Hey KIMYA!

Is My Smart Speaker Spying on Me?

Taking Control of Sensor Privacy Through Isolation and Amnesia

Piet De Vaere and Adrian Perrig
ETH Zürich

Smart speakers come with a paradox.

Require a high-level of trust in vendor honesty & competence

Vendors have repeatedly broken this trust

The Guardian

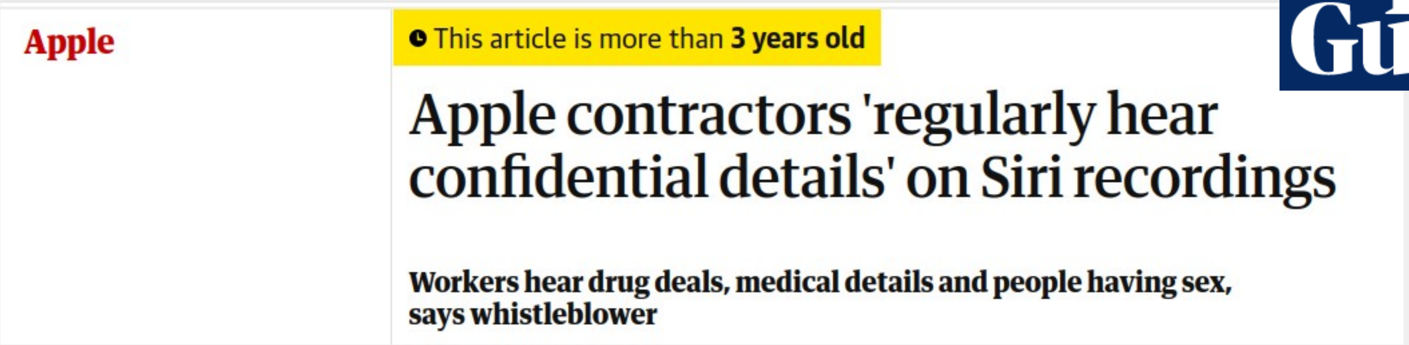

Apple

This article is more than **3 years old**

Apple contractors 'regularly hear confidential details' on Siri recordings

Workers hear drug deals, medical details and people having sex, says whistleblower

Today's status indicators are opaque & (probably) not very secure.

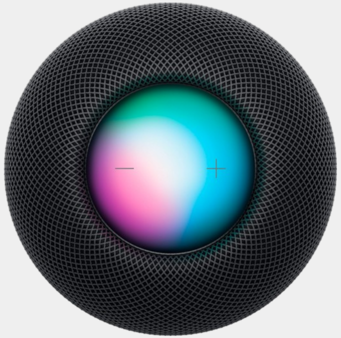

Can we do better?

First attempt: a microphone gateway.

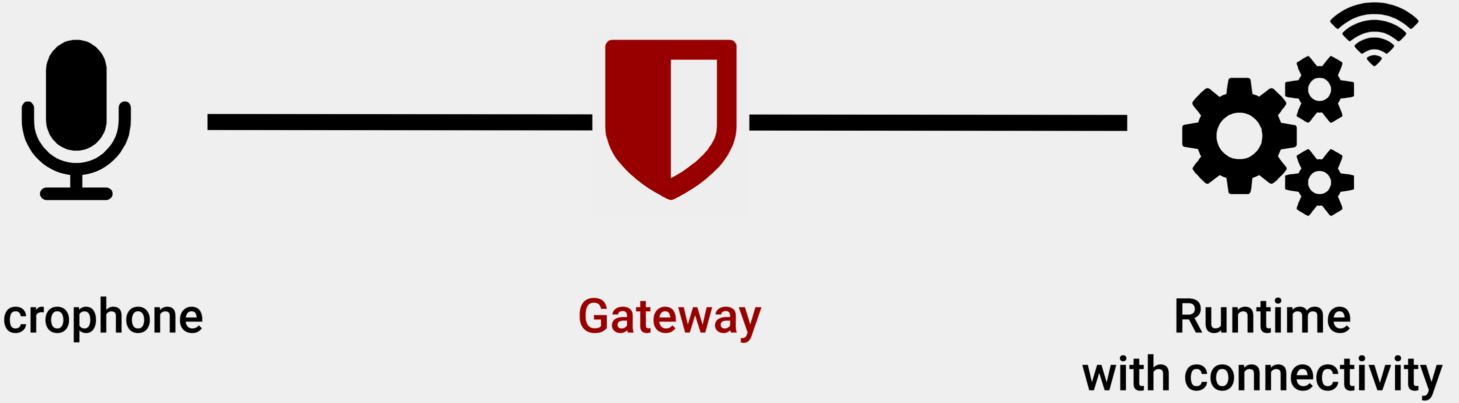

Problem: when to grant access?

