



(51) International Patent Classification:

G10L 15/22 (2006.01) G10L 17/00 (2013.01)
G10L 15/30 (2013.01) G06F 21/64 (2013.01)

(21) International Application Number:

PCT/GB2019/052401

(22) International Filing Date:

28 August 2019 (28.08.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/115,654 29 August 2018 (29.08.2018) US

(71) Applicant: CIRRUS LOGIC INTERNATIONAL
SEMICONDUCTOR LIMITED [GB/GB]; 7B Nightin-

gale Way, Quartermile, Edinburgh, Central Scotland EH3
9EG (GB).

(72) Inventor: LESSO, John Paul; 8 Galachlaw Shot, Edin-
burgh, Lothian EH10 7JF (GB).

(74) Agent: HASELTINE LAKE KEMPNER LLP; Redcliff
Quay, 120 Redcliff Street, Bristol, Bristol BS1 6HU (GB).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,

(54) Title: SPEECH RECOGNITION




Fig. 1

(57) Abstract: A method of performing speech recognition, comprises, at a first device: receiving an audio signal representing speech; performing a first data integrity check operation on the received audio signal; performing a speaker recognition process on the received audio signal; forwarding the received audio signal to a second device, wherein the second device comprises a speech recognition function; and forwarding an output of the first data integrity check operation to the second device. The method further comprises, at the second device: receiving the audio signal forwarded from the first device; receiving the output of the first data integrity check operation forwarded from the first device; performing a second data integrity check operation on the audio signal forwarded from the first device; and using a result of performing the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.



SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

— *with international search report (Art. 21(3))*

SPEECH RECOGNITION

Technical Field

- 5 This invention relates to speech recognition, and in particular to a system that performs both speaker recognition and speech recognition.

In this document, the term speaker recognition is used to refer to a process in which information is obtained about the identity of a speaker. For example, this process may
10 involve determining whether or not the speaker is a specific individual (speaker verification), or may involve identifying the speaker, for example from a group of enrolled speakers (speaker identification).

The term speech recognition is used to refer to a process in which information is
15 obtained about the content of speech, for example in order to be able to determine what the speaker is saying.

Background

20 Systems are known in which a verbal command from a speaker is recognised and processed, subject to the speaker's identity being verified. One such system is disclosed in GB-2515527A, for example, in which, if it is determined that a predetermined trigger phrase has been spoken, a speaker recognition process is performed in a first device. If the speaker recognition process determines that the
25 predetermined trigger phrase was spoken by a specific enrolled user, the signal representing the speech is passed to a speech recognition engine, which is provided in a second device. Typically, such a system may be used to allow the user to issue voice commands, causing the system to perform some action, or retrieve some requested information, for example.

30

One problem that could in theory arise in such a system is an attack, in which a third party attempts to gain unauthorised access to the system, by providing a recording of the enrolled user's speech to the speaker recognition process, but then providing a signal containing their own speech to the speech recognition engine. The speaker
35 recognition process would signal that the enrolled user was speaking, but the

command that the speech recognition process would recognise may be a command that the enrolled user would not want to issue.

Summary

5

According to an aspect of the present invention, there is provided a method of performing speech recognition, comprising:

at a first device:

receiving an audio signal representing speech;

10

performing a first data integrity check operation on the received audio signal;

performing a speaker recognition process on the received audio signal;

forwarding the received audio signal to a second device, wherein the second device comprises a speech recognition function; and

forwarding an output of the first data integrity check operation to the second

15

device; and

at the second device:

receiving the audio signal forwarded from the first device;

receiving the output of the first data integrity check operation forwarded from the first device;

20

performing a second data integrity check operation on the audio signal forwarded from the first device; and

using a result of performing the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

25

The method may further comprise:

using a result of performing the speech recognition function on the audio signal forwarded from the first device only if a result of performing the speaker recognition process on the received audio signal satisfies a predetermined criterion.

30

The method may further comprise:

performing the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, and

35

if the result of performing the speaker recognition process on the received audio signal satisfies the predetermined criterion.

The method may comprise:

- 5 performing the speech recognition function on the audio signal forwarded from the first device;
- determining if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device; and
- 10 using the result of performing the speech recognition function on the audio signal forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

The method may further comprise:

- 15 if the output of the second data integrity check operation does not match the output of the first data integrity check operation forwarded from the first device, deleting the result of performing the speech recognition function on the audio signal forwarded from the first device.
- 20 The method may comprise performing the speech recognition function on the audio signal in the second device at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

- The method may comprise determining in the second device if the output of the second
- 25 data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

- The method may comprise, if it is determined in the second device that the output of
- 30 the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, performing the speech recognition function on the audio signal in the second device at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

The method may comprise performing the first data integrity check operation on the received audio signal only if it is determined that the audio signal represents a predetermined trigger phrase.

- 5 The step of performing the first data integrity check operation may comprise obtaining a predetermined hash function, and wherein the step of performing the second data integrity check operation comprises obtaining said predetermined hash function.

- According to a second aspect of the invention, there is provided a system for
- 10 performing speech recognition, comprising:
- a first device and a second device,
 - wherein the first device comprises:
 - an input for receiving an audio signal representing speech;
 - a processor for performing a first data integrity check operation on the received
 - 15 audio signal, and for performing a speaker recognition process on the received audio signal; and
 - an interface for forwarding the received audio signal to the second device, and for forwarding an output of the first data integrity check operation to the second device; and
 - 20 wherein the second device is configured for performing a speech recognition function, and comprises:
 - an interface for receiving the audio signal forwarded from the first device, and for receiving the output of the first data integrity check operation forwarded from the first
 - 25 device; and
 - a processor for performing a second data integrity check operation on the audio signal forwarded from the first device;
 - wherein the system is configured to use a result of performing the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data
 - 30 integrity check operation forwarded from the first device.

