The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence

Jascha Grübel[1], Tyler Thrash[2], Didier Hélal[3], Robert W. Sumner[1], Christoph Hölscher[1], Victor R. Schinazi[4]

[1]ETH Zürich, Zürich, Switzerland
[2]Miami University, Oxford, OH, USA
[3]Orbiwise, Geneva, Switzerland
[4]Bond University, Robina, Australia

Presentation available here: 10.5281/zenodo.4605350

Good morning, my name is Jascha Grübel and I am a PhD student at ETH Zürich and today I will present our work “the feasibility of dense indoor LoRaWAN towards passively sensing human presence”. [NS]

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

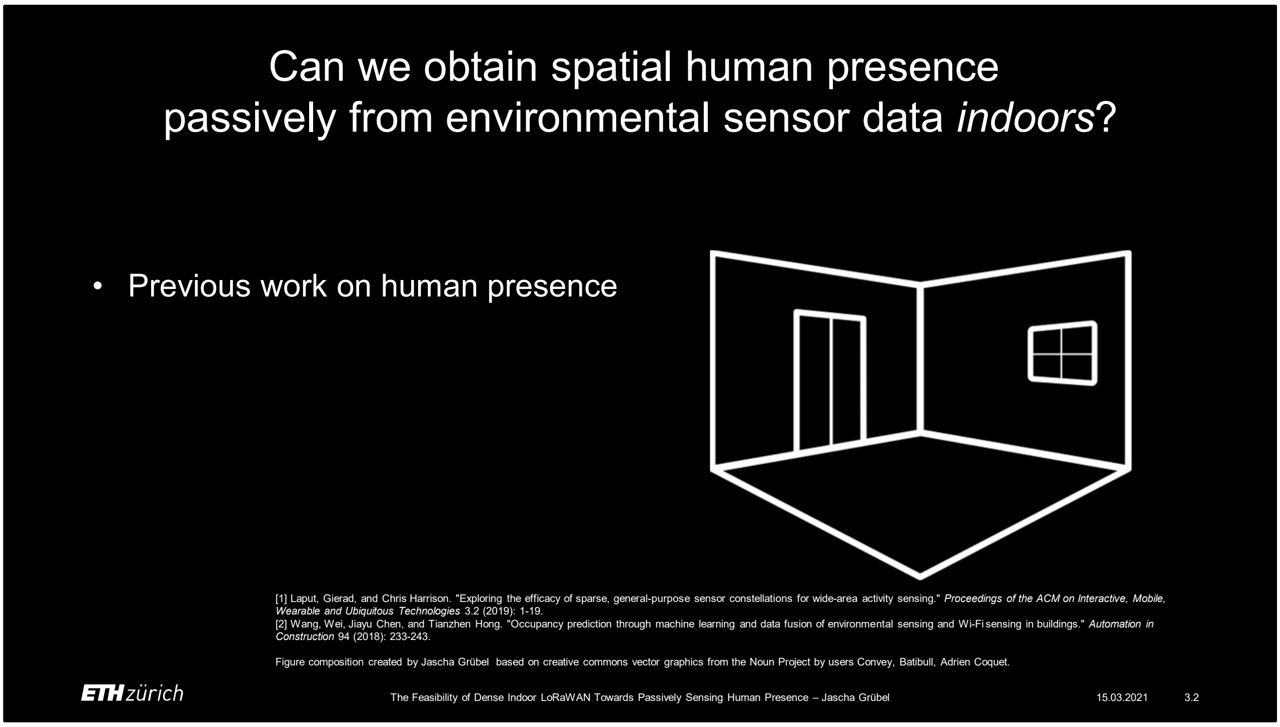

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

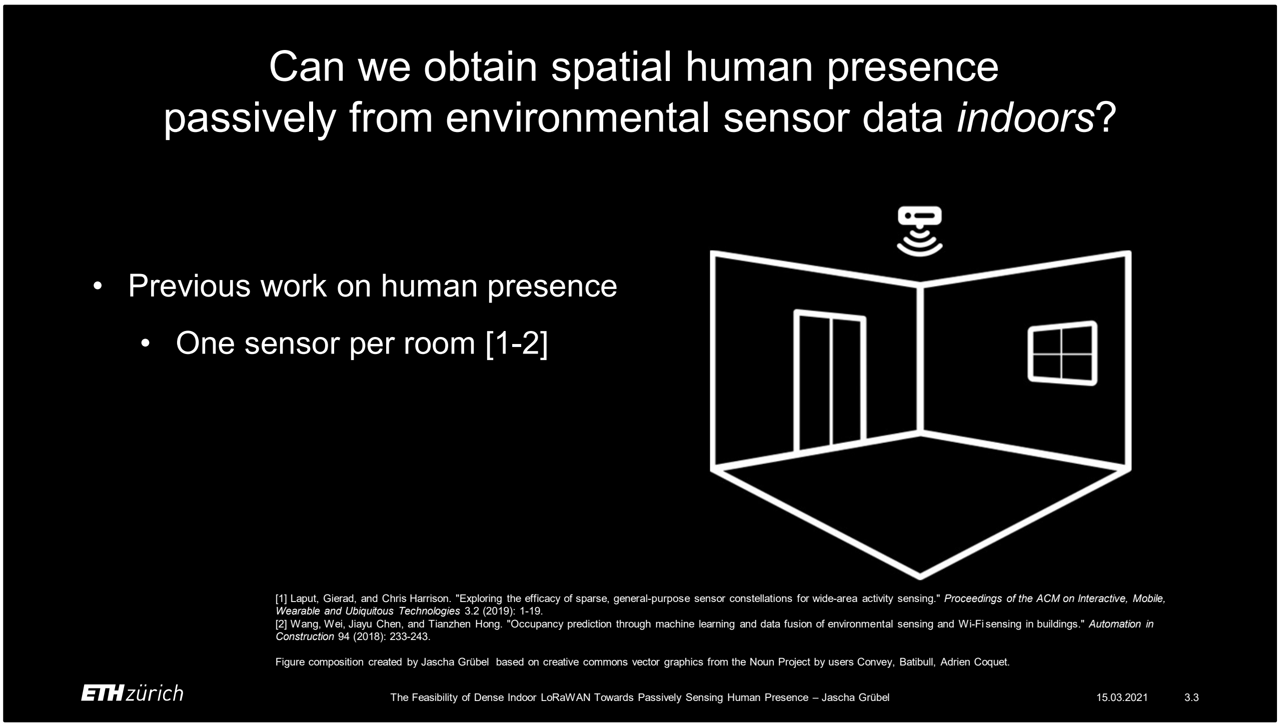

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

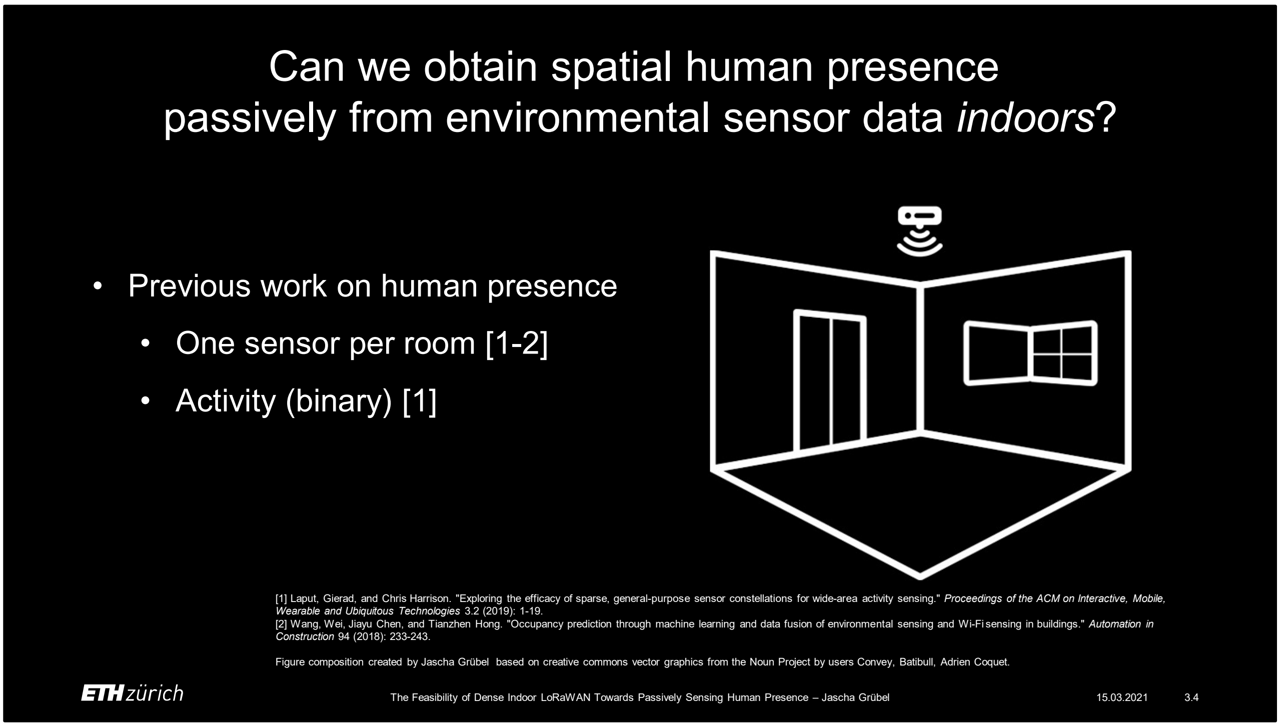

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

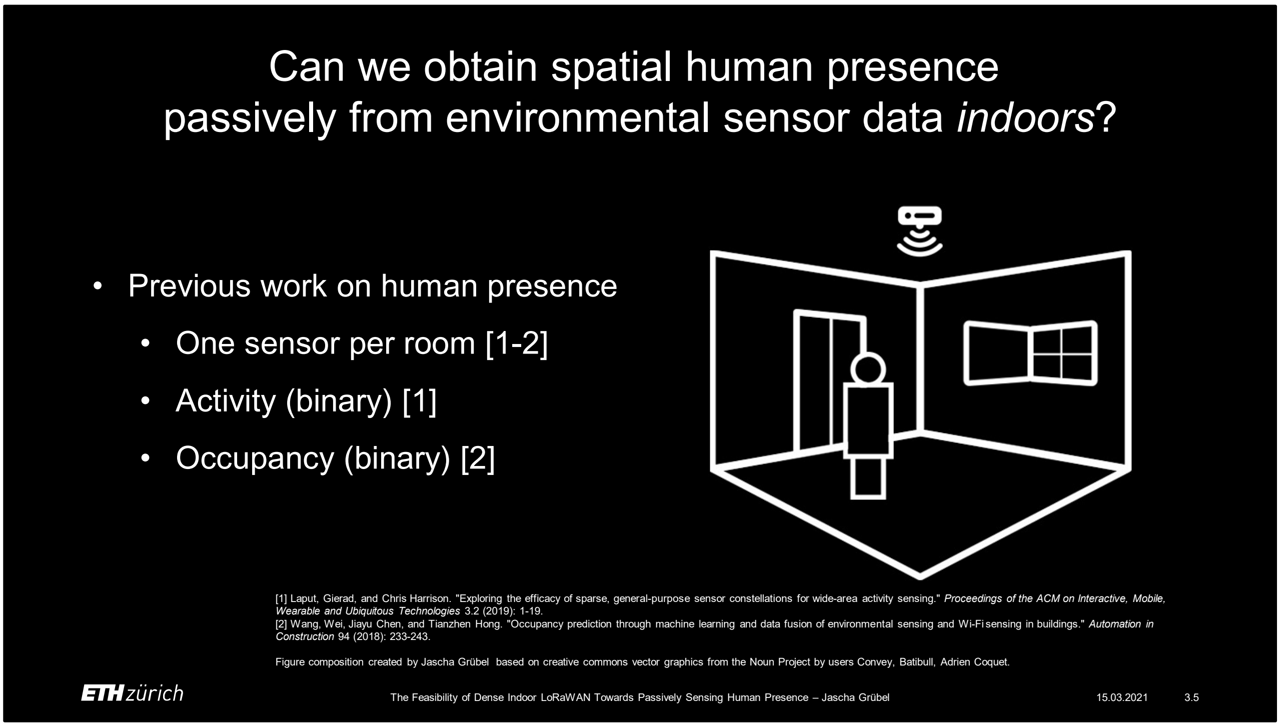

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

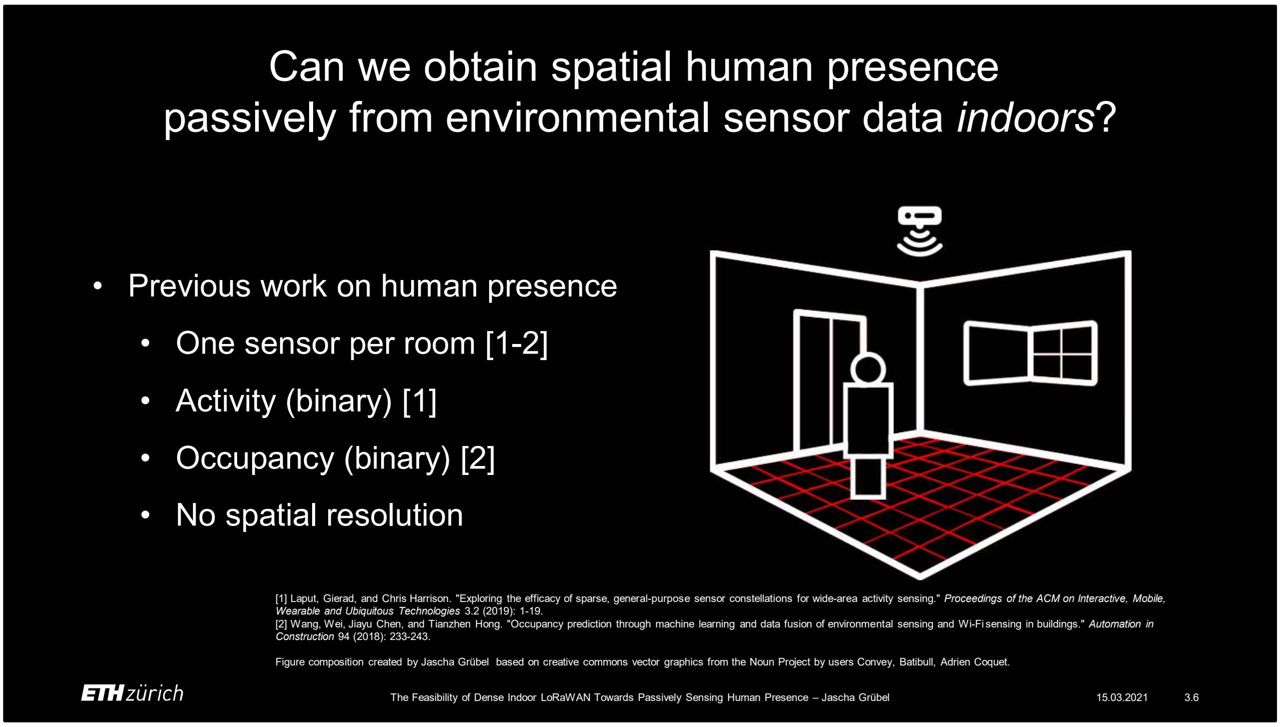

We started with the question whether we can obtain human presence passively from environmental sensor data indoors.
[C] Previous work on human presence has mostly focused on room-level setups.
[C] A single sensor device is installed to measure different forms of presence.
[C] Activity measures binary indicators like an opened window.
[C] Occupancy measures whether anybody at all is in the room.
[C] Neither approaches produce spatial resolution within rooms.

