This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Bäckström, Tom

Speech Coding, Speech Interfaces and IoT - Opportunities and Challenges

Published in:
2018 52nd Asilomar Conference on Signals, Systems, and Computers

DOI:
[10.1109/ACSSC.2018.8645502](https://doi.org/10.1109/ACSSC.2018.8645502)

Published: 01/01/2018

Document Version
Peer reviewed version

Please cite the original version:
Bäckström, T. (2018). Speech Coding, Speech Interfaces and IoT - Opportunities and Challenges. In *2018 52nd Asilomar Conference on Signals, Systems, and Computers* (pp. 1931 - 1935). (Conference Record of the Asilomar Conference on Signals Systems and Computers). IEEE. <https://doi.org/10.1109/ACSSC.2018.8645502>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

SPEECH CODING, SPEECH INTERFACES AND IOT – OPPORTUNITIES AND CHALLENGES

Tom Bäckström

Aalto University, Department of Signal Processing and Acoustics, Espoo, Finland

ABSTRACT

Recent speech and audio coding standards such as 3GPP Enhanced Voice Services match the foreseeable needs and requirements in transmission of speech and audio, when using current transmission infrastructure and applications. Trends in Internet-of-Things technology and development in personal digital assistants (PDAs) however begs us to consider future requirements for speech and audio codecs. The opportunities and challenges are here summarized in three concepts: collaboration, unification and privacy. First, an increasing number of devices will in the future be speech-operated, whereby the ability to focus voice commands to a specific devices becomes essential. We therefore need methods which allows collaboration between devices, such that ambiguities can be resolved. Second, such collaboration can be achieved with a unified and standardized communication protocol between voice-operated devices. To achieve such collaboration protocols, we need to develop distributed speech coding technology for ad-hoc IoT networks. Finally however, collaboration will increase the demand for privacy protection in speech interfaces and it is therefore likely that technologies for supporting privacy and generating trust will be in high demand.

Index Terms— speech coding, acoustic front-end, acoustic sensor networks, privacy

1. INTRODUCTION

A dominant trend in speech processing technology is smart-devices which include speech interfaces. Personal digital assistants such as Siri [1] and Google [2] have become commonplace features of smart-phones, while smart-speakers with speech interfaces such as Amazon Echo [3], Google Home [4] and Apple Home [5] have increased their sales greatly [6]. A central benefit of such speech interfaces is that they provide a unified hands-free access to services, including news and information retrieval as well as control of local actuators and devices. Specifically, in comparison to a tactile interface, requirements to the users physical proximity to the device is relaxed. This is beneficial especially when accessing multiple devices in different locations.

Fig. 1. Illustration of a distributed acoustic sensor network.

In terms of user-interface design, speech interfaces are thus beneficial in the sense that the user does not need to focus on the location of devices, but can concentrate on the primary task. For example, if the user wants to turn off the lights, he can say “*Computer, lights off*”, even if he does not know where the computer is located. In contrast, with a tactile interface, the user first has to locate and reach the physical device, before he can even start to operate it. Thus with speech interfaces, the user’s *attention is shifted from device to task*, which is beneficial since focusing on the device is a distraction with respect to the task. Clearly the benefit of such speech interfaces is emphasized when the number of devices increases.

In a parallel development, Internet of Things (IoT) technologies have recently received much attention [7]. To get the advantage of a large range of devices, they need to be connected and we need to be able to access them through a unified interface. For unification, the user interface does not necessarily have to be a speech interface, but again, a speech interface relaxes the requirement for physical proximity in comparison to tactile interfaces. Speech interfaces are thus often the logical choice for IoT devices.

Further significant benefits of IoT devices can be achieved with devices equipped with microphones. A collection of devices can then act like a (wireless) acoustic sensor network (WASN), to provide efficient pick-up of acoustic signals even in adverse conditions [8] (see Fig. 1). Specifically, it is possible to apply multi-microphone enhancement methods such as noise attenuation, beamforming and source separation, to extract individual spatially separated sound sources [9].