Second attempt: event-detection container.

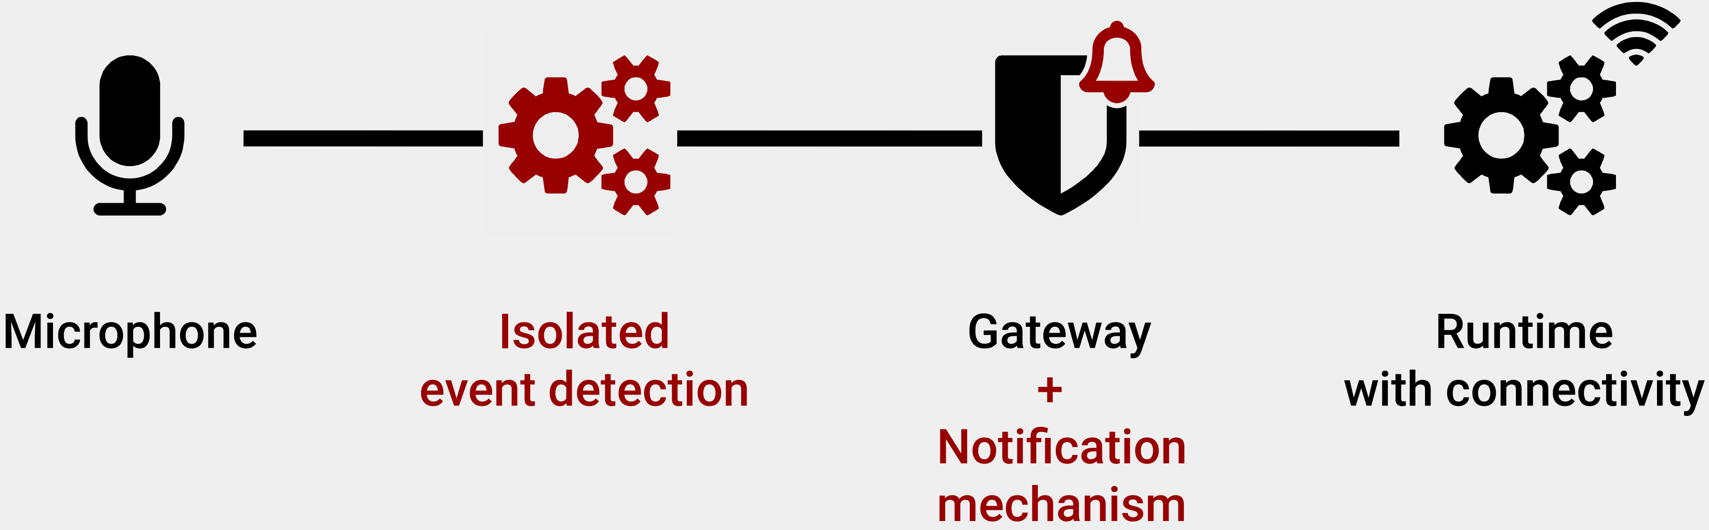

Microphone

Isolated
event detection

Gateway
+
Notification
mechanism

Runtime
with connectivity

Problem: no control over storage

Solution: amnesic event-detection container.

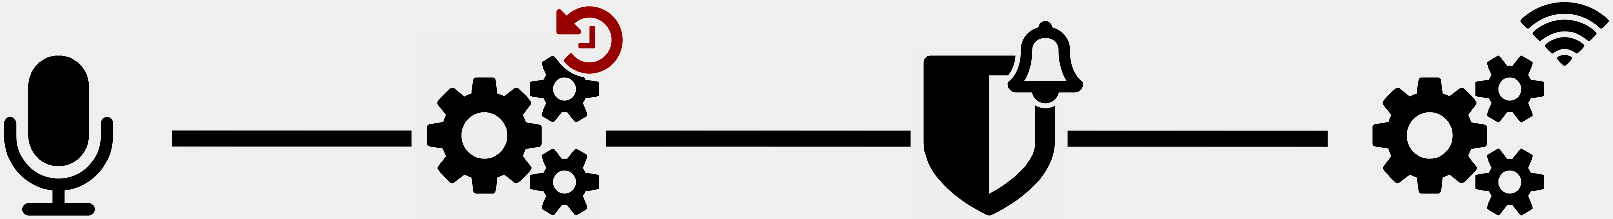

Microphone

Amnesic
& isolated
event detection

Gateway
+
Notification
mechanism

Runtime
with connectivity

KIMYA

Kimya segments the Microcontroller (MCU) into five memory regions.