- The system may be configured to use the result of performing the speech recognition function on the audio signal forwarded from the first device only if a result of performing the speaker recognition process on the received audio signal satisfies a predetermined
- 35 criterion.

The second device may be configured to perform the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, and if the result of performing the speaker recognition
5 process on the received audio signal satisfies the predetermined criterion.

The system may be configured for:

- performing the speech recognition function on the audio signal forwarded from the first device;
- 10 determining if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device; and
- using the result of performing the speech recognition function on the audio signal forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from
15 the first device.

The system may be further configured for:

- if the output of the second data integrity check operation does not match the output of the first data integrity check operation forwarded from the first device, deleting
20 the result of performing the speech recognition function on the audio signal forwarded from the first device.

The system may be further configured for performing the speech recognition function on the audio signal in the second device at least partly in parallel with performing the
25 speaker recognition process on the received audio signal in the first device.

The second device may be further configured for determining if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, at least partly in parallel with
30 performing the speaker recognition process on the received audio signal in the first device.

The second device may be further configured for, if it is determined that the output of the second data integrity check operation matches the output of the first data integrity
35 check operation forwarded from the first device, performing the speech recognition function on the audio signal in the second device at least partly in parallel with

performing the speaker recognition process on the received audio signal in the first device.

The first device may be configured for performing the first data integrity check operation on the received audio signal only if it is determined that the audio signal represents a predetermined trigger phrase.

Performing the first data integrity check operation may comprise obtaining a predetermined hash function, and wherein performing the second data integrity check operation may comprise obtaining said predetermined hash function.

The first device and the second device may be separate integrated circuits in an electronic device.

The first device may be an electronic device, and the second device may be a remotely located computing device.

The second device may be an electronic device, and the first device may be an accessory for use with the second device.

In any of these examples, the electronic device may be a smartphone or other communications device, a smart speaker, a tablet or laptop computer, a games console, a home control system, a home entertainment system, an in-vehicle entertainment system, or a domestic appliance.

According to a further aspect, there is provided a method of operation of a first device in a speech recognition system, the method comprising:

- receiving an audio signal representing speech;
- performing a first data integrity check operation on the received audio signal;
- performing a speaker recognition process on the received audio signal;
- forwarding the received audio signal to a second device, wherein the second device comprises a speech recognition function; and
- forwarding an output of the first data integrity check operation to the second device.

According to a further aspect, there is provided a first device configured to operate in accordance with the method.

5 The first device may be an integrated circuit. Alternatively, the first device may be an electronic device, for example a smartphone or other communications device, a smart speaker, a tablet or laptop computer, a games console, a home control system, a home entertainment system, an in-vehicle entertainment system, or a domestic appliance.

10 According to a further aspect, there is provided a method of operation of a second device in a speech recognition system, wherein the second device comprises a speech recognition function, the method comprising:

receiving from a first device an audio signal representing speech;
receiving from the first device an output of a first data integrity check operation;
performing a second data integrity check operation on the audio signal forwarded
15 from the first device; and
using a result of performing the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

20

The method may further comprise:

using the result of performing the speech recognition function on the audio signal forwarded from the first device only in response to receiving from the first device an indication that a result of performing a speaker recognition process on the received
25 audio signal satisfies a predetermined criterion.

According to a further aspect, there is provided a second device configured to operate in accordance with the method.

30 The second device may be an integrated circuit. Alternatively, the second device may be an electronic device, for example a smartphone or other communications device, a smart speaker, a tablet or laptop computer, a games console, a home control system, a home entertainment system, an in-vehicle entertainment system, or a domestic appliance.

35

According to a further aspect, there is provided a method for providing an authenticated/identified speech recognition output, the method comprising the steps of:

- receiving audio including speech;
- providing the received audio to a speaker recognition module, to
- 5 authenticate/identify a speaker, wherein the step of providing the received audio to a speaker recognition module comprises performing a first data integrity check operation on the received audio signal;
- providing the received audio to a speech recognition module, to recognise speech in the received audio, wherein the step of providing the received audio to the
- 10 speech recognition module comprises performing a second data integrity check operation on the received audio signal;
- comparing the output of the first and second data integrity check operations to determine the audio provided to the speaker recognition module is the same as the audio provided to the speech recognition module,
- 15 wherein the outputs of the speaker recognition module and the speech recognition module are combined to provide an authenticated/identified speech recognition output when said comparison step determines that the audio provided to the speaker recognition module is the same as the audio provided to the speech recognition module.

20

According to a further aspect, there is provided a non-transitory storage medium having stored thereon software code which, when run on a suitable processor, performs any of the methods described above.

25

This has the advantage that the speech that is provided to the speech recognition function can be seen to be the same as the speech that was provided to the speaker recognition process, ensuring that the attack described above cannot be successful.

30 **Brief Description of Drawings**

Figure 1 illustrates a system in accordance with an aspect of the invention.

Figure 2 illustrates an example of a first device in the system of Figure 1.

35

Figure 3 illustrates an example of a second device in the system of Figure 1.

Figure 4 illustrates an example of a system in accordance with Figure 1.

Figure 5 illustrates an alternative example of a system in accordance with Figure 1.

5

Figure 6 is a flow chart, illustrating a method in accordance with an aspect of the invention.

10 Figure 7 illustrates an embodiment of a system in accordance with an aspect of the invention.

Figure 8 illustrates an alternative embodiment of a system in accordance with an aspect of the invention.