Communication between such devices requires speech and audio coding techniques, which is a field which has recently seen several advances. Namely, the 3GPP has stan-

standardized the Enhanced Voice Services (EVS) for telecommunications, while MPEG has standardized Unified Speech and Audio coding for speech and audio broadcasting, as well as 3D-Audio for multichannel broadcasting, and at the time of writing, Bluetooth is about to publish their new codec for low-delay speech and audio communication [10–13]. These codecs are state-of-the-art technology and provide the best known compromise between quality and available resources. However, none of these codecs are suitable for the IoT scenario. The available codecs are for communication between single devices, whereas for IoT, we need simultaneous communication between multiple devices.

In communication between multiple independent devices in ad-hoc networks, conventional single-device codecs give a suboptimal performance, because each device can then in theory be sending the same signal, which is redundant and thus inefficient. Methods of distributed source coding can then be used to avoid such redundancy [14–18]. It is then in principle possible to achieve optimal compromise between rate and distortion, while simultaneously optimizing resource consumption between devices. Available methods for distributed source coding have however not been able to provide competitive performance in a single channel setting.

In contrast, we have recently presented methods which extend single device codecs with features of distributed coding, with the objective of improving quality when multiple devices are available, that is, when we can use distributed coding [19–21]. The initial approach here is to decorrelate quantization error between devices, such that they can be removed by post-filtering methods [22, 23]. While our current methods do not yet reach the full benefit of distributed source coding, they do retain competitive quality in single-device scenarios while enabling additional devices to improve quality.

Given this *scientific* background, the purpose of the current work is to highlight some problems in state-of-the-art *implementations* in commercial products, and to propose a path to resolving these issues to the benefit of all actors. In short, this paper argues for a standard communication protocol between IoT devices, a unified speech and audio codec for both telecommunication and digital services, which is designed to ensure the privacy of the users. Indeed, the presented arguments point to a system where minimization of the amount data transmitted provides best resource efficiency, minimizes hardware requirements and provides best privacy.

The challenges with state-of-the-art implementations are described through two use cases, the cafeteria and the bedroom, described in the following section, which have been chosen to represent typical scenarios for a large percentage of the users. As a solution to such problems, in Section 3, we propose a systems design where devices are allowed to collaborate when in the same acoustic space. In the discussion section we further comment also on the ethical and legal landscape of speech processing in multi-device scenarios.

2. USE CASES

2.1. Cafeteria

Suppose two friends, Alice and Bob, meet at their favorite cafeteria. Alice happens to prefer products from the Apple brand, while Bob uses an Android phone. The cafeteria also provides access to Amazon Echo devices. Alice and Bob then have access to *three* different speech interfaces. Which interface should they use?

I am in no position to prescribe preference among products, but the point is that Alice and Bob are *forced to focus on devices* to choose their speech interface. In contrast, a central promise of speech interfaces was to shift attention away from devices to the task at hand. Current products give no solution to that issue. If Alice and Bob want to call their mutual friend Steven, choosing a device for that task is a distraction, which reduces their user experience.

This unsatisfactory experience occurs because current products which provide speech interfaces are *not interoperable* across manufacturers and service providers. Though the argument for speech interfaces was centered around the promise for a shift of attention from device to task, current products are unable to fulfill this promise because users have to choose the device and service jointly. Specifically, the available products are stellar products when observed in isolation, but when devices of multiple manufacturers are used in parallel, performance is unsatisfactory.

Conversely, to achieve the shift of attention from device to task, devices should collaborate, which means that they should be interoperable. Collaboration across different manufacturers is probably best achieved with a common application programming interface (API). To maximize the number of manufacturers which support such an API, it seems obvious that the API should be open or at least free-to-use. Several different standardization organizations already have the necessary procedures for adopting such technology, whereby standardization is a viable alternative.

2.2. Bedroom

When I am in my bedroom and say “*Computer, lights off*,” the lights are turned off. This is an extremely private event; it provides insight into my habits in the bedroom. Nobody, with the exception of my preferred romantic partner, has access to my bedroom, not even my best friends. It is therefore inconceivable to me how someone would provide access to his or her bedroom to some company providing a cloud service.