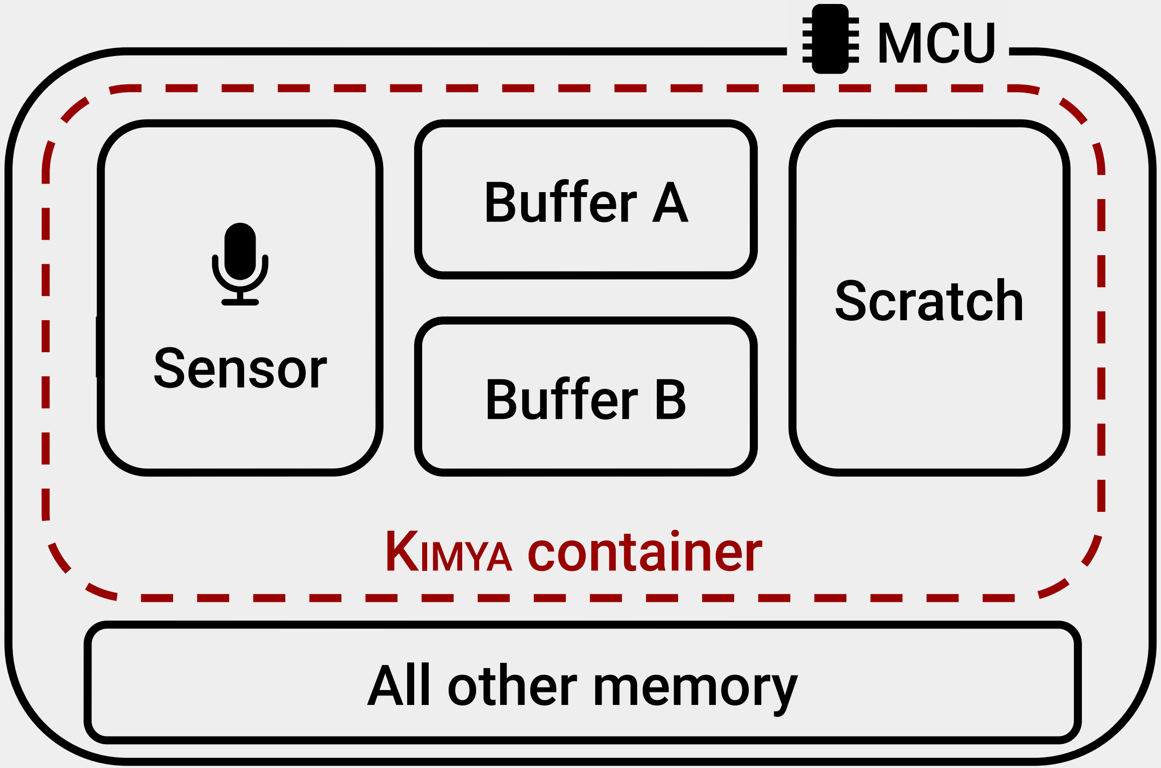

KIMYA introduces 4 different MCU phases.

IDLE

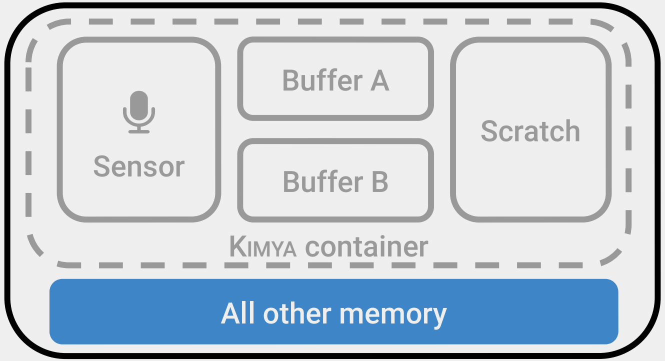

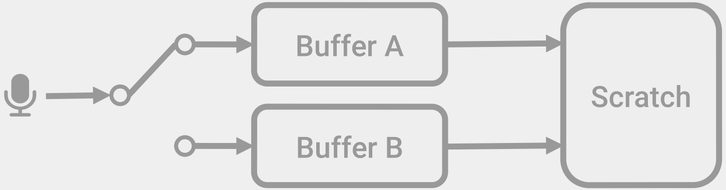

Read + Write

KIMYA introduces 4 different MCU phases.