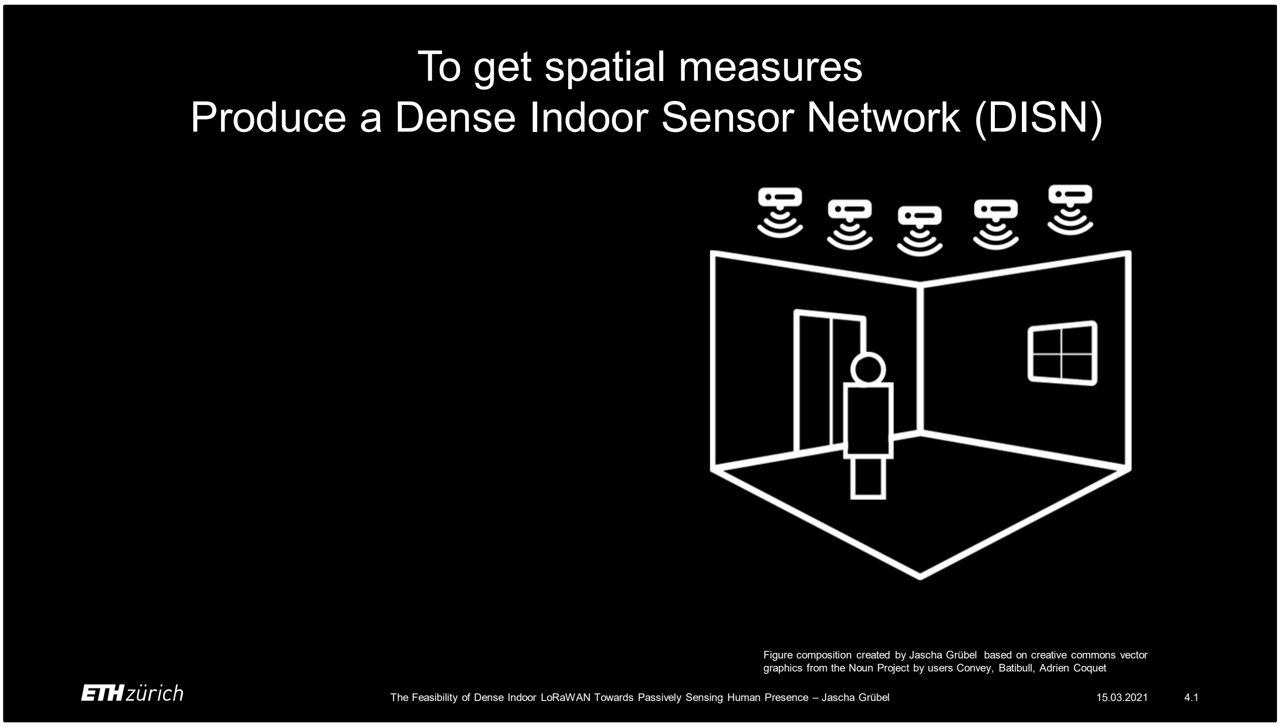

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

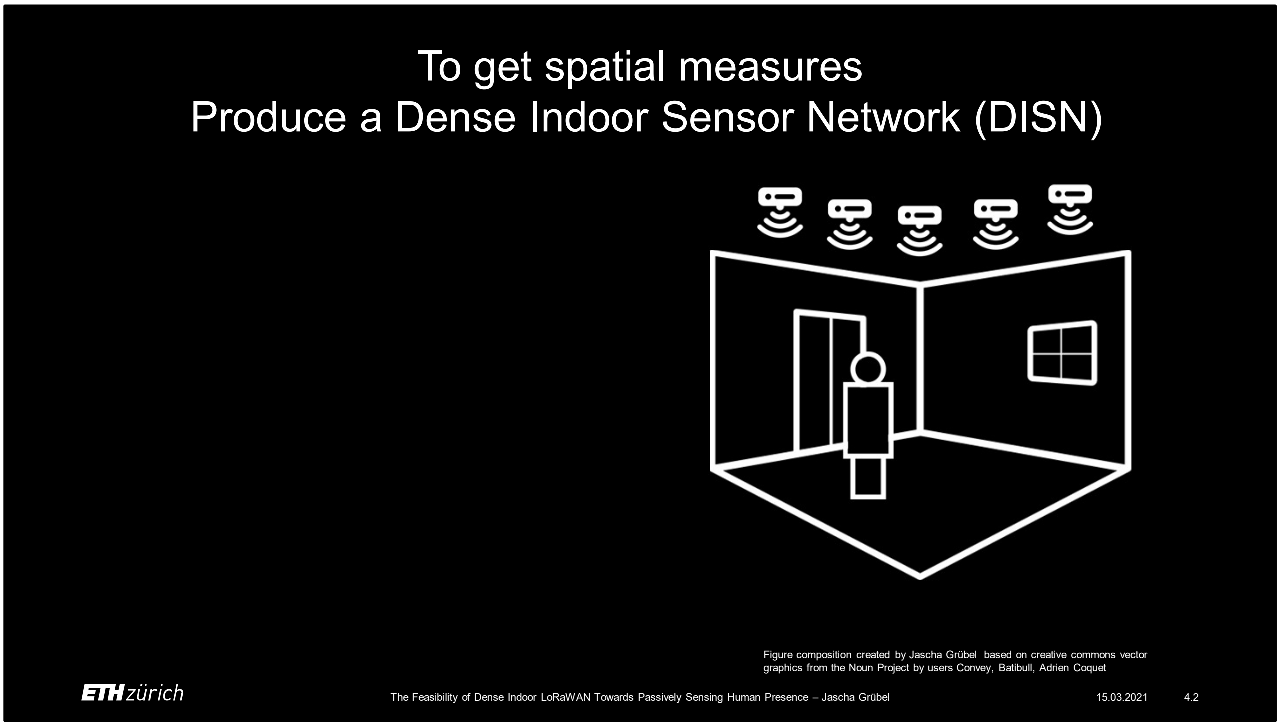

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

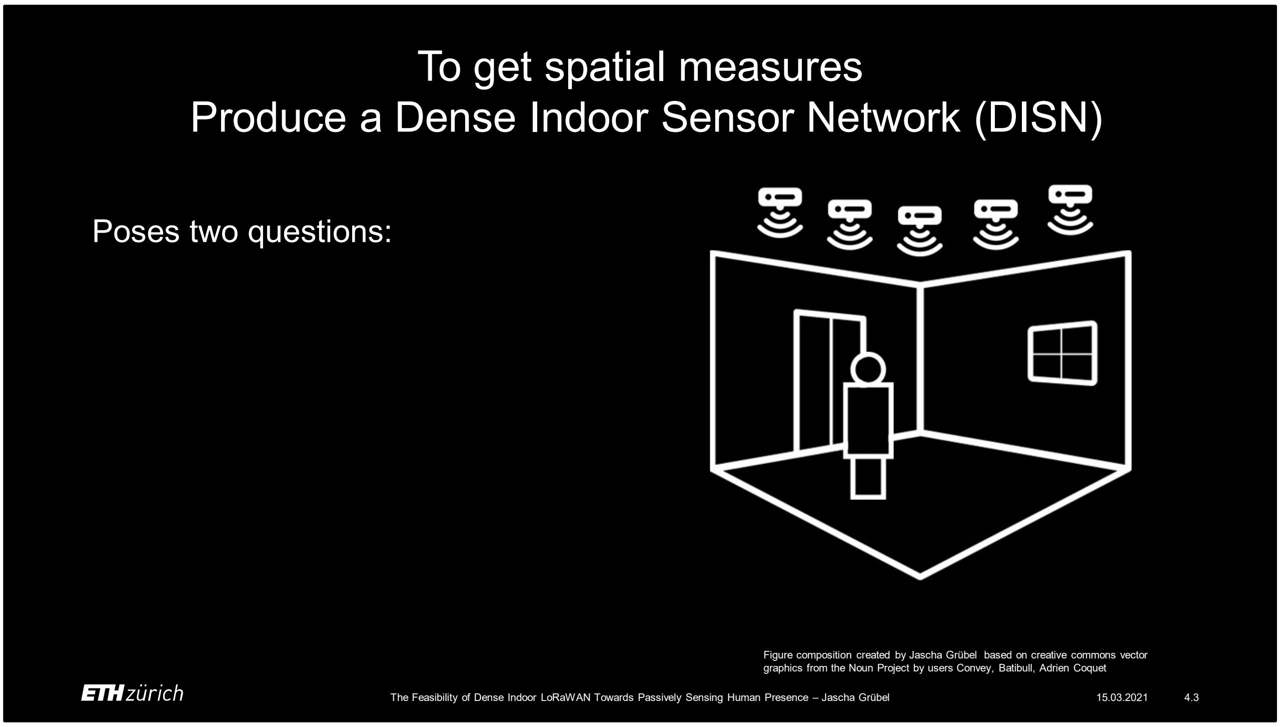

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

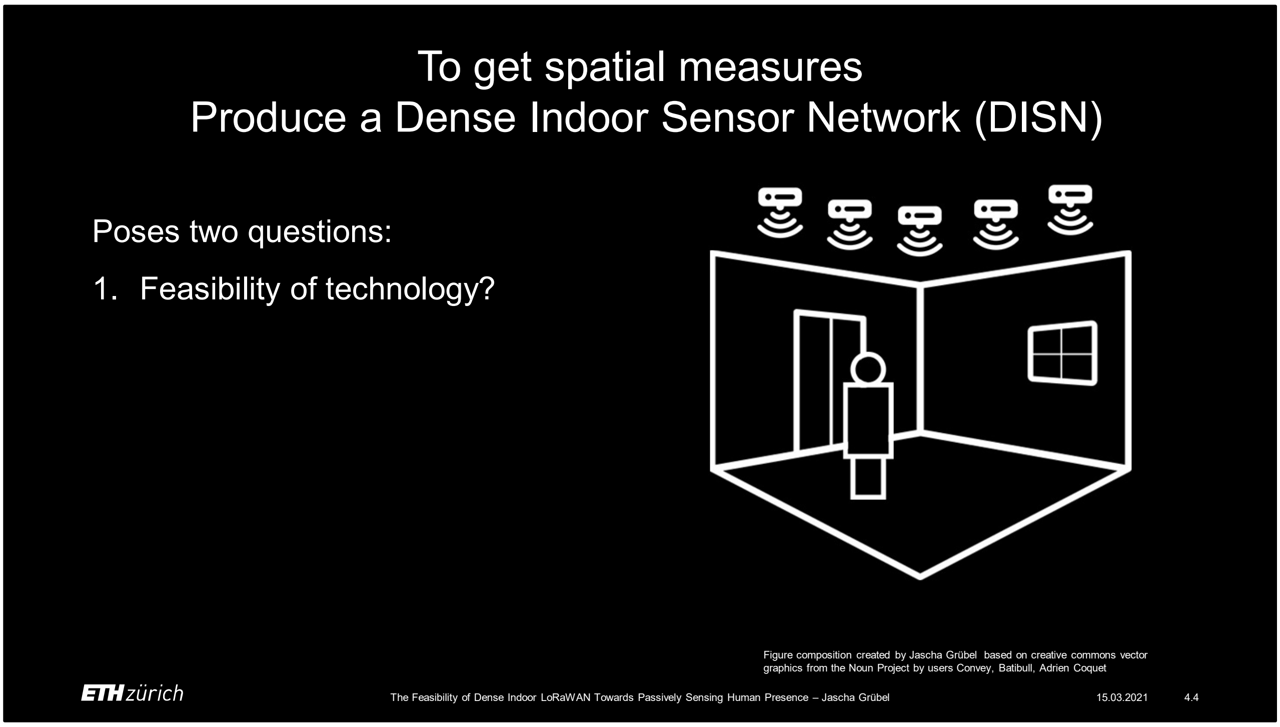

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

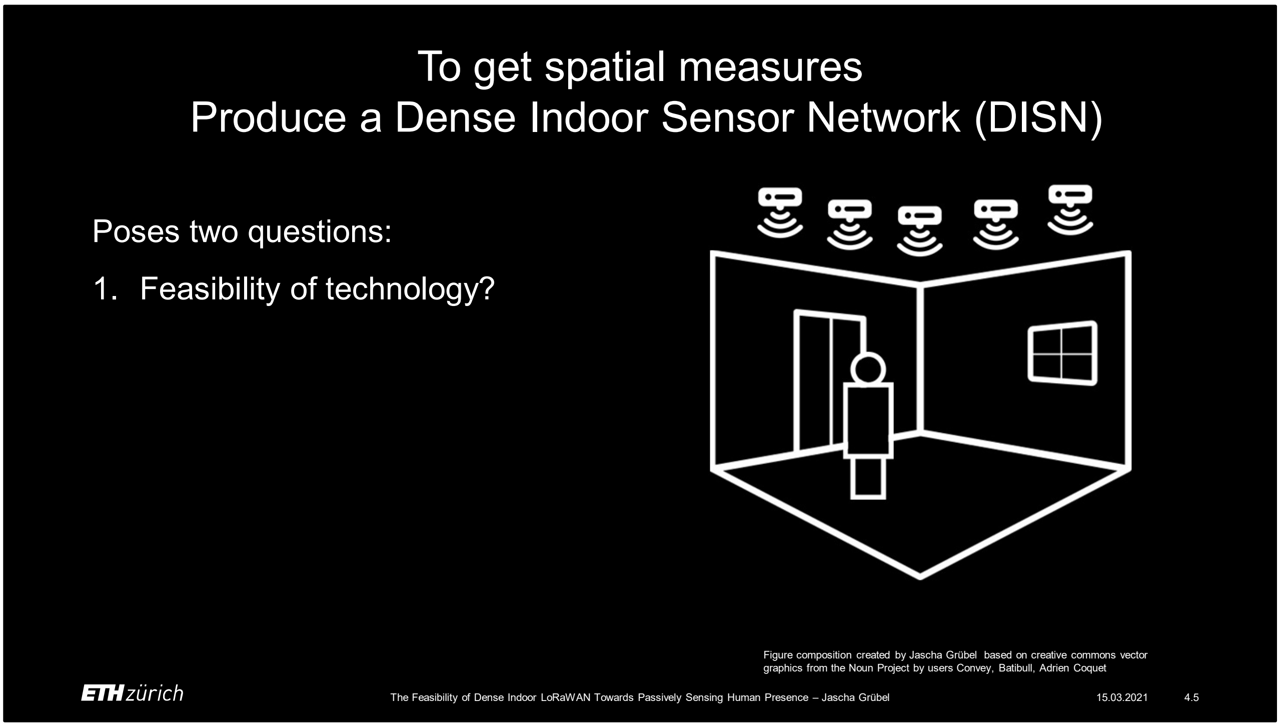

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

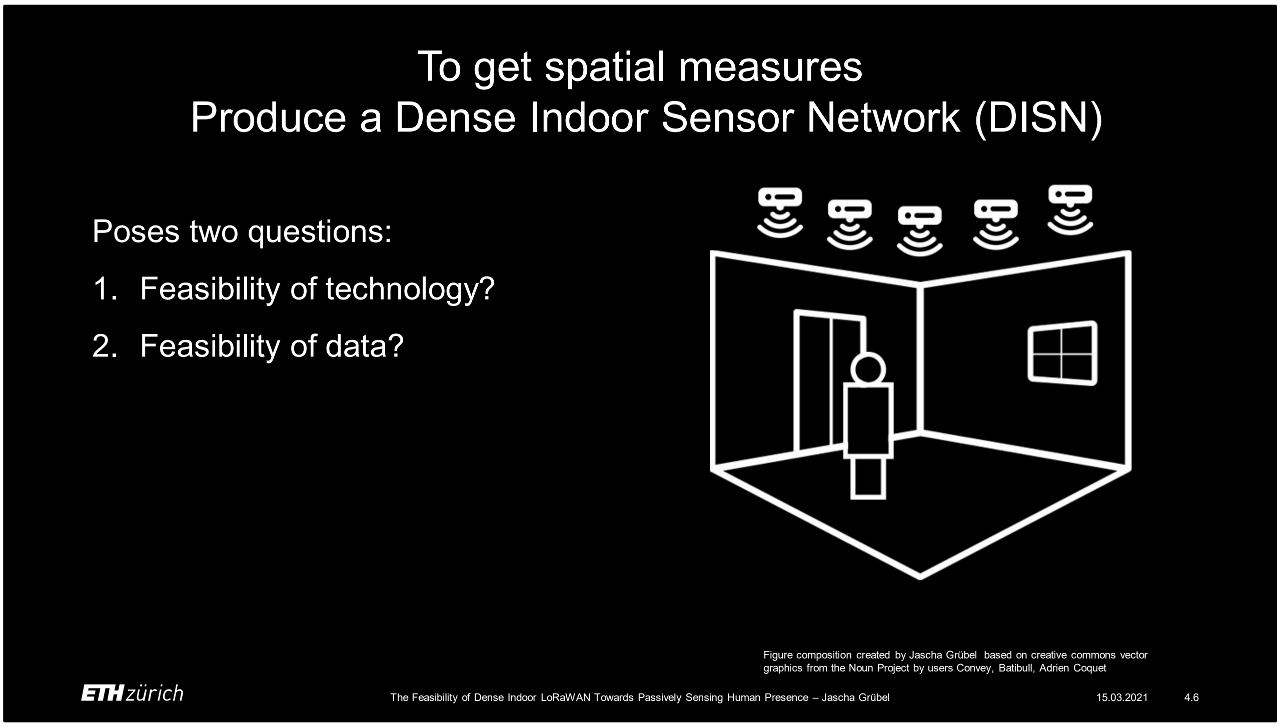

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

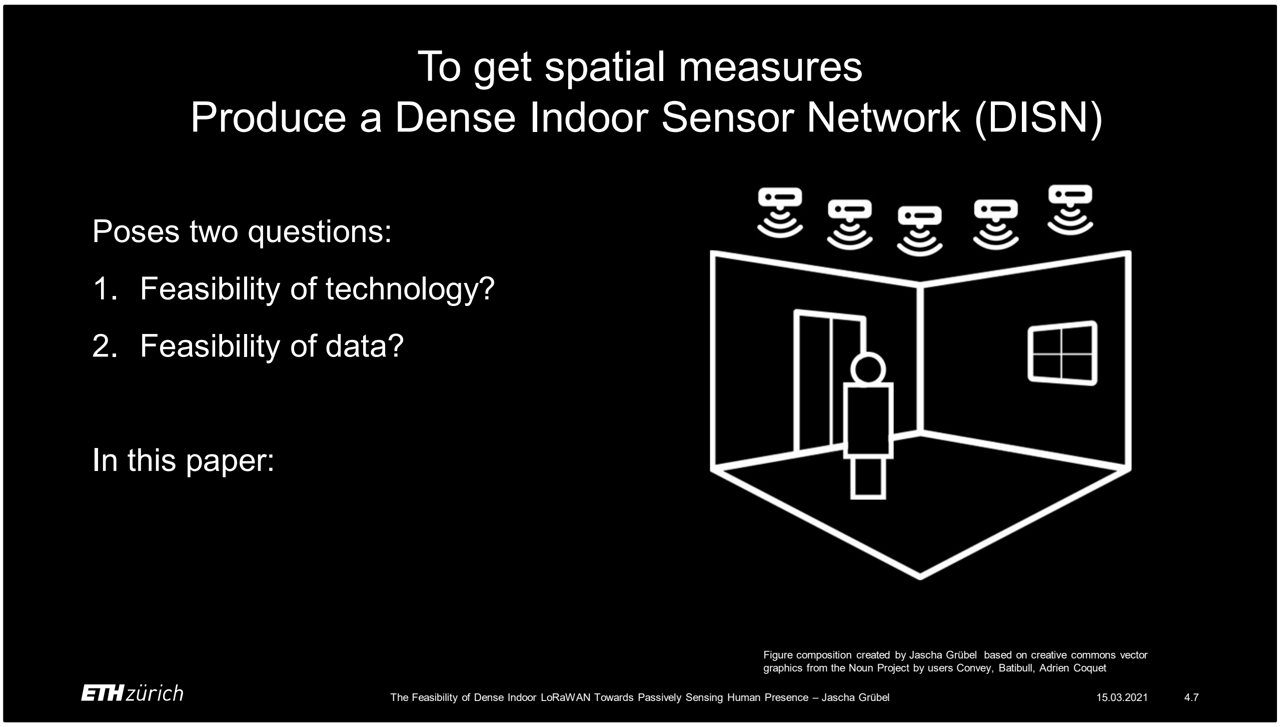

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

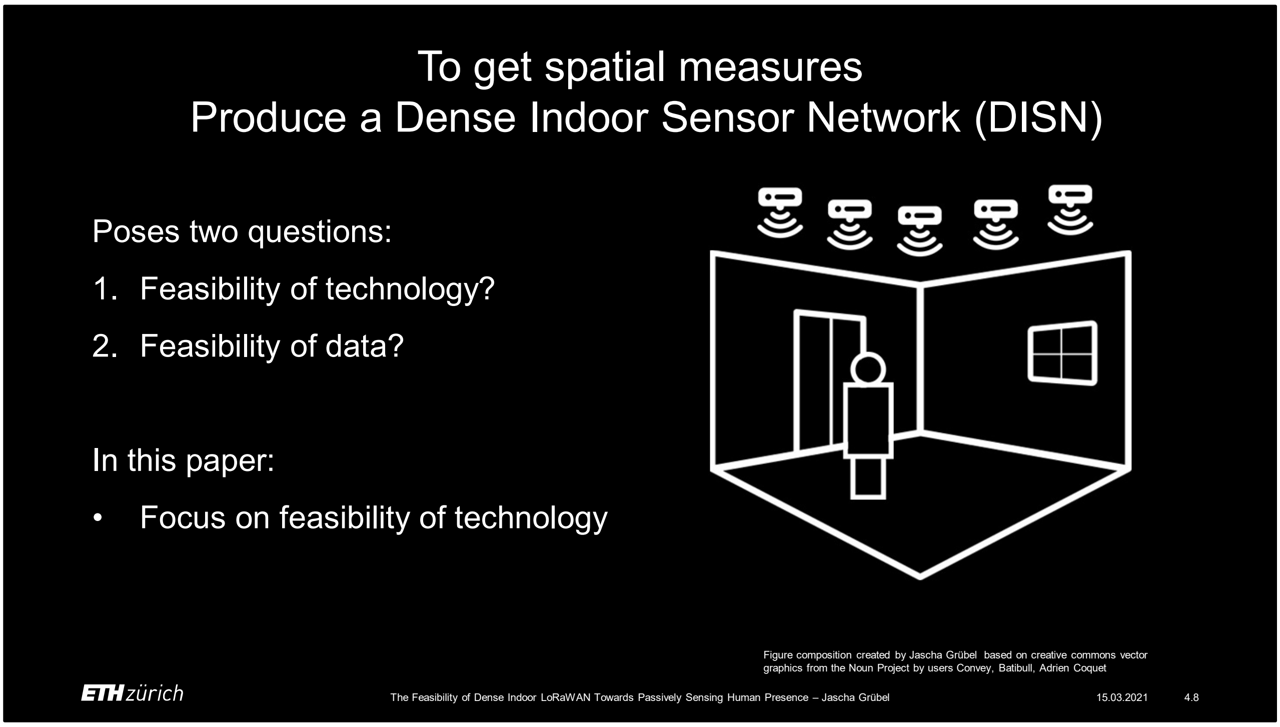

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

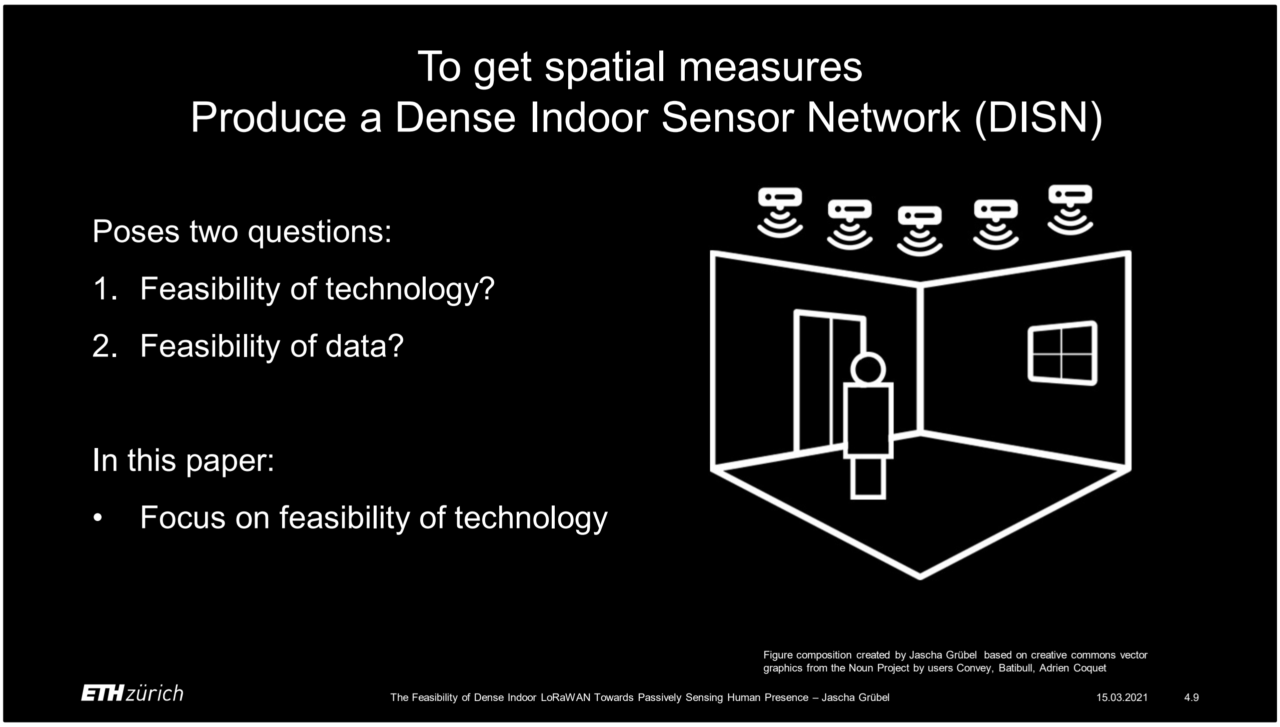

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

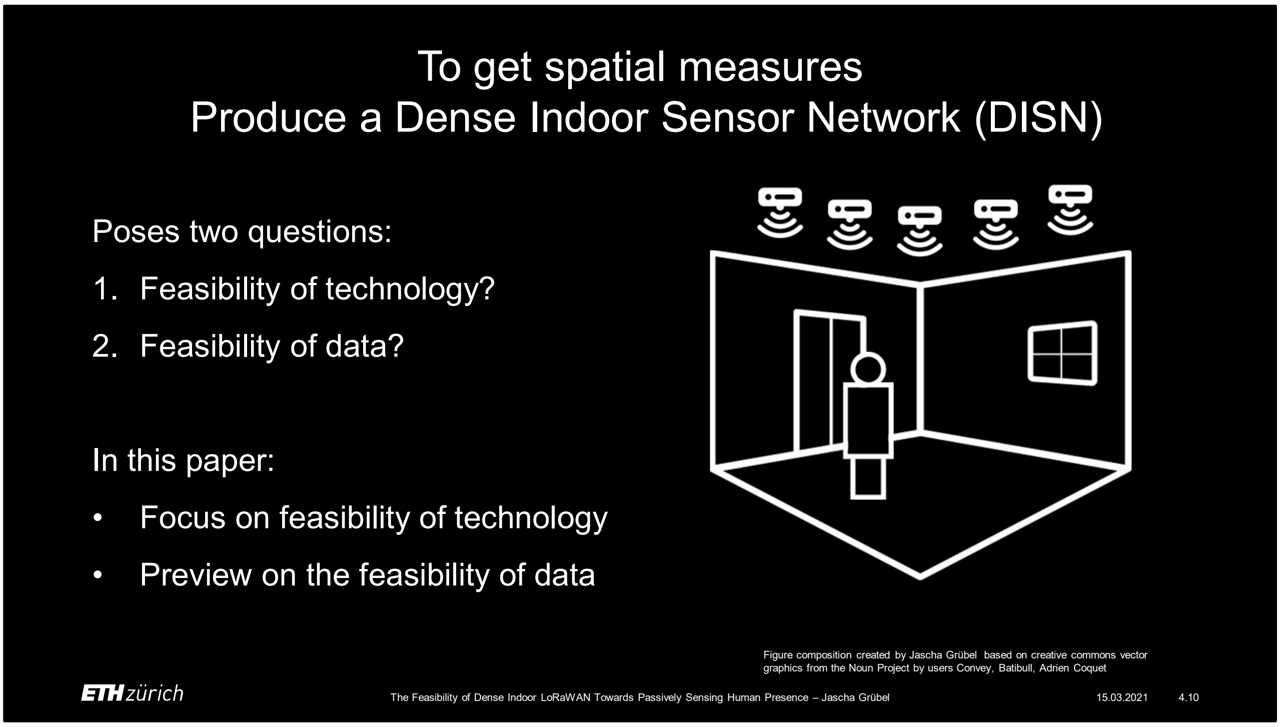

To get spatial measures we need to produce a dense indoor sensor network or D.I.S.N.
[C] This poses two major questions.
[C] First, how feasible is this technologically.
[C] Second, how feasible is the data to derive spatial measures.
[C] In this paper we focus on the feasibility of the technology.
[C] And only provide a preview of the feasibility of the data

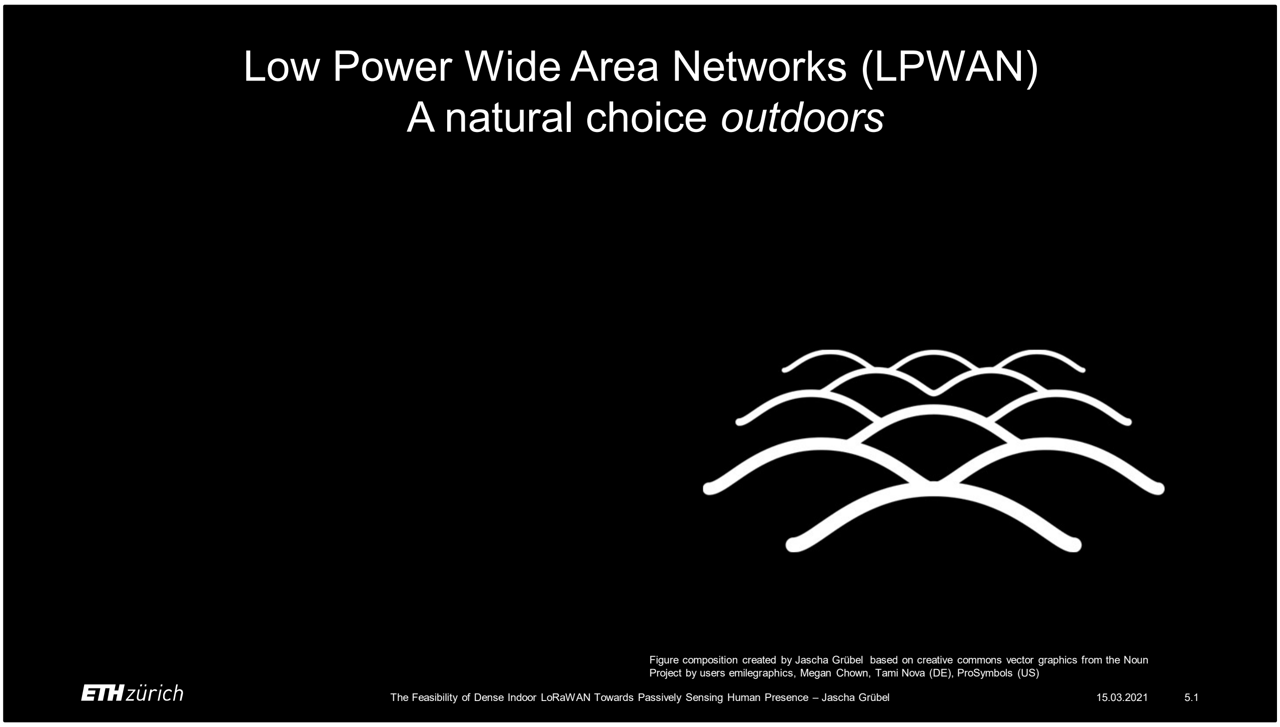

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

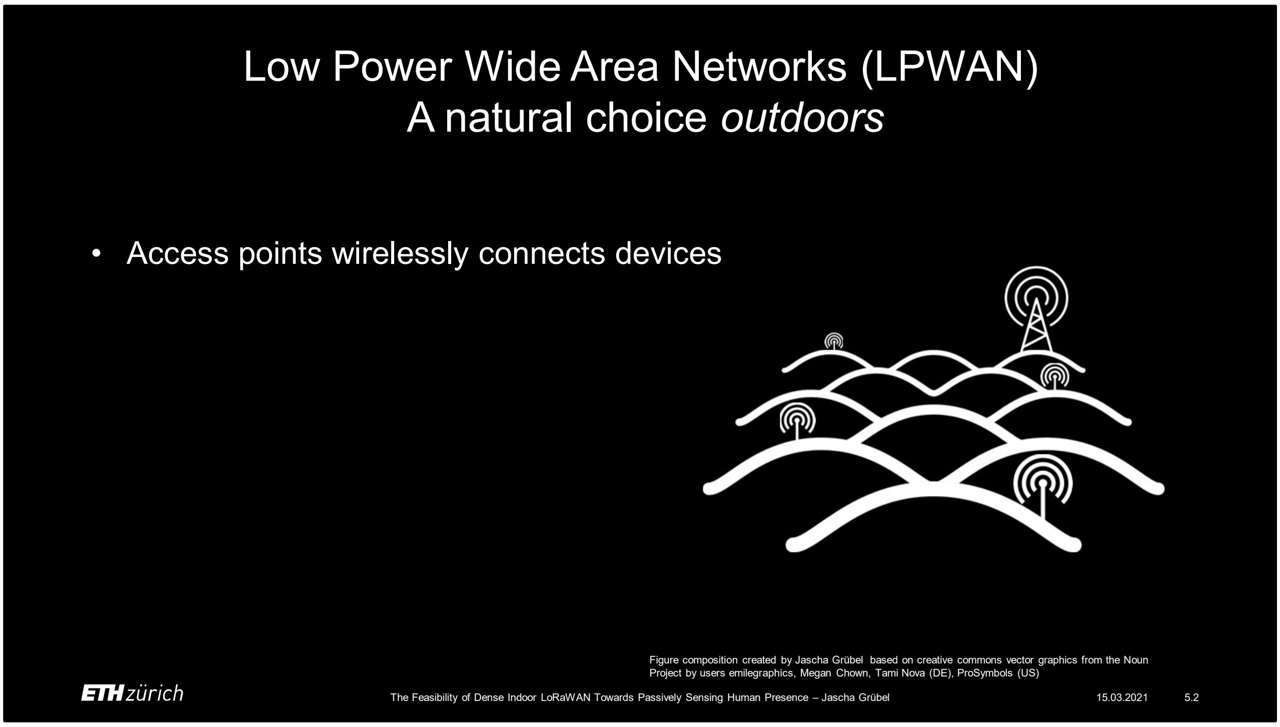

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

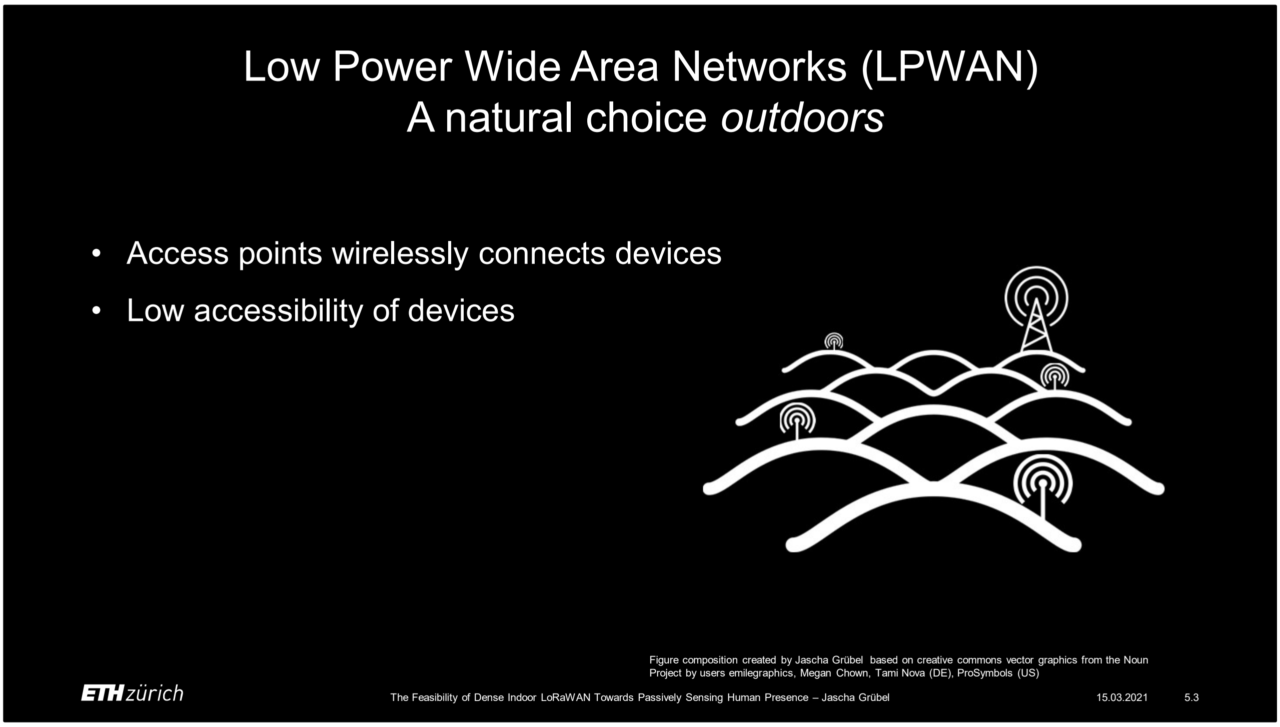

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

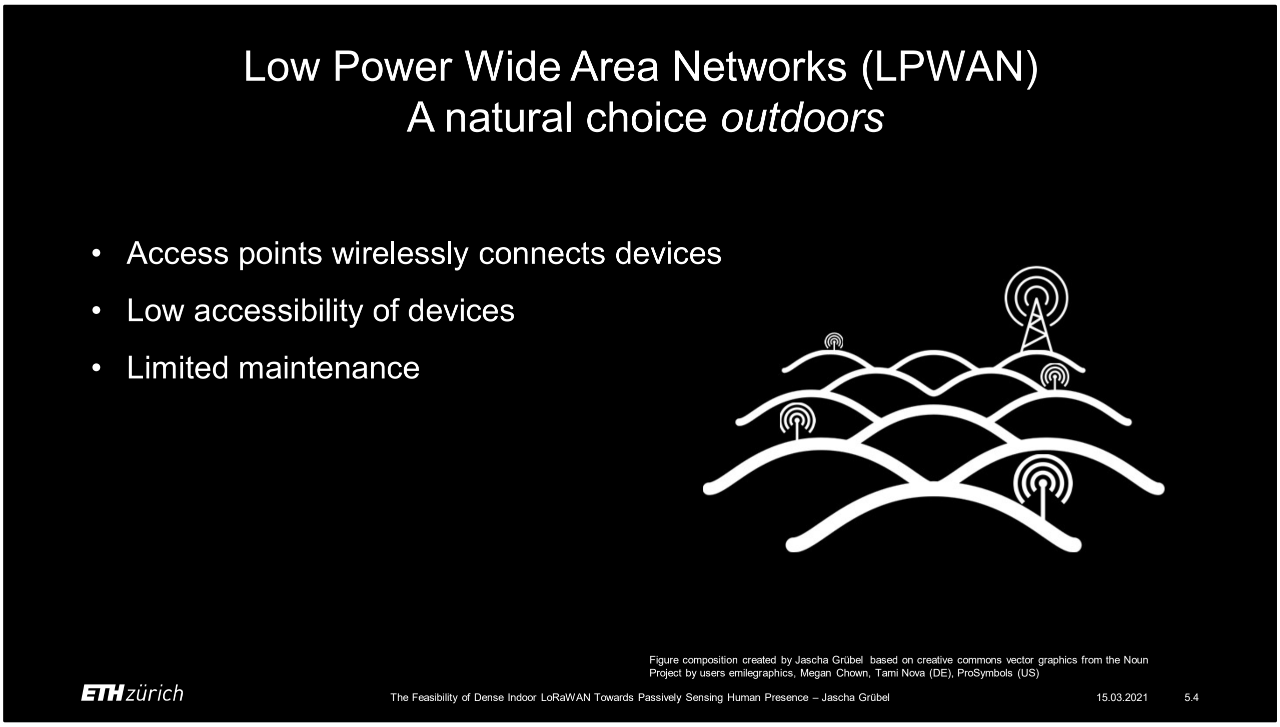

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

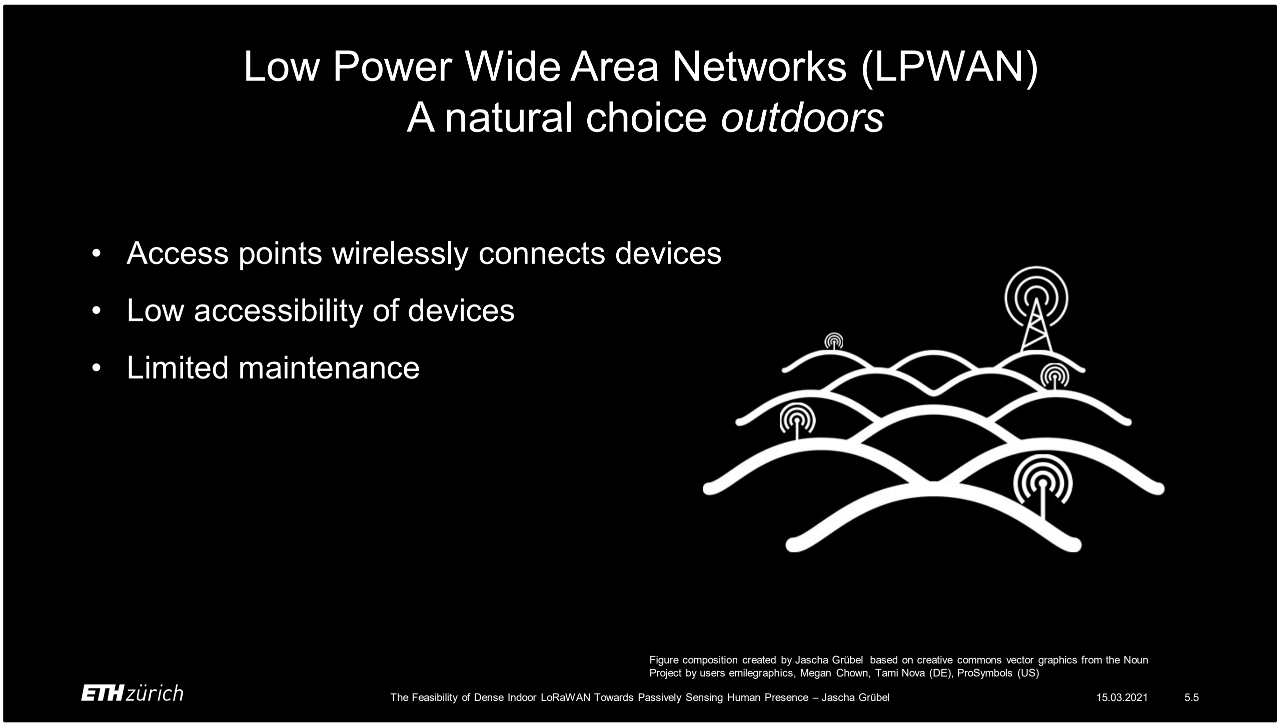

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

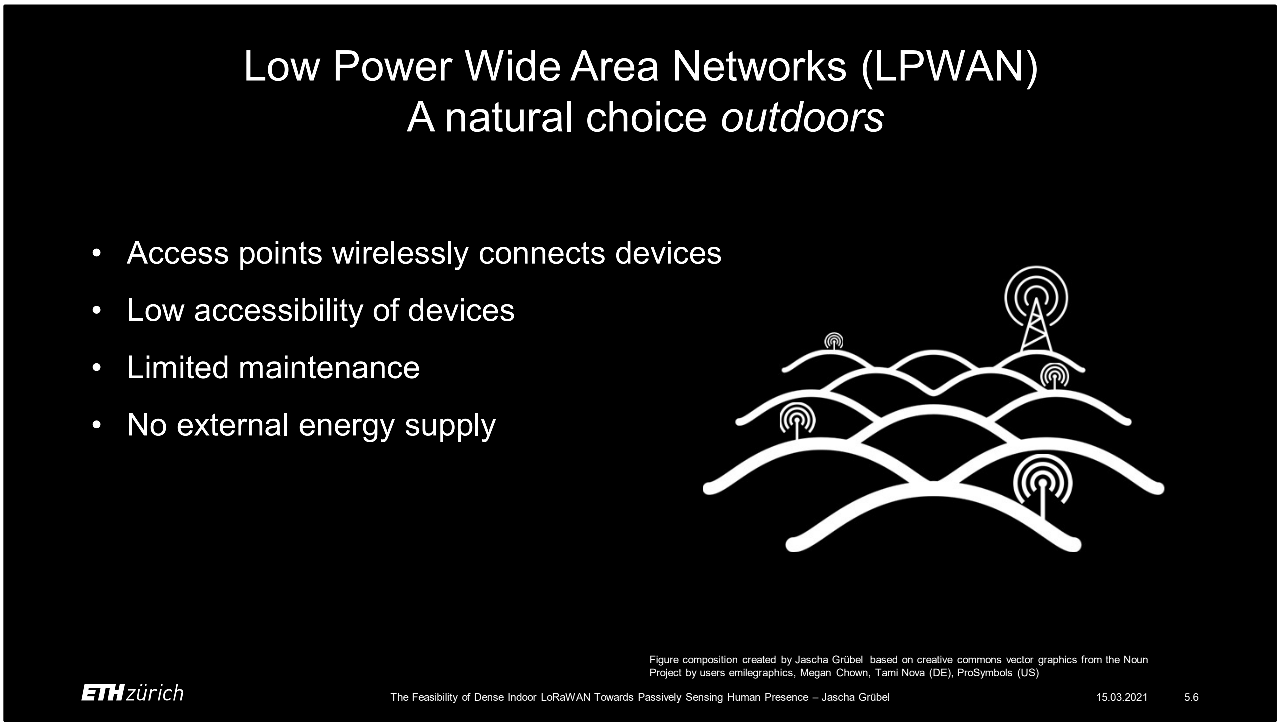

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

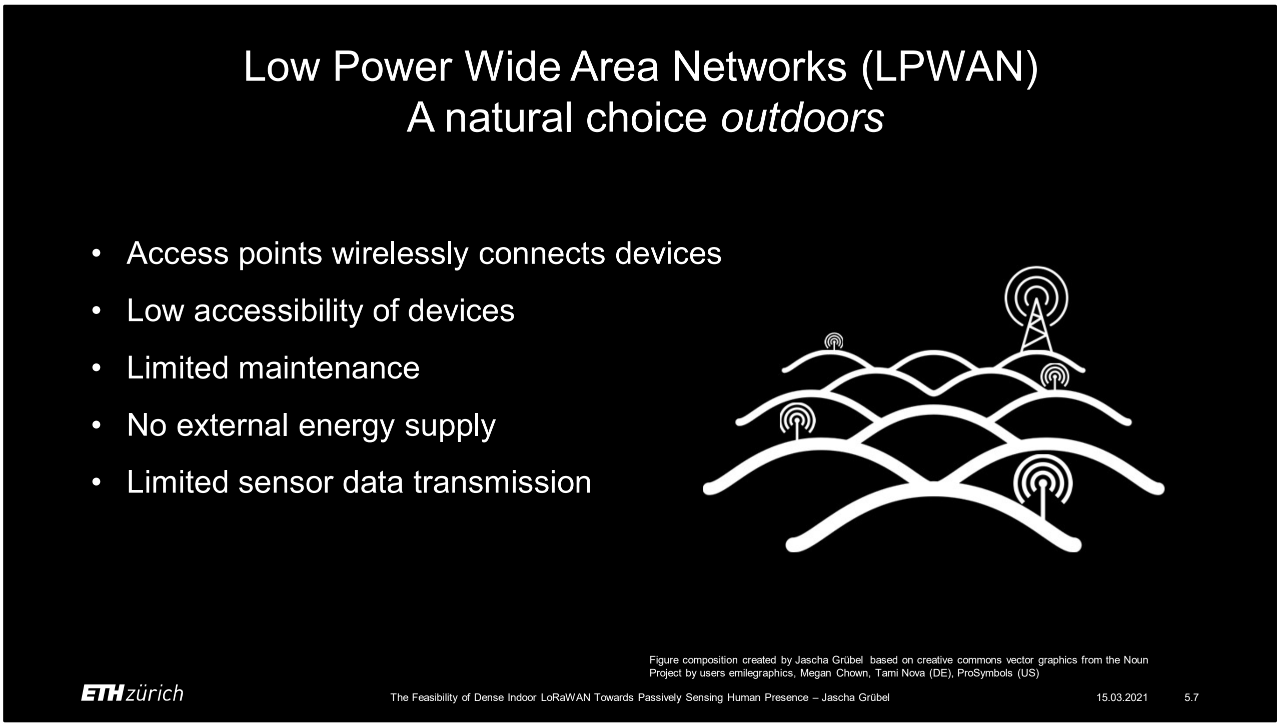

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

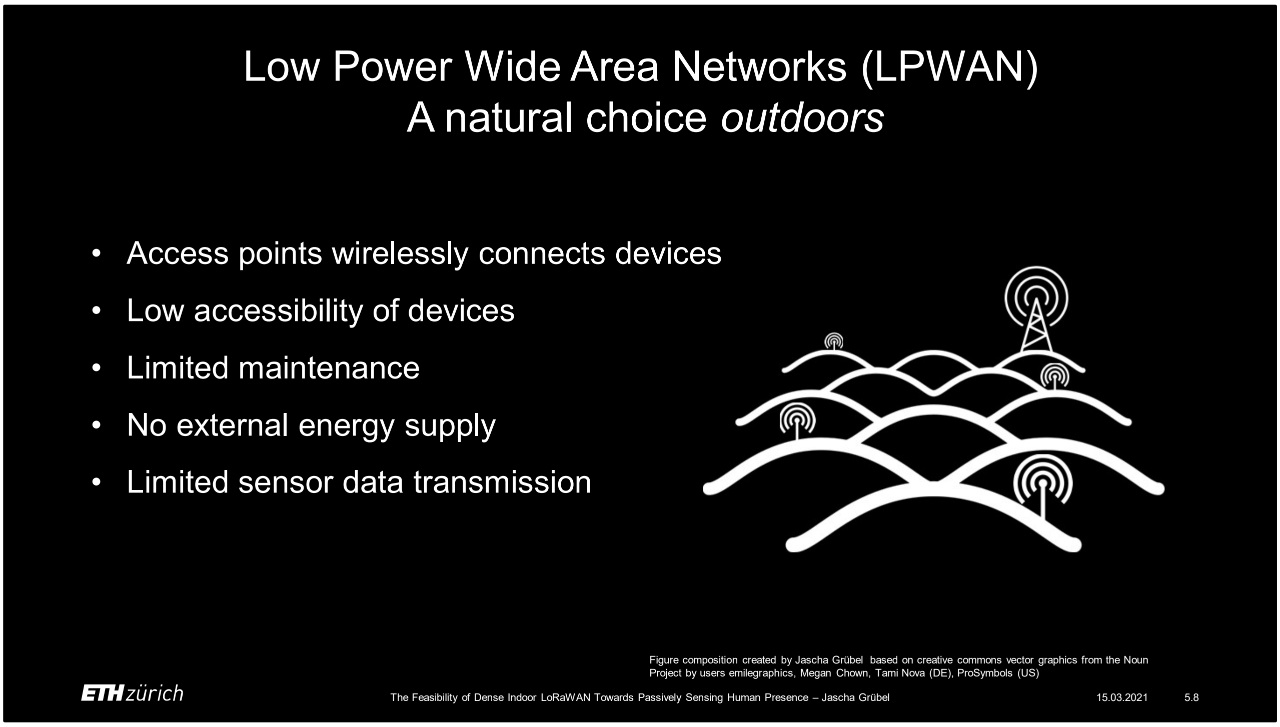

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

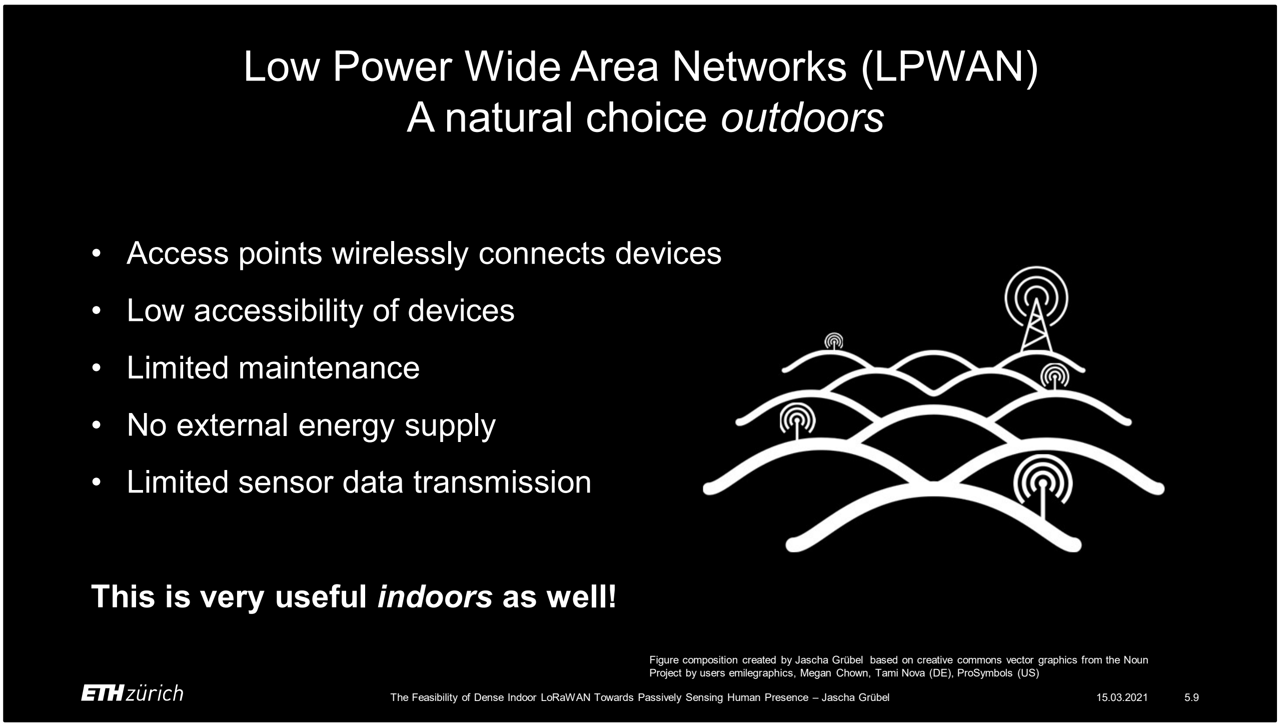

Low Power Wide Area Networks (or LPWAN) are a natural choice in outdoor scenarios.
[C] An access point is used to wirelessly connect devices across large distances.
[C] They can handle devices with low accessibility such as in forests, underground or distributed in cities.
[C] Furthermore, they can work even with limited maintenance because of the robust design.
[C] Usually they require very little power for transmission and therefore can run without external energy supply for years on batteries.
[C] Environmental sensor data is usually tiny and does not require large bandwidth, so limited bandwidth poses no issue.
[C] Interestingly, these characteristics are also very useful for indoor applications.

