Researchers explore inaudible voice command attack

September 8 2017, by Nancy Owano

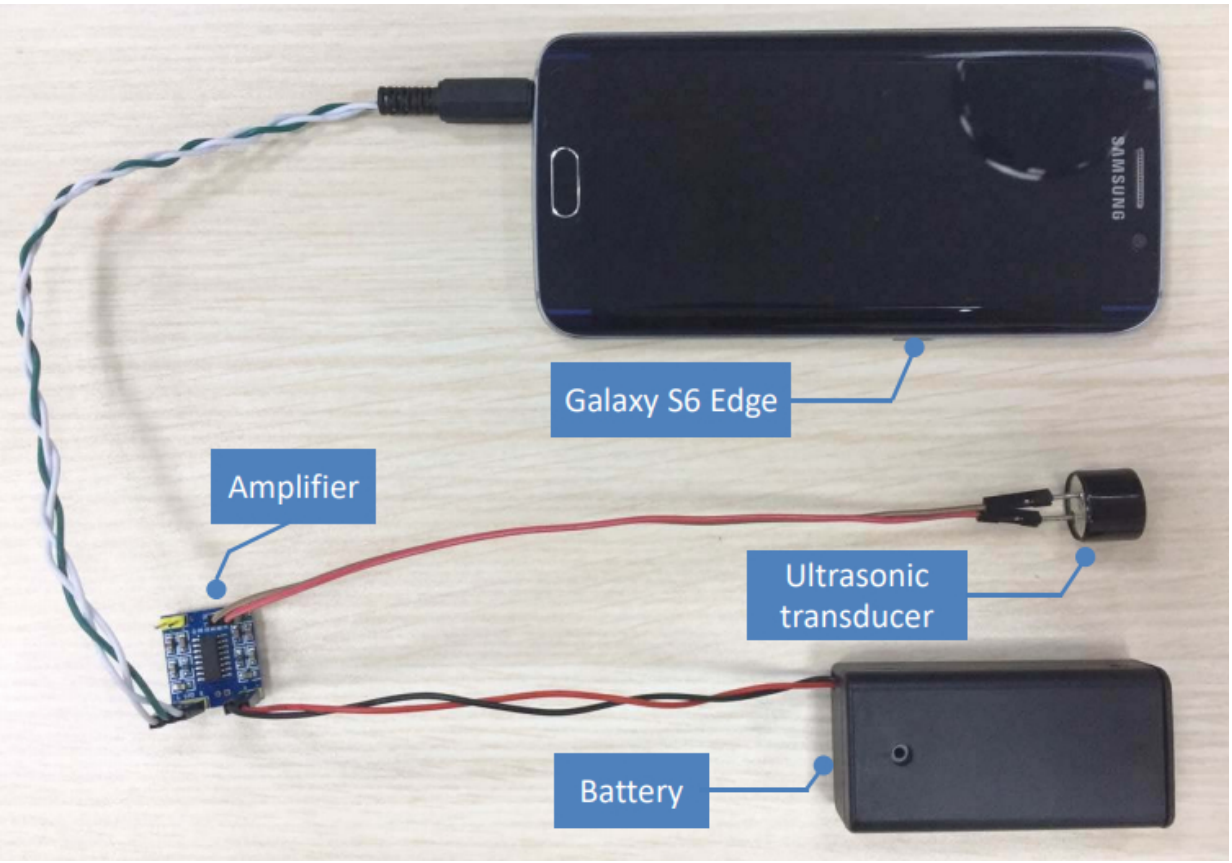

Portable attack implementation with a Samsung Galaxy S6 Edge smartphone, an ultrasonic transducer and a low-cost amplifier. The total price for the amplifier, the ultrasonic transducer plus the battery is less than $3. Credit: arXiv:1708.09537 [cs.CR]

(Tech Xplore)—The novelty has worn off; we are now quite accustomed to the technology advances leveraged by leading vendors that allow the daily use of human-device interactions via voice.

Well, six researchers, Guoming Zhang, Chen Yan, Xiaoyu Ji, Taimin Zhang, Tianchen Zhang, Wenyuan Xuhave, have an important wake-up call. They have determined voice assistants that can be attacked by inaudible commands. They are discussing what they call the DolphinAttack.