15 Figure 9 illustrates a further alternative embodiment of a system in accordance with an aspect of the invention.

Detailed Description

20

The description below sets forth example embodiments according to this disclosure. Further example embodiments and implementations will be apparent to those having ordinary skill in the art. Further, those having ordinary skill in the art will recognize that various equivalent techniques may be applied in lieu of, or in conjunction with, the
25 embodiments discussed below, and all such equivalents should be deemed as being encompassed by the present disclosure.

Figure 1 shows a system in accordance with an aspect of the invention. Specifically, Figure 1 shows a system 10, which includes a first device 12 and a second device 14.
30 An input signal can be supplied to the first device 12 on an input line 16. The first device 12 has a direct or indirect connection 18 to the second device. The second device 14 provides an output on an output connection 20.

The first device 12 includes a speaker recognition block 22. When the input signal
35 supplied to the first device 12 on the input line 16 represents speech, the speaker

recognition block 22 may for example determine whether or not the speaker is a specific individual, or may identify the speaker from a group of enrolled speakers.

5 The second device 14 includes a speech recognition block 24. When an input signal supplied to the second device 14 represents speech, the speech recognition block 24 may obtain information about the content of speech, for example in order to be able to determine the content of a command that was uttered by the speaker.

Figure 2 shows one possible form of the first device 12. In this example, the first
10 device 12 includes an interface 30, for connecting to other devices; a memory 32, for storing data and program instructions; and a processor 34, for performing operations in accordance with program instructions stored in the memory 32. Thus, with reference to Figure 1, the speaker recognition block 22 may be implemented by suitable program instructions stored in the memory 32, causing the processor 34 to perform the speaker
15 recognition functionality.

Figure 3 shows one possible form of the second device 14. In this example, the second device 14 includes an interface 40, for connecting to other devices; a memory
20 42, for storing data and program instructions; and a processor 44, for performing operations in accordance with program instructions stored in the memory 42. Thus, with reference to Figure 1, the speech recognition block 24 may be implemented by suitable program instructions stored in the memory 42, causing the processor 44 to perform the speech recognition functionality.

25 In the system 10 shown in Figure 1, the first device 12 and the second device 14 are typically physically separate devices, with some direct or indirect physical connection between them. For example, the first device 12 and the second device 14 may be separate integrated circuit devices, which may form part of the same product, or may be provided in separate products.

30 For example, the first device 12 and the second device 14 may be separate integrated circuits within a device such as a smartphone, or smart speaker, or another (portable) electronic device with some sort of voice operability, for example a tablet or laptop computer, a games console, a home control system, a home entertainment system, an
35 in-vehicle entertainment system, a domestic appliance, or the like.

In other examples, the first device 12 and the second device 14 may be separate (portable) electronic devices.

Figure 4 shows one example, where the first device 12 takes the form of a smartphone 60, or another (portable) electronic device with some sort of voice operability, for example a smart speaker, a tablet or laptop computer, a games console, a home control system, a home entertainment system, an in-vehicle entertainment system, a domestic appliance, or the like.

10 The smartphone 60 includes a speaker recognition block. Thus, the smartphone 60 also includes at least one microphone and, when it is determined that a sound detected by the microphone(s) contains speech, the speaker recognition block may for example determine whether or not the speaker is a specific individual, or may identify the speaker from a group of enrolled speakers.

15

The second device 14 is located remotely from the smartphone 60. For example, the second device 14 may take the form of a server 62 located in the cloud 64, accessible by the first device over a wireless communications network. The server 62 includes a speech recognition block. When a signal representing speech is supplied to the server 20 62, the speech recognition block may obtain information about the content of the speech, for example in order to be able to determine the content of a command.

Figure 5 shows another example, where the first device 12 takes the form of an accessory, for example a headset 80, which has a connection to a smartphone 82, or 25 another (portable) electronic device with some sort of voice operability, for example a smart speaker, a tablet or laptop computer, a games console, a home control system, a home entertainment system, an in-vehicle entertainment system, a domestic appliance, or the like. The accessory 80 may have a connection to the host device 82 over a short-range communications link, such as a Bluetooth or WiFi connection.

30

The accessory 80 includes a speaker recognition block. Thus, the accessory 80 includes at least one microphone and, when it is determined that a sound detected by the microphone(s) contains speech, the speaker recognition block may for example determine whether or not the speaker is a specific individual, or may identify the 35 speaker from a group of enrolled speakers.

The host device 82 includes a speech recognition block. When a signal representing speech is supplied to the host device 82, the speech recognition block may obtain information about the content of the speech, for example in order to be able to determine the content of a command.

5

Figure 6 is a flow chart, illustrating an example of a method of performing speech recognition, in a system as illustrated in Figure 1, 4, or 5, for example.

In step 100, a first device receives an audio signal representing speech.

10

At step 102, in the first device, a first data integrity check operation is performed on the received audio signal.

Any suitable data integrity check operation may be performed. For example, the step of performing a data integrity check operation may comprise generating a checksum value, or generating a Message Authentication Code based on the relevant part of the received audio signal. Typically, the relevant part of the received audio signal is a part that is intended to be forwarded to the second device. In some embodiments, the step of performing a data integrity check operation comprises obtaining a predetermined hash function from the relevant part of the received audio signal.

20

At step 104, in the first device, a speaker recognition process is performed on the received audio signal.

The speaker recognition process may be any suitable process, for example a voice biometric process, in which features are extracted from the received signal, and compared with a previously obtained model, or voiceprint, of the speech of one or more enrolled user. Depending on a required degree of security of the operation, the difference between the features and the model can be used to assess whether the comparison provides the required degree of confidence that the person speaking is the enrolled user.

30

When the audio signal is received, a voice activity detection process may be used to determine whether the signal represents speech. Only those parts of a received signal that do represent speech may be forwarded for subsequent processing. Once it has been determined that at least a part of the signal represents speech, a voice keyword

35

detection process may be performed to determine whether the signal represents a predetermined trigger phrase. In that case, the first data integrity check operation and/or the speaker recognition process may be performed on the received audio signal only if it is determined that the audio signal represents the predetermined trigger
5 phrase.