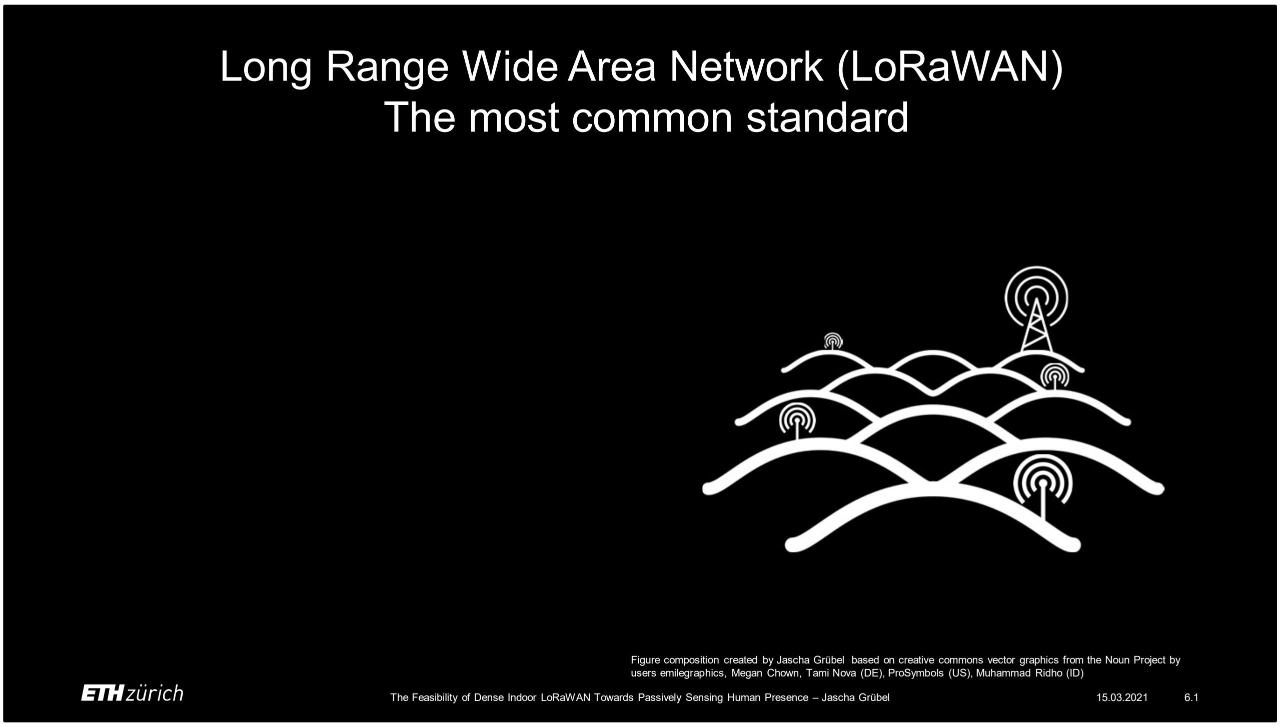

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

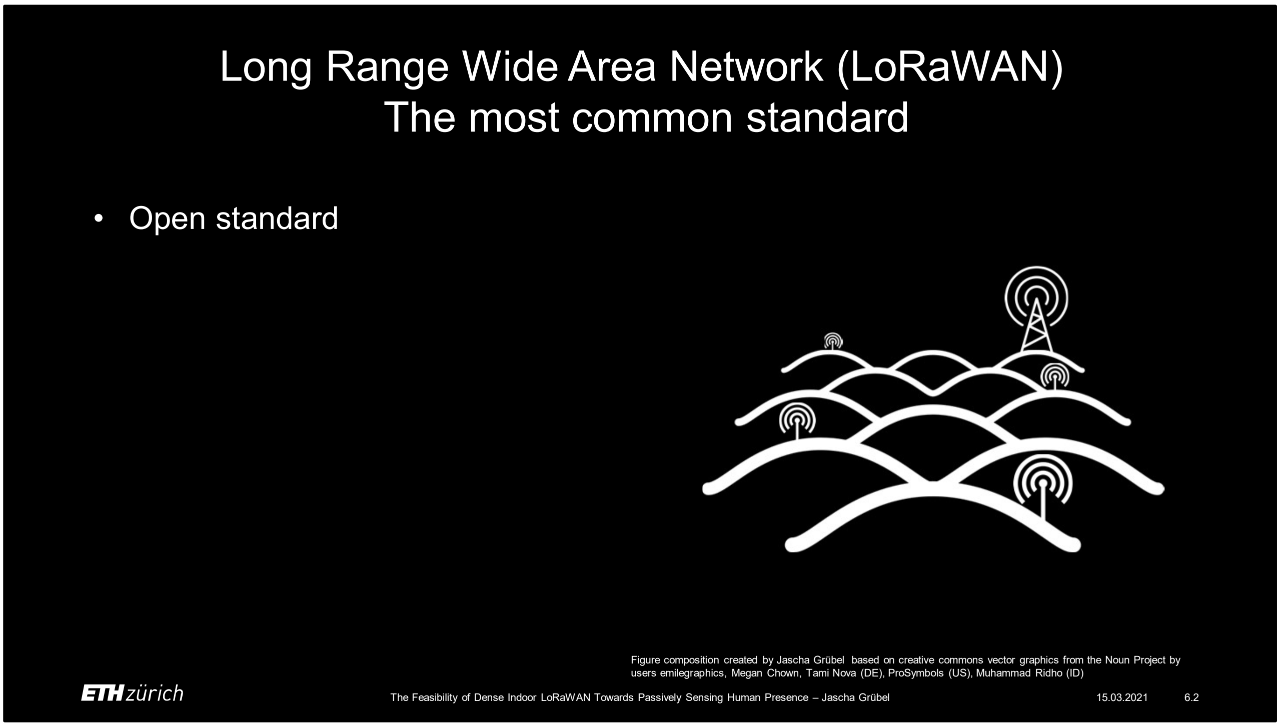

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

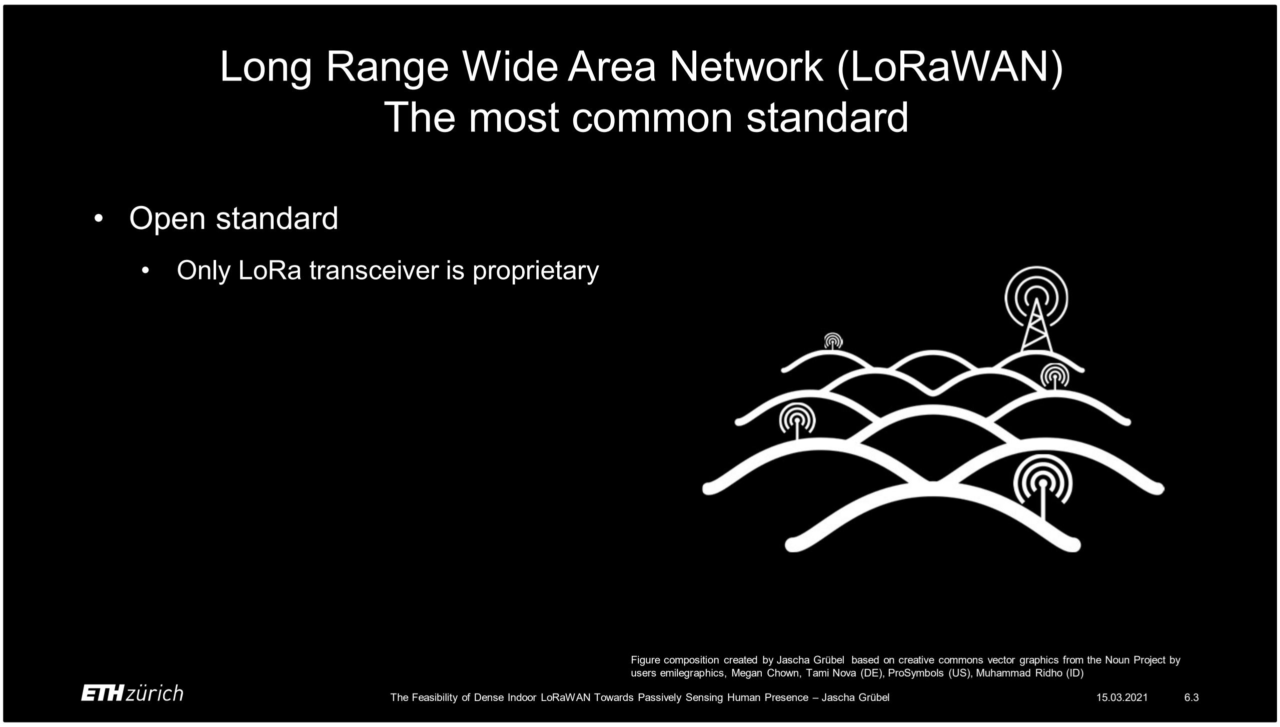

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

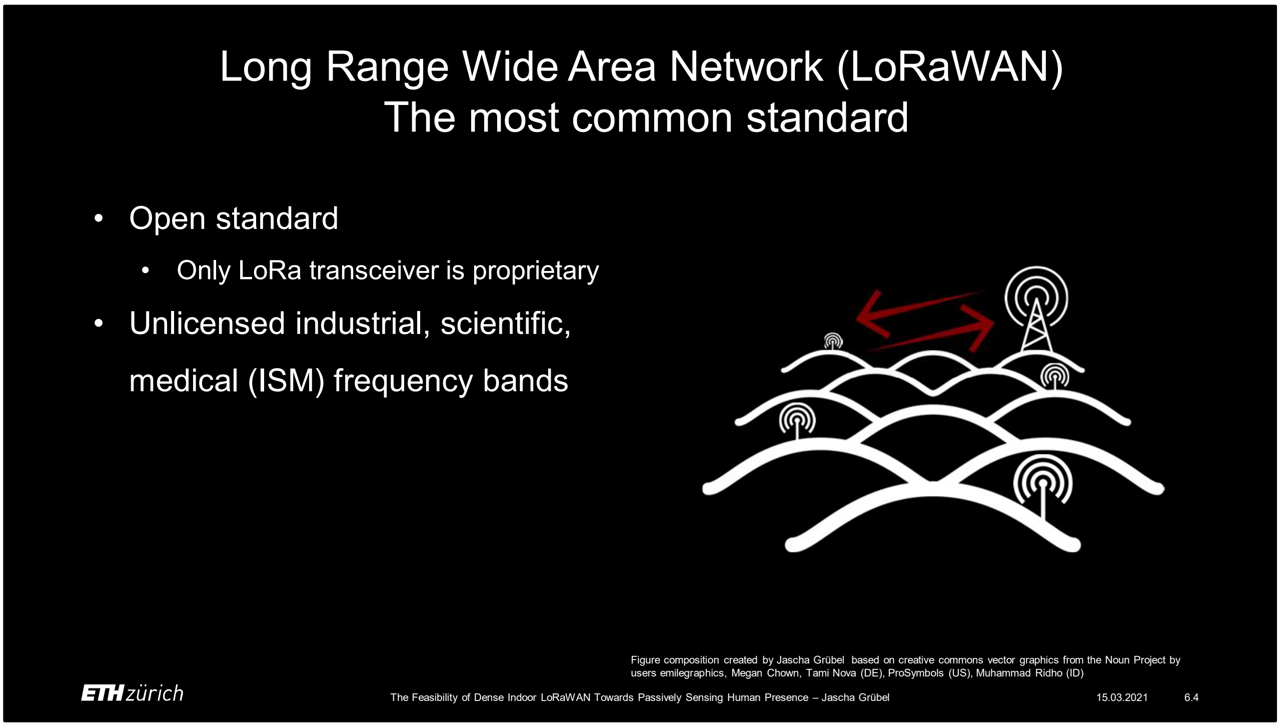

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

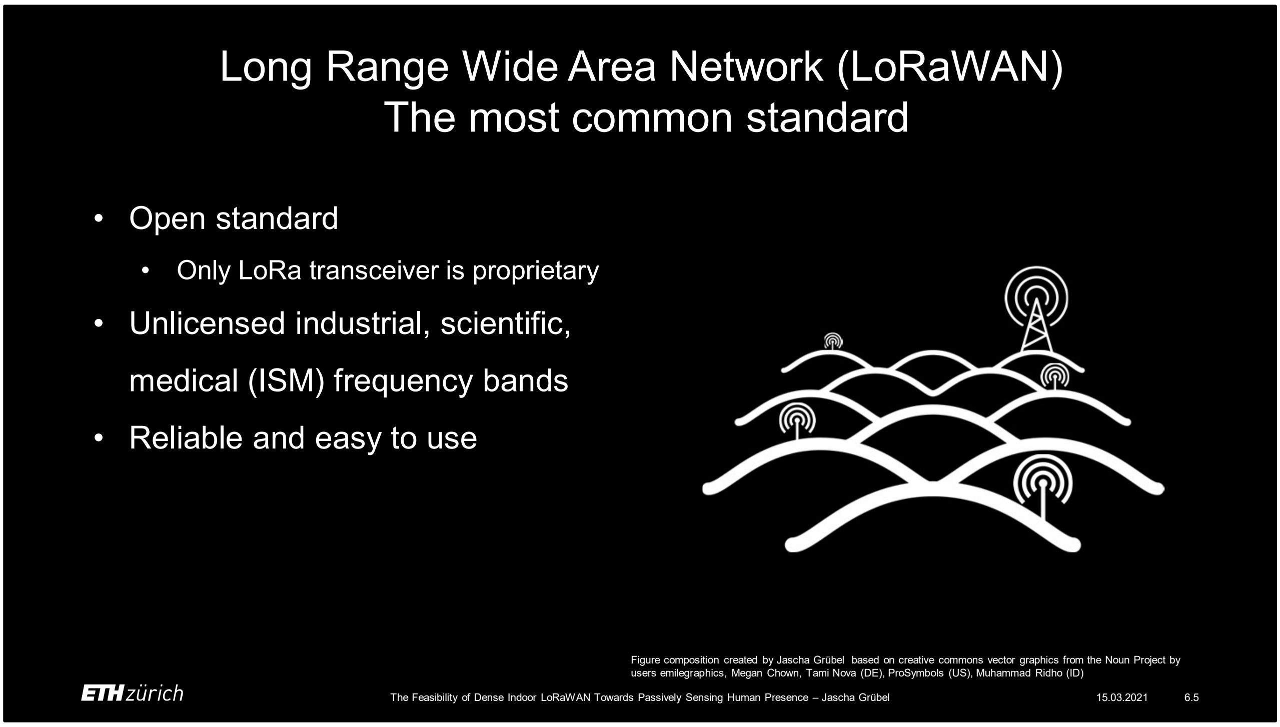

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

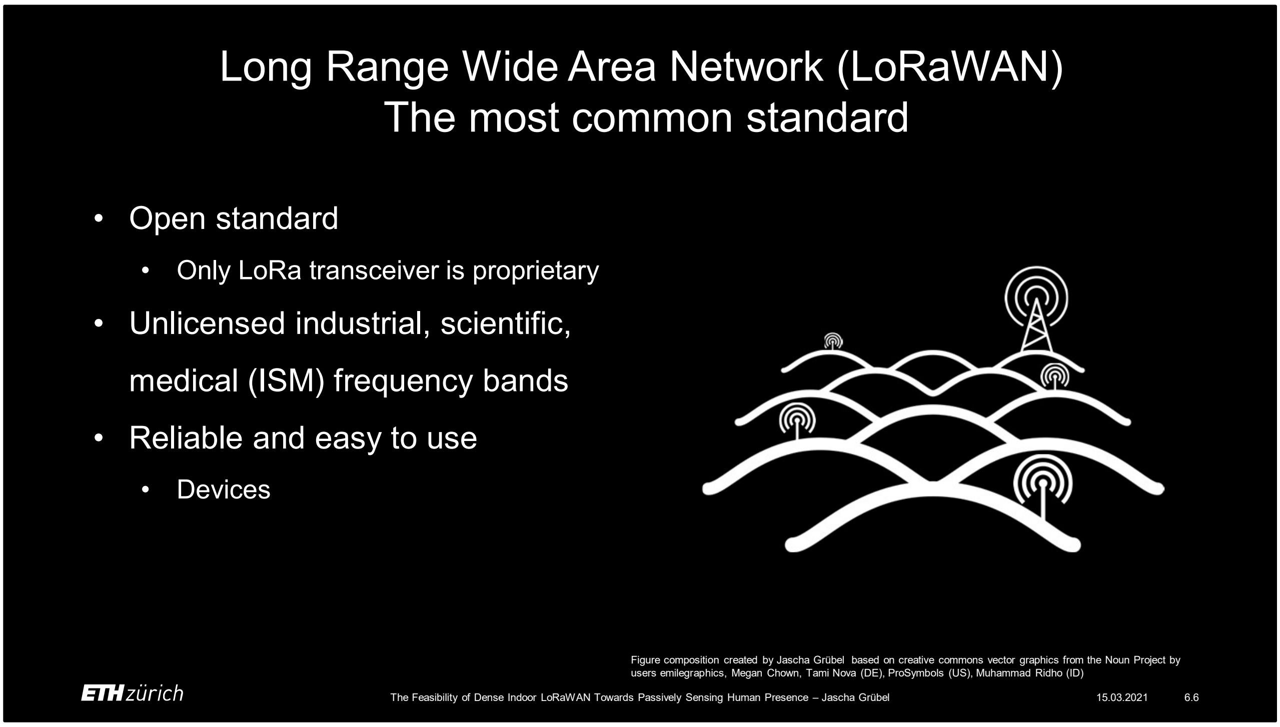

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

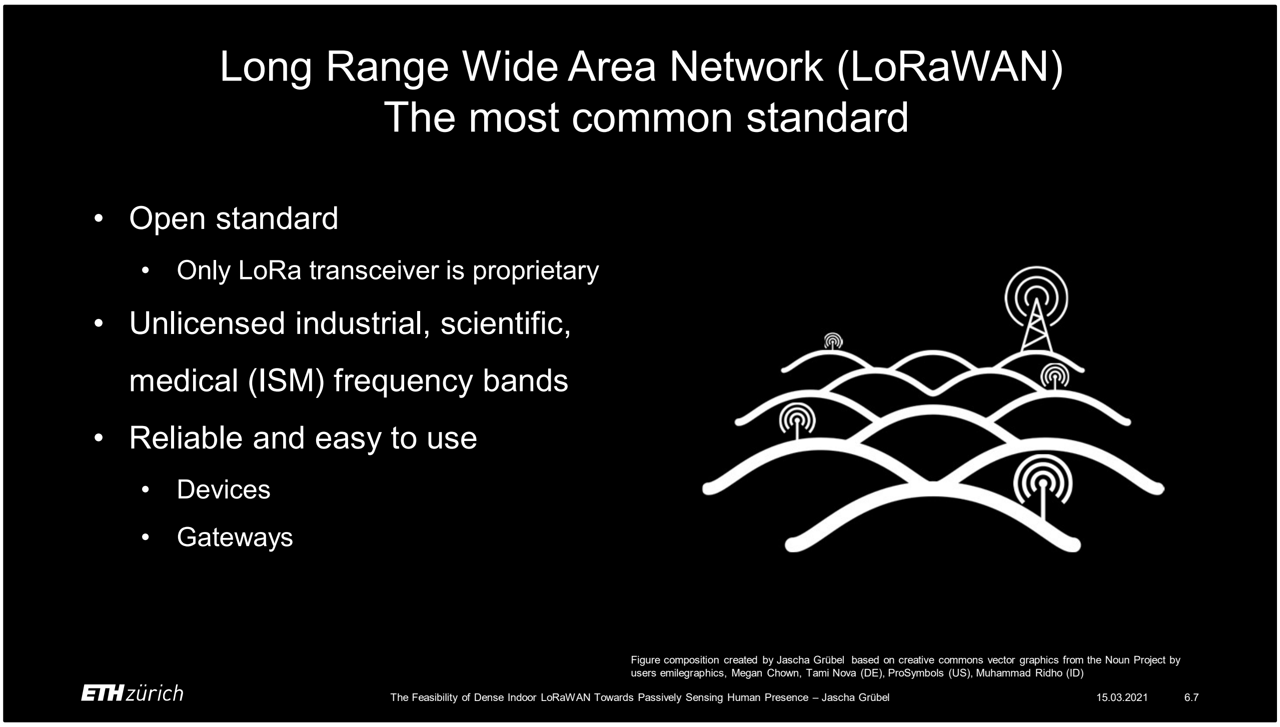

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

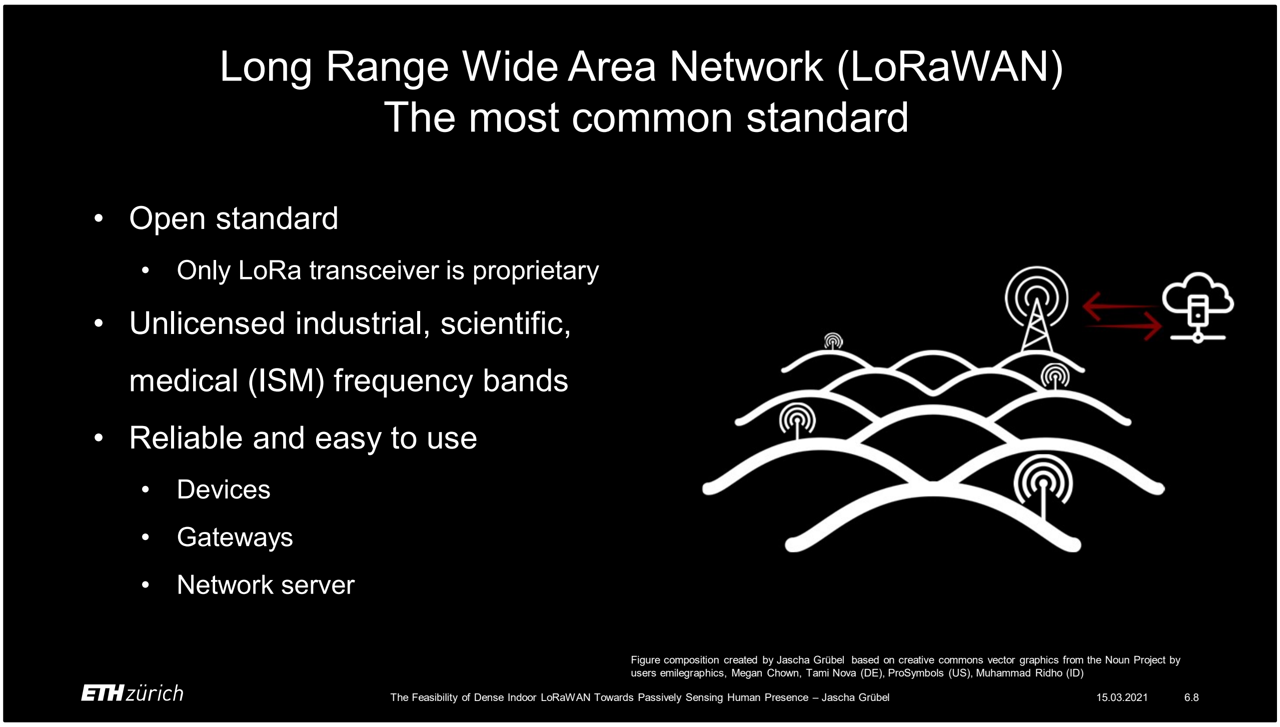

Amongst LPWAN implementations the Long Range Wide Area Network (or LoRaWAN) is the most commonly used standard exemplified by the rise of The Things Network.
LoRaWAN offers three advantages:
[C] It is an open standard backed by the industry-based LoRaWAN Alliance with
[C] the LoRa transceiver being the only proprietary technology
[C] It used the unlicensed industrial, scientific, medical (or ISM) frequency bands keeping costs low as no licensing is required.
[C] From an end user perspective, LoRaWAN is reliable and easy to use as it only requires
[C] devices,
[C] a gateways and
[C] a network server.

Before we dive into our implemtation, let me give you a quick recap of LoRaWAN and LoRa. Many of you may be experts but I would like to take everybody along.

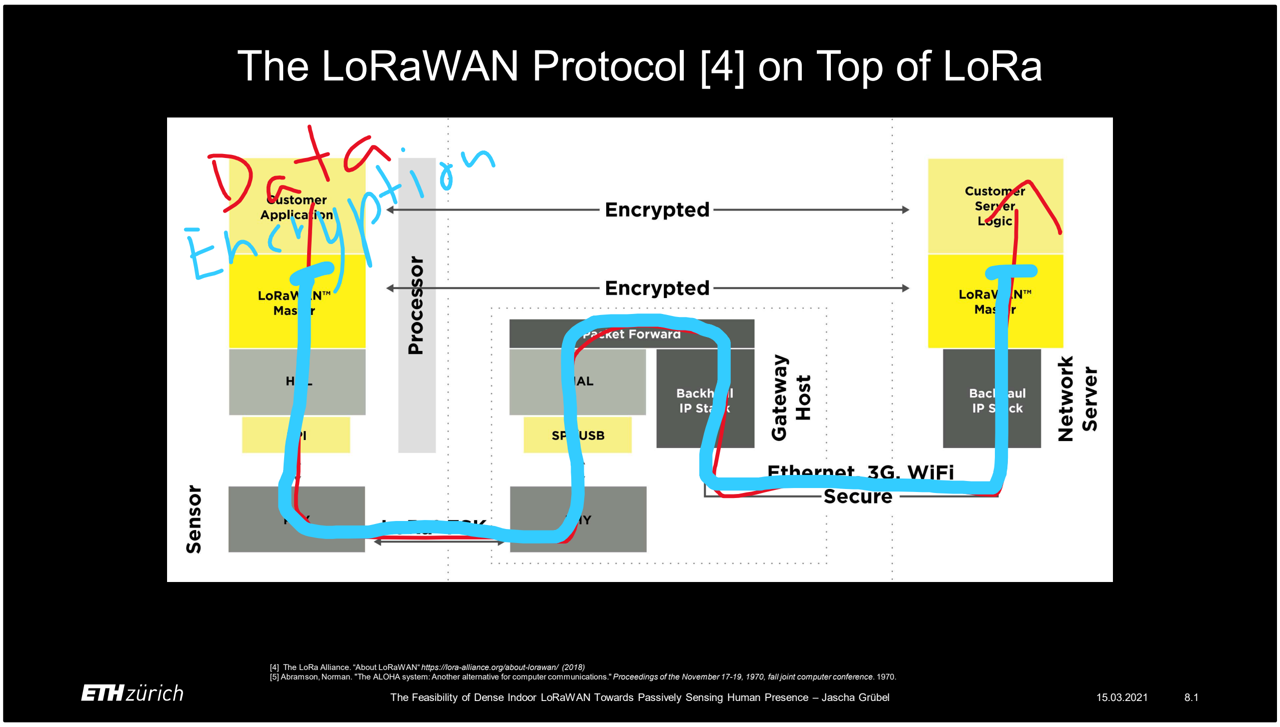

In the LoRaWAN protocol,
[C] the encrypted data from a sensor device flows to some remote application on a server.
The device uses a LoRa transceiver to send a message to the gateway.
The gateway transforms the message to the Internet Protocol (IP) to send it to the network server.
[C] The transmission is also encrypted on a second level.
The payload is encrypted between device and network server enabling a separation between transmission and content underpinning services such as The Things Network.
[C] Most commonly, LoRaWAN class A devices are used.
To conserve the most energy possible, they have an ALOHA-like receive time window after every transmission. That means the device can sleep most of the time and only expects transmission in these selected time windows.

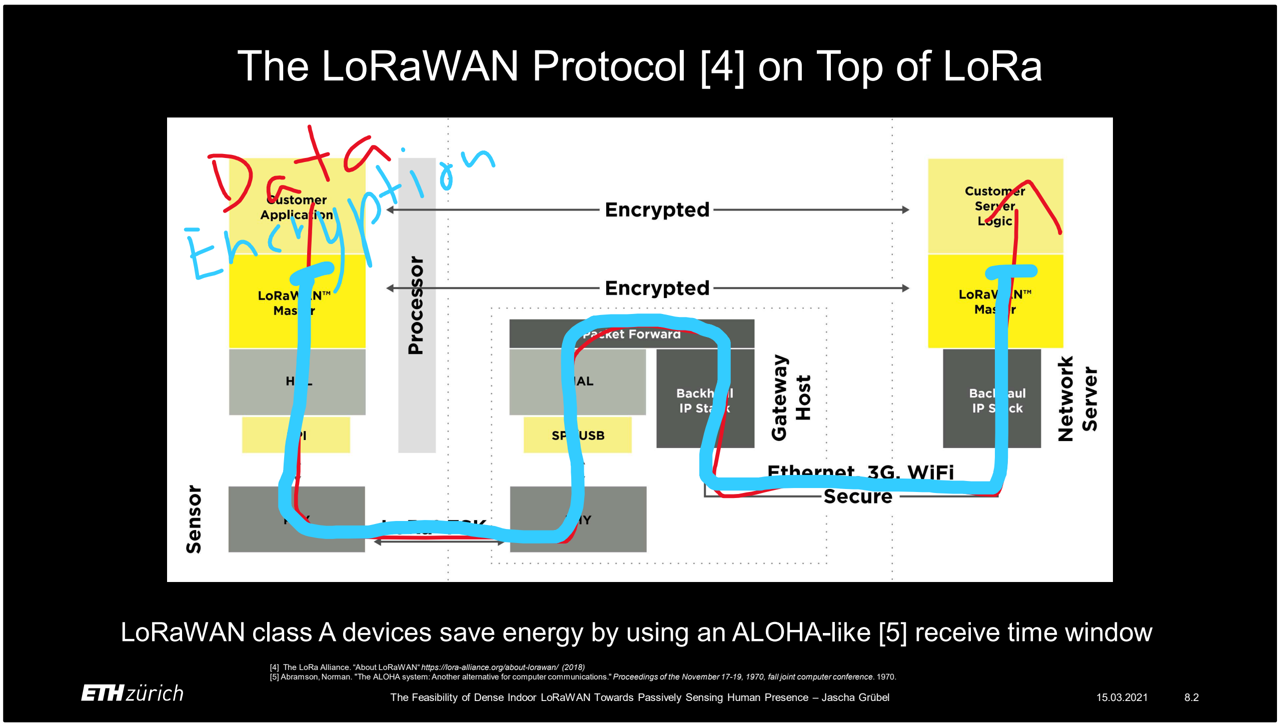

In the LoRaWAN protocol,
[C] the encrypted data from a sensor device flows to some remote application on a server.
The device uses a LoRa transceiver to send a message to the gateway.
The gateway transforms the message to the Internet Protocol (IP) to send it to the network server.
[C] The transmission is also encrypted on a second level.
The payload is encrypted between device and network server enabling a separation between transmission and content underpinning services such as The Things Network.
[C] Most commonly, LoRaWAN class A devices are used.
To conserve the most energy possible, they have an ALOHA-like receive time window after every transmission. That means the device can sleep most of the time and only expects transmission in these selected time windows.

