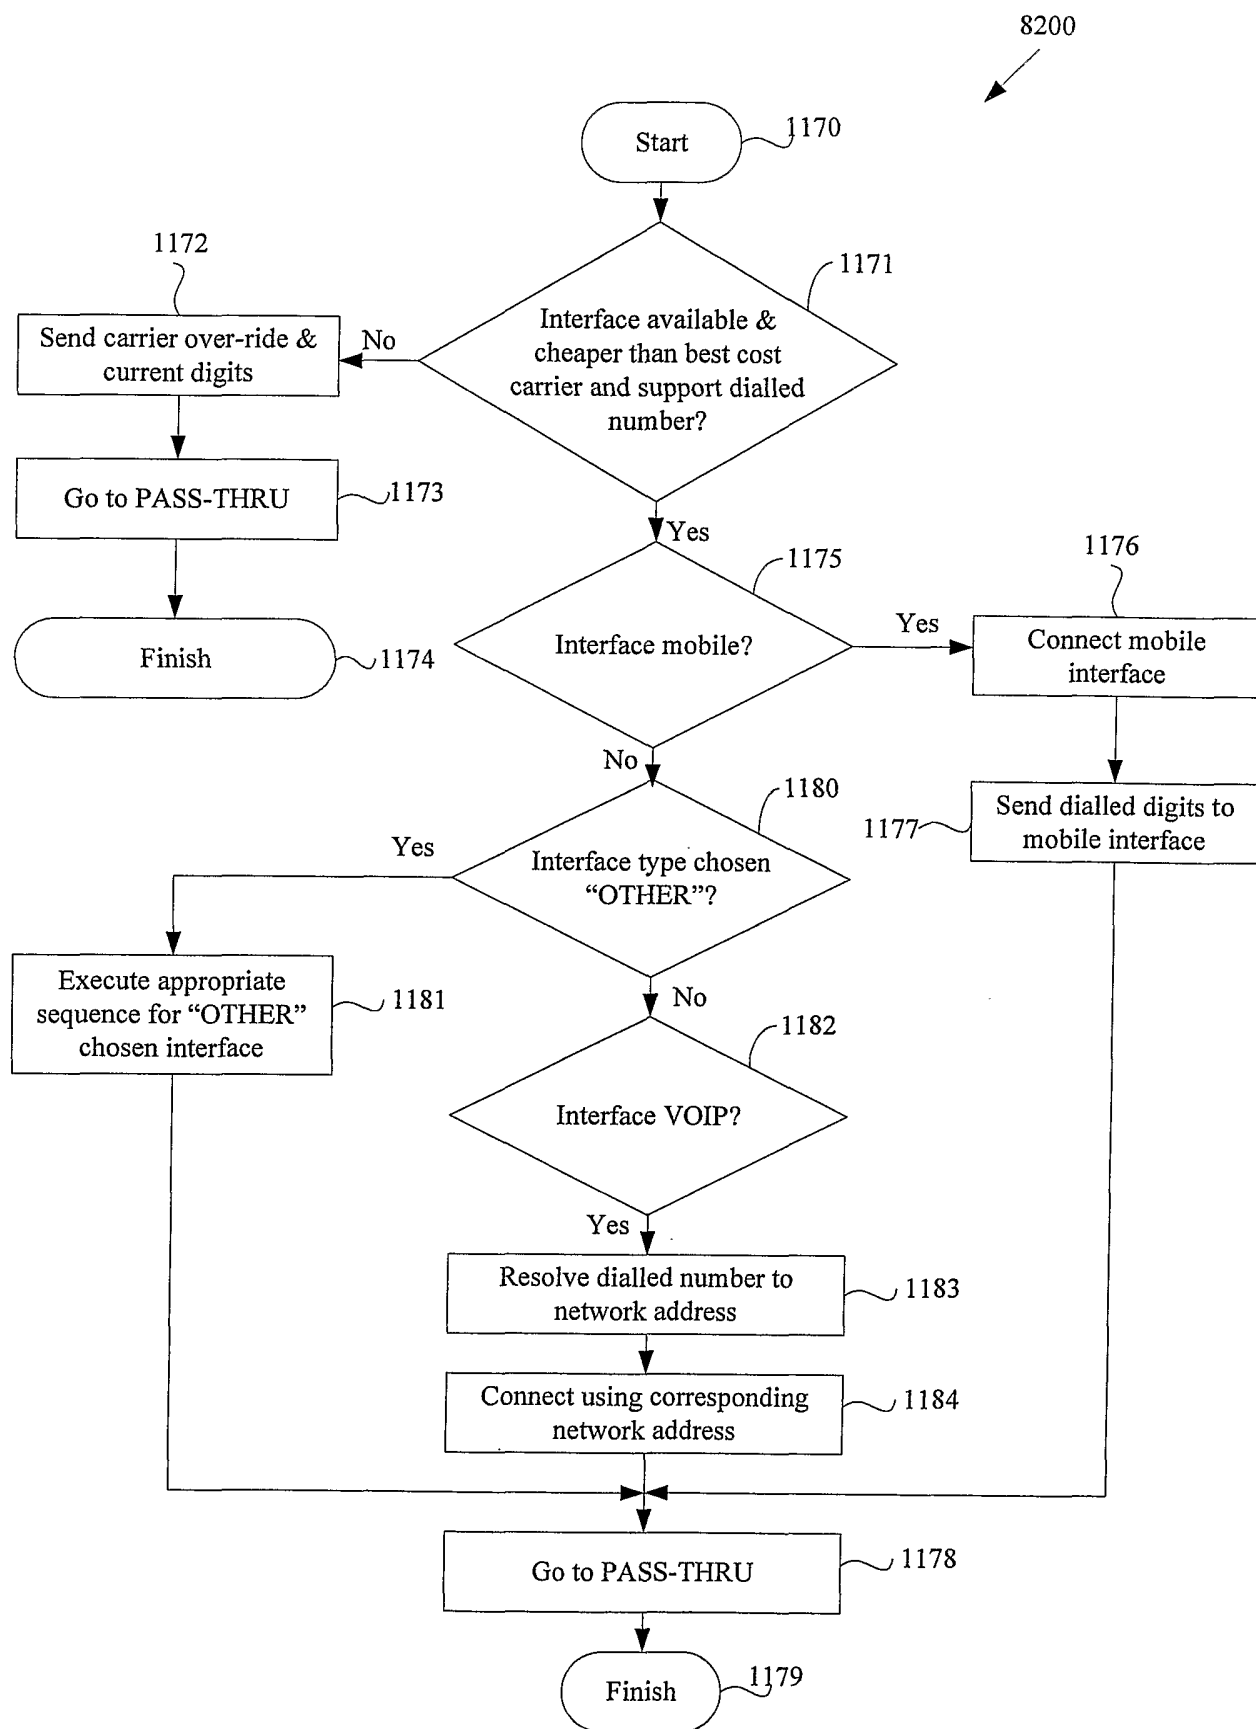


Figure 32

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2005/000922

A. CLASSIFICATION OF SUBJECT MATTER

Int. Cl. ⁷: G06F 9/445, H04Q 7/00

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
DWPI, PCT, USPTO, Google Scholar (firmware, upgrade, update, wireless, assemble, reassemble, application, cheapest, low cost, carrier etc.)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/0181415 A1 (WEST et al), 5 December 2002 the whole document (paragraphs 153-160 in particular)	1-29
X	WO 2003/025742 A2 (SOMA NETWORKS, INC.), 27 March 2003 the whole document (page 22, lines 16-29 in particular)	1-29
X	US 2003/0022663 A1 (RAJARAM et al), 30 January 2003 the whole document	1-29
X, P	WO 2004/073332 A1 (UNISYS CORPORATION), 26 August 2004 the whole document	1-29