At step 106, the first device forwards the received audio signal to a second device, which has a speech recognition function.

10 Again, the first device may forward the received audio signal to the second device only if it has been determined that at least a part of the signal represents speech, and/or if it is determined that the audio signal represents the predetermined trigger phrase.

In addition, at step 108, the first device forwards an output of the first data integrity
15 check operation to the second device

At step 110, the second device receives the audio signal forwarded from the first device.

20 In addition, at step 112, the second device receives the output of the first data integrity check operation forwarded from the first device.

At step 114, the second device performs a second data integrity check operation on the audio signal forwarded from the first device. As described with reference to step 102,
25 any suitable data integrity check operation may be performed. For example, the step of performing a data integrity check operation may comprise generating a checksum value, or generating a Message Authentication Code based on the received audio signal. Generally, the second data integrity check operation performed in the second device is the same as the first data integrity check operation performed in the first
30 device. For example, when the step of performing a data integrity check operation in the first device comprises obtaining a predetermined hash function from the relevant part of the received audio signal, the step of performing the second data integrity check operation comprises performing the same predetermined hash function on the signal received in the second device.

The data integrity check operation is such that, provided that the operation is performed on the same input (i.e. the relevant part of the received audio signal) in steps 102 and 114, it will provide the same output. However, if it is performed on different inputs, for example because an attacker has inserted a rogue signal into the
5 second device, the data integrity check operation will provide different outputs.

At step 116, a result of performing a speech recognition function in the second device on the audio signal forwarded from the first device is used, only if an output of the second data integrity check operation performed in the second device matches the
10 output of the first data integrity check operation forwarded to the second device from the first device.

In addition, the result of performing the speech recognition function on the audio signal forwarded from the first device may be used only if a result of performing the speaker
15 recognition process on the received audio signal in the first device satisfies a predetermined criterion.

If, in step 114, the output of the second data integrity check operation does not match the output of the first data integrity check operation forwarded from the first device, this
20 may be due to an attack on the system, or it may be due to an unacceptably large Bit Error Rate (BER) on the communications link between the first device and the second device. If the speech recognition function has already been performed on the audio signal forwarded from the first device when it is determined that the output of the second data integrity check operation does not match the output of the first data
25 integrity check operation forwarded from the first device, the result of performing the speech recognition function on the audio signal forwarded from the first device may be deleted.

Thus, if an attacker inserts a rogue signal into the second device, but provides a signal
30 that passes the voice biometric identity check to the first device, the data integrity check operation will provide different outputs, and the positive result of the voice biometric identity check will not result in the rogue signal being taken as a genuine signal.

35 Although Figure 6 shows various steps being performed sequentially in the first device, and then various other steps being performed sequentially in the second device, it

should be noted that steps may be performed in parallel in some cases, and/or the order in which the steps are performed may differ from that shown in Figure 6. In particular, the order in which the steps are performed in either one of the devices may differ from that described. In addition, or alternatively, one or more of the steps that are performed in the second device may be performed before one or more of the steps that are performed in the first device.

Figure 7 is a block diagram, illustrating a first system for performing the method of Figure 6.

10

The system includes a first device 130 and a second device 132. The system also includes a microphone 134 for detecting sounds in the vicinity. In embodiments such as those shown in Figures 4 and 5, the microphone 134 may be provided on the first device 130.

15

The system also includes a buffer 136, which performs initial processing on the signal generated by the microphone 134, and stores the result. For example, speaker recognition and speech recognition processes typically operate on signals that have been divided into frames having a duration of 10-30 ms, and so the buffer 136 may perform this division. Specifically, the buffer 136 may divide the signal into frames, and may include a voice activity detection block, configured to determine which frames contain speech. In embodiments such as those shown in Figures 4 and 5, the buffer 136 may be provided in the first device 130.

25 In the system illustrated in Figure 7, the system also includes an optional voice keyword detection (VKD) block 138. This may be configured to detect whether the signal represents a predetermined keyword, or trigger phrase. The VKD block 138 may act only on frames that are determined as containing speech.

30 In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains speech. In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains the predetermined keyword.

35 The audio signal, or at least part of the audio signal, is passed to a voice biometric (VBio) or speaker recognition block 140. The speaker recognition block 140 performs

a speaker recognition process as described with reference to step 104 above, and determines whether or not the speaker is a specific individual (in the case of speaker verification), or identifies the speaker from a group of enrolled speakers (in the case of speaker identification). Thus, the output of the speaker recognition block 140 may be a
5 signal indicating whether or not the speaker is a specific individual, or may be a signal identifying the speaker.

The audio signal, or at least part of the audio signal, is also passed to a first data integrity check block 142. In this illustrated embodiment, the first data integrity check
10 block 142 calculates a hash value from the received audio signal. The first data integrity check block 142 may determine a hash value, for example according to the known SHA-256 algorithm as will be understood by skilled in the art, although other hash functions may also be appropriate.

15 The audio signal, or the same part of the audio signal that is passed to the speaker recognition block 140 and the first data integrity check block 142, is also passed to the second device 132.

Specifically, the audio signal, or the part of the audio signal, is passed to an automatic
20 speech recognition (ASR) function 144.

In addition, the audio signal, or the part of the audio signal, is passed to a second data integrity check block 146. In this illustrated embodiment, the second data integrity check block 146 calculates a hash value from the received audio signal, using the
25 same method that is used by the first data integrity check block 142.