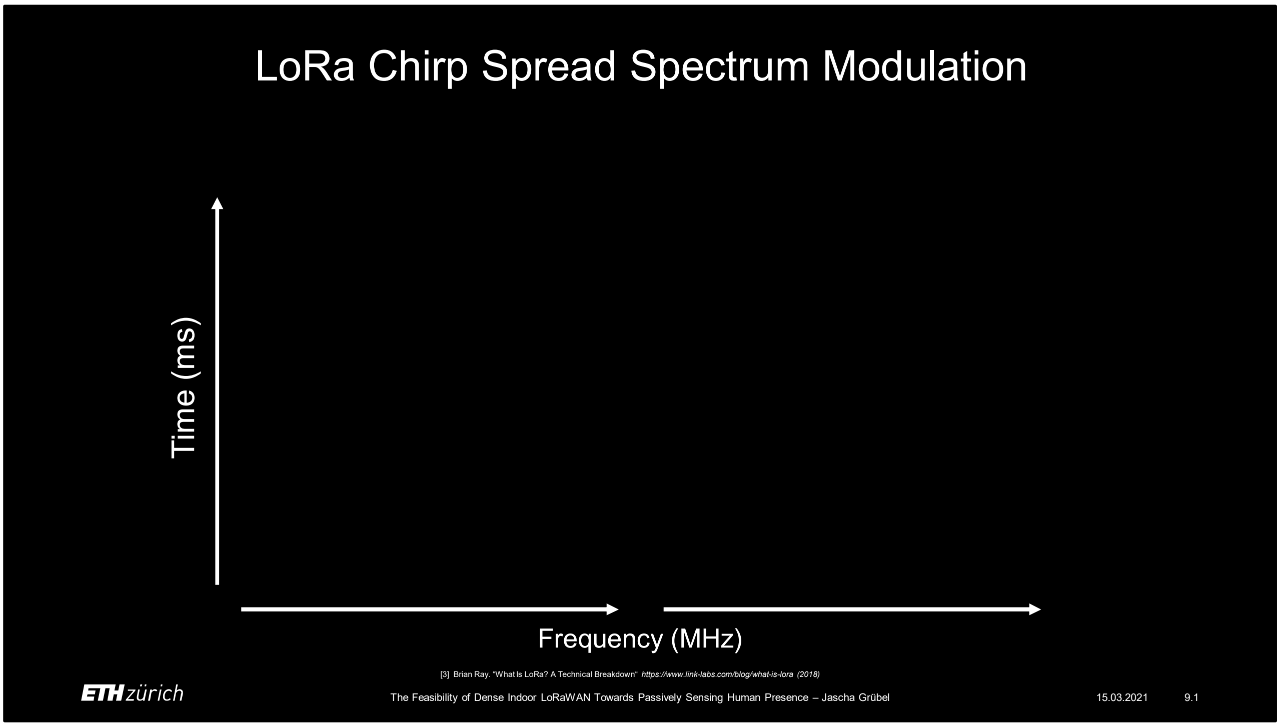

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

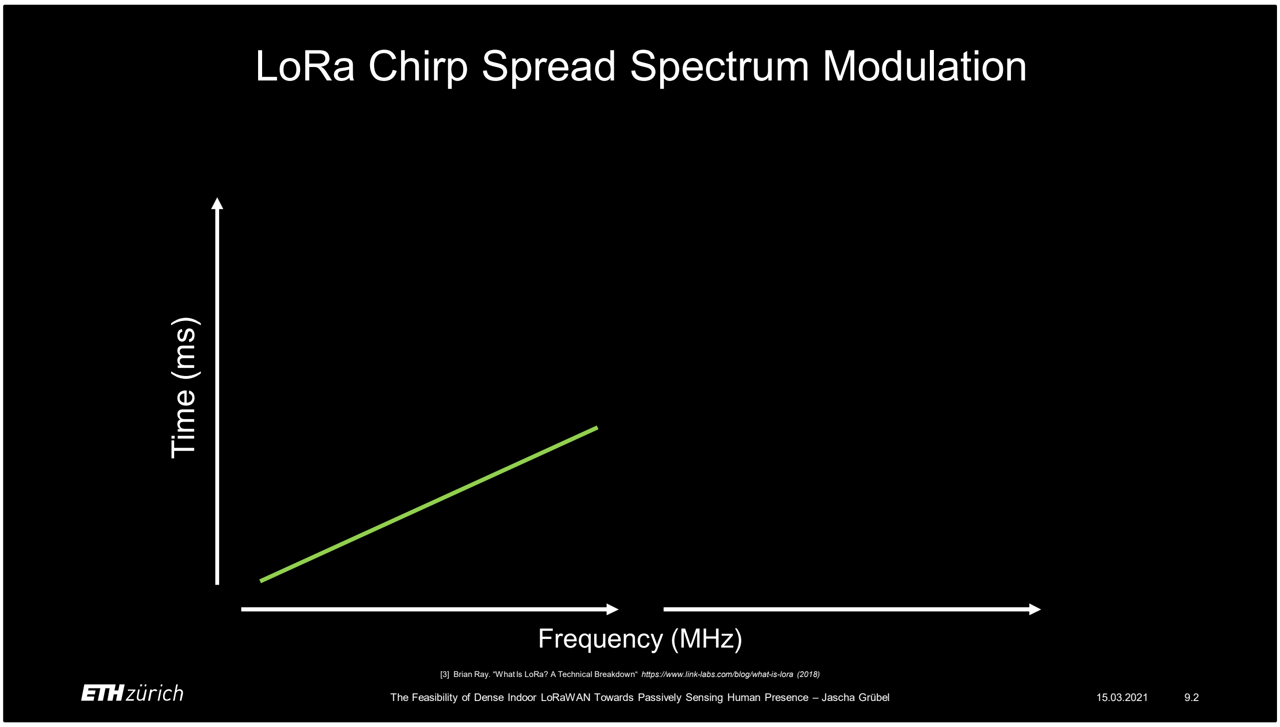

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

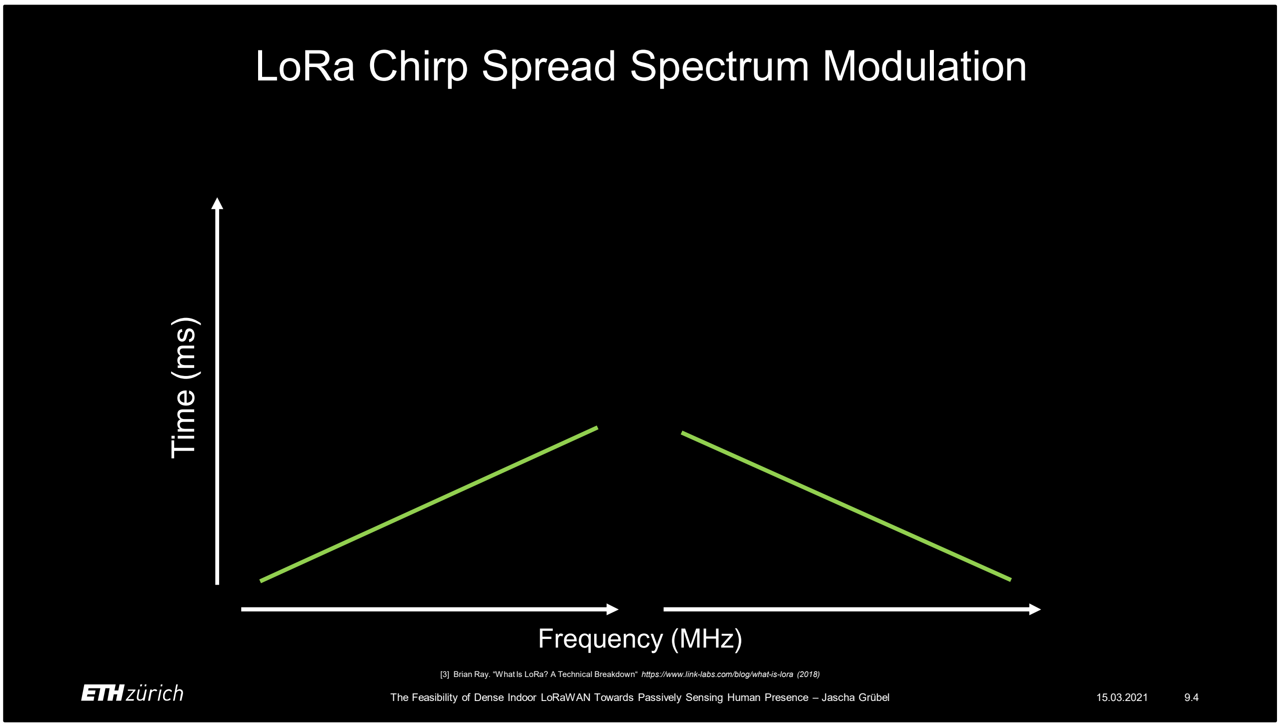

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

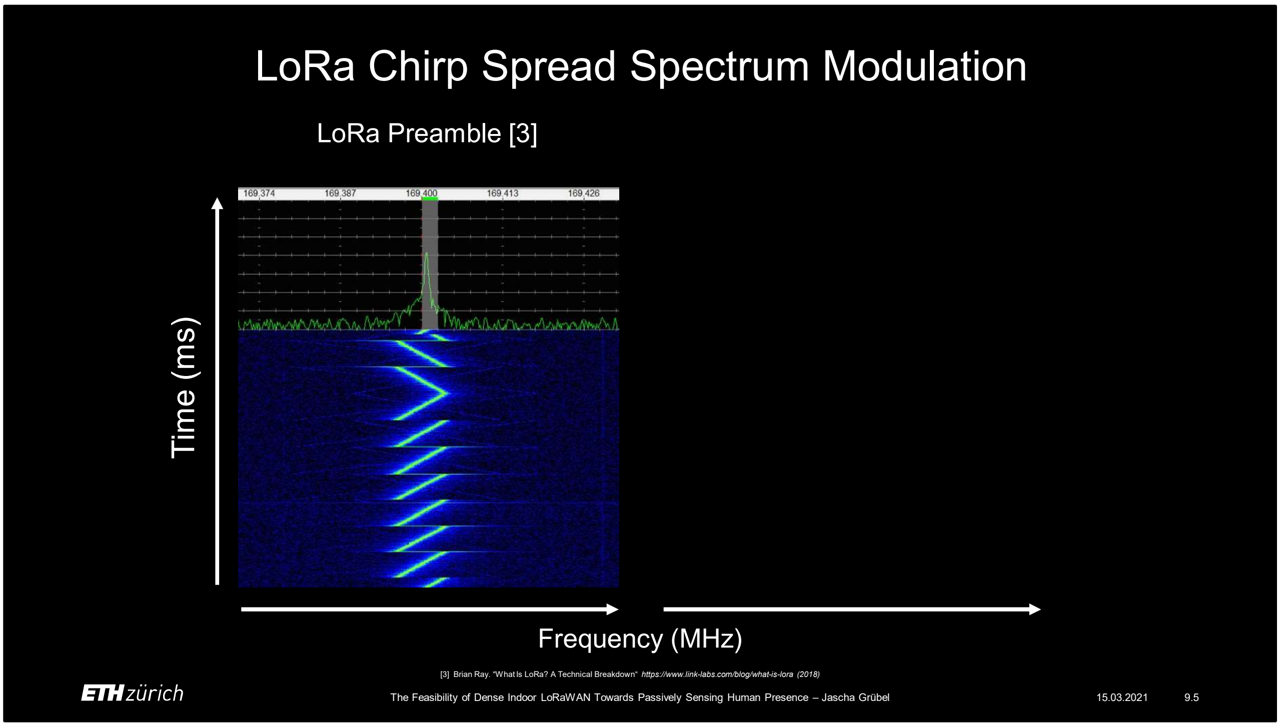

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

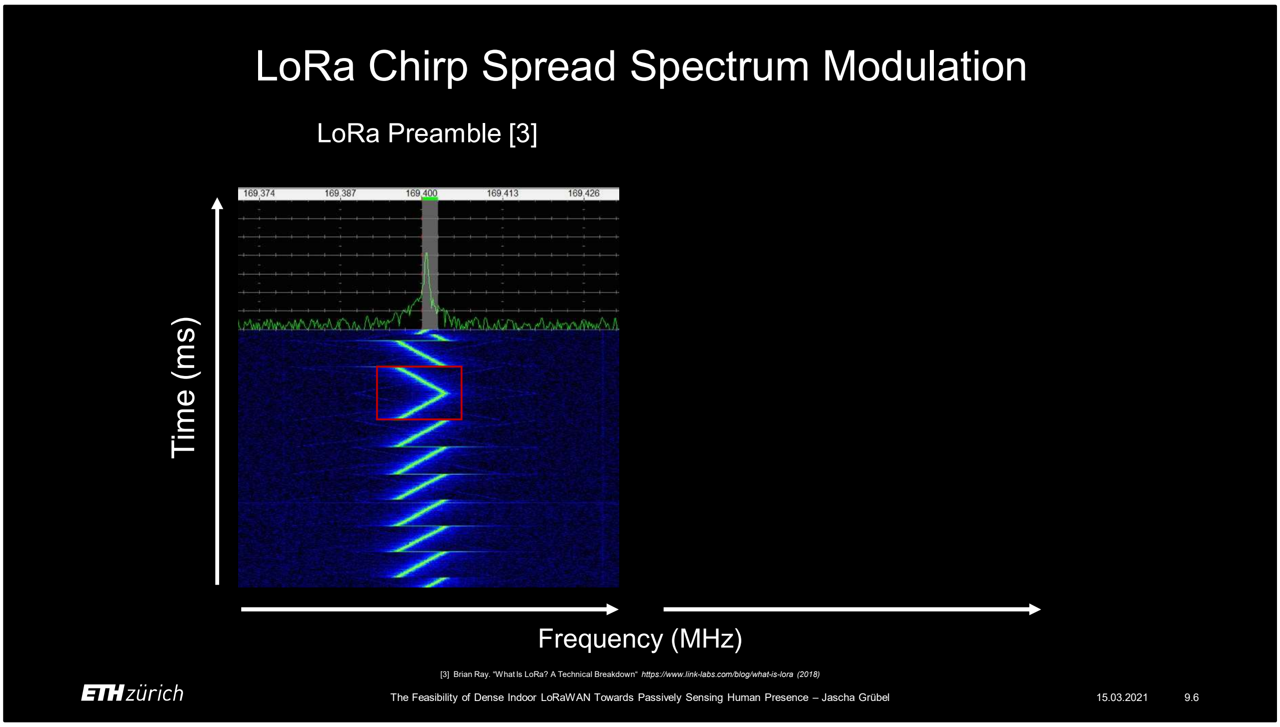

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

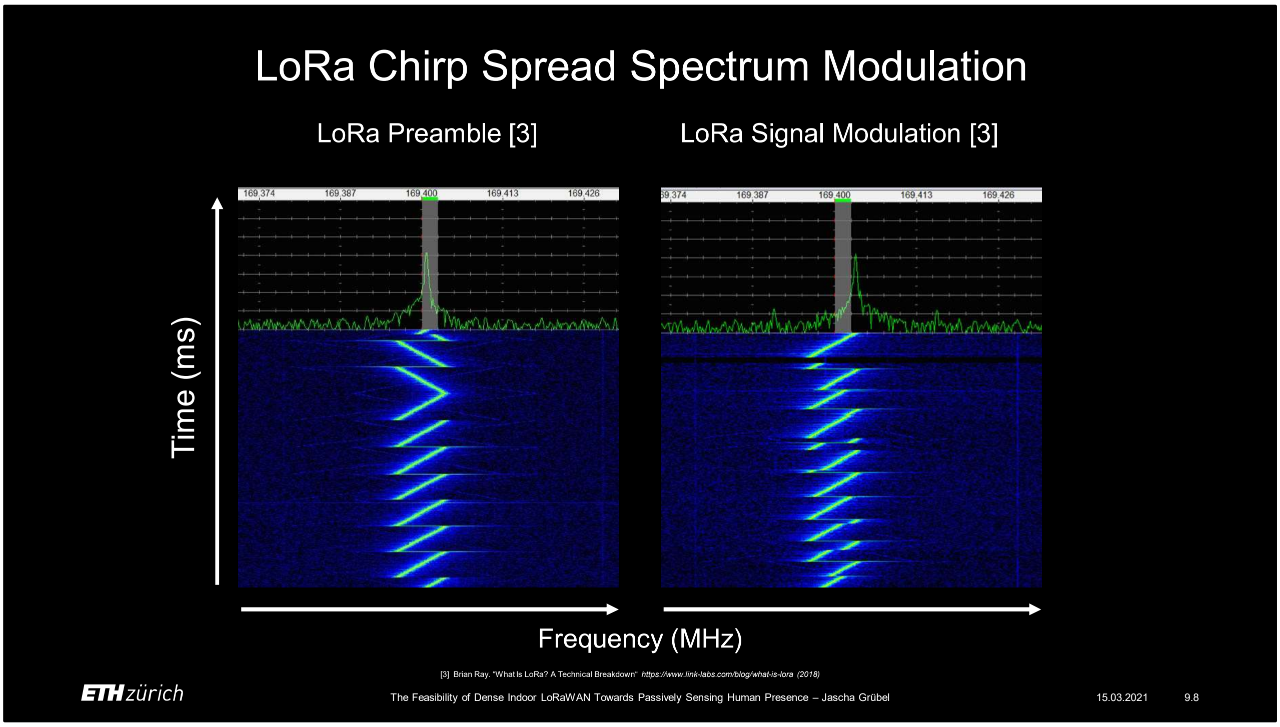

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

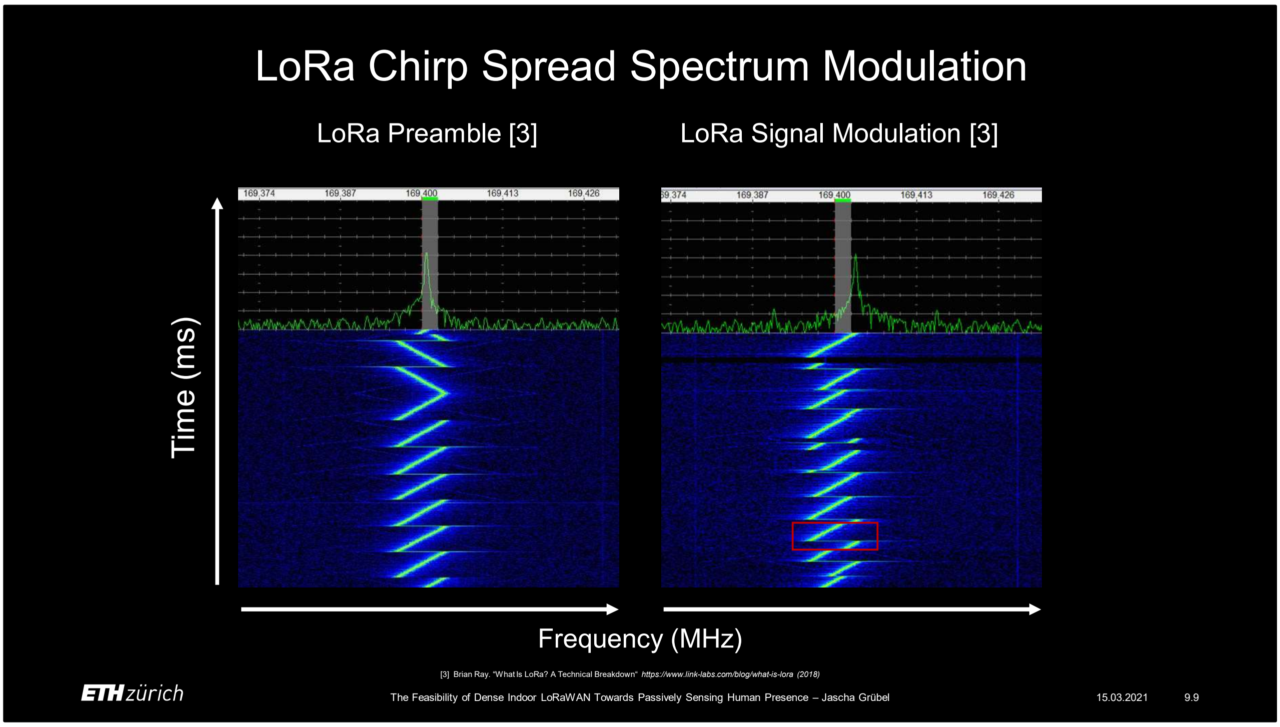

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

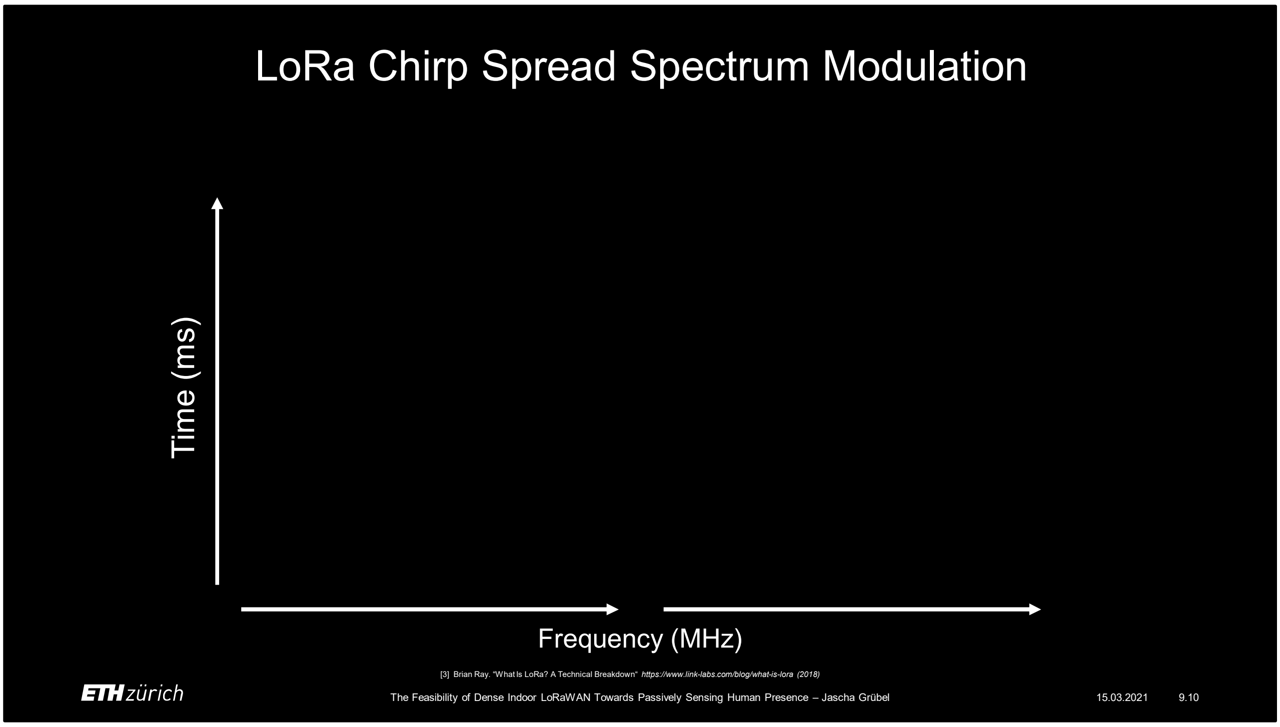

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

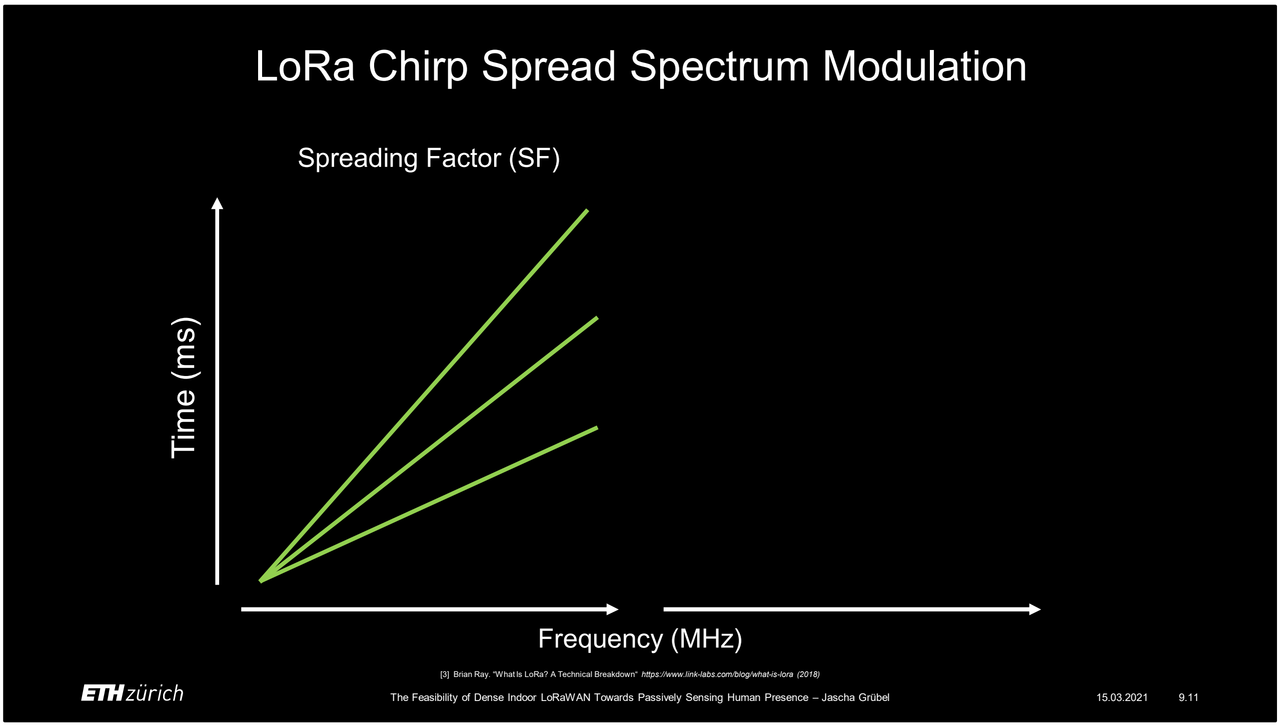

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

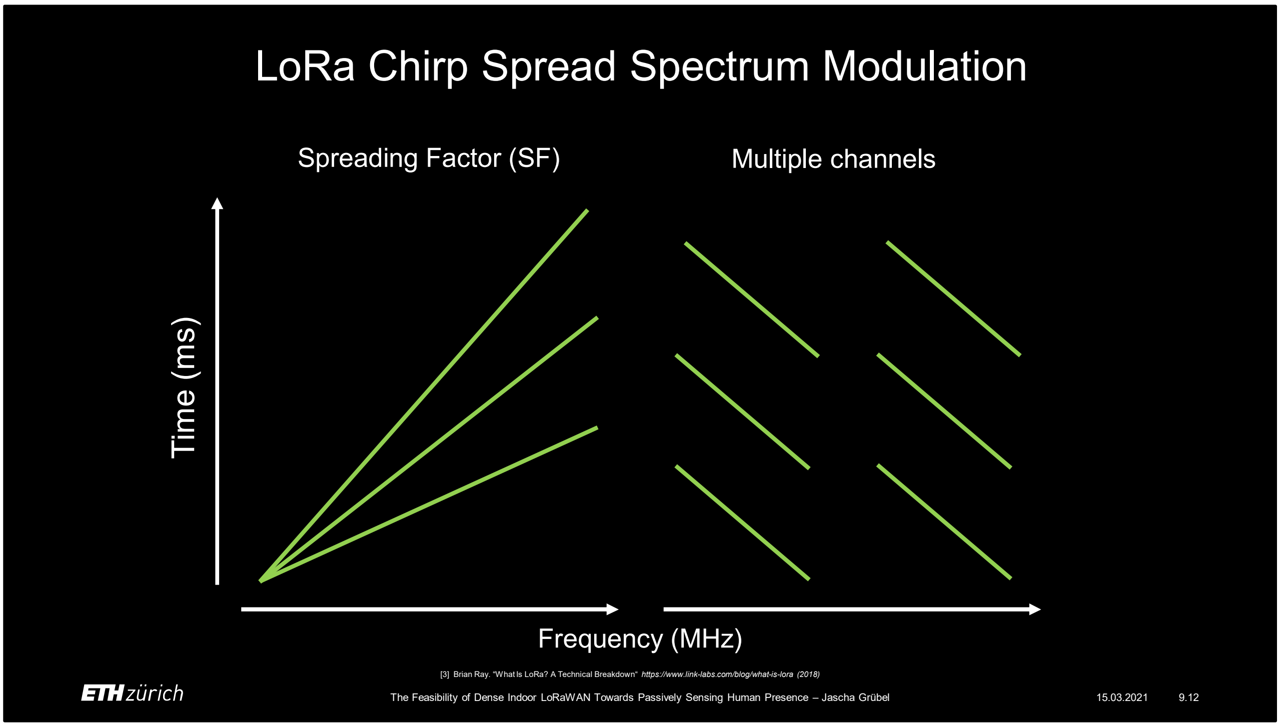

LoRa is based on a chirp from
[C] a lower to higher frequency or
[C] vice versa.
This chirp produces robustness because of its duration and pattern.
[C] To start communication, the preamble is sent. This allows listening devices to record the signal [C] after the inverted chirp.
[C] The symbols are encoded by modulation.
[C] One symbol still consists of a chirp but the starting point is cyclically shifted.
[C] The spreading factor defines the time that a chirp requires and indirectly the data rate that can be achieved.
This entails multiple trade-offs between the data rate (how much we can send), the integration time (here a longer signal is better) and the number of transmissions (here a shorter signal is better).
[C] Using multiple channels, i.e. base frequencies and multiple spreading factors (i.e. chirp durations) allows for robust concurrent transmission of messages because each chirp can be uniquely identified as long as they do not occur on the same channel at the same spreading factor and time.
Generally, the larger distance, the higher the required SF must be to

reconstruct the signal.

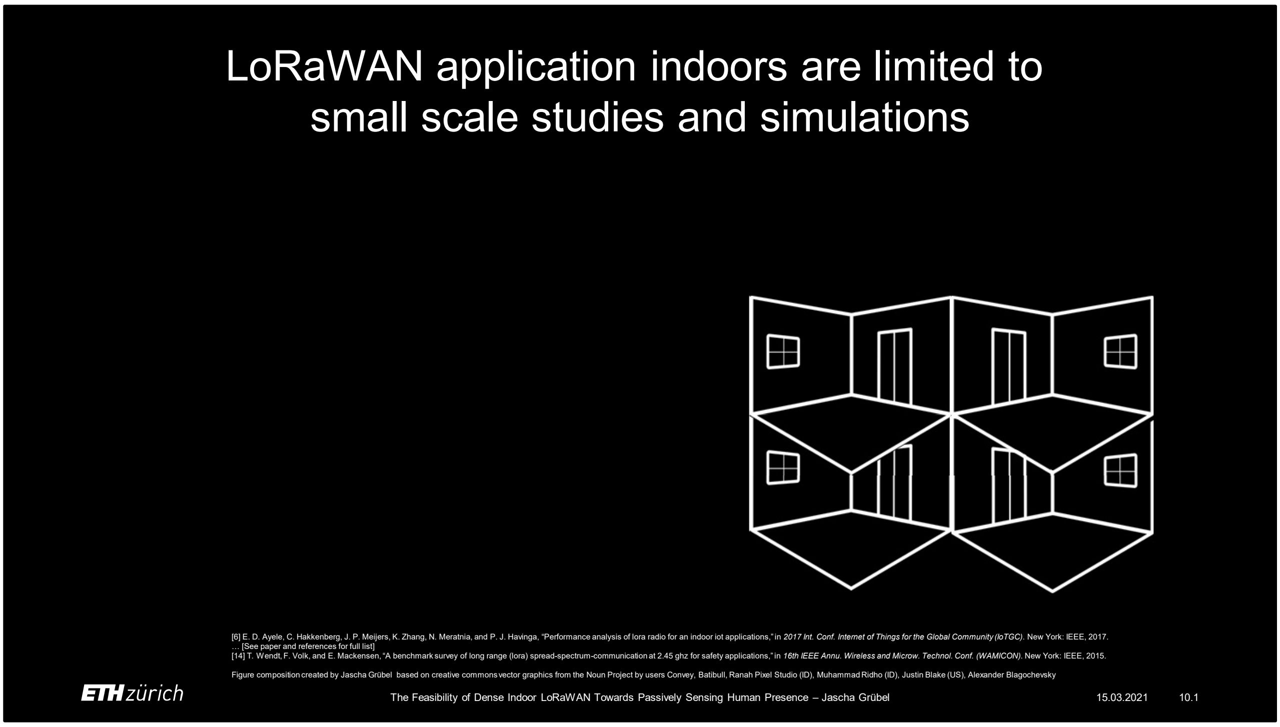

We are not the first to apply this outdoor long-distance technology indoors but previous studies have been small scale or simulation-based.
[C] Usually, a study consisted of a minimal viable design (one device, one gateway, and one network server).
[C] often moving the device around to get different measurements.
[C] Hundreds of devices or multiple gateways have only been used in simulations.
[C] Even when data was gathered, usually very few transmissions were used to provide some descriptive data.

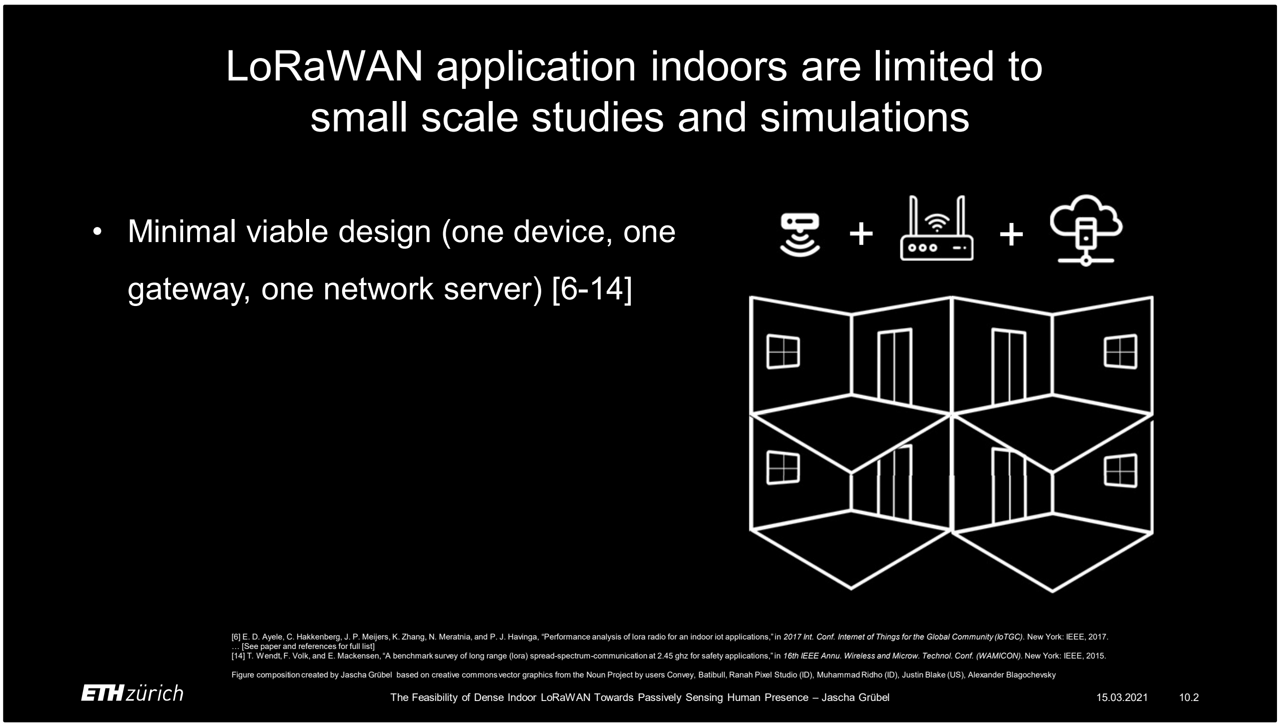

We are not the first to apply this outdoor long-distance technology indoors but previous studies have been small scale or simulation-based.
[C] Usually, a study consisted of a minimal viable design (one device, one gateway, and one network server).
[C] often moving the device around to get different measurements.
[C] Hundreds of devices or multiple gateways have only been used in simulations.
[C] Even when data was gathered, usually very few transmissions were used to provide some descriptive data.

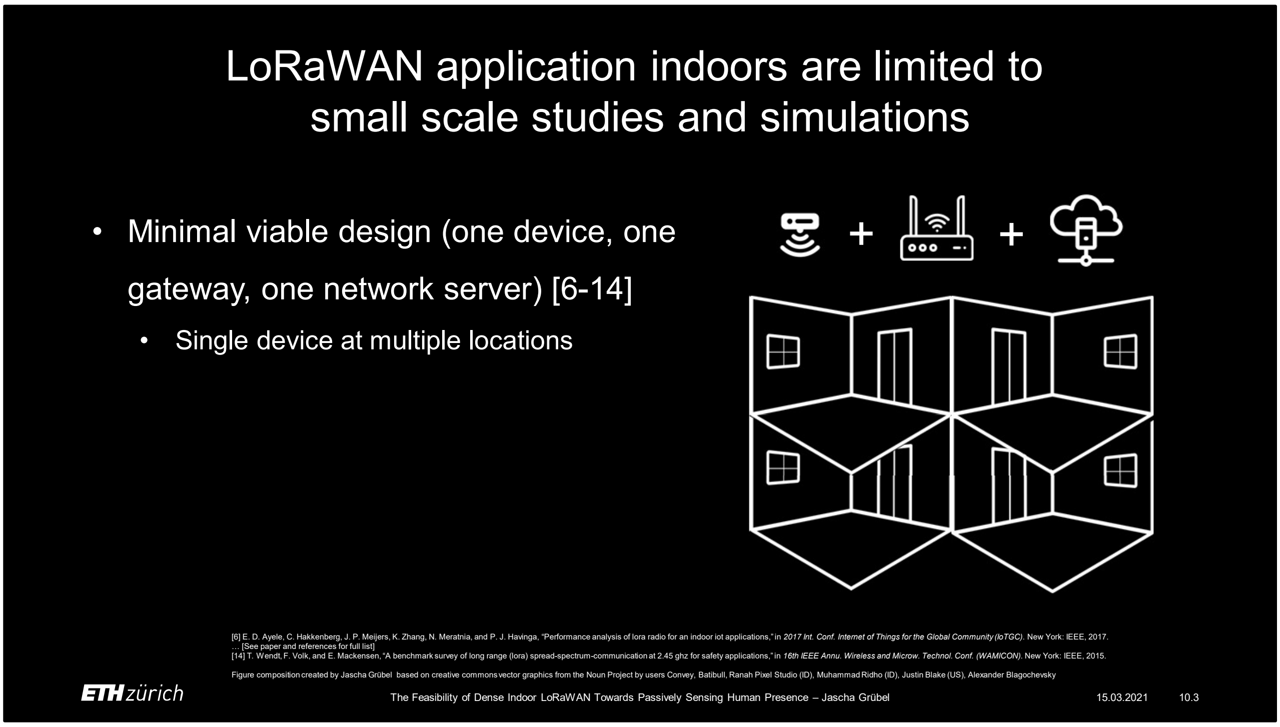

We are not the first to apply this outdoor long-distance technology indoors but previous studies have been small scale or simulation-based.
[C] Usually, a study consisted of a minimal viable design (one device, one gateway, and one network server).
[C] often moving the device around to get different measurements.
[C] Hundreds of devices or multiple gateways have only been used in simulations.
[C] Even when data was gathered, usually very few transmissions were used to provide some descriptive data.

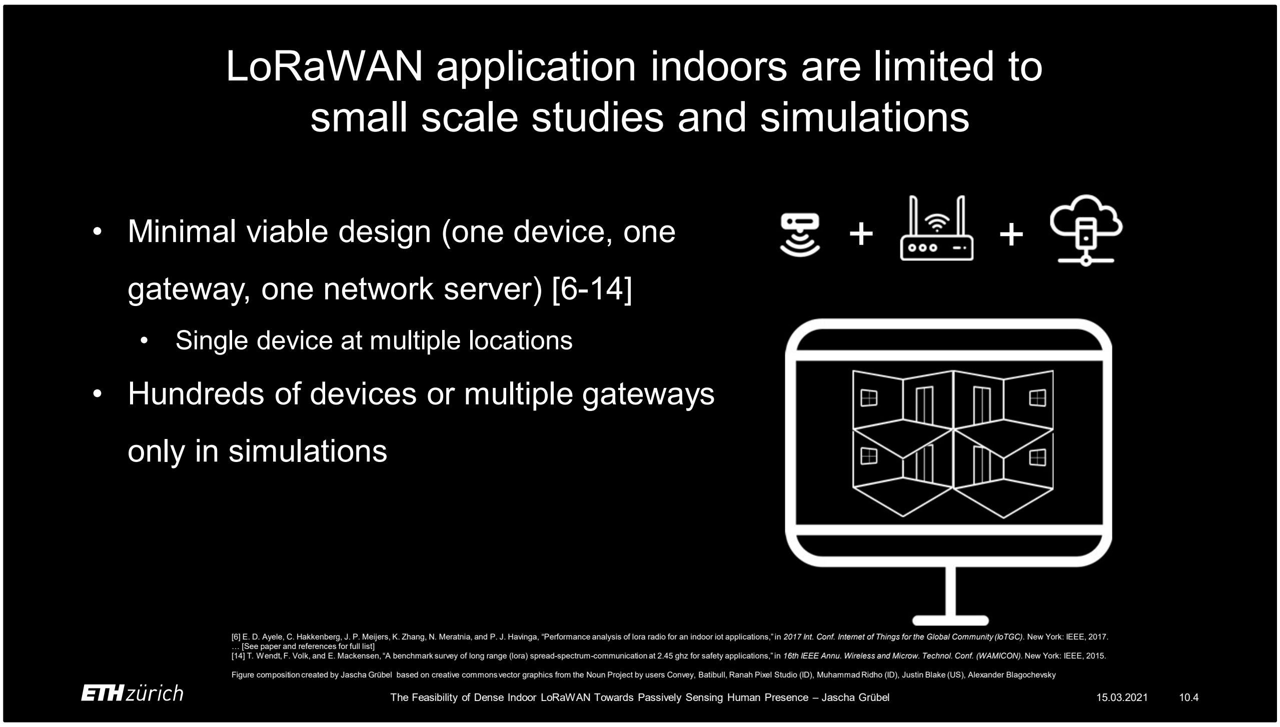

We are not the first to apply this outdoor long-distance technology indoors but previous studies have been small scale or simulation-based.
[C] Usually, a study consisted of a minimal viable design (one device, one gateway, and one network server).
[C] often moving the device around to get different measurements.
[C] Hundreds of devices or multiple gateways have only been used in simulations.
[C] Even when data was gathered, usually very few transmissions were used to provide some descriptive data.

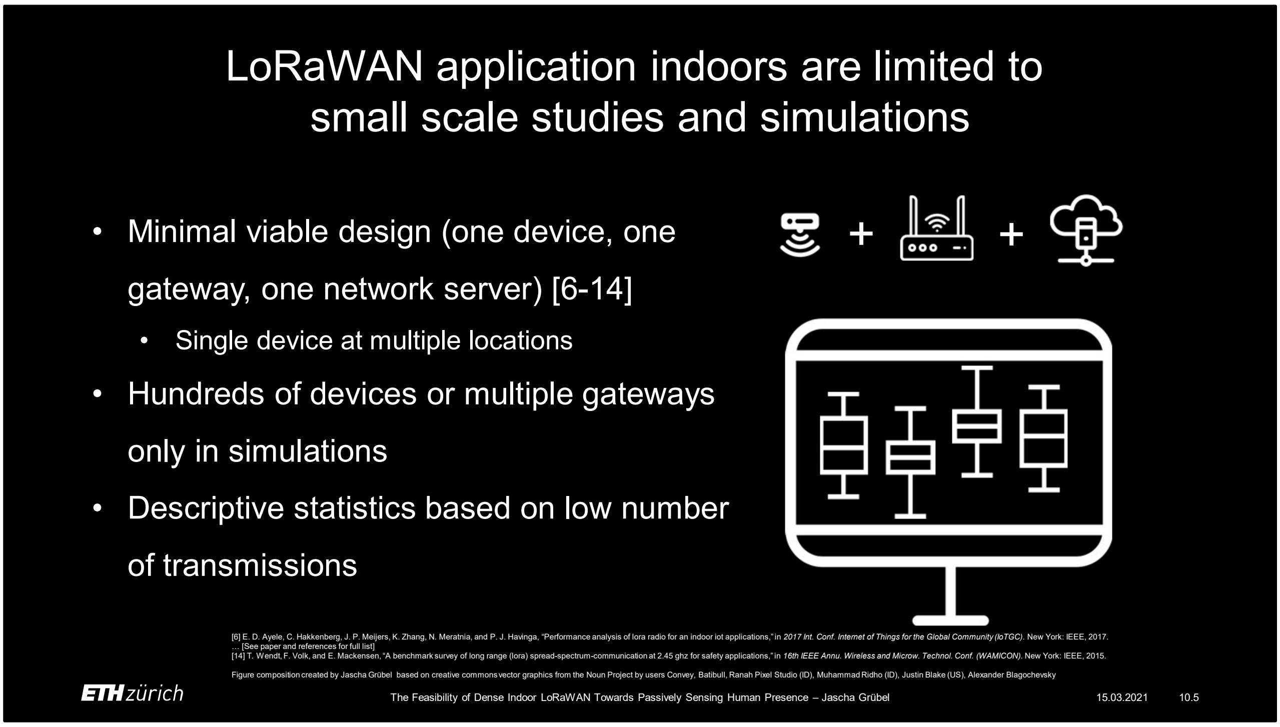

We are not the first to apply this outdoor long-distance technology indoors but previous studies have been small scale or simulation-based.
[C] Usually, a study consisted of a minimal viable design (one device, one gateway, and one network server).
[C] often moving the device around to get different measurements.
[C] Hundreds of devices or multiple gateways have only been used in simulations.
[C] Even when data was gathered, usually very few transmissions were used to provide some descriptive data.

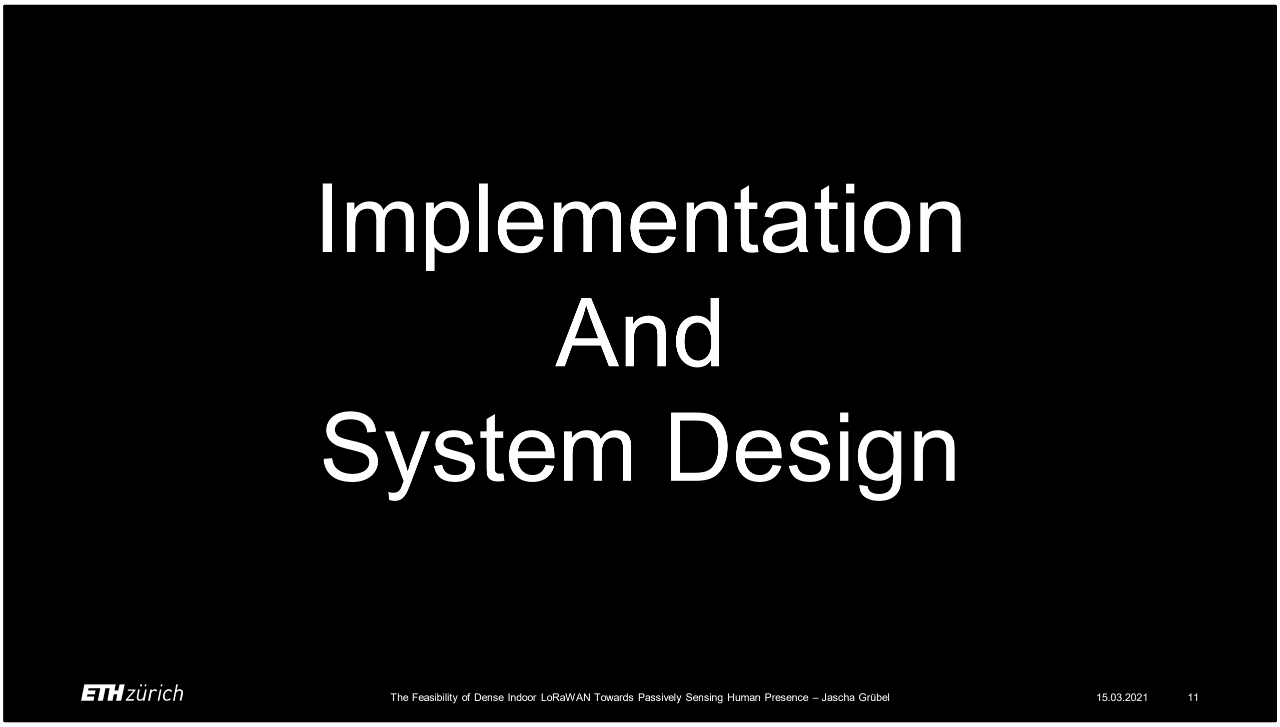

We wanted to design and implement a real dense indoor sensor network that goes beyond simulations to test what happens in the real world.