☒ Further documents are listed in the continuation of Box C☒ See patent family annex

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

Date of the actual completion of the international search
13 October 2005Date of mailing of the international search report
21 OCT 2005Name and mailing address of the ISA/AU
AUSTRALIAN PATENT OFFICE
PO BOX 200, WODEN ACT 2606, AUSTRALIA
E-mail address: pct@ipaustralia.gov.au
Facsimile No. (02) 6285 3929

Authorized officer

MATTHEW HOLLINGWORTH
Telephone No : (02) 6283 2024

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2005/000922

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2003/036933 A1 (RAMASUBRAMANIAN), 1 May 2003 the whole document	30-32
X	US 6,240,169 B1 (HASKINS et al), 29 May 2001 the whole document	30-32
X	US 6,381,315 B1 (NHAISSI), 30 April 2002 the whole document	30-32
X	US 5,862,203 A (WULKAN et al), 19 January 1999 the whole document	30-32
X	US 6,169,791 B1 (POKRESS), 2 January 2001 the whole document	30-32
A	US 2002/0120874 A1 (SHU et al), 29 August 2002 the whole document	1-29

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2005/000922

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a)

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:
See extra sheet.

1. ☒ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
- ☒ No protest accompanied the payment of additional search fees.

Supplemental Box

(To be used when the space in any of Boxes I to VIII is not sufficient)

Continuation of Box No: III

The international application does not comply with the requirements of unity of invention because it does not relate to one invention or to a group of inventions so linked as to form a single general inventive concept. In coming to this conclusion the International Searching Authority has found that there are two inventions:

- Claims 1-29. The formation of an application from file segments received by a communications receiver is a first "special technical feature."
- Claims 30-32. The selection of a communications service based on the destination address of a telecommunications terminal is a second "special technical feature."

Since the abovementioned groups of claims do not share any technical features, a "technical relationship" between the inventions, as defined in PCT rule 13.2 does not exist. Accordingly, the international application does not relate to one invention or to a single inventive concept.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2005/000922

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report				Patent Family Member			
US	2002181415	EP	1382128	US	6842433	WO	02087094
WO	03025742	CA	2357382	CA	2498648	CN	1585926
		EP	1461694	MX	PA04002527	US	2005055595
US	2003022663	CN	1406686	CN	1535418	CN	1535419
		CN	1535420	CN	1535421	CN	1535422
		CN	1535423	CN	1535529	CN	1537272
		CN	1537276	CN	1537397	EP	1410188
		EP	1410189	EP	1410190	EP	1410191
		EP	1410192	EP	1410193	EP	1410209
		EP	1410665	EP	1423959	EP	1425894
		US	6860315	US	6918108	US	2003019603
		US	2003022665	US	2003023964	US	2003033525
		US	2003033599	US	2003064717	US	2003066064
		US	2003110479	US	2003110480	US	2004205746
		US	2004214559	US	2004214560	US	2004214561
		US	2005010917	US	2005026603	US	2005064847
		WO	03010656	WO	03010658	WO	03010662
		WO	03010663	WO	03010664	WO	03010668
		WO	03010932	WO	03010942	WO	03012639
		WO	03013103	WO	2005029891		
WO	2004073332	AU	2003225439				
WO	03036933	SG	111921				
US	6240169						
US	6381315	AU	22995/01	GB	2372901	GB	2395862
		GB	2395863	GB	2395864	GB	2395865
		GB	2396078	WO	0141420		
US	5862203	AU	17705/99	AU	37826/97	AU	63683/96
		CA	2227569	EP	0865706	US	5764741
		US	5799072	US	6078652	WO	9705749
		WO	9808335	WO	9933252		
US	6169791	AU	85880/98	CA	2297976	EP	0998830
		WO	9905872				
US	2002120874	US	2002080888	US	2003084020	WO	02052787
		WO	02062032	WO	02065292		

Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.

END OF ANNEX

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2005/000922