ACQUIRE

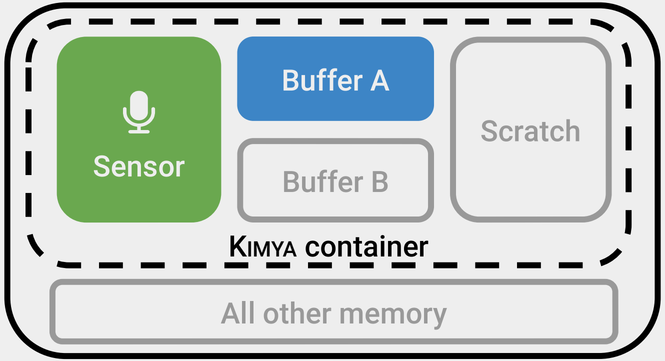

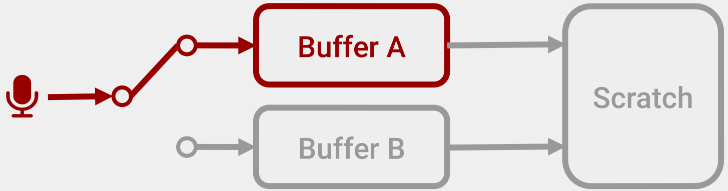

🟦 Read + Write

🟩 Read only

KIMYA introduces 4 different MCU phases.

PROCESS

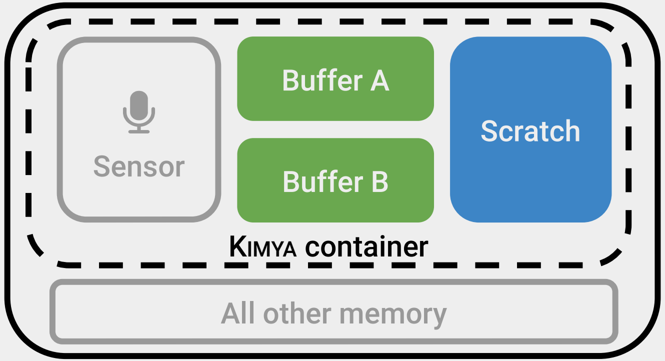

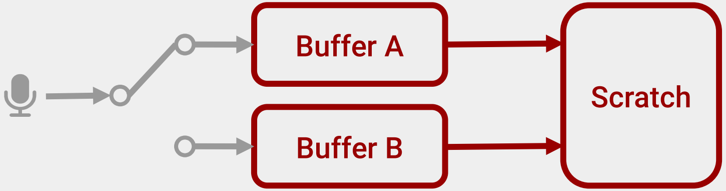

Read + Write

Read only

If an event has been detected, the MCU is TRIGGERED

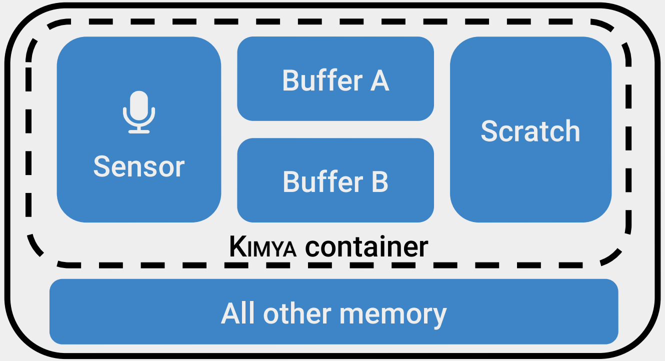

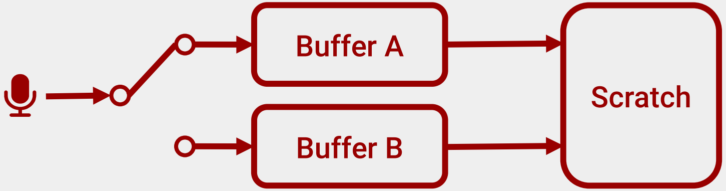

🟦 Read + Write

🟩 Read only

🔔 Notification generated

Each 0.5 T_{lifetime}, buffers are alternated and wiped.