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

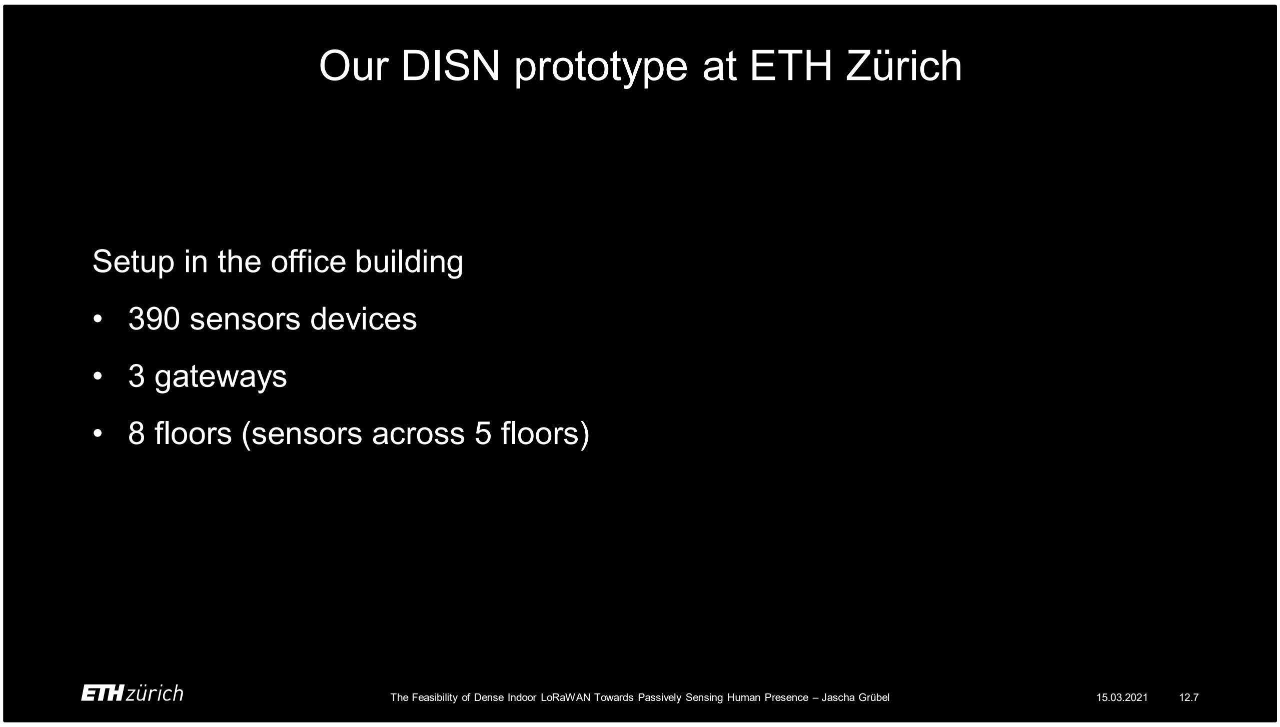

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

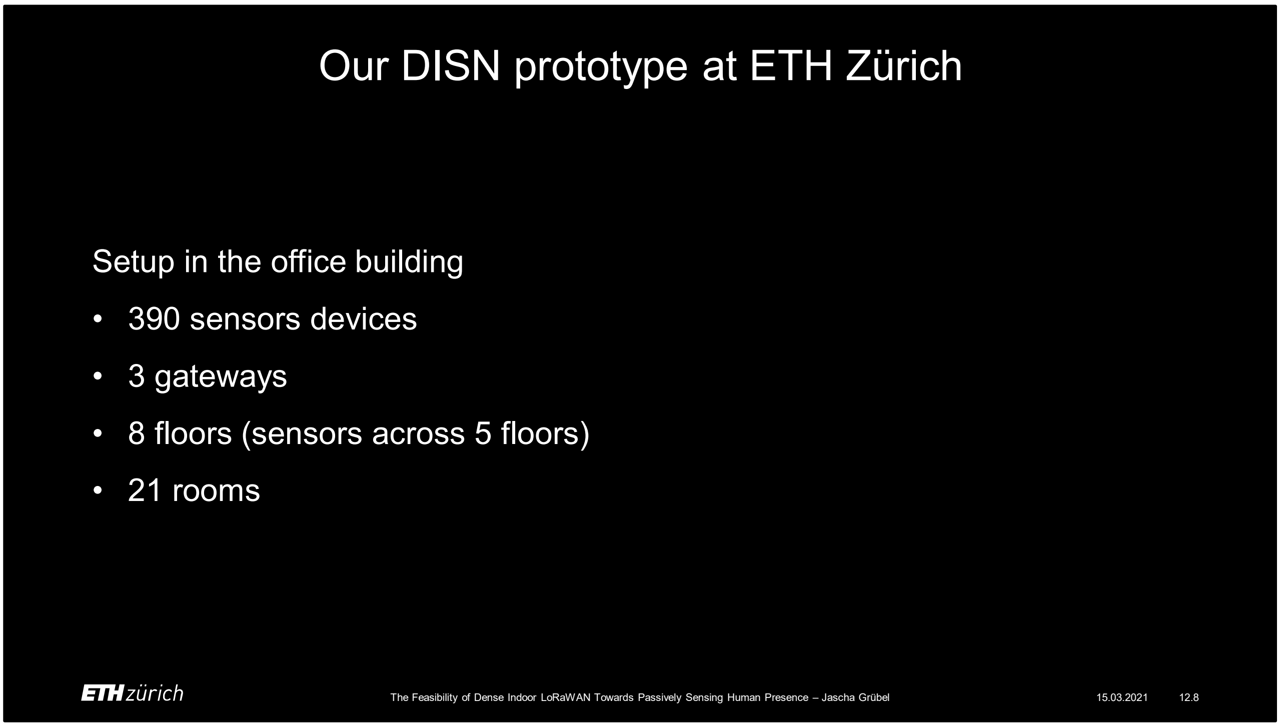

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

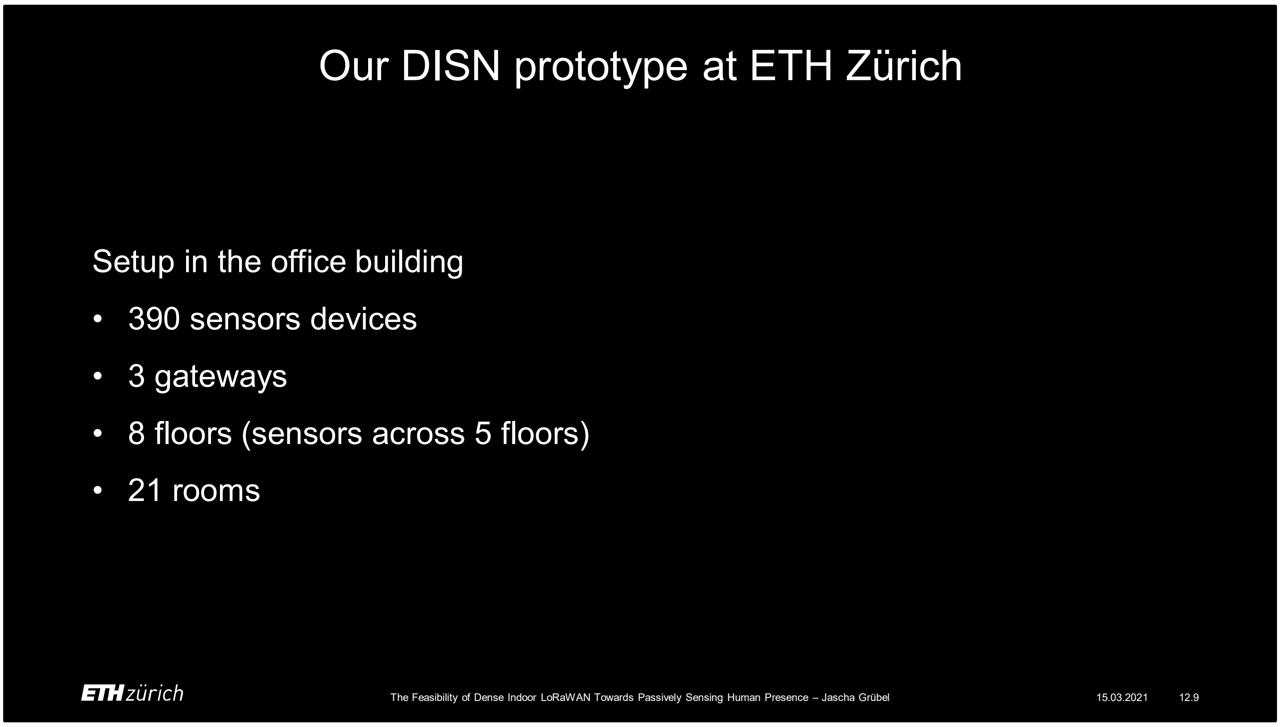

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

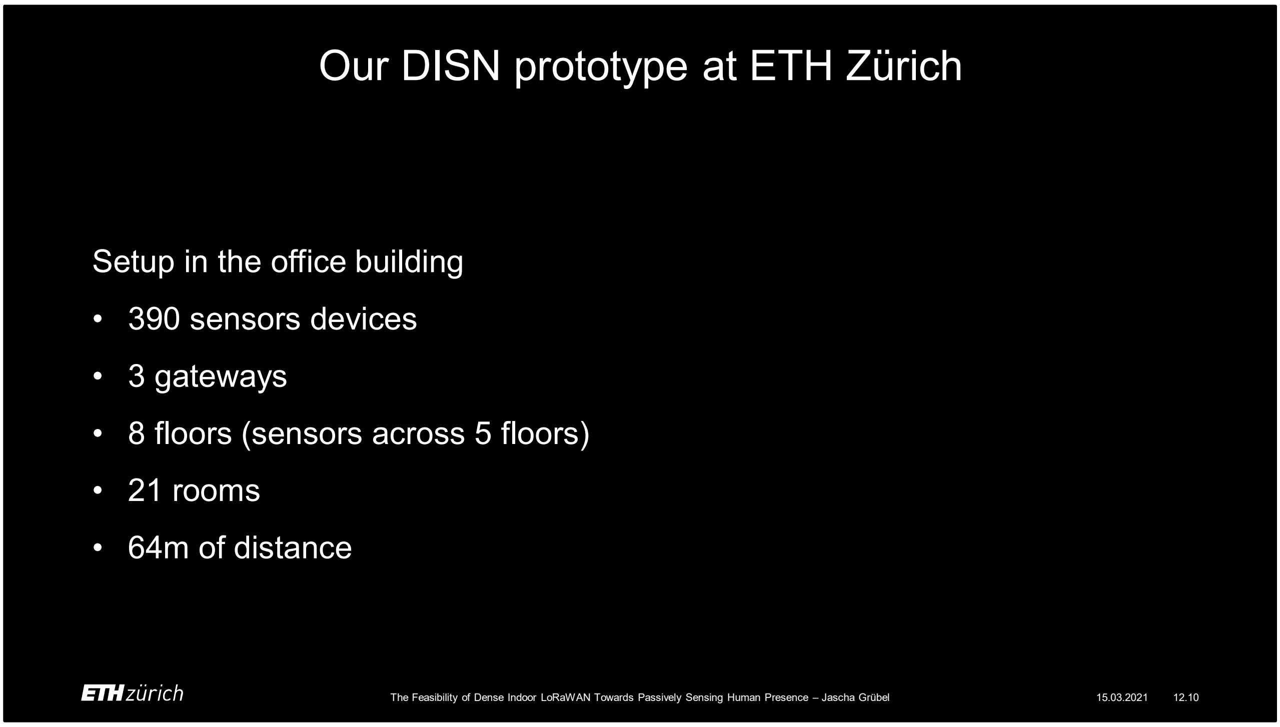

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

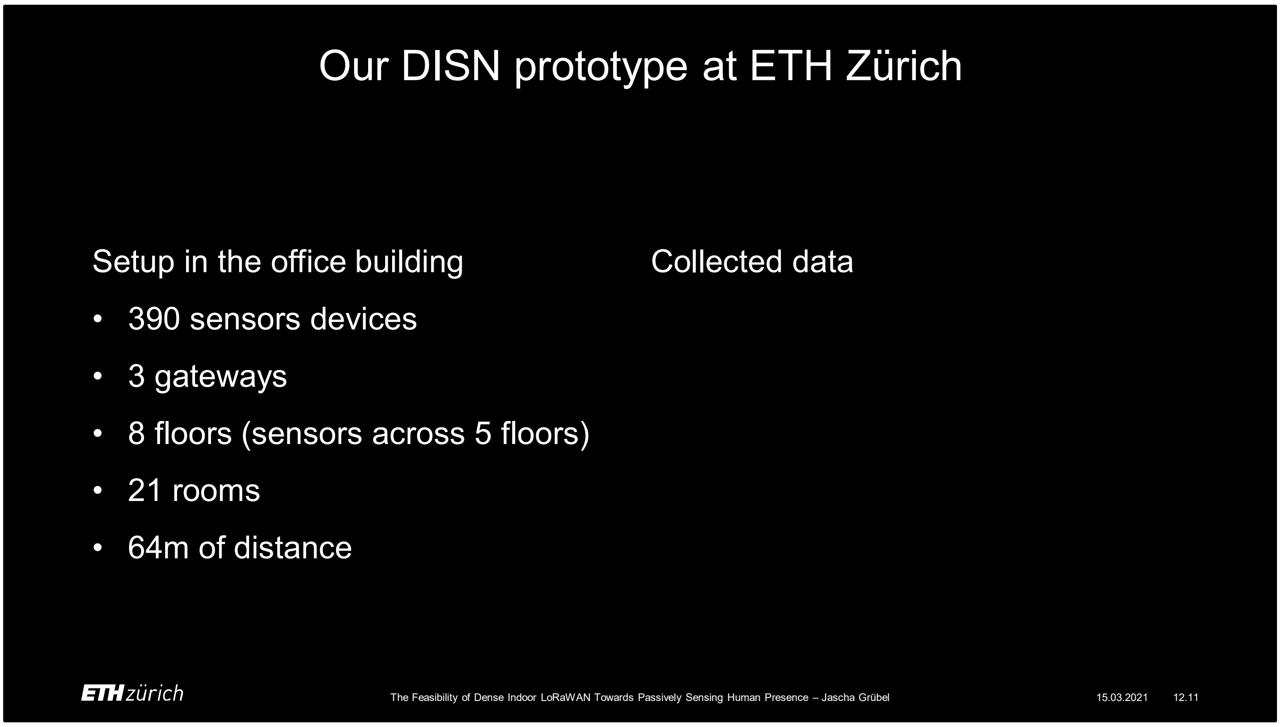

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

Our DISN prototype at ETH Zürich

Setup in the office building

- 390 sensors devices
- 3 gateways
- 8 floors (sensors across 5 floors)
- 21 rooms
- 64m of distance

Collected data

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

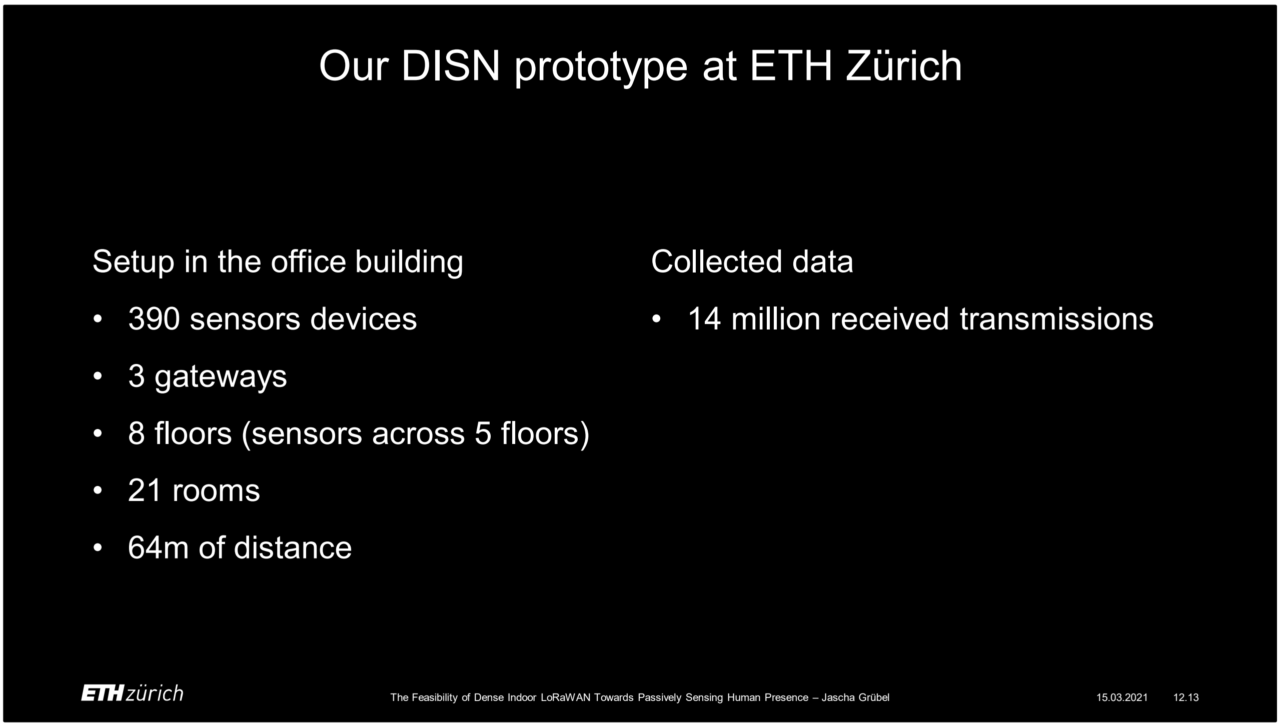

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

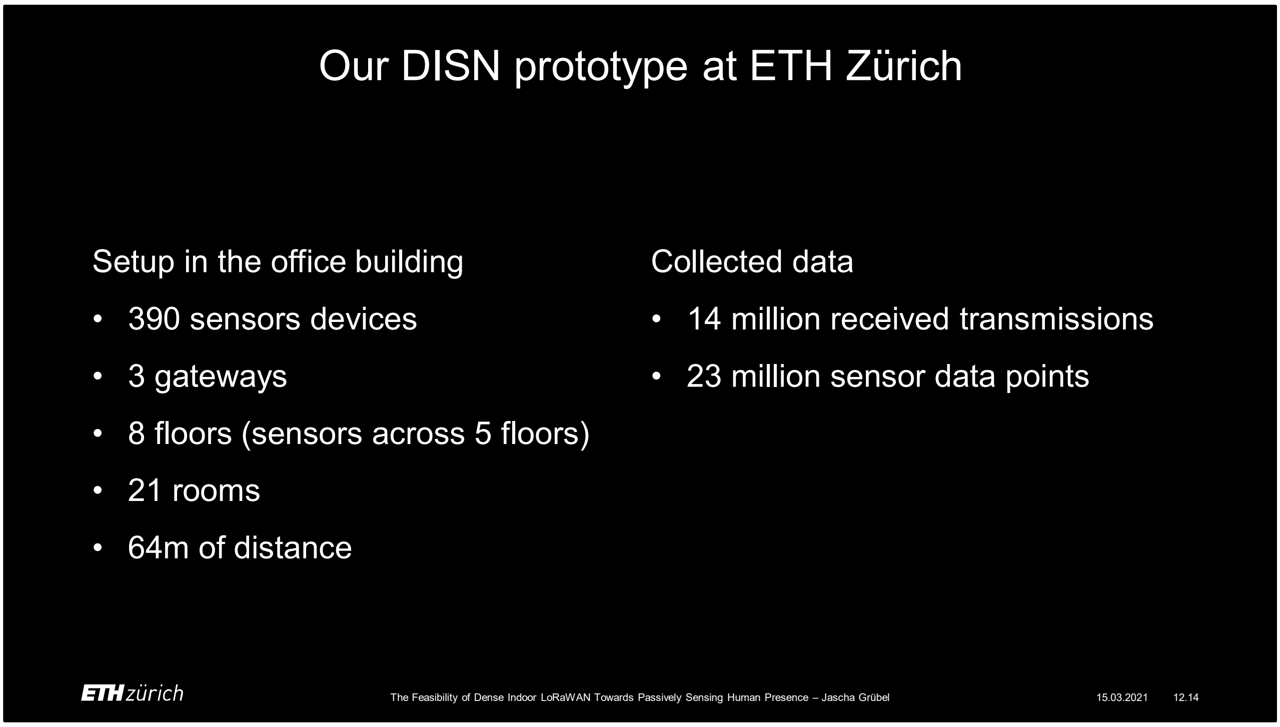

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

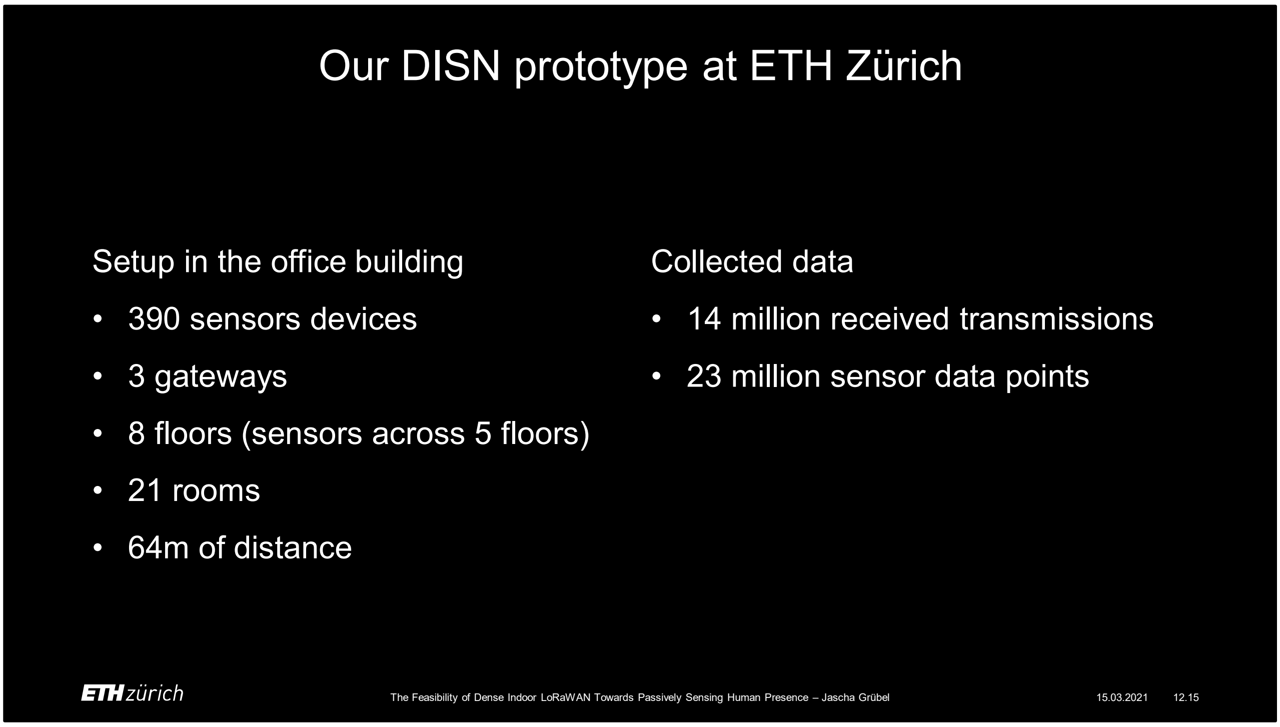

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

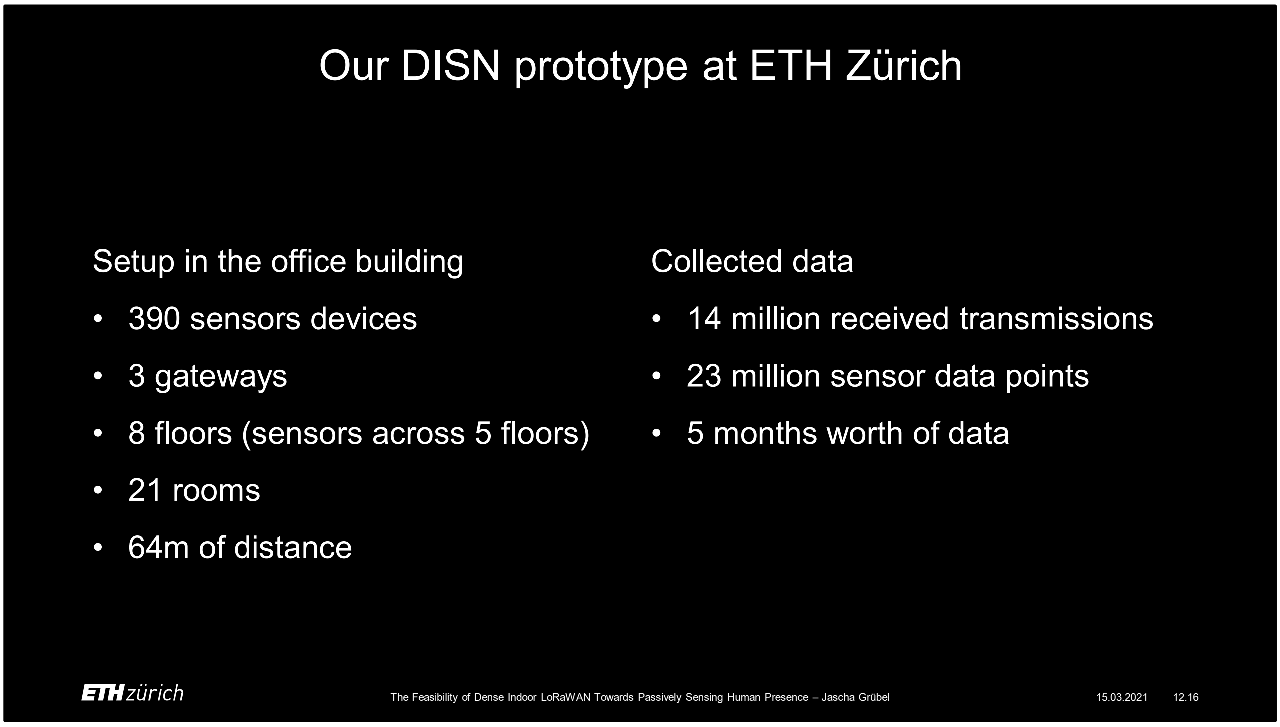

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

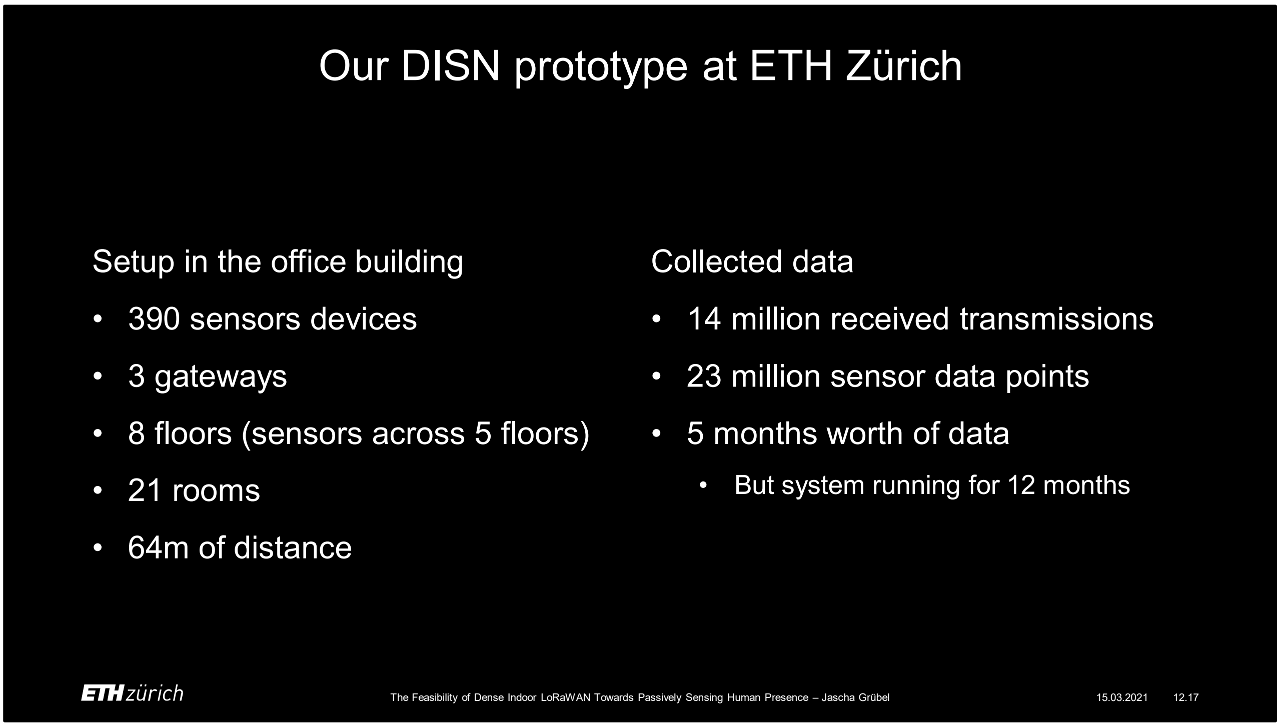

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

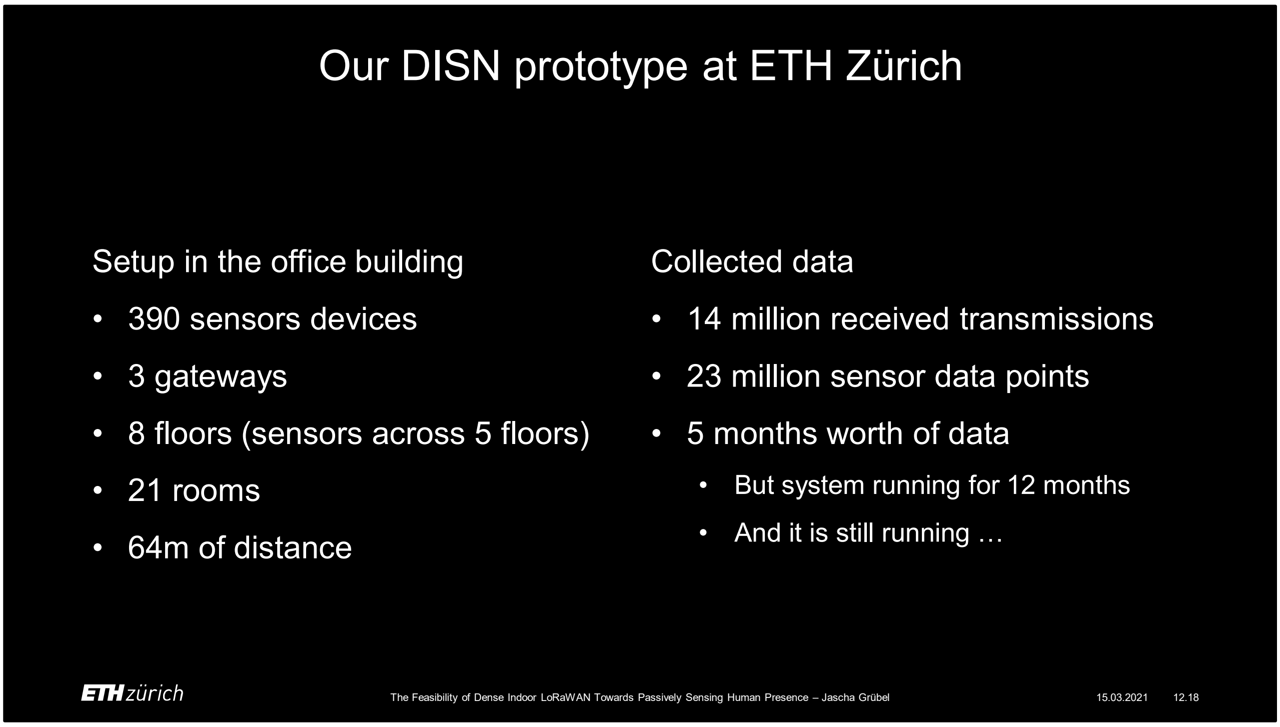

[C] Let me start by giving you an overview of the key data of our system.
[C] We setup our DISN in an office building at ETH Zürch.
[C] We installed 390 sensors and
[C] 3 gateways across
[C] 8 floors and
[C] 21 rooms.
[C] The maximal distance between sensors and gateways was 64 meters.
[C] Our collected data consists of 14 million received transmission across all gateways.
[C] And a total of 23 million sensor data points were transmitted.
[C] This constitutes 5 months worth of data.
[C] Because our goal is to study human presence, the system has been collecting data for 12 months now
[C] and is still running.

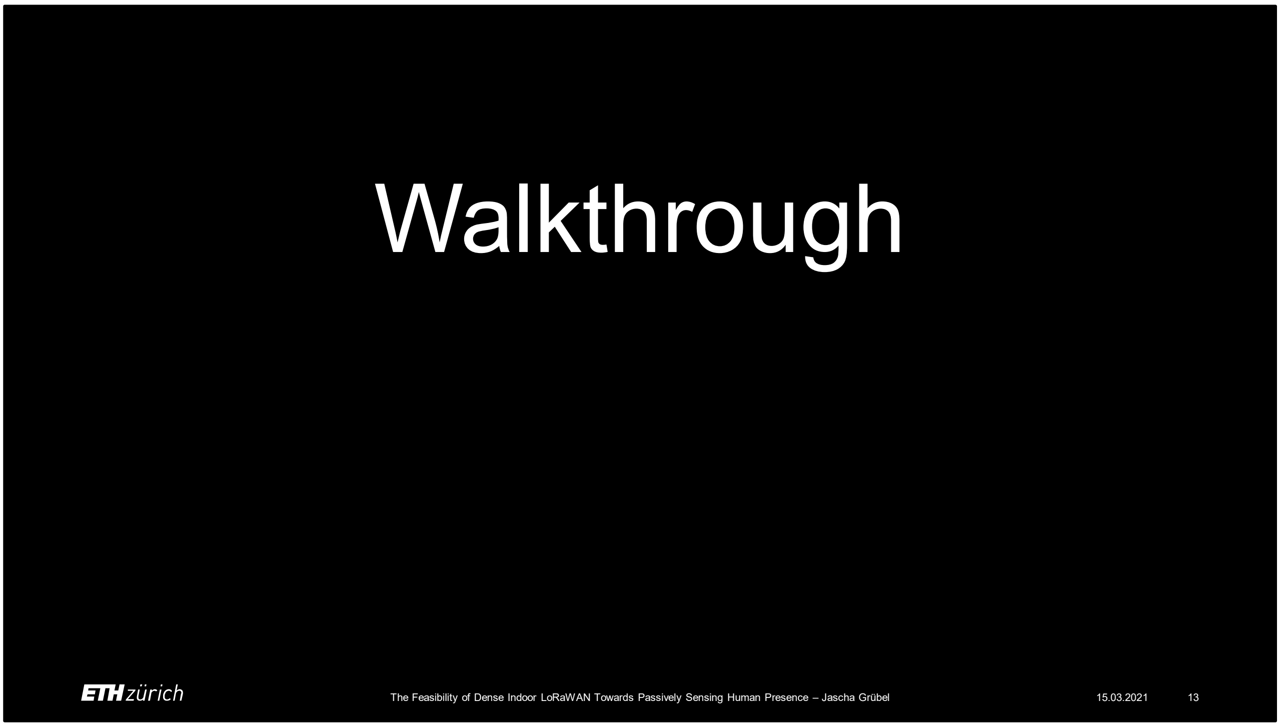

Let me first give a walkthrough the building, setup and implementation.