I would personally contend that the privacy provided by current products is not a happy state of affairs, but objectively, increasing collaboration across devices between different manufacturers, will increase privacy concerns exponentially. Speech interfaces are often served through cloud servers and if all speech input would be shared among the cloud servers of all nearby devices, the weakest link would

define the level of privacy of all users. It would be very hard for the user to determine what the actual level of privacy is and consequently, it would be extremely hard for the user to trust the system.

Privacy becomes a concern when information is transmitted or stored. In the most basic form of privacy protection, users need to be able to protect their own information to be transmitted or stored without their consent. Some of the main issues are that users often do not know what information is transmitted, to whom, or how information can be used. Moreover, companies which provide services and manufacture devices generally have more expertise than the average user, whereby they companies both know better what information is valuable but also may know how to extract information without alerting the average users' suspicion.

Among privacy and security advocates it is widely agreed that the main solution cannot be education of users [24]. Companies and potential criminals always have superior expertise compared to average users and it is unreasonable to expect that we could educate users to such an extent that the difference in expertise would be overcome. Instead, technology should be designed to protect the users. Collaboration among speech interfaces should be designed such that it minimizes information transfer to the minimum amount required to fulfill the task at hand. In particular, to follow an ethically aligned design paradigm [25], services and information transfer should be activated only through informed and meaningful consent.

3. LOCAL COLLABORATION

A solution which fulfills the above requirements is a systems structure where only local collaboration is permitted by default (see Fig. 2). All devices which are within earshot, that is, which can receive my speech with their microphones, already have access to the speech signal. Collaboration between such local devices, $\mathcal{A}$ and $\mathcal{B}$, does thus not expose the user to new threats to privacy. Only remote devices $\mathcal{C}$, remote in space or time, expose the user's privacy, whereby transmission to such devices require separate and specific authorization. By remote in space or time, respectively, we refer to telecommunications and storage applications.

The benefit of local collaboration is that it allows resource-optimization among local devices. Each device could in theory run its own noise attenuation, source separation, voice activity detection, keyword spotting and speech recognition algorithms and more, but those algorithms would then be redundant, since multiple devices run the same tasks. For improved efficiency and accuracy, we can instead use, for example, a distributed voice activity detector among local devices [26], distributed beamforming [16] or distributed speech recognition [27].

It is here evident that if a local device, such as $\mathcal{A}$ and $\mathcal{B}$, has been compromised to allow access to an illegitimate

Fig. 2. Illustration of local collaboration between devices in the same acoustic space, where connections to the outside are forbidden without specific authorization.

user, then that device can leak microphone signals to the outside. The current proposal provides no tools for avoiding such leaks. Device-level security however falls outside the current scope and in contrast, this proposal is only concerned with information flows between legitimate users and services.

Specifically, for example at home, my family members have access to the same network and we trust each other. However, if I go into the study and close the room behind me, that is a clear sign that I want to have privacy. Whereas all the devices of family members are generally allowed to collaborate, when I close the door behind me, I limit access to only those devices which are with me in the same room.

To achieve such functionality, devices have to be able to sense whether they are in the same room or acoustic space, and such sensing must not reveal any private information. Since the objective is to determine access rights to an acoustic signal, it is easiest to use the acoustic signal also as the basis of access right management. Specifically, each device can extract data from the microphone signal, apply error correction to obtain a fingerprint which is associated to that unique acoustic environment and use that fingerprint to cipher communication among local devices [28]. Since the fingerprint is used as a key for encryption, it is never transmitted between devices and no sensitive data is exchanged unencrypted. Moreover, since only such devices which reside in the same acoustic space can know the key, only local devices can decrypt the messages and participate in the collaboration.

4. SYSTEMS DESIGN

Conceptually, the proposed functionality can be interpreted as a new abstraction layer, which lumps together all local devices, to only give out the jointly estimated speech signal (see Fig. 3). Such an abstraction layer comprising local devices can be named the *fog* as a parallel to the abstraction of remote servers as a *cloud*. Remote services do not need to know the particulars of local devices – the cloud needs access to only the output sound – whereby it is well-warranted to hide such information from the cloud. Any local device can then act as a merge-node where the desired speech signal is estimated from distributed sensor nodes (see Fig. 1).