The output of the speaker recognition block 140 and the output of the first data integrity check block 142 are both passed to a block 148 which determines whether to proceed with automatic speech recognition (ASR Go?). The output of the second data integrity
30 check block 146 is also passed to the block 148.

The block 148 compares the output of the first data integrity check block 142 with the output of the second data integrity check block 146.

35 The block 148 then controls the operation of the speech recognition function in the block 144, such that the speech recognition function is performed on the audio signal

forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

- 5 In addition, in this embodiment, the speech recognition function is performed on the audio signal forwarded from the first device only if the result of performing the speaker recognition process on the received audio signal satisfies a predetermined criterion, for example if the speaker is identified as a specific individual, or if the speaker is identified as a person who is authorised to issue a command to the device.

10

As described above, it is assumed that the output of the speaker recognition block 140 may be a signal indicating whether or not the speaker is a specific individual, or may be a signal identifying the speaker. In other embodiments, the output of the speaker recognition block 140 may for example be a signal indicating a likelihood of the speaker
15 being a specific individual, and the block 148 may determine based on criteria set in the second device 132 whether this likelihood is high enough to proceed with automatic speech recognition.

- If the block 144 is controlled such that the speech recognition function is performed on
20 the audio signal forwarded from the first device, then the output of that block 144 is forwarded for subsequent processing. For example, the audio signal may contain a spoken command relating to a function of the first device 130, in which case the output of the speech recognition block 144 may be passed to the first device 130.

- 25 **Figure 8** is a block diagram, illustrating a second system for performing the method of Figure 6.

- The system includes a first device 160 and a second device 162. The system also includes a microphone 164 for detecting sounds in the vicinity. In embodiments such
30 as those shown in Figures 4 and 5, the microphone 164 may be provided on the first device 160.

- The system also includes a buffer 166, which performs initial processing on the signal generated by the microphone 164, and stores the result. For example, speaker
35 recognition and speech recognition processes typically operate on signals that have been divided into frames having a duration of 10-30 ms, and so the buffer 136 may

perform this division. Specifically, the buffer 166 may divide the signal into frames, and may include a voice activity detection block, configured to determine which frames contain speech. In embodiments such as those shown in Figures 4 and 5, the buffer 166 may be provided in the first device 160.

5

In the system illustrated in Figure 8, the system also includes an optional voice keyword detection (VKD) block 168. This may be configured to detect whether the signal represents a predetermined keyword, or trigger phrase. The VKD block 168 may act only on frames that are determined as containing speech.

10

In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains speech. In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains the predetermined keyword.

15

The audio signal, or at least part of the audio signal, is passed to a voice biometric (VBio) or speaker recognition block 170. The speaker recognition block 170 performs a speaker recognition process as described with reference to step 104 above, and determines whether or not the speaker is a specific individual (in the case of speaker verification), or identifies the speaker from a group of enrolled speakers (in the case of speaker identification). Thus, the output of the speaker recognition block 170 may be a signal indicating whether or not the speaker is a specific individual, or may be a signal identifying the speaker.

20

25 The audio signal, or at least part of the audio signal, is also passed to a first data integrity check block 172. In this illustrated embodiment, the first data integrity check block 172 calculates a hash value from the received audio signal. The first data integrity check block 172 may determine a hash value, for example according to the known SHA-256 algorithm as will be understood by skilled in the art, although other
30 hash functions may also be appropriate.

The audio signal, or the same part of the audio signal that is passed to the speaker recognition block 170 and the first data integrity check block 172, is also passed to the second device 162.

35

Specifically, the audio signal, or the part of the audio signal, is passed to an automatic speech recognition (ASR) function 174. In this embodiment, the automatic speech recognition function may be performed as soon as the audio signal is received by the second device 162, without deliberate delay. Thus, for example, the speech
5 recognition function may be performed by the block 174 at least partly in parallel with the speaker recognition function being performed by the block 170, in order to minimize overall processing time.

In addition, the audio signal, or the part of the audio signal, is passed to a second data
10 integrity check block 176. In this illustrated embodiment, the second data integrity check block 176 calculates a hash value from the received audio signal, using the same method that is used by the first data integrity check block 172.

The output of the first data integrity check block 172 is passed to a block 178 which
15 determines whether to use the result of automatic speech recognition (ASR OK?). The output of the second data integrity check block 176 is also passed to the block 178.

The block 178 compares the output of the first data integrity check block 172 with the
20 output of the second data integrity check block 176.

The block 178 then controls the operation of the output of the speech recognition
function in the block 174, such that the output of the speech recognition function is
forwarded for subsequent processing only if the output of the second data integrity
check operation matches the output of the first data integrity check operation forwarded
25 from the first device.

If it is determined that the output of the second data integrity check operation matches
the output of the first data integrity check operation forwarded from the first device, the
result of the automatic speech recognition is forwarded for subsequent processing only
30 if the result of performing the speaker recognition process on the received audio signal
satisfies a predetermined criterion, for example if the speaker is identified as a specific
individual, or if the speaker is identified as a person who is authorised to issue a
command to the device.

In this embodiment, the output of the speaker recognition block 170 is passed to a
35 block 180 (Go/No Go?) which determines whether the received audio signal satisfies

the predetermined criterion. The block 180 is shown as being in the second device 162, but may be in the first device 160.

As described above, it is assumed that the output of the speaker recognition block 170
5 may be a signal indicating whether or not the speaker is a specific individual, or may be a signal identifying the speaker. In other embodiments, the output of the speaker recognition block 170 may for example be a signal indicating a likelihood of the speaker being a specific individual, and the block 180 may determine based on criteria set in the first device 160 or the second device 162 whether this likelihood is high enough to
10 use the result of the automatic speech recognition. The criteria set in the first device 160 or the second device 162 may depend on the result of the automatic speech recognition itself. For example, a command to a smart speaker device to play a particular piece of music may be acted upon even if the system has relatively low confidence that the speaker is a specified individual, because the adverse
15 consequences of wrongly playing that music are not great. However, a command to a banking app on a smartphone to make a (large) payment may be acted upon only if the system has relatively high confidence that the speaker is the specified individual, because the adverse consequences of wrongly making that payment are much greater.