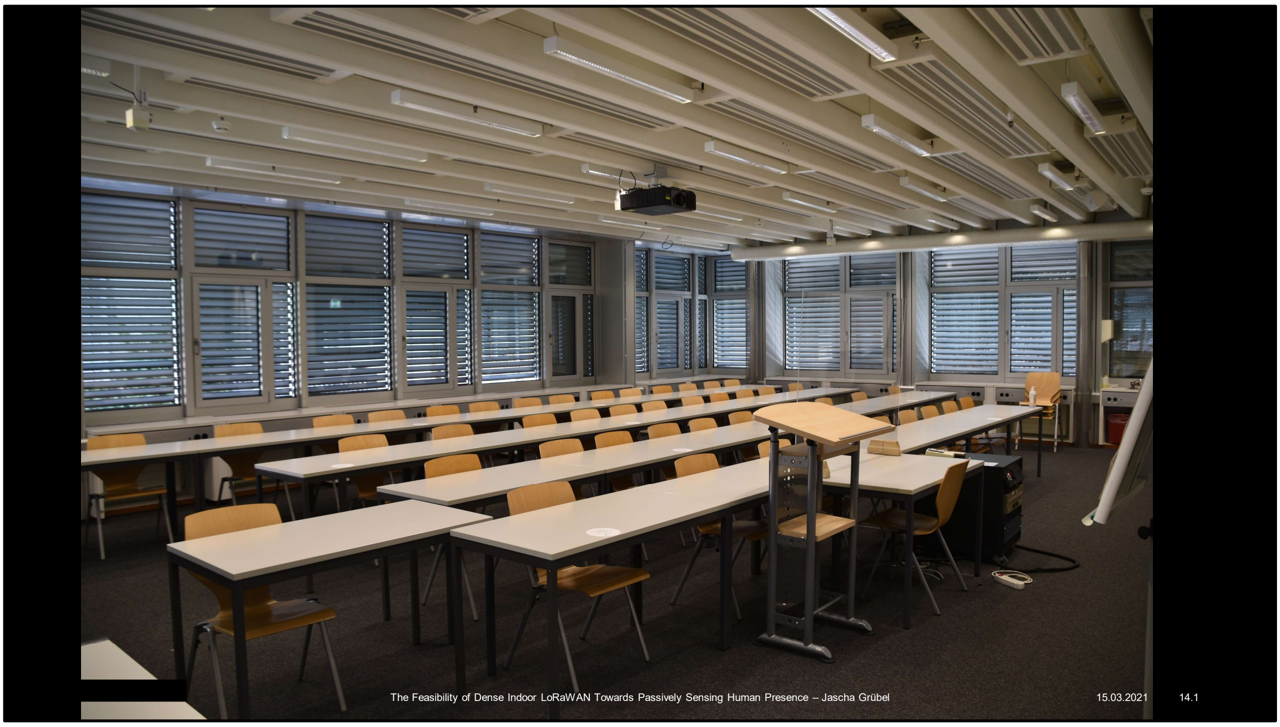

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence.
Previous research has connected
[C] levels of CO_2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

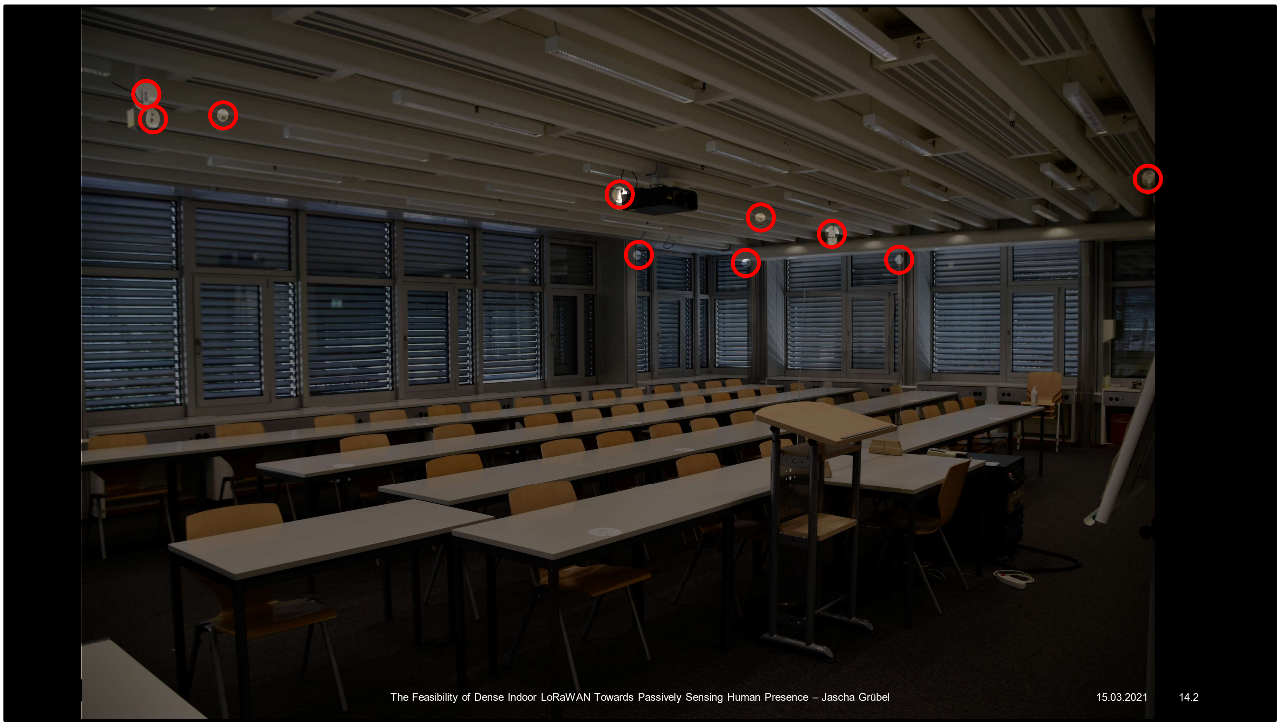

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence. Previous research has connected
[C] levels of CO_2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

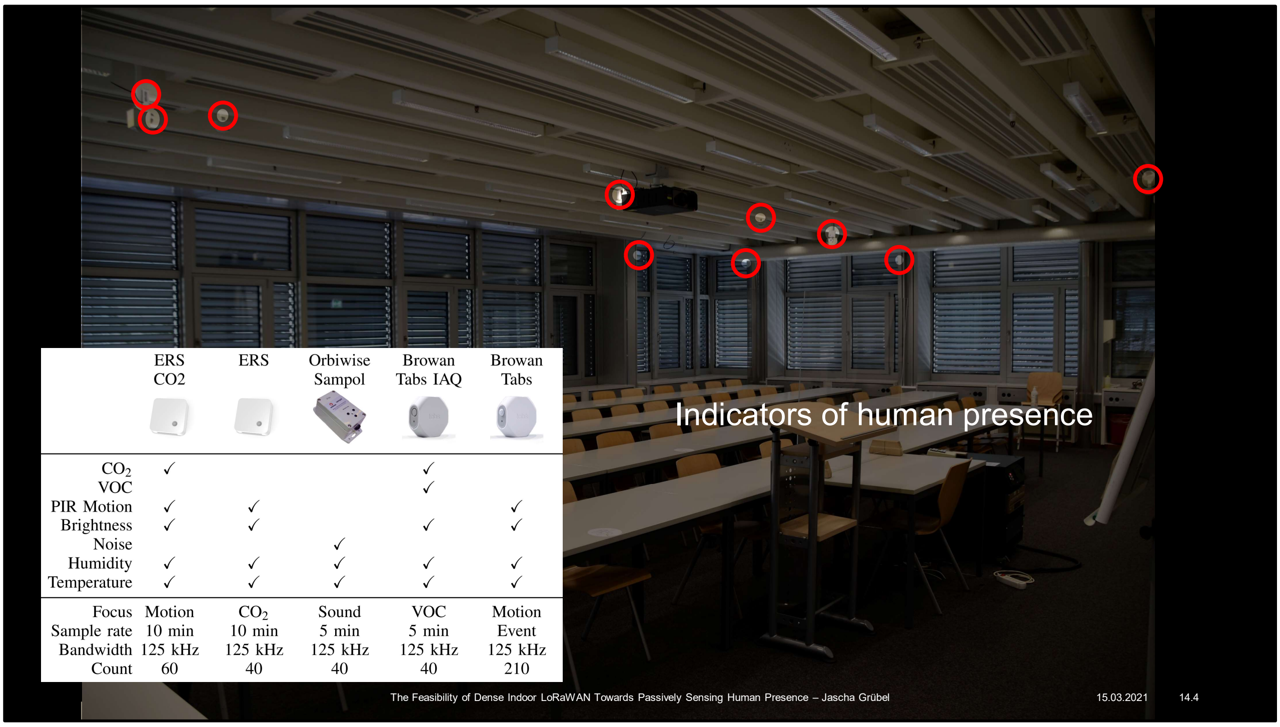

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence. Previous research has connected
[C] levels of CO2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

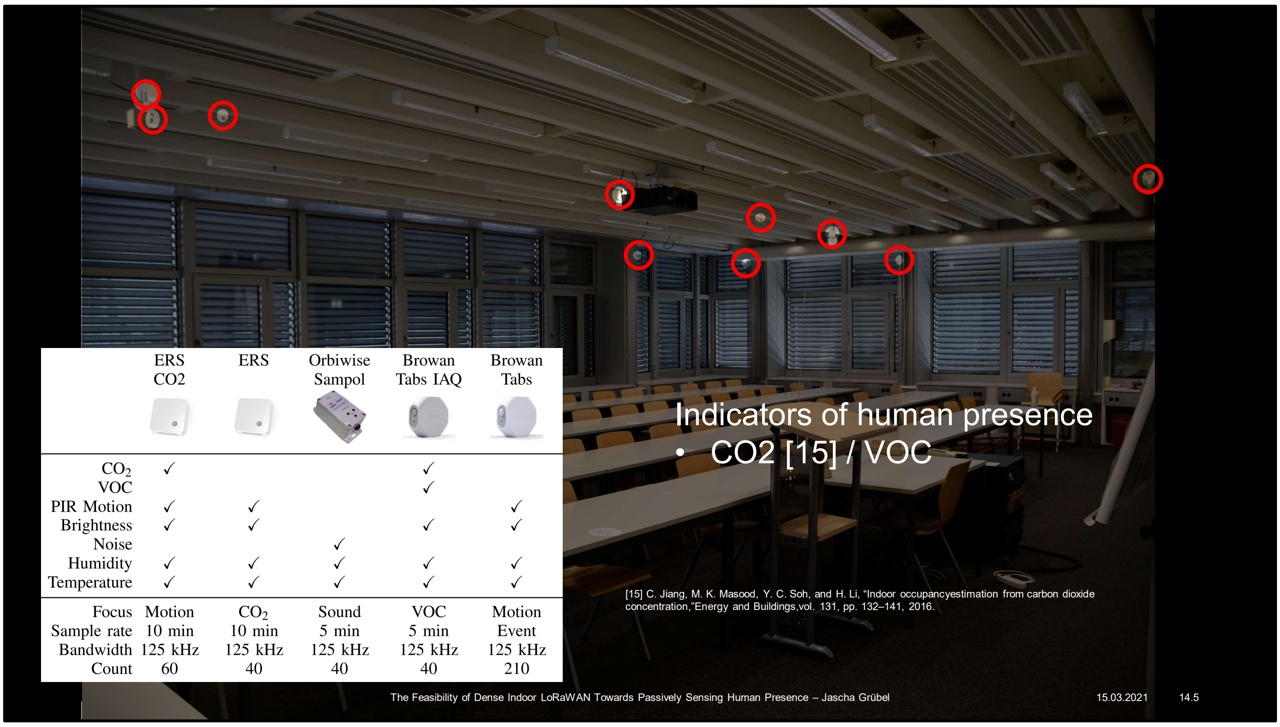

	ERS CO2	ERS	Orbiwise Sampol	Browan Tabs IAQ	Browan Tabs
CO_2	✓			✓	
VOC				✓	
PIR Motion	✓	✓			✓
Brightness	✓	✓		✓	✓
Noise			✓		
Humidity	✓	✓	✓	✓	✓
Temperature	✓	✓	✓	✓	✓
Focus	Motion	CO_2	Sound	VOC	Motion
Sample rate	10 min	10 min	5 min	5 min	Event
Bandwidth	125 kHz	125 kHz	125 kHz	125 kHz	125 kHz
Count	60	40	40	40	210

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence. Previous research has connected
[C] levels of CO2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

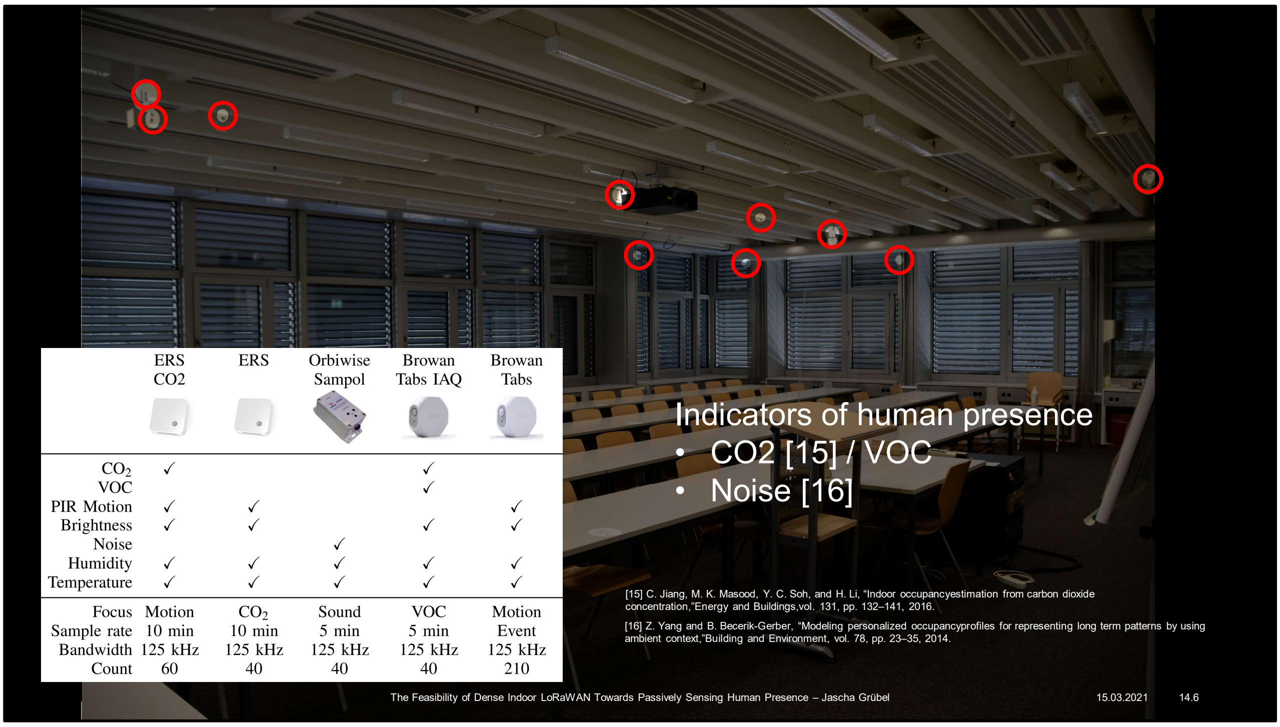

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence. Previous research has connected
[C] levels of CO2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

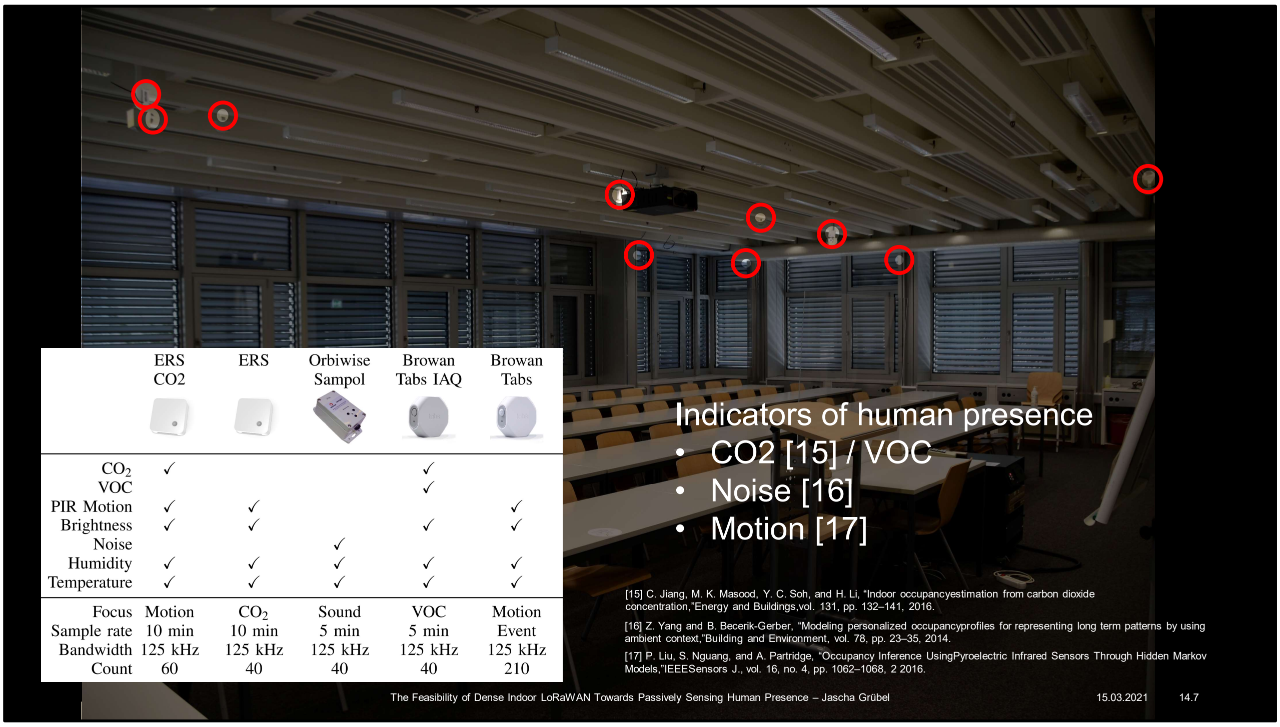

	ERS CO2	ERS	Orbiwise Sampol	Browan Tabs IAQ	Browan Tabs
CO_2	✓			✓	
VOC				✓	
PIR Motion	✓	✓			✓
Brightness	✓	✓		✓	✓
Noise			✓		
Humidity	✓	✓	✓	✓	✓
Temperature	✓	✓	✓	✓	✓
Focus	Motion	CO_2	Sound	VOC	Motion
Sample rate	10 min	10 min	5 min	5 min	Event
Bandwidth	125 kHz	125 kHz	125 kHz	125 kHz	125 kHz
Count	60	40	40	40	210

We start in a lecture hall in the building.
[C] 5 types of LoRaWAN sensors are placed out of sight in all public spaces of the building.
[C] We chose the types based on their ability to indicate human presence. Previous research has connected
[C] levels of CO2,
[C] levels of Noise, and
[C] Motion detection in particular to human presence.
[C] We transition to virtual representation of the lecture hall in the digital twin.
[NS]

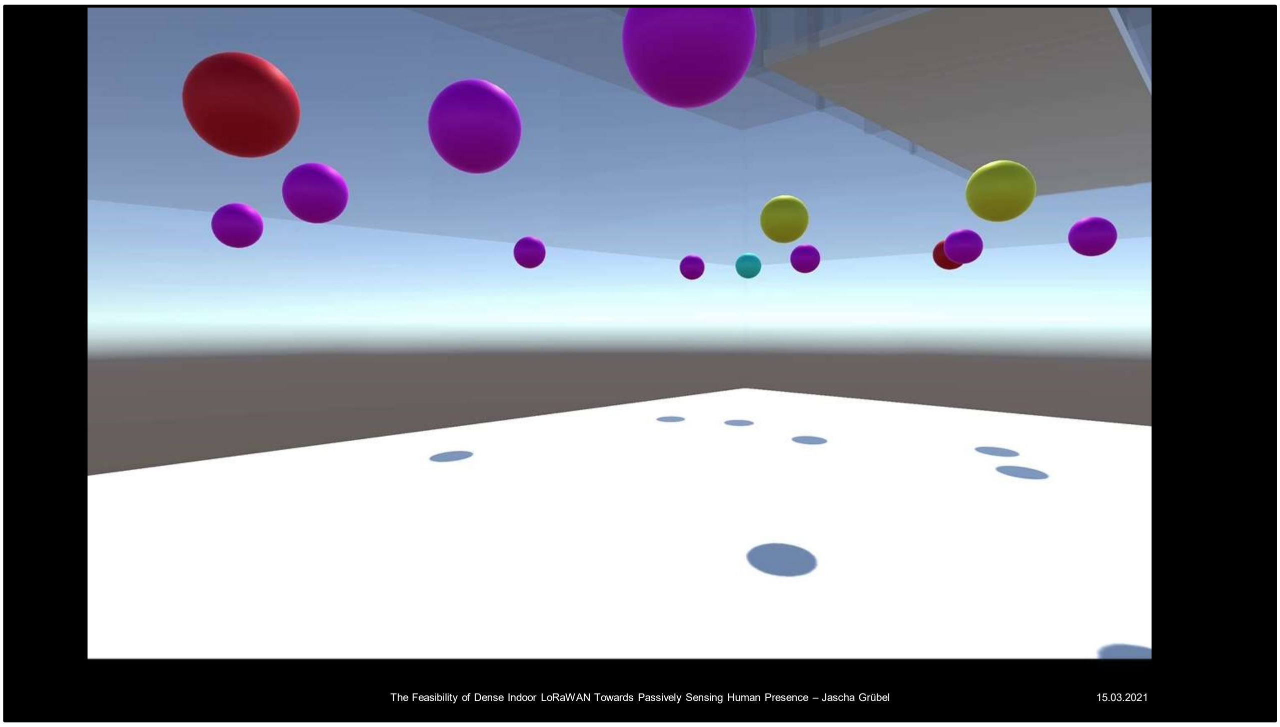

The digital twin is employed to manage the data and visualize the 390 sensors.
Please note particularly the green nodes representing gateways. One is now appearing from the left in the non-public part of the building 3 floors below.
We can also see the 5 sensor types color-coded in cyan, yellow, blue, purple, and red.

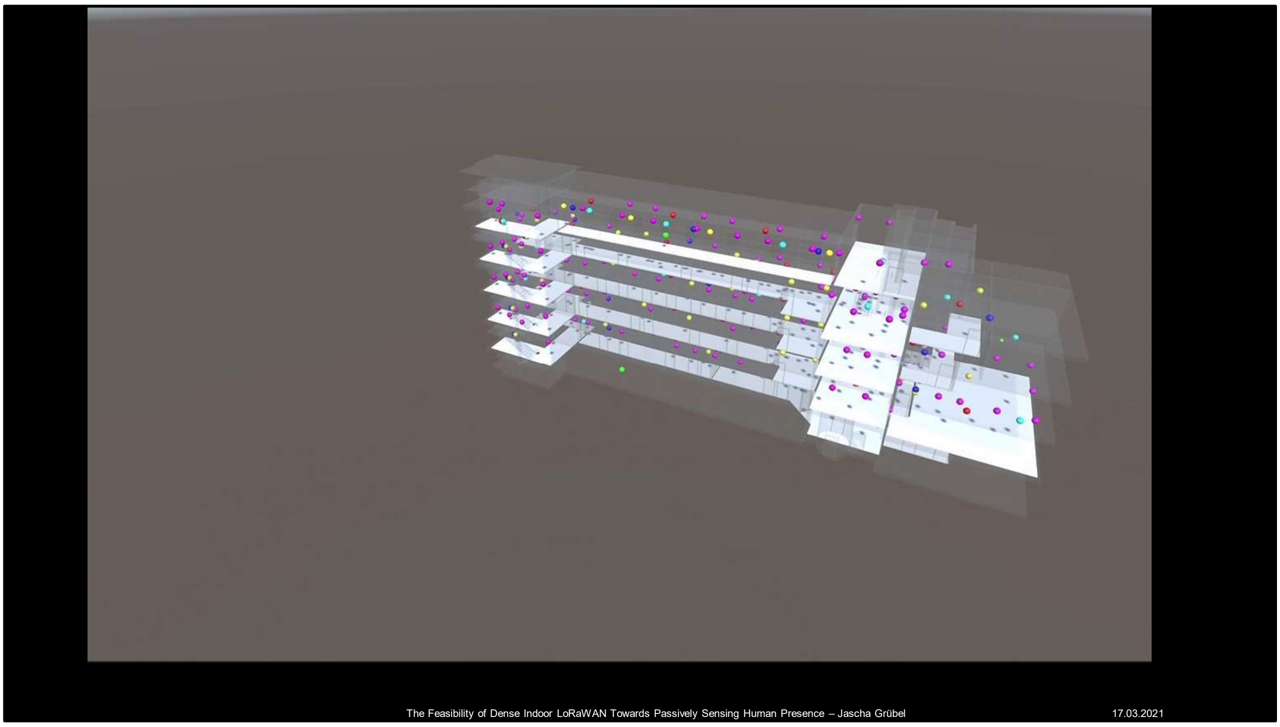

The digital twin is employed to manage the data and visualize the 390 sensors.
Please note particularly the green nodes representing gateways. One is now appearing from the left in the non-public part of the building 3 floors below.
We can also see the 5 sensor types color-coded in cyan, yellow, blue, purple, and red.

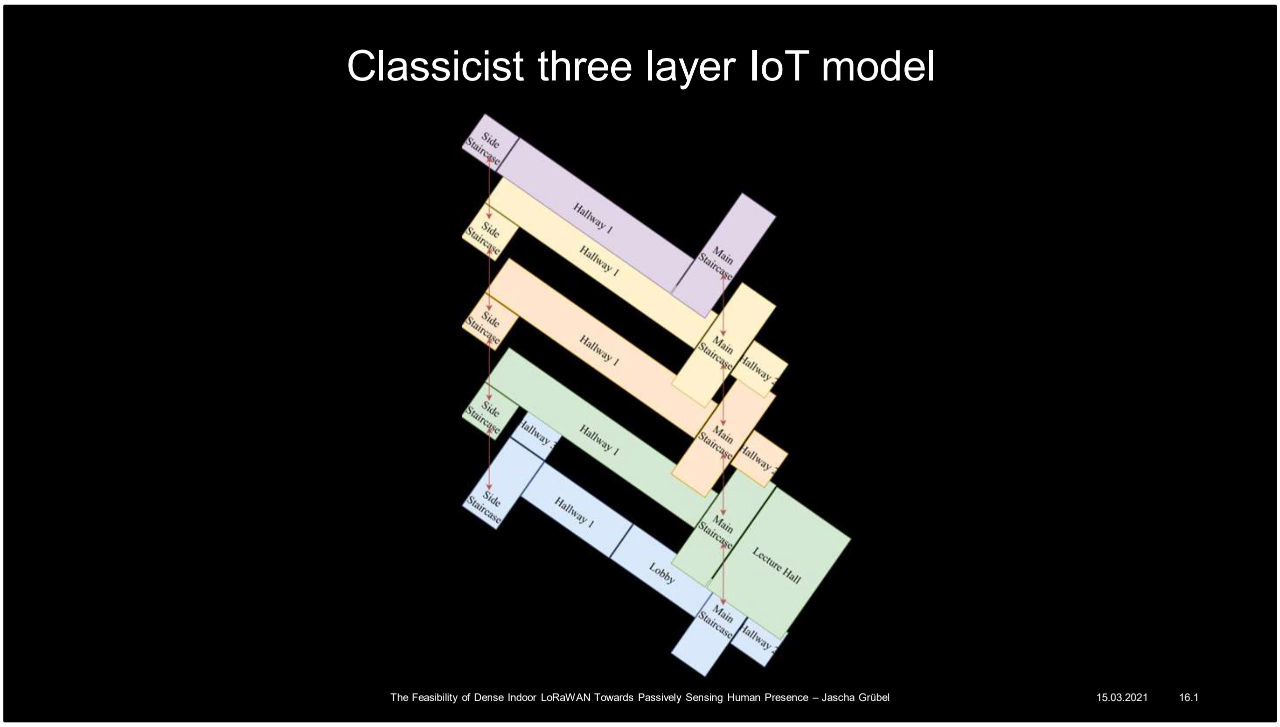

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

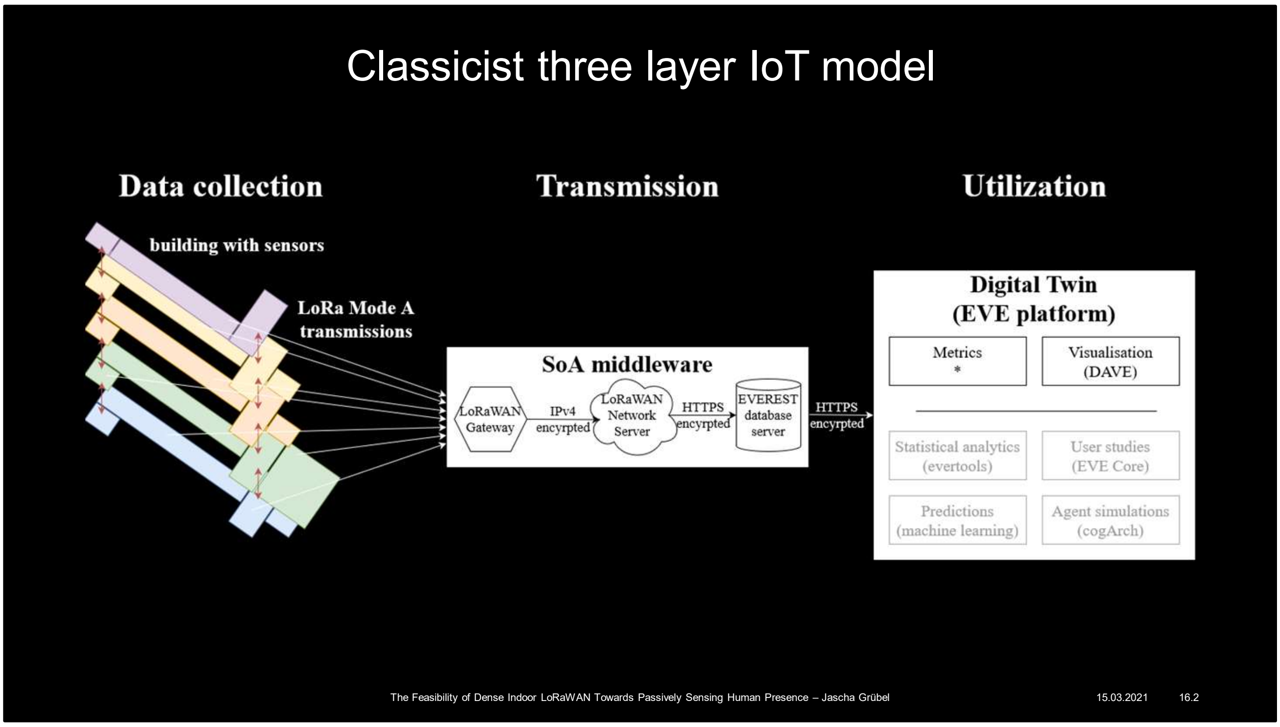

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

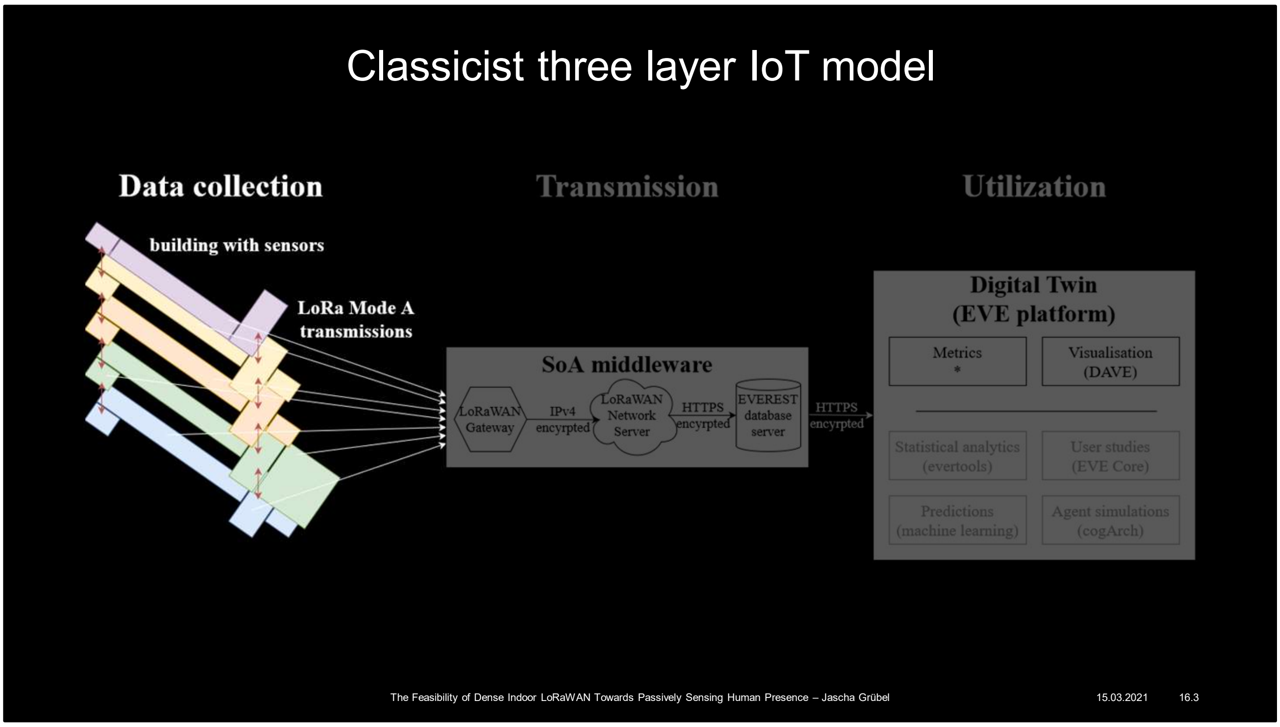

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

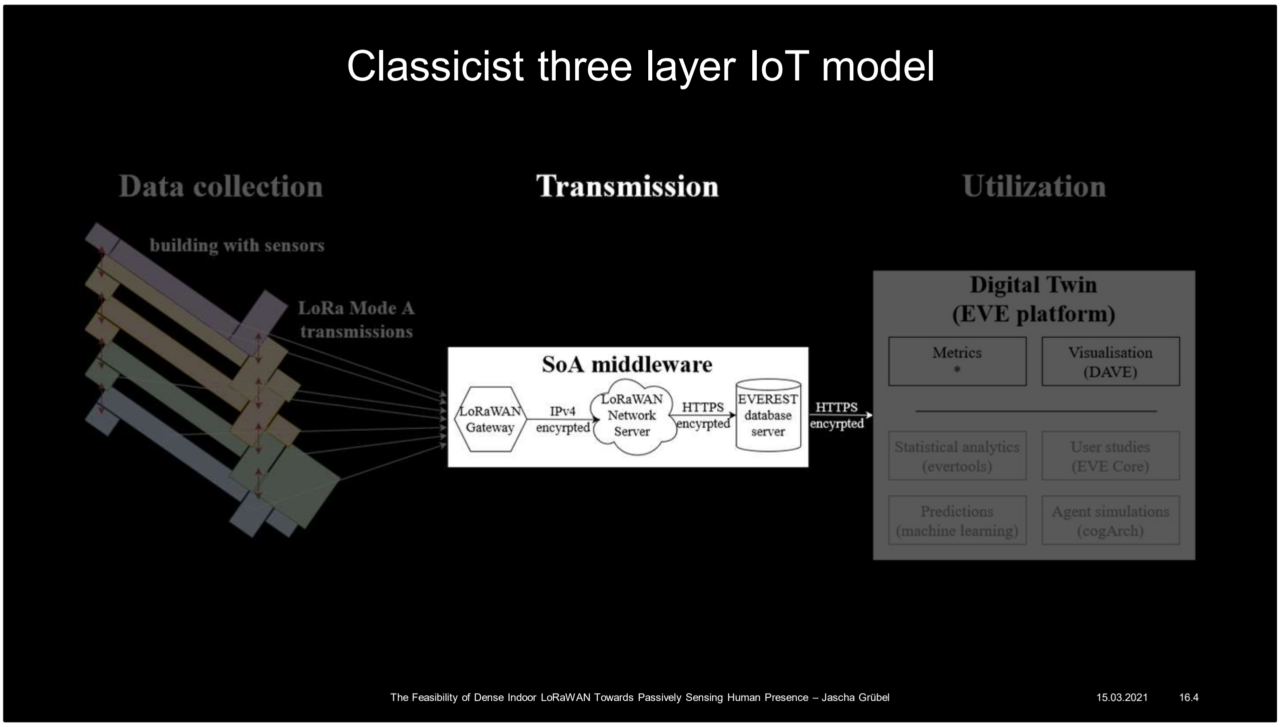

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

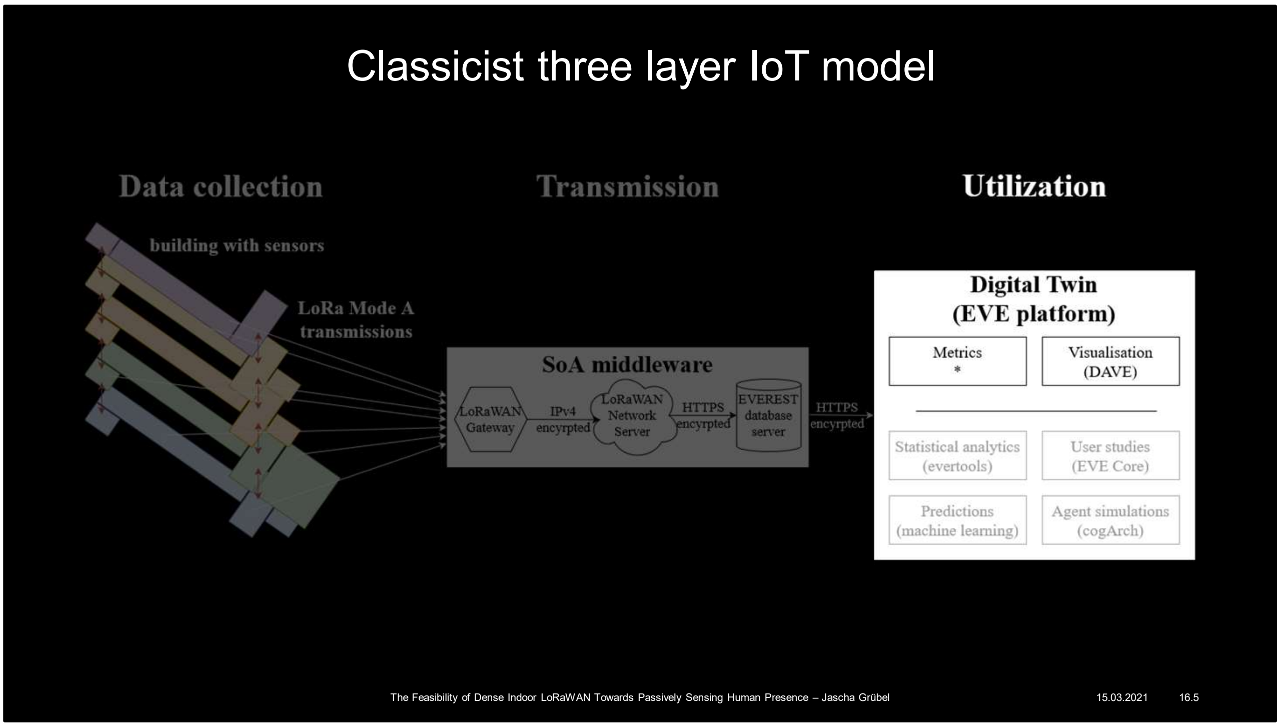

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

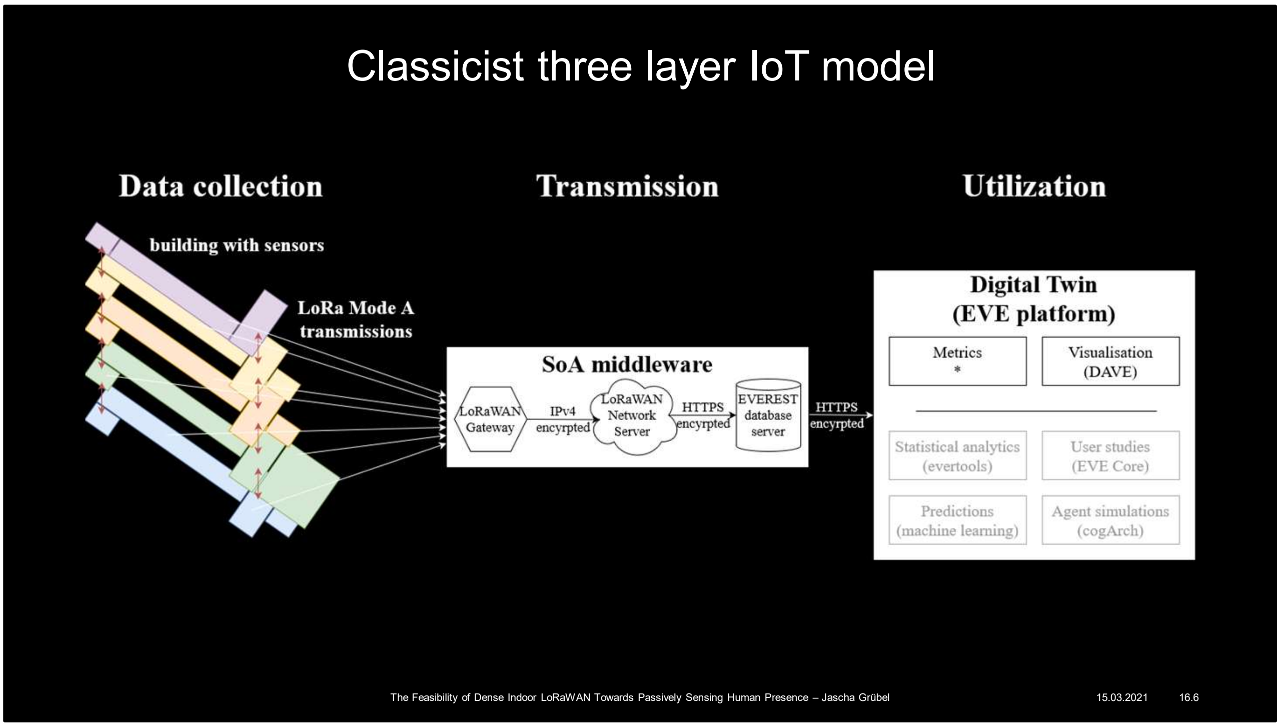

We transition to this schematic view of the building where we show the 21 rooms across the 5 floors involved in our setup.
[C] To structure our system, a classicist three layer IoT model is applied.
[C] The data collection consists of a sensor network in the building.
All 390 sensors use battery-powered LoRa class A and are evenly distributed across floor space.
[C] The transmission is managed by a service-oriented architecture middleware.
It consists of 3 independent modular components.
First, the LoRaWAN gateways receive the sensor data and forwards it to the LoRaWAN network server for meta-data augmentation, decoding and decryption.
Then our own database server receives the sensor data and stores them in a PostgreSQL database with extensions to manage spatial and temporal data, PostGIS and TimescaleDB, respectively.
Each component is independent and could be replaced with another implementation.
[C] The end user utilization is enabled through the digital twin, we saw a preview just now in the fly-through.