Fig. 3. Illustration of the proposed systems design for a distributed acoustic speech interface acting as an abstraction layer to unify access to local devices.

To transmit speech to remote recipients then requires that local devices reach a consensus on authorization [29]. Such authorization mechanisms must prevent spoofing attacks from any single device which implies that the authorization mechanism should be distributed and immutable. In other words, authorization likely requires distributed ledger technologies, which are colloquially known as block-chain technologies [30].

5. DISCUSSION AND CONCLUSIONS

The presented use-cases lead to the conclusion that, for voices services in IoT devices, it would be beneficial to develop a new speech and audio coding standard applicable to multi-device scenarios. In comparison to prior standards, the novelty of such a standard would be that it incorporates methods of both distributed source coding and built-in provisions for protecting the privacy of its users. Central benefits of such a standard include that 1. users can access services and devices without the necessity of physical proximity to the device, thus releasing attention from the device to the task at hand, 2. microphones in several devices can be used to jointly estimate the speech signal for better signal quality, 3. available hardware and software resources can be optimized to minimize costs in infrastructure and energy consumption and 4. a unified and transparent approach to privacy helps to build trust to technology among users as well as allows third-party audits of privacy.

Recent work on privacy has highlighted the need for the users to be able to manage their own information [31] and the general data protection regulation recently legislated by the European Union [32] enforces such a model in practice. In comparison, speech information is by nature a dialogue between two or more parties and it seems intuitively clear that ownership of and access rights to such a dialogue is always shared among the participants. The proposed system takes a first step in that direction by enabling a distributed access management system, where access to data is granted through a consensus among participants.

This approach to privacy is in stark difference to current products, where access management is in some cases governed by a single primary user who owns the device. This exposes secondary users to a breach in trust. For example, if I have a smart speaker at home and I could remotely access the voice command history, then while away from home, I could listen to the commands other family members make at home. Though I do not have the expertise to determine whether this is legal, certainly this situation seems ethically unreasonable.

In conclusion, smart speakers and other smart devices have recently become commonplace products and they provide useful hands-free access to services through voice operated user-interfaces. The current conclusion is that the usability, signal quality and resource efficiency would improve if all IoT devices could collaborate and jointly provide a speech user-interface. The proposed approach is a unified multi-device speech and audio codec for telecommunication and digital services. An essential part of such a codec would be provisions for enforcing and maximizing users' privacy.

6. REFERENCES

- [1] "Siri homepage," <https://www.apple.com/siri/>, Accessed 30.10.2018.
- [2] "Google homepage," <https://www.google.com/intl/en-GB/landing/now/>, Accessed 30.10.2018.
- [3] "Amazon echo homepage," <https://www.amazon.com/echo>, Accessed 30.10.2018.
- [4] "Google home homepage," https://store.google.com/au/product/google_home, Accessed 30.10.2018.
- [5] "Apple home homepage," <https://www.apple.com/homepod>, Accessed 30.10.2018.
- [6] Canals, "Global smart speaker shipments grew 187% year on year in Q2 2018, with China the fastest-growing market," <https://www.canalys.com/newsroom/global-smart-speaker-shipments-grew-187-year-on-year-in-q2-2018-with-china-the-fastest-growing-market>, Accessed 30.10.2018.
- [7] D Miorandi, S Sicari, F De Pellegrini, and I Chlamtac, "Internet of things: Vision, applications and research challenges," *Ad Hoc Networks*, vol. 10, no. 7, pp. 1497–1516, 2012.
- [8] A Bertrand, "Applications and trends in wireless acoustic sensor networks: A signal processing perspective," in *Communications and Vehicular Technology in the*