20 If the block 180 determines that the output of the speech recognition function that is performed on the audio signal forwarded from the first device should be used, then the output of that block 174 is forwarded for subsequent processing. For example, the audio signal may contain a spoken command relating to a function of the first device 160, in which case the output of the speech recognition block 174 may be passed to
25 the first device 160. As already noted, the block 180 may itself be located in the first device 160.

If the block 180 determines that the output of the speech recognition function that is performed on the audio signal forwarded from the first device should not be used, then
30 the output of the speech recognition block 174 may be explicitly deleted from the second device 162, and may additionally or alternatively be deleted from the first device 160 if it has been sent there. Similarly, if the block 178 determines that the output of the second data integrity check operation does not match the output of the first data integrity check operation forwarded from the first device, then the output of
35 the speech recognition block 174 may be explicitly deleted from the second device 162,

and may additionally or alternatively be deleted from the first device 160 if it has been sent there.

Figure 9 is a block diagram, illustrating a third system for performing the method of Figure 6.

The system includes a first device 190 and a second device 192. The system also includes a microphone 194 for detecting sounds in the vicinity. In embodiments such as those shown in Figures 4 and 5, the microphone 194 may be provided on the first device 190.

The system also includes a buffer 196, which performs initial processing on the signal generated by the microphone 194, and stores the result. For example, speaker recognition and speech recognition processes typically operate on signals that have been divided into frames having a duration of 10-30 ms, and so the buffer 196 may perform this division. Specifically, the buffer 196 may divide the signal into frames, and may include a voice activity detection block, configured to determine which frames contain speech. In embodiments such as those shown in Figures 4 and 5, the buffer 196 may be provided in the first device 190.

In the system illustrated in Figure 9, the system also includes an optional voice keyword detection (VKD) block 198. This may be configured to detect whether the signal represents a predetermined keyword, or trigger phrase. The VKD block 198 may act only on frames that are determined as containing speech.

In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains speech. In some embodiments, the signal is only passed for subsequent processing if it is determined that the signal contains the predetermined keyword.

The audio signal, or at least part of the audio signal, is passed to a voice biometric (VBio) or speaker recognition block 200. The speaker recognition block 200 performs a speaker recognition process as described with reference to step 104 above, and determines whether or not the speaker is a specific individual (in the case of speaker verification), or identifies the speaker from a group of enrolled speakers (in the case of speaker identification). Thus, the output of the speaker recognition block 200 may be a

signal indicating whether or not the speaker is a specific individual, or may be a signal identifying the speaker.

The audio signal, or at least part of the audio signal, is also passed to a first data integrity check block 202. In this illustrated embodiment, the first data integrity check block 202 calculates a hash value from the received audio signal. The first data integrity check block 202 may determine a hash value, for example according to the known SHA-256 algorithm as will be understood by skilled in the art, although other hash functions may also be appropriate.

The audio signal, or the same part of the audio signal that is passed to the speaker recognition block 200 and the first data integrity check block 202, is also passed to the second device 192.

In the second device 192, the audio signal, or the part of the audio signal, is passed to a second data integrity check block 204. In this illustrated embodiment, the second data integrity check block 204 calculates a hash value from the received audio signal, using the same method that is used by the first data integrity check block 202.

The output of the first data integrity check block 202 is passed to a block 206 which determines whether to proceed with automatic speech recognition (ASR OK?). The output of the second data integrity check block 204 is also passed to the block 206.

The block 206 compares the output of the first data integrity check block 202 with the output of the second data integrity check block 204.

The block 206 then controls the operation of the speech recognition function in an automatic speech recognition (ASR) block 208 in the second device 192, such that the speech recognition function is performed on the audio signal forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

In this embodiment, the comparison of the output of the first data integrity check block 202 with the output of the second data integrity check block 204 automatic speech recognition function may be performed as soon as the audio signal is received by the second device 192, without deliberate delay. Thus, for example, the comparison may

be performed by the block 206 at least partly in parallel with the speaker recognition function being performed by the block 200, in order to minimize overall processing time. Further, if the comparison test is passed, the automatic speech recognition function may be performed by the ASR function 208 without additional deliberate delay.

- 5 Thus, for example, the speech recognition function may be performed by the block 208 at least partly in parallel with the speaker recognition function being performed by the block 170, in order to minimize overall processing time.

- 10 If it is determined that the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, the automatic speech recognition process is performed, but the result of the automatic speech recognition is forwarded for subsequent processing only if the result of performing the speaker recognition process on the received audio signal satisfies a predetermined criterion, for example if the speaker is identified as a specific individual,
- 15 or if the speaker is identified as a person who is authorised to issue a command to the device.

- In this embodiment, the output of the speaker recognition block 200 is passed to a block 210 (Go/No Go?) which determines whether the received audio signal satisfies
- 20 the predetermined criterion. The block 210 is shown as being in the second device 192, but may be in the first device 190.

- As described above, it is assumed that the output of the speaker recognition block 200 may be a signal indicating whether or not the speaker is a specific individual, or may be
- 25 a signal identifying the speaker. In other embodiments, the output of the speaker recognition block 200 may for example be a signal indicating a likelihood of the speaker being a specific individual, and the block 210 may determine based on criteria set in the first device 190 or the second device 192 whether this likelihood is high enough to use the result of the automatic speech recognition. The criteria set in the first device
- 30 190 or the second device 192 may depend on the result of the automatic speech recognition itself. For example, a command to a smart speaker device to play a particular piece of music may be acted upon even if the system has relatively low confidence that the speaker is a specified individual, because the adverse consequences of wrongly playing that music are not great. However, a command to a
- 35 banking app on a smartphone to make a (large) payment may be acted upon only if the

system has relatively high confidence that the speaker is the specified individual, because the adverse consequences of wrongly making that payment are much greater.