We base the digital twin on a development version of our open-source platform for virtual reality data collection and analysis EVE.
Before we move on, let me draw the attention to two particular aspects of the implementation in the SoA Middleware.

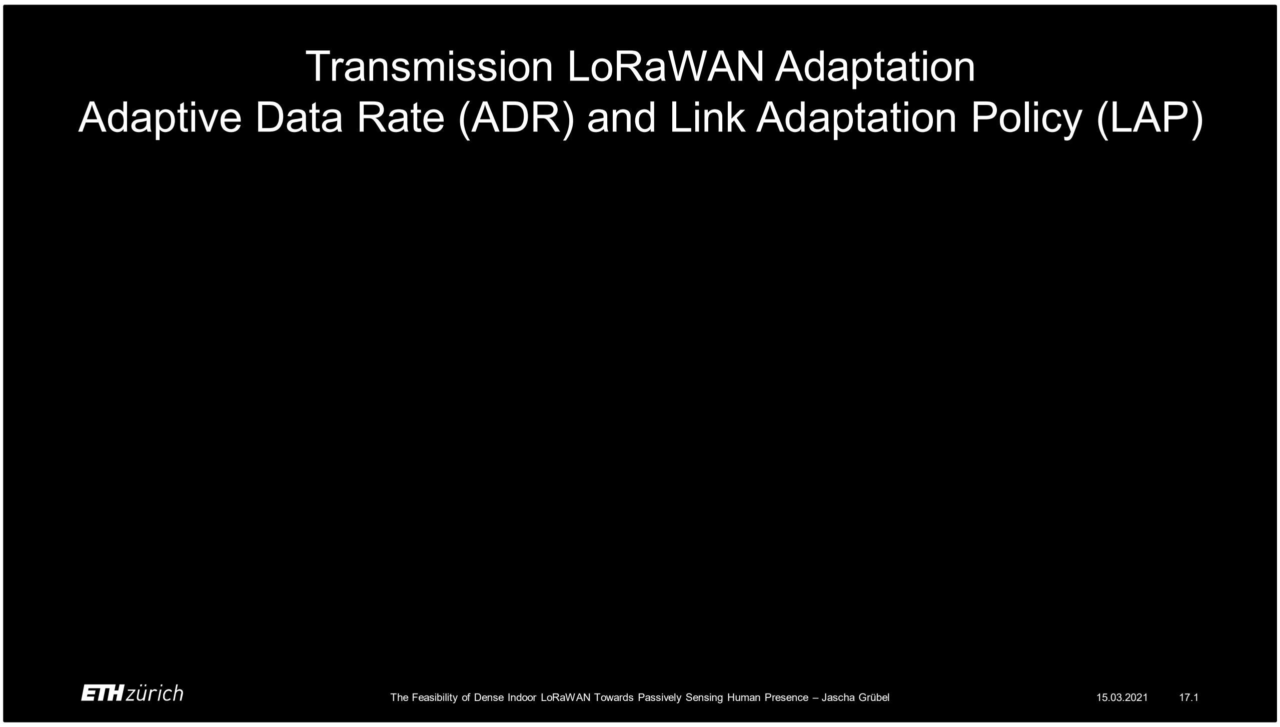

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

Transmission LoRaWAN Adaptation
Adaptive Data Rate (ADR) and Link Adaptation Policy (LAP)

Two-step optimization on network server:

ETH zürich

The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence – Jascha Grübel

15.03.2021 17.2

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

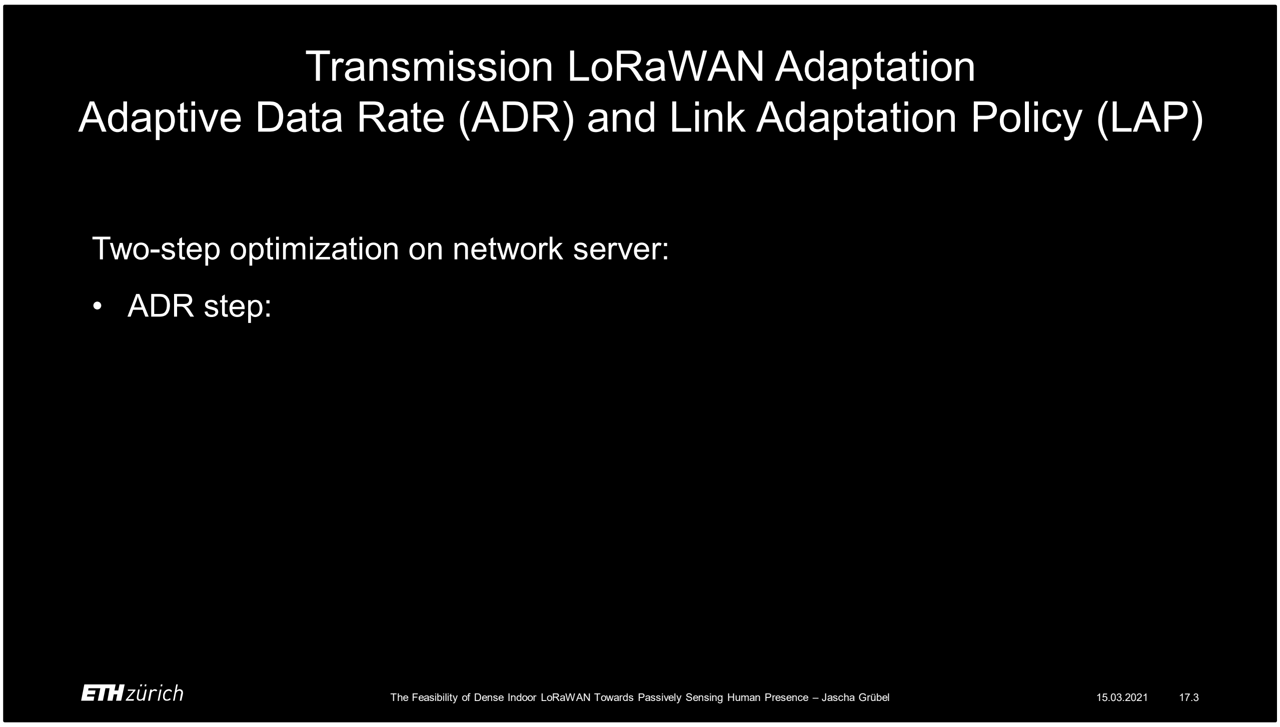

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

Transmission LoRaWAN Adaptation
Adaptive Data Rate (ADR) and Link Adaptation Policy (LAP)

Two-step optimization on network server:

- ADR step:

ETH zürich The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence – Jascha Grübel 15.03.2021 17.4

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

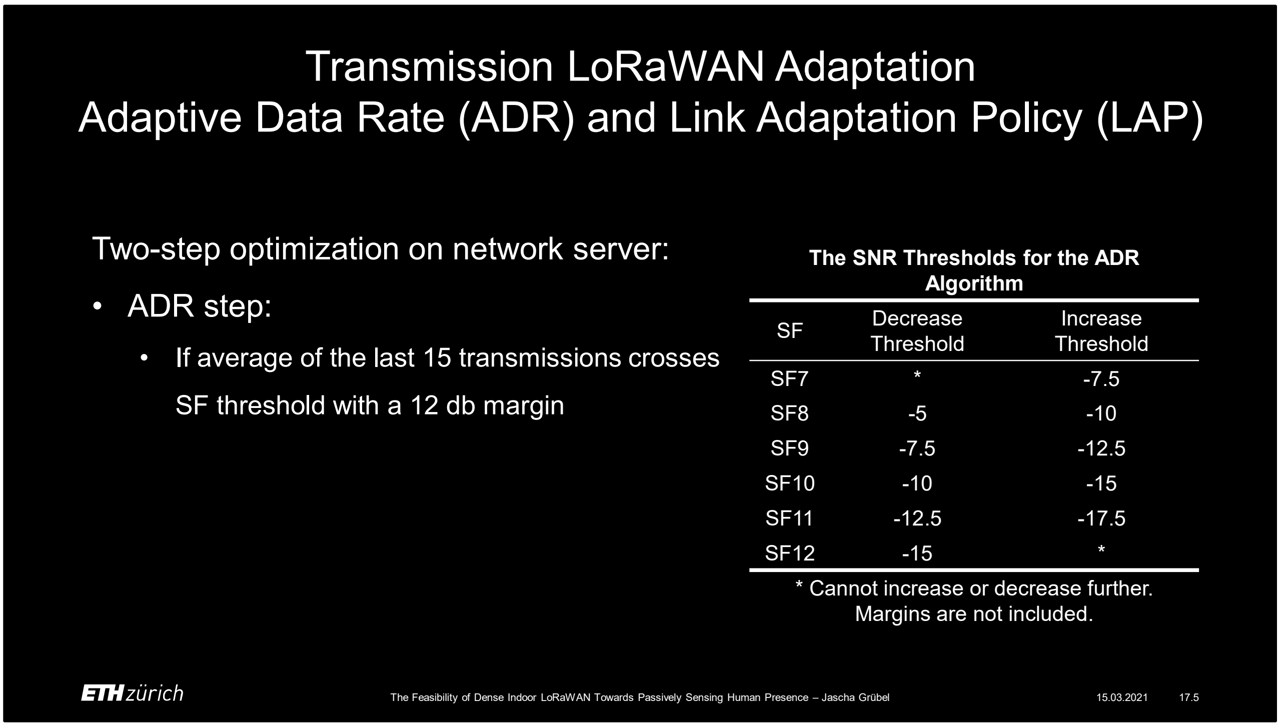

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

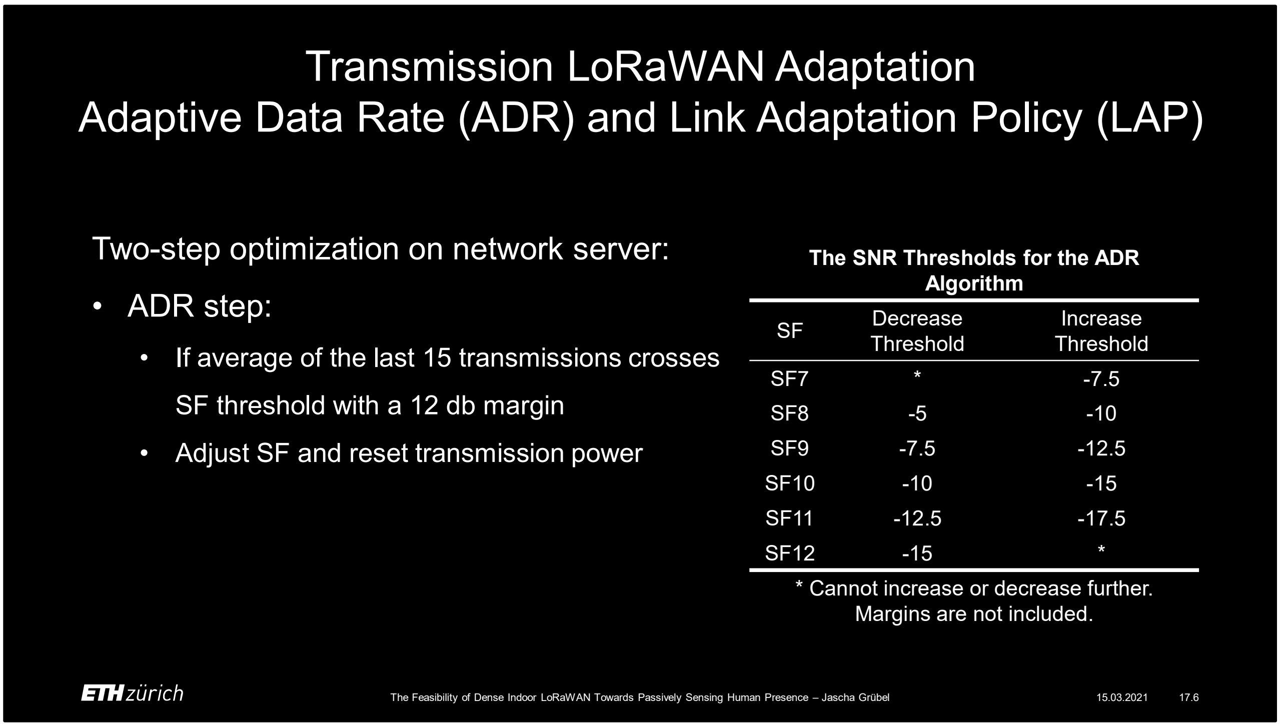

SF	Decrease Threshold	Increase Threshold
SF7	*	-7.5
SF8	-5	-10
SF9	-7.5	-12.5
SF10	-10	-15
SF11	-12.5	-17.5
SF12	-15	*

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

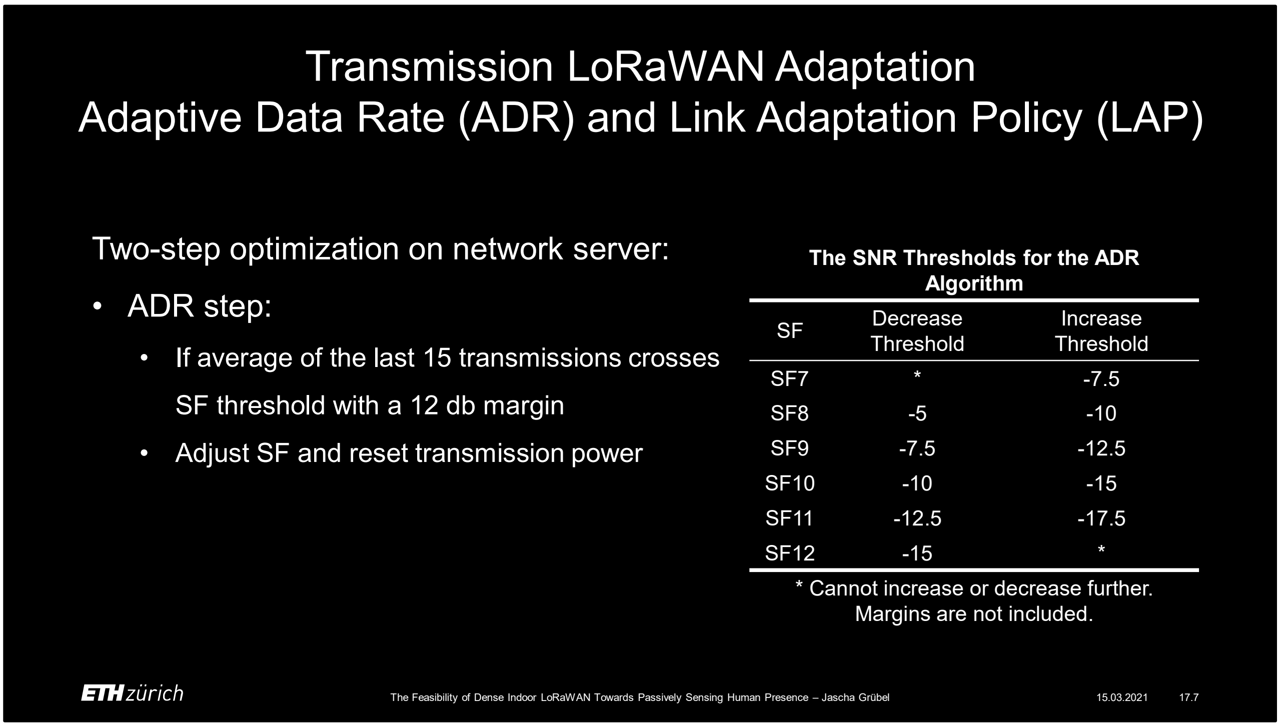

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

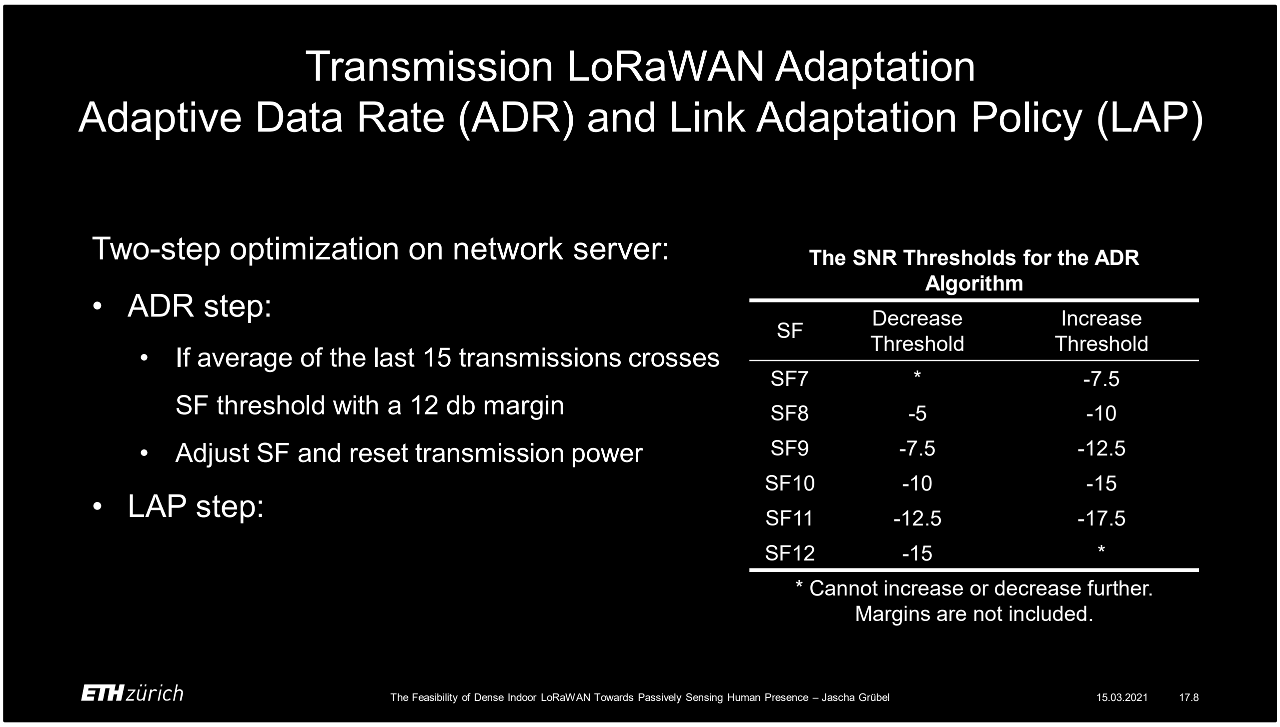

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

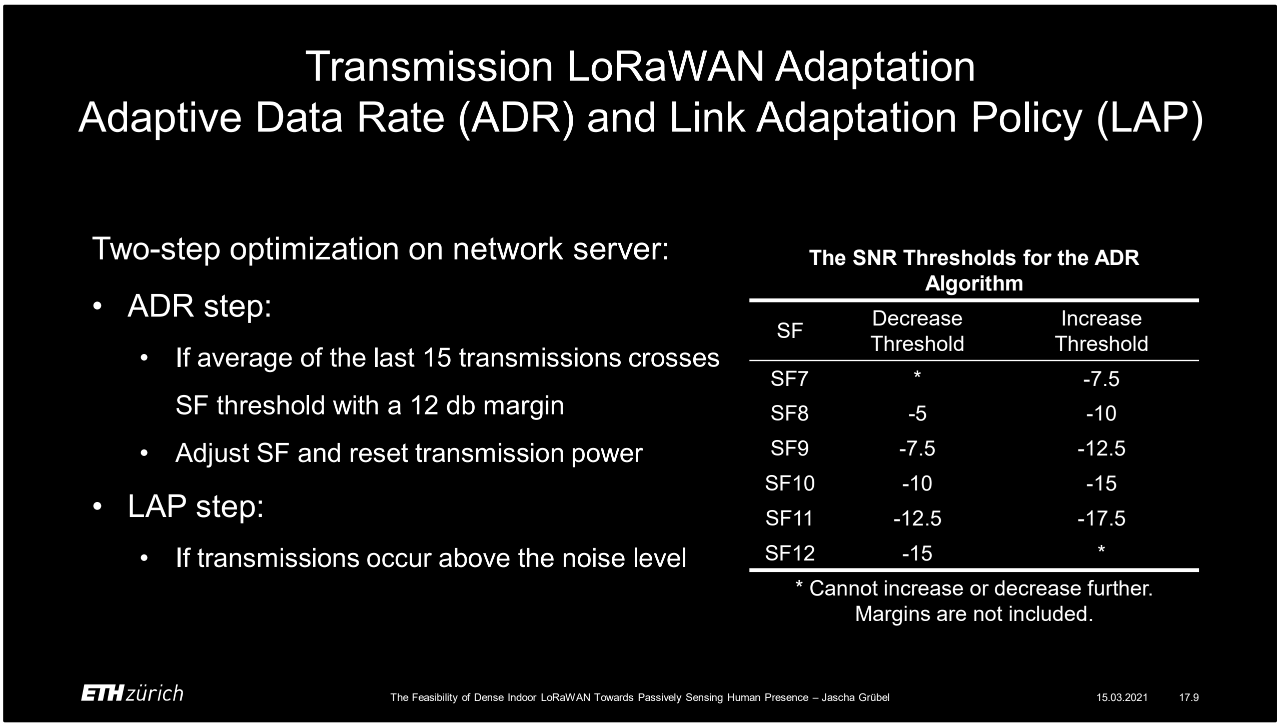

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

Transmission LoRaWAN Adaptation
Adaptive Data Rate (ADR) and Link Adaptation Policy (LAP)

Two-step optimization on network server:

- ADR step:
  - If average of the last 15 transmissions crosses SF threshold with a 12 db margin
  - Adjust SF and reset transmission power
- LAP step:
  - If transmissions occur above the noise level
  - Reduce transmission power

The SNR Thresholds for the ADR Algorithm

SF	Decrease Threshold	Increase Threshold
SF7	*	-7.5
SF8	-5	-10
SF9	-7.5	-12.5
SF10	-10	-15
SF11	-12.5	-17.5
SF12	-15	*

* Cannot increase or decrease further. Margins are not included.

The first implementation details is about conserving energy with so many devices.
We use a variation of Adapative Data Rate (or ADR) and Link Adaptation Policy (or LAP).
[C] The network server performs a two-step optimization.
[C] In the ADR step,
[C] if the average of the last 15 transmissions crosses the SF threshold with a 12db margin,
[C] the SF is adjusted and the transmission power reset to maximal power.
[C] In the LAP step,
[C] as long as the transmission occurs above the noise level,
[C] transmission power on the current SF is reduced.
In practise, the second step mostly happens on SF7 as there is too little room for maneouver on higher SF

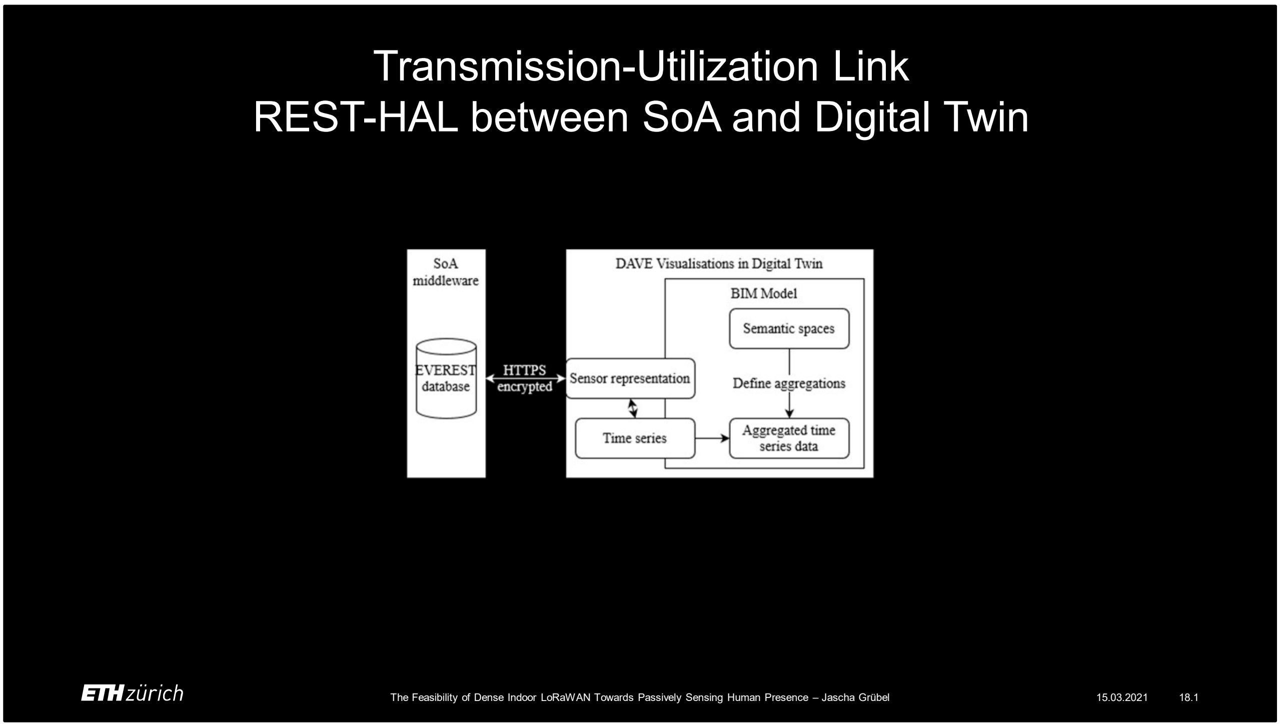

The second implementation detail is how to make utilization independent from the SoA middleware implementation.
[C] We employ Representational State Transfer Hypertext Application Language (or REST-HAL) for communication between the database server in the SoA and Digital Twin.
This expansion of REST adds relational links and embeddings between entries.
[C] For example, here we see the meta information on sensor 1.
[C] The type field is a reference to another REST entry.
Instead of having to retrieve it manually,
[C] the embedding provides it below.
The content on the database can thus be machine-readable explored.
New sensor types can easily be added and used with without having to expand the codebase elsewhere.

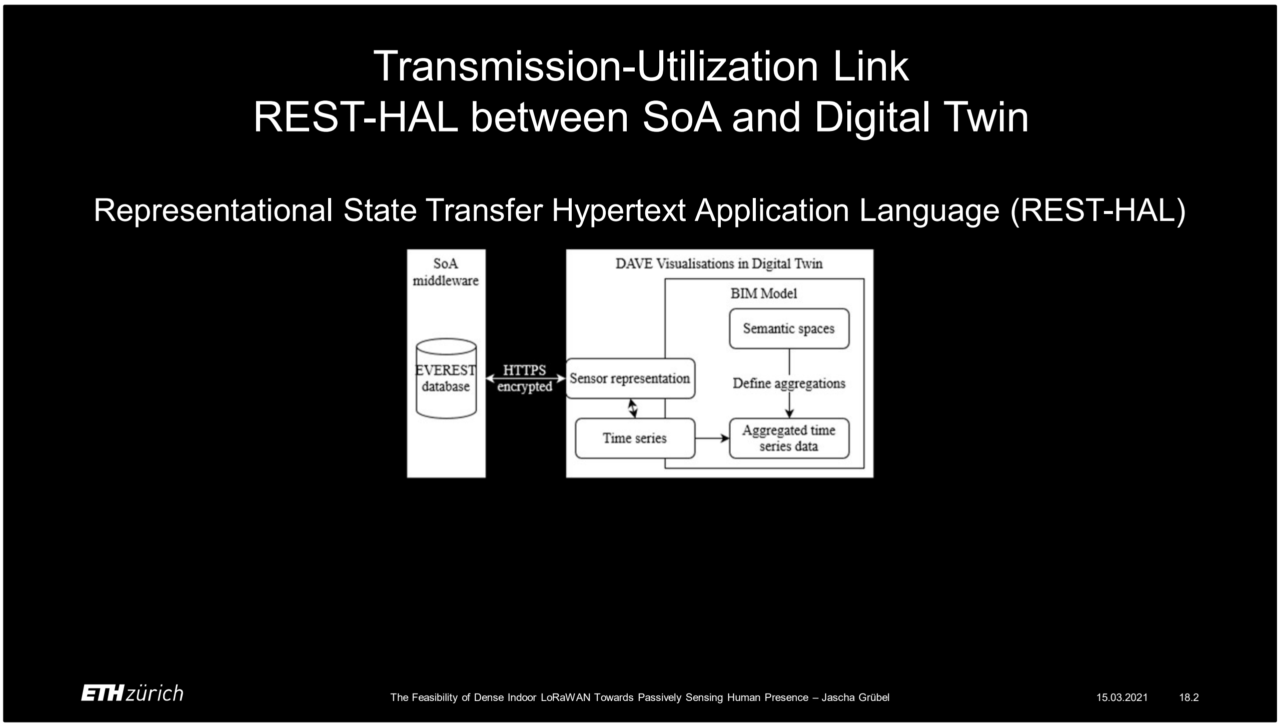

The second implementation detail is how to make utilization independent from the SoA middleware implementation.
[C] We employ Representational State Transfer Hypertext Application Language (or REST-HAL) for communication between the database server in the SoA and Digital Twin.
This expansion of REST adds relational links and embeddings between entries.
[C] For example, here we see the meta information on sensor 1.
[C] The type field is a reference to another REST entry.
Instead of having to retrieve it manually,
[C] the embedding provides it below.
The content on the database can thus be machine-readable explored.
New sensor types can easily be added and used with without having to expand the codebase elsewhere.

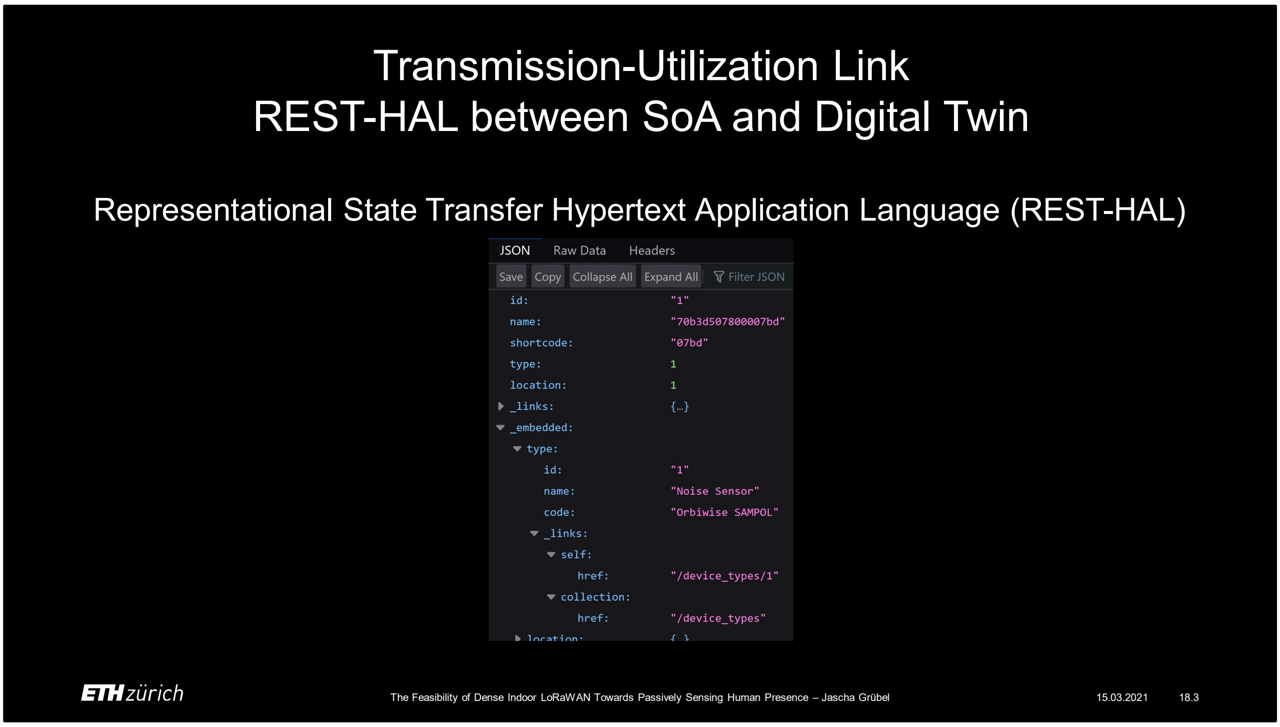

The second implementation detail is how to make utilization independent from the SoA middleware implementation.
[C] We employ Representational State Transfer Hypertext Application Language (or REST-HAL) for communication between the database server in the SoA and Digital Twin.
This expansion of REST adds relational links and embeddings between entries.
[C] For example, here we see the meta information on sensor 1.
[C] The type field is a reference to another REST entry.
Instead of having to retrieve it manually,
[C] the embedding provides it below.
The content on the database can thus be machine-readable explored.
New sensor types can easily be added and used with without having to expand the codebase elsewhere.

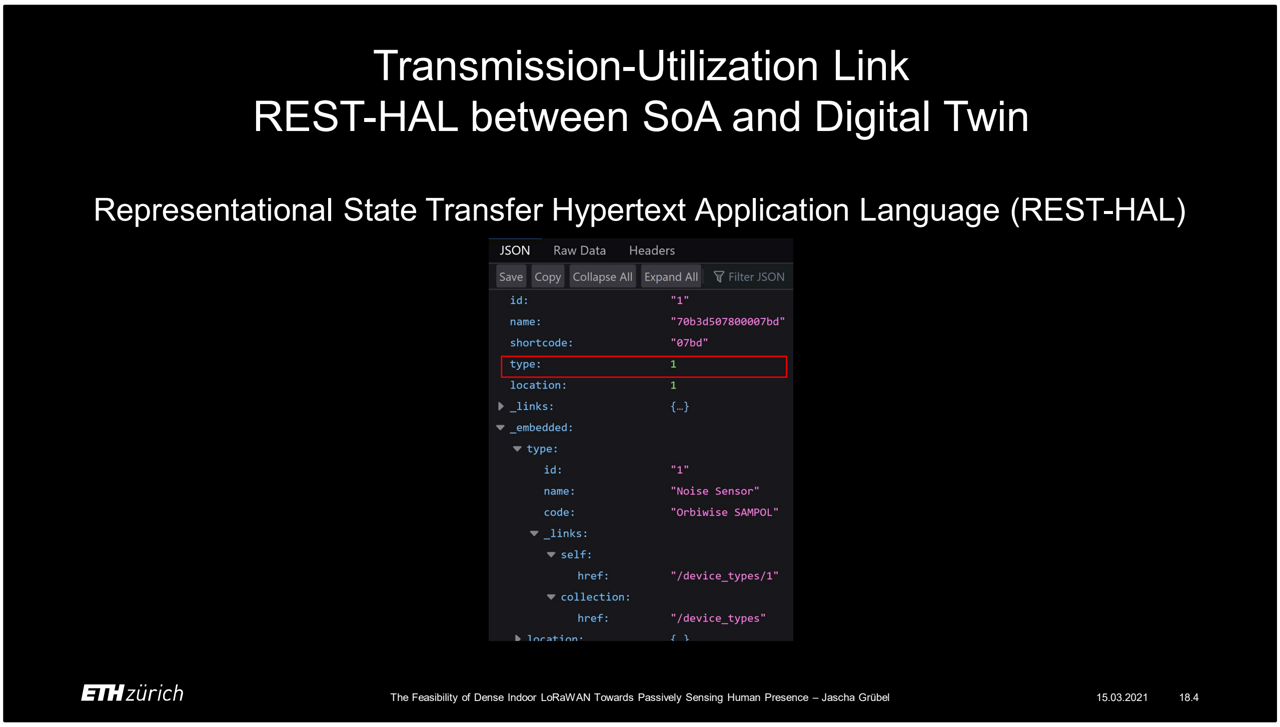

The second implementation detail is how to make utilization independent from the SoA middleware implementation.
[C] We employ Representational State Transfer Hypertext Application Language (or REST-HAL) for communication between the database server in the SoA and Digital Twin.
This expansion of REST adds relational links and embeddings between entries.
[C] For example, here we see the meta information on sensor 1.
[C] The type field is a reference to another REST entry.
Instead of having to retrieve it manually,
[C] the embedding provides it below.
The content on the database can thus be machine-readable explored.
New sensor types can easily be added and used with without having to expand the codebase elsewhere.