- Benelux (SCVT), 2011 18th IEEE Symposium on*, 2011, pp. 1–6.
- [9] J Benesty, M Sondhi, and Y Huang, *Springer Handbook of Speech Processing*, Springer, 2008.
 - [10] 3GPP, *TS 26.445, EVS Codec Detailed Algorithmic Description; 3GPP Technical Specification (Release 12)*, 2014.
 - [11] ISO/IEC 23003–3:2012, “MPEG-D (MPEG audio technologies), Part 3: Unified speech and audio coding,” 2012.
 - [12] J Herre, J Hilpert, A Kuntz, and J Plogsties, “MPEG-H Audio – The new standard for universal spatial/3D audio coding,” *Journal of the Audio Engineering Society*, vol. 62, no. 12, pp. 821–830, 2015.
 - [13] T Bäckström, *Speech Coding with Code-Excited Linear Prediction*, Springer, 2017.
 - [14] Z Xiong, A D Liveris, and S Cheng, “Distributed source coding for sensor networks,” *IEEE Signal Process. Mag.*, vol. 21, no. 5, pp. 80–94, 2004.
 - [15] H Dong, J Lu, and Y Sun, “Distributed audio coding in wireless sensor networks,” in *Computational Intelligence and Security, 2006 International Conference on*. IEEE, 2006, vol. 2, pp. 1695–1699.
 - [16] G Barriac, R Mudumbai, and U Madhow, “Distributed beamforming for information transfer in sensor networks,” in *Information Processing in Sensor Networks, 2004. IPSN 2004. Third International Symposium on*. IEEE, 2004, pp. 81–88.
 - [17] Y Jia, Y Luo, Y Lin, and I Kozintsev, “Distributed microphone arrays for digital home and office,” in *Proc. ICASSP*, 2006, vol. 5.
 - [18] A Majumdar, K Ramchandran, and I Kozintsev, “Distributed coding for wireless audio sensors,” in *Applications of Signal Processing to Audio and Acoustics, IEEE Workshop on*, 2003, pp. 209–212.
 - [19] T Bäckström, F Ghido, and J Fischer, “Blind recovery of perceptual models in distributed speech and audio coding,” in *Proc. Interspeech*, 2016, pp. 2483–2487.
 - [20] T Bäckström, J Fischer, and S Das, “Dithered quantization for frequency-domain speech and audio coding,” in *Proc. Interspeech*, 2018, pp. 3533–3537.
 - [21] T Bäckström and J Fischer, “Fast randomization for distributed low-bitrate coding of speech and audio,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 26, no. 1, January 2018.
 - [22] S Das and T Bäckström, “Postfiltering with complex spectral correlations for speech and audio coding,” *Proc. Interspeech 2018*, pp. 3538–3542, 2018.
 - [23] S Das and T Bäckström, “Postfiltering using log-magnitude spectrum for speech and audio coding,” *Proc. Interspeech 2018*, pp. 3543–3547, 2018.
 - [24] B Schneier, “Stop trying to fix the user,” *IEEE Security & Privacy*, vol. 14, no. 5, pp. 96–96, 2016.
 - [25] The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, *Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*, IEEE, version 2 edition, 2017.
 - [26] V Berisha, H Kwon, and A Spanias, “Real-time implementation of a distributed voice activity detector,” in *Sensor Array and Multichannel Processing, 2006. Fourth IEEE Workshop on*. IEEE, 2006, pp. 659–662.
 - [27] ETSI, *ES 201 108 standard – Distributed speech recognition (V1.1.3)*, 2003.
 - [28] D Schürmann and S Sigg, “Secure communication based on ambient audio,” *IEEE Trans. Mobile Comput.*, vol. 12, no. 2, pp. 358–370, 2013.
 - [29] G F Coulouris, J Dollimore, and T Kindberg, *Distributed systems: concepts and design*, Pearson education, 2005.
 - [30] A Deshpande, K Stewart, L Lepetit, and S Gunashekar, “Distributed ledger technologies/blockchain: Challenges, opportunities and the prospects for standards,” *Overview report The British Standards Institution (BSI)*, 2017.
 - [31] K Kuikkaniemi, A Poikola, and O Pitkänen, “My data fueling wellbeing applications,” in *Proceedings of the 2014 Workshops on Advances in Computer Entertainment Conference*. ACM, 2014, p. 18.
 - [32] European Parliament, “Directive 95/46/EC general data protection regulation,” 2016.