If the block 210 determines that the output of the speech recognition function that is performed on the audio signal forwarded from the first device should be used, then the output of that block 208 is forwarded for subsequent processing. For example, the audio signal may contain a spoken command relating to a function of the first device 190, in which case the output of the speech recognition block 208 may be passed to the first device 190. As already noted, the block 210 may itself be located in the first device 190.

If the block 210 determines that the output of the speech recognition function that is performed on the audio signal forwarded from the first device should not be used, then the output of the speech recognition block 208 may be explicitly deleted from the second device 192, and may additionally or alternatively be deleted from the first device 190 if it has been sent there.

Thus, in all of the illustrated embodiments, a result of performing the speech recognition function on the audio signal forwarded from the first device is used only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device. This may be achieved either by performing the speech recognition function on the audio signal forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, or by pre-emptively performing the speech recognition function on the audio signal forwarded from the first device and using the result thereof only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

30

The skilled person will recognise that some aspects of the above-described apparatus and methods, for example the discovery and configuration methods may be embodied as processor control code, for example on a non-volatile carrier medium such as a disk, CD- or DVD-ROM, programmed memory such as read only memory (Firmware), or on a data carrier such as an optical or electrical signal carrier. For many

applications, embodiments will be implemented on a DSP (Digital Signal Processor), ASIC (Application Specific Integrated Circuit) or FPGA (Field Programmable Gate Array). Thus the code may comprise conventional program code or microcode or, for example code for setting up or controlling an ASIC or FPGA. The code may also
5 comprise code for dynamically configuring re-configurable apparatus such as re-programmable logic gate arrays. Similarly the code may comprise code for a hardware description language such as Verilog™ or VHDL (Very high speed integrated circuit Hardware Description Language). As the skilled person will appreciate, the code may be distributed between a plurality of coupled components in communication with one
10 another. Where appropriate, the embodiments may also be implemented using code running on a field-(re)programmable analogue array or similar device in order to configure analogue hardware.

It should be noted that the above-mentioned embodiments illustrate rather than limit
15 the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. The word “comprising” does not exclude the presence of elements or steps other than those listed in a claim, “a” or “an” does not exclude a plurality, and a single feature or other unit may fulfil the functions of several units recited in the claims. Any reference
20 numerals or labels in the claims shall not be construed so as to limit their scope.

CLAIMS

1. A method of performing speech recognition, comprising:
at a first device:
5 receiving an audio signal representing speech;
performing a first data integrity check operation on the received audio signal;
performing a speaker recognition process on the received audio signal;
forwarding the received audio signal to a second device, wherein the second
device comprises a speech recognition function; and
10 forwarding an output of the first data integrity check operation to the second
device; and
at the second device:
receiving the audio signal forwarded from the first device;
receiving the output of the first data integrity check operation forwarded from the
15 first device;
performing a second data integrity check operation on the audio signal forwarded
from the first device; and
using a result of performing the speech recognition function on the audio signal
forwarded from the first device only if an output of the second data integrity check
20 operation matches the output of the first data integrity check operation forwarded from
the first device.
2. A method according to claim 1, further comprising:
using a result of performing the speech recognition function on the audio signal
25 forwarded from the first device only if a result of performing the speaker recognition
process on the received audio signal satisfies a predetermined criterion.
3. A method according to claim 2, comprising:
performing the speech recognition function on the audio signal forwarded from
30 the first device only if an output of the second data integrity check operation matches
the output of the first data integrity check operation forwarded from the first device, and
if the result of performing the speaker recognition process on the received audio signal
satisfies the predetermined criterion.
- 35 4. A method according to claim 1, comprising:

performing the speech recognition function on the audio signal forwarded from the first device;

determining if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device; and

5 using the result of performing the speech recognition function on the audio signal forwarded from the first device only if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

10 5. A method according to claim 4, further comprising:

if the output of the second data integrity check operation does not match the output of the first data integrity check operation forwarded from the first device, deleting the result of performing the speech recognition function on the audio signal forwarded from the first device.

15

6. A method according to claim 4, comprising performing the speech recognition function on the audio signal in the second device at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

20

7. A method according to claim 1, comprising determining in the second device if the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

25

8. A method according to claim 7, comprising, if it is determined in the second device that the output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device, performing the speech recognition function on the audio signal in the second device at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

30

9. A method according to claim 1, comprising performing the first data integrity check operation on the received audio signal only if it is determined that the audio signal represents a predetermined trigger phrase.

35

10. A method according to claim 1, wherein the step of performing the first data integrity check operation comprises obtaining a predetermined hash function, and wherein the step of performing the second data integrity check operation comprises
5 obtaining said predetermined hash function.

11. A system for performing speech recognition, comprising:
a first device and a second device,
wherein the first device comprises:
10 an input for receiving an audio signal representing speech;
a processor for performing a first data integrity check operation on the received audio signal, and for performing a speaker recognition process on the received audio signal; and
an interface for forwarding the received audio signal to the second device, and
15 for forwarding an output of the first data integrity check operation to the second device;
and
wherein the second device is configured for performing a speech recognition function, and comprises:
an interface for receiving the audio signal forwarded from the first device, and for
20 receiving the output of the first data integrity check operation forwarded from the first device; and
a processor for performing a second data integrity check operation on the audio signal forwarded from the first device;
wherein the system is configured to use a result of performing the speech
25 recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output of the first data integrity check operation forwarded from the first device.