"While we might not hear the bad guys talking, our computers clearly can," said *Fast Company*'s Mark Wilson in *CO.DESIGN*.

(Dolphins use sounds like whistles and squeaks and clicks. They can hear high frequencies that humans cannot hear.)

"DolphinAttack: Inaudible Voice Commands" is now on *arXiv*. The paper was accepted to the 24th ACM Conference on Computer and Communications Security taking place next month in Dallas.

In brief, these researchers have discussed an attack on popular voice assistants "by commanding these assistants using speech that has been shifted to ultrasonic ranges," said Cory Doctorow in *Boing Boing* on Thursday.

Vishal Laul in *Neowin* also said that the researchers managed to modulate the frequency of a voice command from a human into ultrasonic frequencies. Laul said that while we cannot hear it, the microphones on a number of consumer devices can hear it.

They said by "leveraging the nonlinearity of the microphone circuits, the modulated low frequency audio commands can be successfully demodulated, recovered, and more importantly interpreted by the speech

recognition systems."

They said that "We validate DolphinAttack on popular speech recognition systems, including Siri, Google Now, Samsung S Voice, Huawei HiVoice, Cortana and Alexa."

The team managed to not only activate assistants but perform commands. They could initiate a call on a phone, switch a phone to airplane mode and manipulate a car's navigation system, depending on the voice assistant they were using.

Their paper stated the attack distances varied from 2 cm to a maximum value of 175 cm and showed a great variation across devices.

Meanwhile, said India Ashok in *IBTimes*, DolphinAttack's effectiveness would depend on the device. For instance, using the attack technique to command an Amazon Echo to open a door would not be feasible as it would require the attacker to already be inside the target's house.

"From a practical standpoint, few users should fear a DolphinAttack today. The attacker needs to be in near proximity, know the device you are using and be in an environment with little background noise," said Bret Kinsella in *Voicebot*. Kinsella added: "However, the researchers have performed a great service in bringing this issue to light now and offering practical suggestions for manufacturers to patch this vulnerability."

So what might be solutions in the form of safeguards?

Can voice controllable systems avoid inaudible voice command attacks?

Wilson, in *CO.DESIGN*, quoted an industrial designer, Gadi Amit, who said that the design of such microphones makes them difficult to secure

from this type of attack. "Microphones' components themselves vary in type, but most use air pressures that probably cannot be blocked from ultrasounds," Amit said. Many popular mics "transform turbulent air–or sound waves–into electrical waves."

The authors said they proposed hardware and software defense solutions. They said they validated it was feasible to detect DolphinAttack by classifying audios using supported vector machine (SVM). They suggested to re-design voice-controllable systems "to be resilient to inaudible voice command attacks."

More information: DolphinAtack: Inaudible Voice Commands, arXiv:1708.09537 [cs.CR] , DOI: 10.1145/3133956.3134052 , arxiv.org/abs/1708.09537

Abstract
Speech recognition (SR) systems such as Siri or Google Now have become an increasingly popular human-computer interaction method, and have turned various systems into voice controllable systems(VCS). Prior work on attacking VCS shows that the hidden voice commands that are incomprehensible to people can control the systems. Hidden voice commands, though hidden, are nonetheless audible. In this work, we design a completely inaudible attack, DolphinAttack, that modulates voice commands on ultrasonic carriers (e.g., f > 20 kHz) to achieve inaudibility. By leveraging the nonlinearity of the microphone circuits, the modulated low frequency audio commands can be successfully demodulated, recovered, and more importantly interpreted by the speech recognition systems. We validate DolphinAttack on popular speech recognition systems, including Siri, Google Now, Samsung S Voice, Huawei HiVoice, Cortana and Alexa. By injecting a sequence of inaudible voice commands, we show a few proof-of-concept attacks, which include activating Siri to initiate a FaceTime call on iPhone, activating Google Now to switch the phone to the airplane mode, and

even manipulating the navigation system in an Audi automobile. We propose hardware and software defense solutions. We validate that it is feasible to detect DolphinAttack by classifying the audios using supported vector machine (SVM), and suggest to re-design voice controllable systems to be resilient to inaudible voice command attacks.