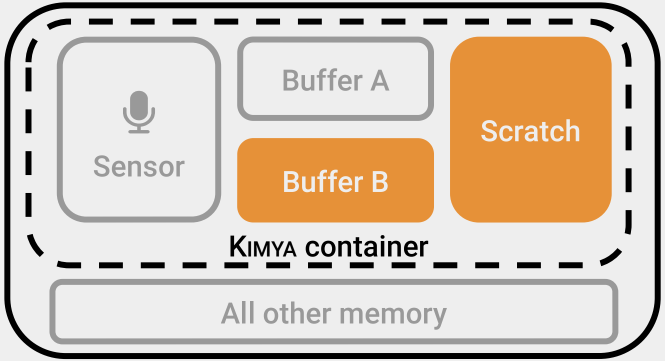

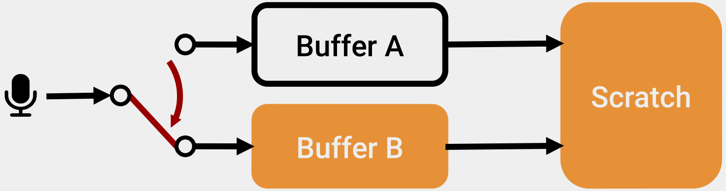

⇒ Maximum data age: T_{lifetime}

Read + Write

Read only

Zero'd out

In the paper:
Implementation with TrustZone on Cortex-M.

What is an “interaction”?

When does it start, when does it end?

How to enforce isolation …

… in a way that’s compatible with existing OSes

… without {timing, cache, peripheral, …} covert channels

Kimya introduces 1 ms of latency

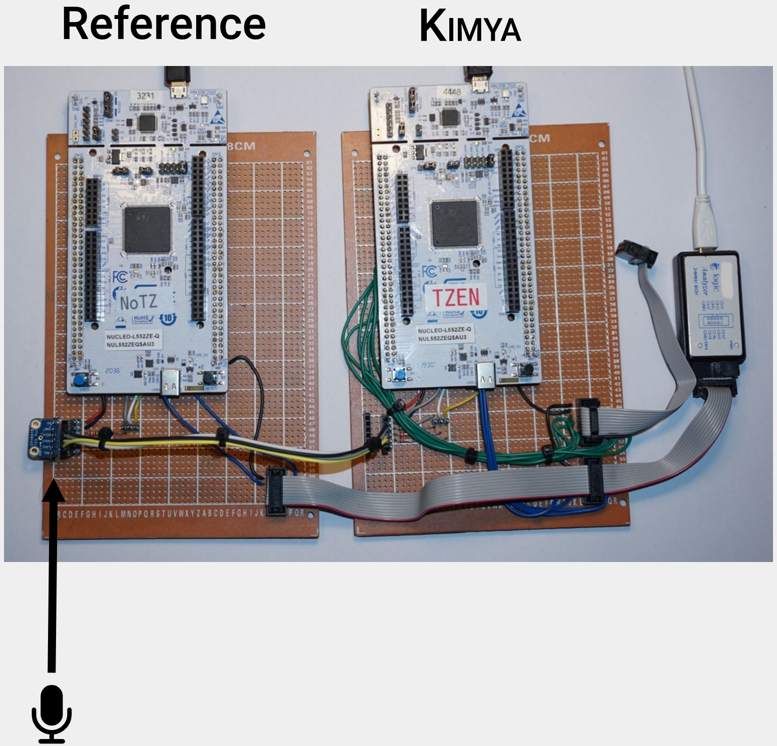

No additional HW required

Evaluated using on-chip keyword-spotting pipeline
(mel spectrum + CNN)

only 1.19 ms of latency
(spread = 0.03 ms)

Detailed benchmarks in paper

Conclusion

Smart speakers come with a paradox
& current protections are insufficient

KIMYA provides an isolated and amnesic event-detection container that

- Introduces low overhead,
- does not restrict which algorithms can be used,
- and is independent of crypto.

Questions?
Job offers?
➔ piet@devae.re