12. A system according to claim 11, wherein the system is configured to use the
30 result of performing the speech recognition function on the audio signal forwarded from the first device only if a result of performing the speaker recognition process on the received audio signal satisfies a predetermined criterion.

13. A system according to claim 12, wherein the second device is configured to
35 perform the speech recognition function on the audio signal forwarded from the first device only if an output of the second data integrity check operation matches the output

of the first data integrity check operation forwarded from the first device, and if the result of performing the speaker recognition process on the received audio signal satisfies the predetermined criterion.

- 5 14. A system according to claim 11, wherein the system is configured for:
performing the speech recognition function on the audio signal forwarded from
the first device;
determining if an output of the second data integrity check operation matches the
output of the first data integrity check operation forwarded from the first device; and
10 using the result of performing the speech recognition function on the audio signal
forwarded from the first device only if the output of the second data integrity check
operation matches the output of the first data integrity check operation forwarded from
the first device.

- 15 15. A system according to claim 14, wherein the system is further configured for:
if the output of the second data integrity check operation does not match the
output of the first data integrity check operation forwarded from the first device, deleting
the result of performing the speech recognition function on the audio signal forwarded
from the first device.

- 20 16. A system according to claim 14, wherein the system is further configured for
performing the speech recognition function on the audio signal in the second device at
least partly in parallel with performing the speaker recognition process on the received
audio signal in the first device.

- 25 17. A system according to claim 11, wherein the second device is further configured
for determining if the output of the second data integrity check operation matches the
output of the first data integrity check operation forwarded from the first device, at least
partly in parallel with performing the speaker recognition process on the received audio
30 signal in the first device.

18. A system according to claim 17, wherein the second device is further configured
for, if it is determined that the output of the second data integrity check operation
matches the output of the first data integrity check operation forwarded from the first
35 device, performing the speech recognition function on the audio signal in the second

device at least partly in parallel with performing the speaker recognition process on the received audio signal in the first device.

19. A system according to claim 11, wherein the first device is configured for
5 performing the first data integrity check operation on the received audio signal only if it is determined that the audio signal represents a predetermined trigger phrase.

20. A system according to claim 11, wherein performing the first data integrity check
operation comprises obtaining a predetermined hash function, and wherein performing
10 the second data integrity check operation comprises obtaining said predetermined hash function.




Fig. 1

2/8




Fig. 2




Fig. 3




Fig. 4




Fig. 5

5/8

First Device

Second Device




Fig. 6




Fig. 7




Fig. 8




Fig. 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2019/052401

A. CLASSIFICATION OF SUBJECT MATTER

INV. G10L15/22 G10L15/30
ADD. G10L17/00 G06F21/64

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)

G10L G06F

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y A	US 2016/118044 A1 (BONDAREV PAVEL [CA] ET AL) 28 April 2016 (2016-04-28) paragraphs [0077], [0080] - [0092]	1-3, 10-13,20 4-9, 14-19
Y A	US 2016/019894 A1 (TSUJIKAWA MISAKI [JP] ET AL) 21 January 2016 (2016-01-21) paragraphs [0059], [0081], [0082], [0087], [0089], [0090], [0091], [0092], [0100], [0101], [0105]; figures 1A,3B,4	1-3, 10-13,20 4-9, 14-19
A	GB 2 363 236 A (MOTOROLA INC [US]) 12 December 2001 (2001-12-12) abstract; figure 3	1-20
A	EP 2 669 889 A2 (SAMSUNG ELECTRONICS CO LTD [KR]) 4 December 2013 (2013-12-04) claim 1; figure 2	9,19



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

"E" earlier application or patent but published on or after the international filing date

"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)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"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

"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

"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

"&" document member of the same patent family

Date of the actual completion of the international search

9 October 2019

Date of mailing of the international search report

17/10/2019

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Krembel, Luc

INTERNATIONAL SEARCH REPORT

International application No.
PCT/GB2019/052401

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 additional 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 fees, this Authority did not invite payment of additional fees.
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 and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-20

Distributed speech recognition with integrity checking.

1.1. claims: 4-8, 14-18

Reducing overall latency

1.2. claims: 9, 19

Trigger phrase based processing

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2019/052401

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016118044 A1	28-04-2016	CA 2869245 A1	27-04-2016
		US 2016118044 A1	28-04-2016

US 2016019894 A1	21-01-2016	JP 2016029468 A	03-03-2016
		US 2016019894 A1	21-01-2016
		US 2019228775 A1	25-07-2019

GB 2363236 A	12-12-2001	AT 336778 T	15-09-2006
		AU 8383901 A	17-12-2001
		AU 2001283839 B2	16-03-2006
		CA 2407791 A1	13-12-2001
		CN 1437745 A	20-08-2003
		DE 60122327 T2	13-09-2007
		EA 200201298 A1	24-04-2003
		EP 1292942 A1	19-03-2003
		ES 2269443 T3	01-04-2007
		GB 2363236 A	12-12-2001
		JP 5122716 B2	16-01-2013
		JP 2003536108 A	02-12-2003
		US 2004039569 A1	26-02-2004
		WO 0195311 A1	13-12-2001
		ZA 200208371 B	26-01-2004

EP 2669889 A2	04-12-2013	CN 103456306 A	18-12-2013
		CN 106297802 A	04-01-2017
		CN 106448678 A	22-02-2017
		EP 2669889 A2	04-12-2013
		EP 3001414 A1	30-03-2016
		KR 20130133629 A	09-12-2013
		US 2013325484 A1	05-12-2013
		US 2017162198 A1	08-06-2017
		US 2019035399 A1	31-01-2019