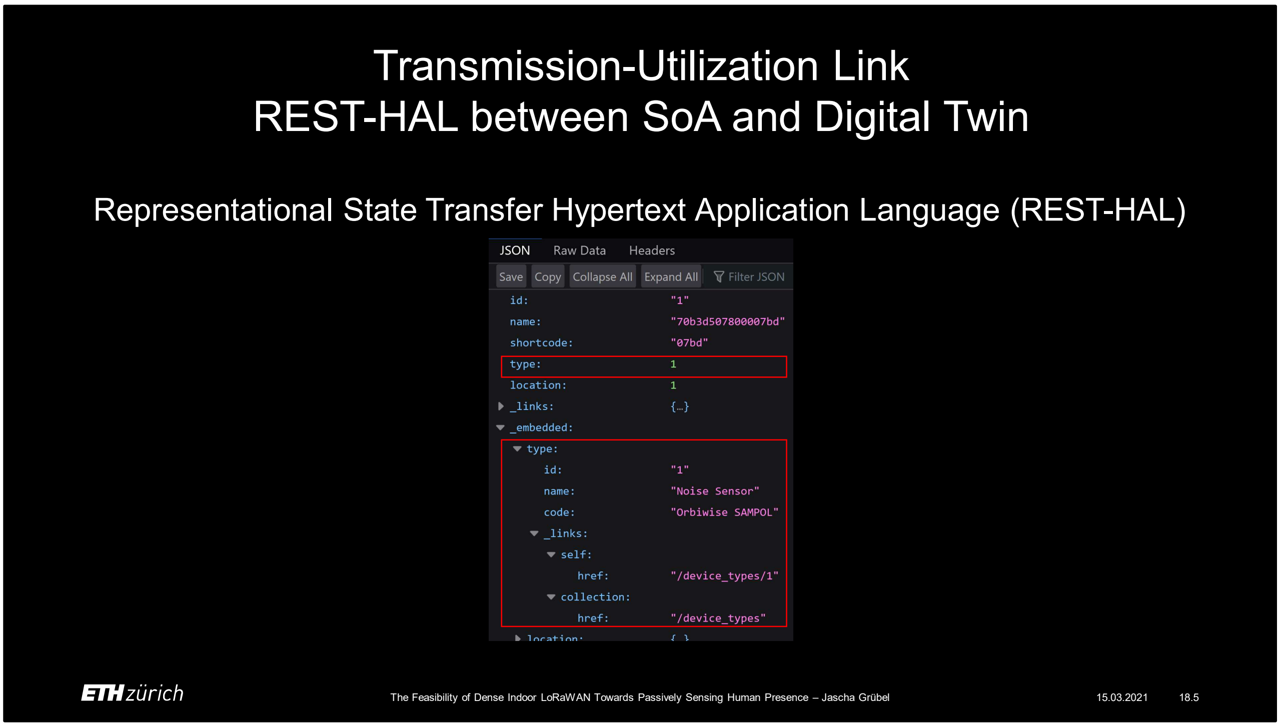

The second implementation detail is how to make utilization independent from the SoA middleware implementation.
[C] We employ Representational State Transfer Hypertext Application Language (or REST-HAL) for communication between the database server in the SoA and Digital Twin.
This expansion of REST adds relational links and embeddings between entries.
[C] For example, here we see the meta information on sensor 1.
[C] The type field is a reference to another REST entry.
Instead of having to retrieve it manually,
[C] the embedding provides it below.
The content on the database can thus be machine-readable explored.
New sensor types can easily be added and used with without having to expand the codebase elsewhere.

With a detailed overview of the system.
Let’s first look at the descriptives of the data to understand what we have at our hand.

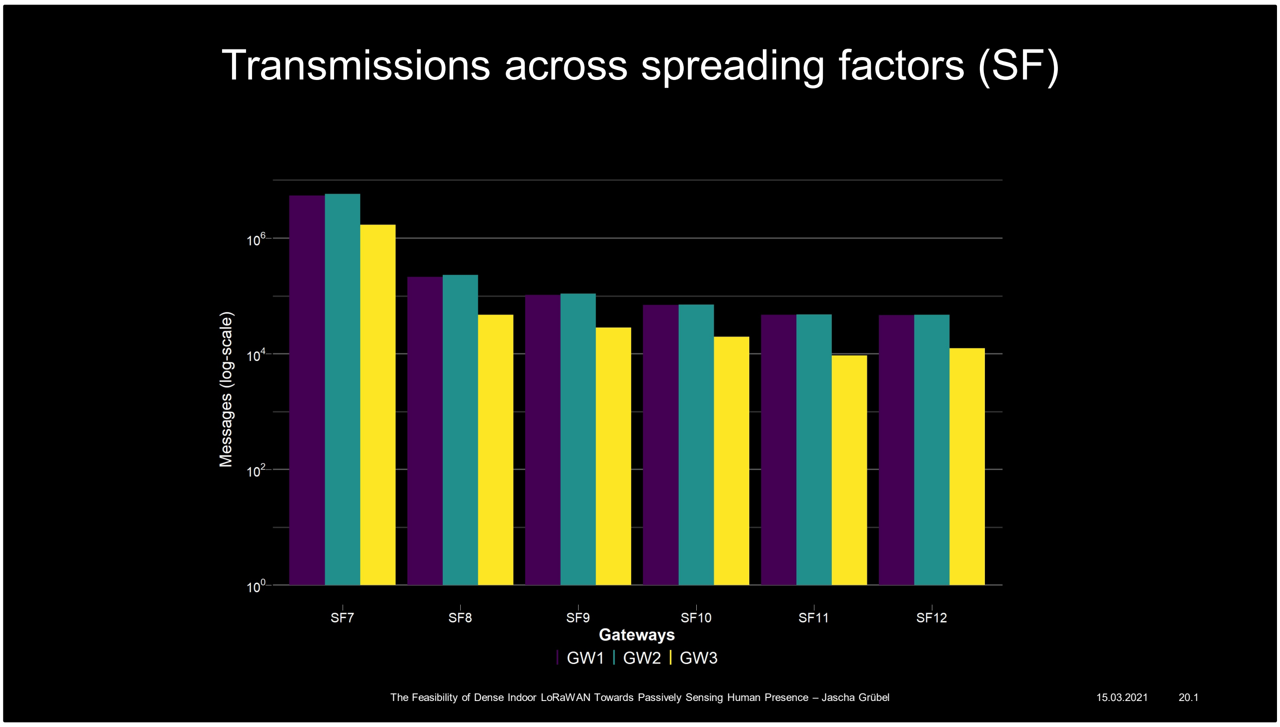

Let's start with understanding how our data was transmitted.
Overall, our 14 million transmissions are unevenly spread across spreading factor (SF).
[C] Each bar color represents a gateway and they are group by SF.
More than 90% of transmissions occur on SF7.
[C] A log scale as used to make the decrease in higher SF visible.
Given the short distances involved, SF7 would be the best choice for transmission.
We want to gain a deeper understanding why our system used higher SF.

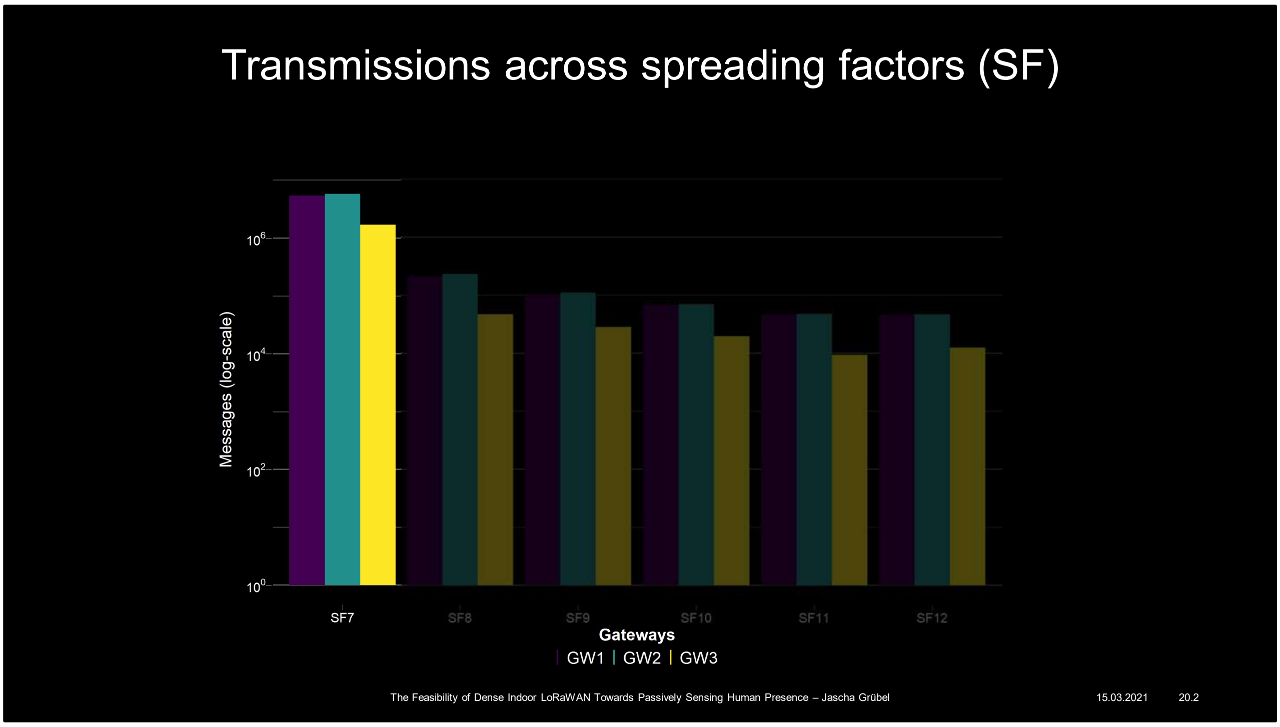

Let’s start with understanding how our data was transmitted.
Overall, our 14 million transmissions are unevenly spread across spreading factor (SF).
[C] Each bar color represents a gateway and they are group by SF.
More than 90% of transmissions occur on SF7.
[C] A log scale as used to make the decrease in higher SF visible.
Given the short distances involved, SF7 would be the best choice for transmission.
We want to gain a deeper understanding why our system used higher SF.

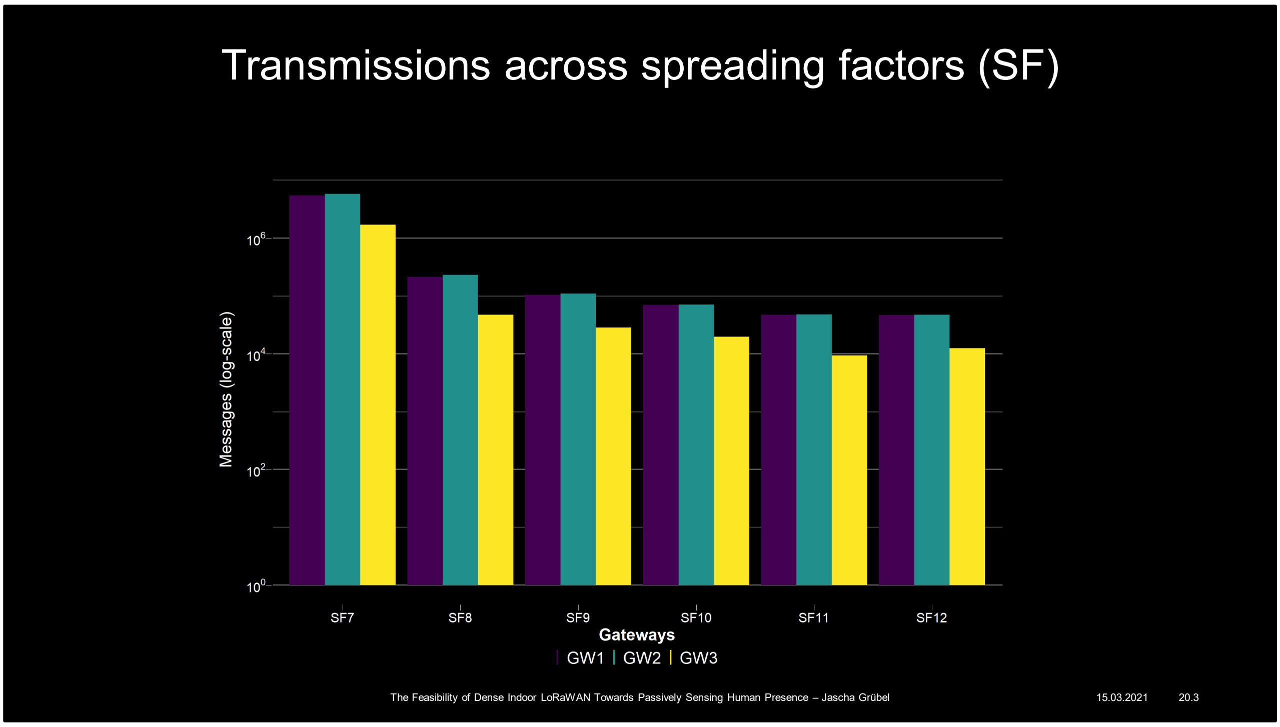

Let’s start with understanding how our data was transmitted.
Overall, our 14 million transmissions are unevenly spread across spreading factor (SF).
[C] Each bar color represents a gateway and they are group by SF.
More than 90% of transmissions occur on SF7.
[C] A log scale as used to make the decrease in higher SF visible.
Given the short distances involved, SF7 would be the best choice for transmission.
We want to gain a deeper understanding why our system used higher SF.

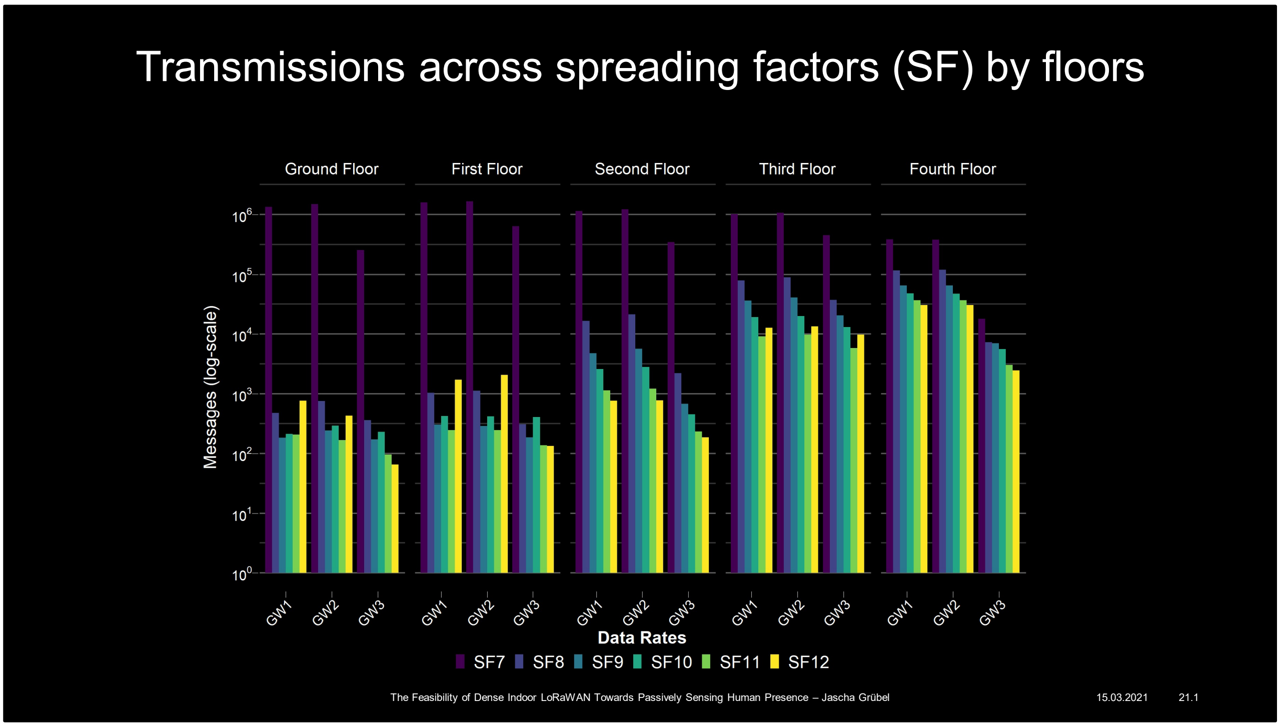

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

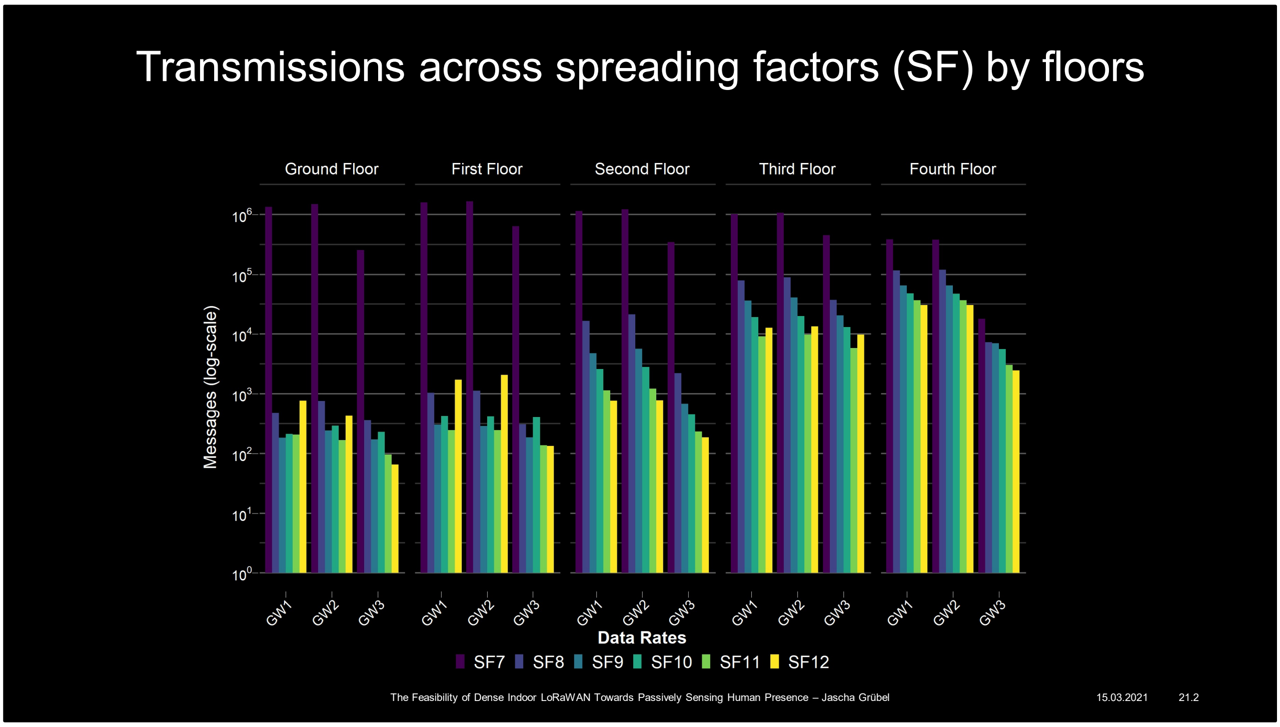

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

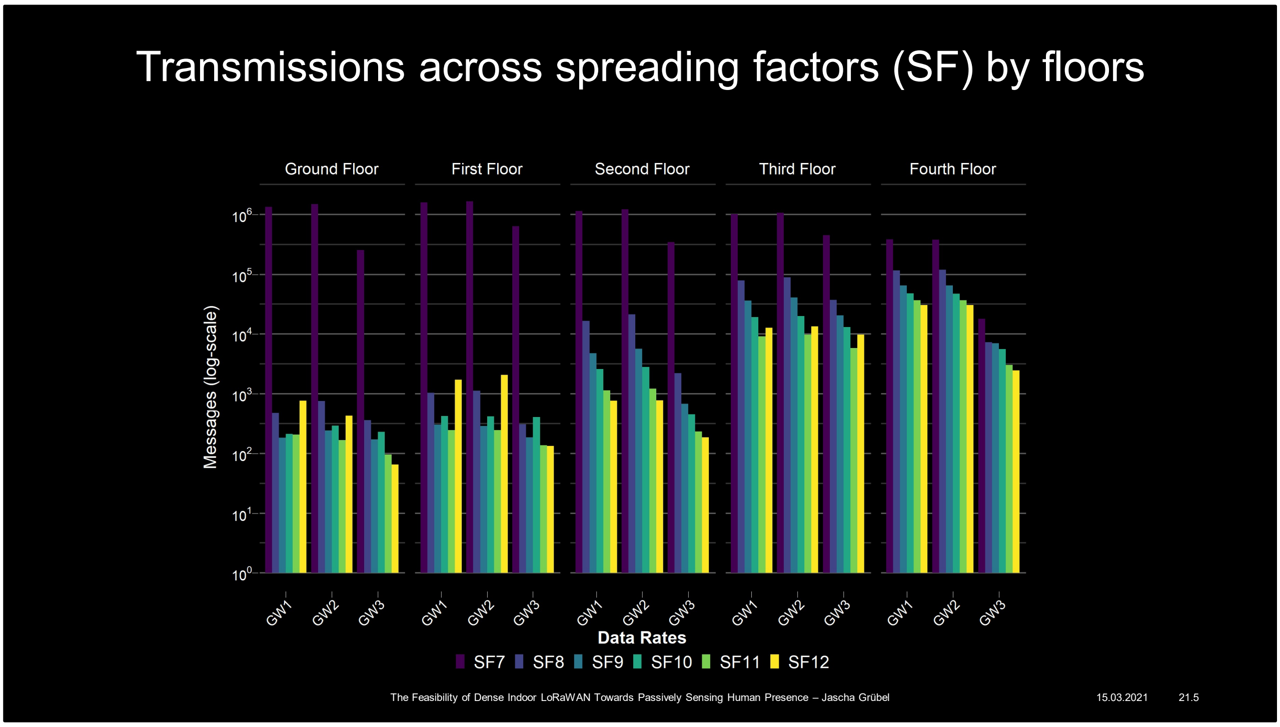

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

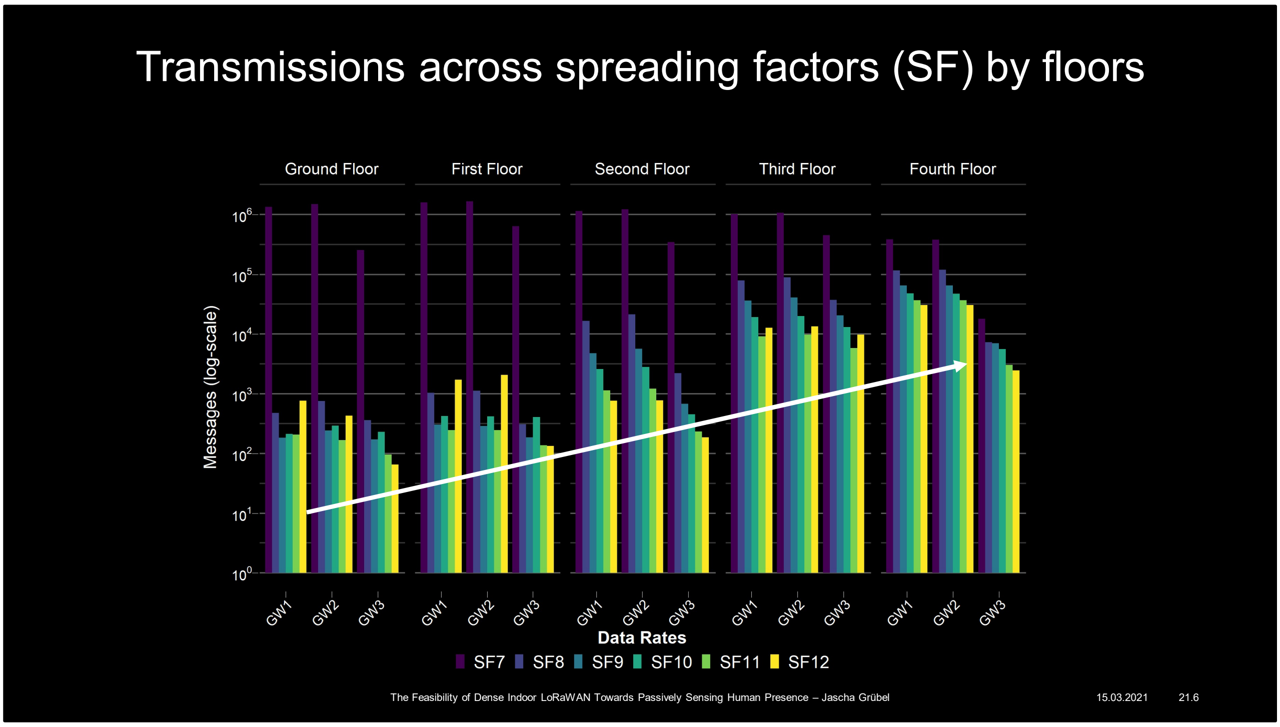

Dividing the transmissions by floor gives us more insight into the use of higher SF.
[C] Here the bar color represents the spreading factor and is
[C] grouped by gateways and
[C] facetted by floor.
[C] Higher SF are used more often for transmission, the further we are from the gateways in the ground floor and below.
This indicates that SF7 is limited indoors in its transmission capacity even at short distance below 64 meters.

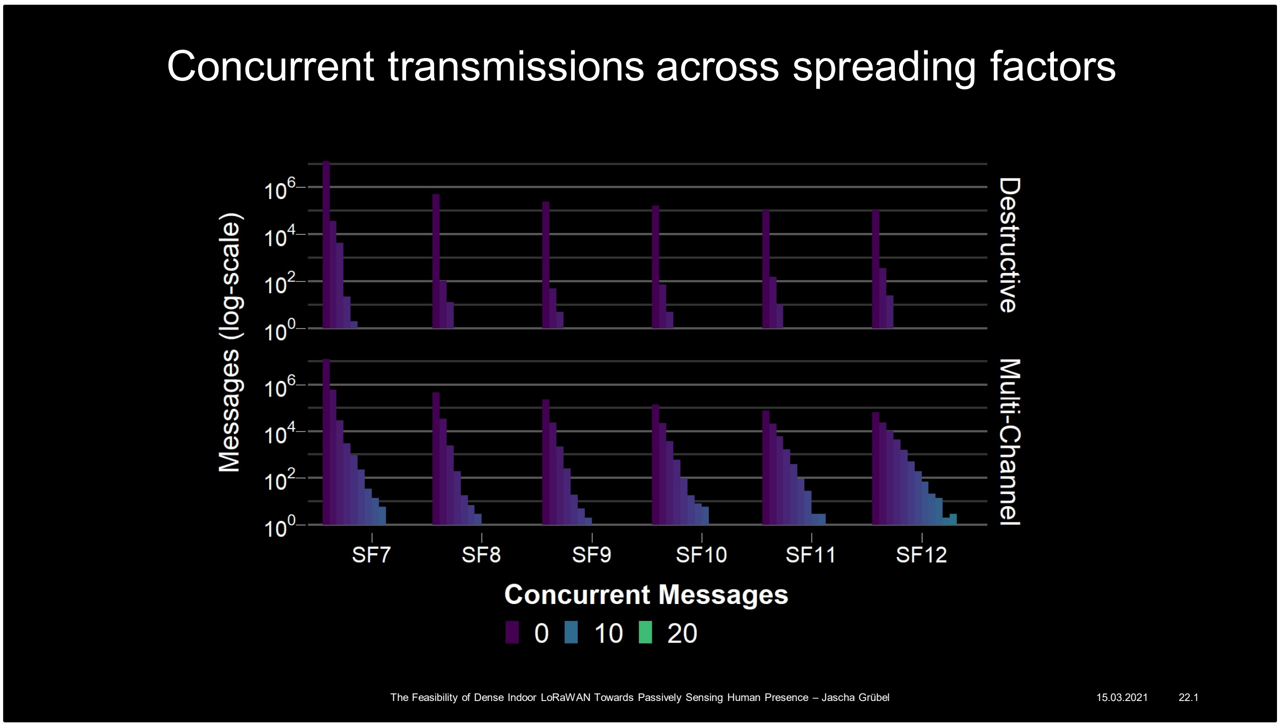

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

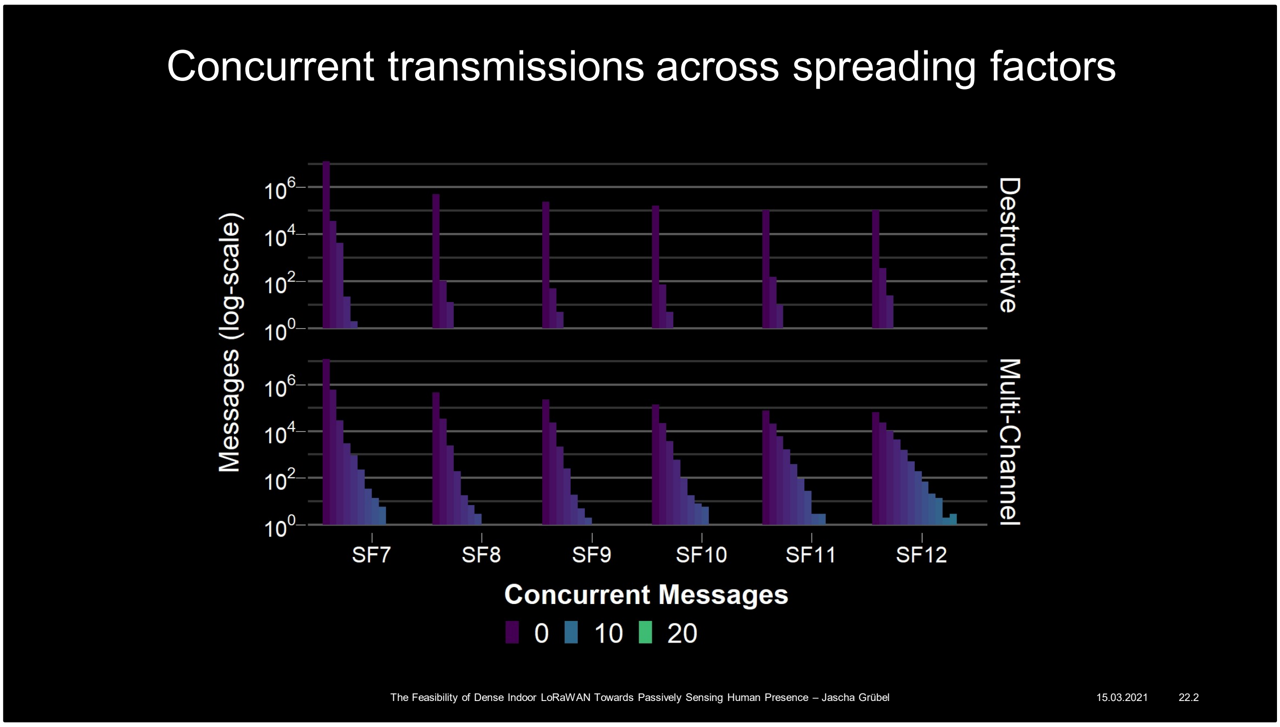

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

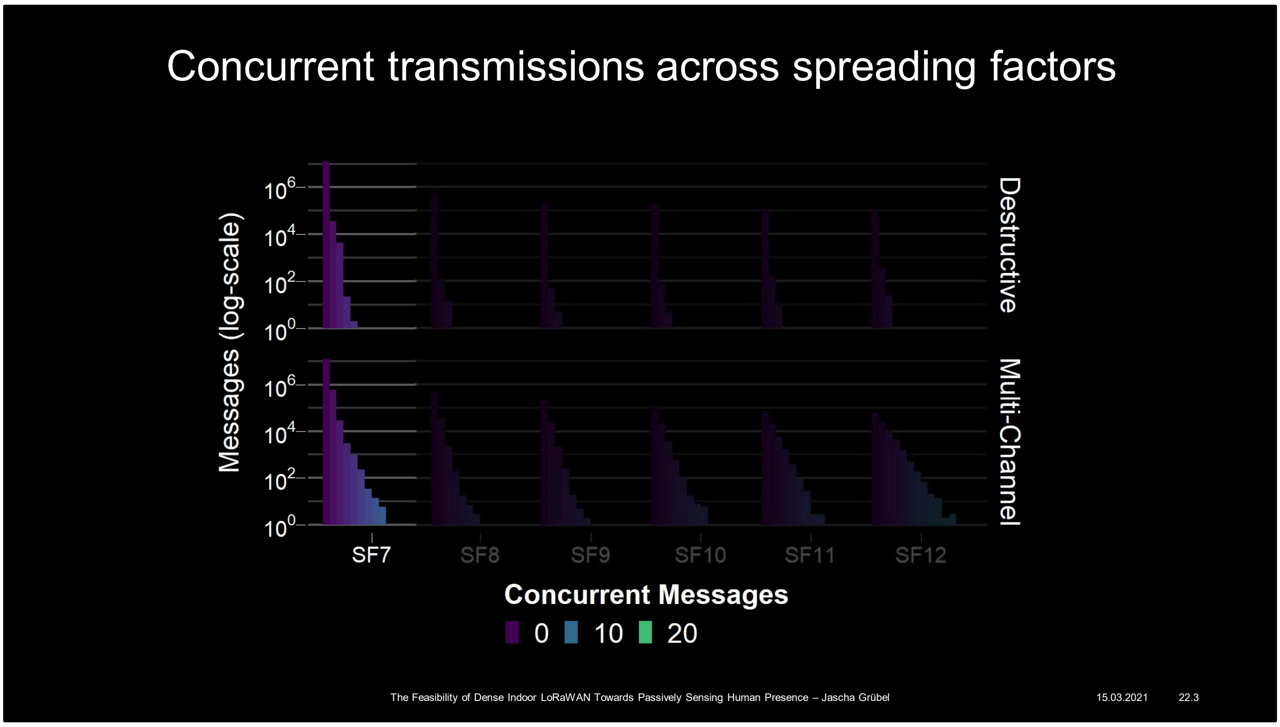

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

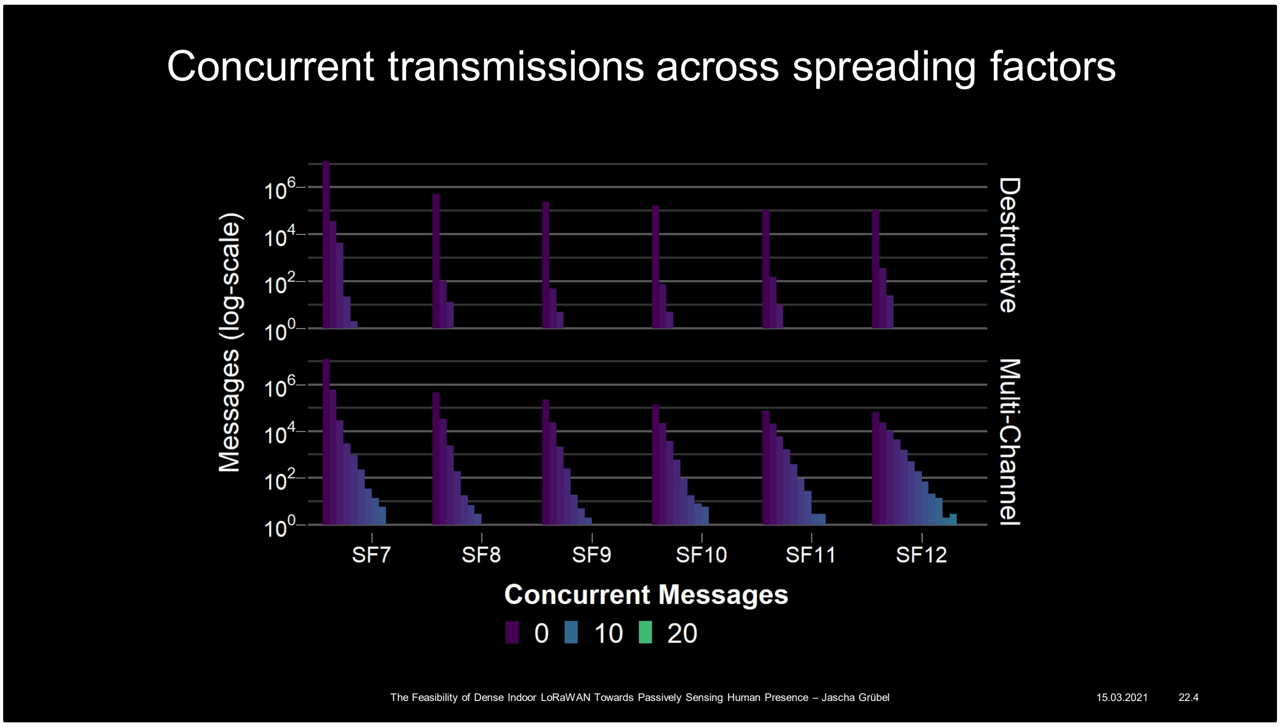

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

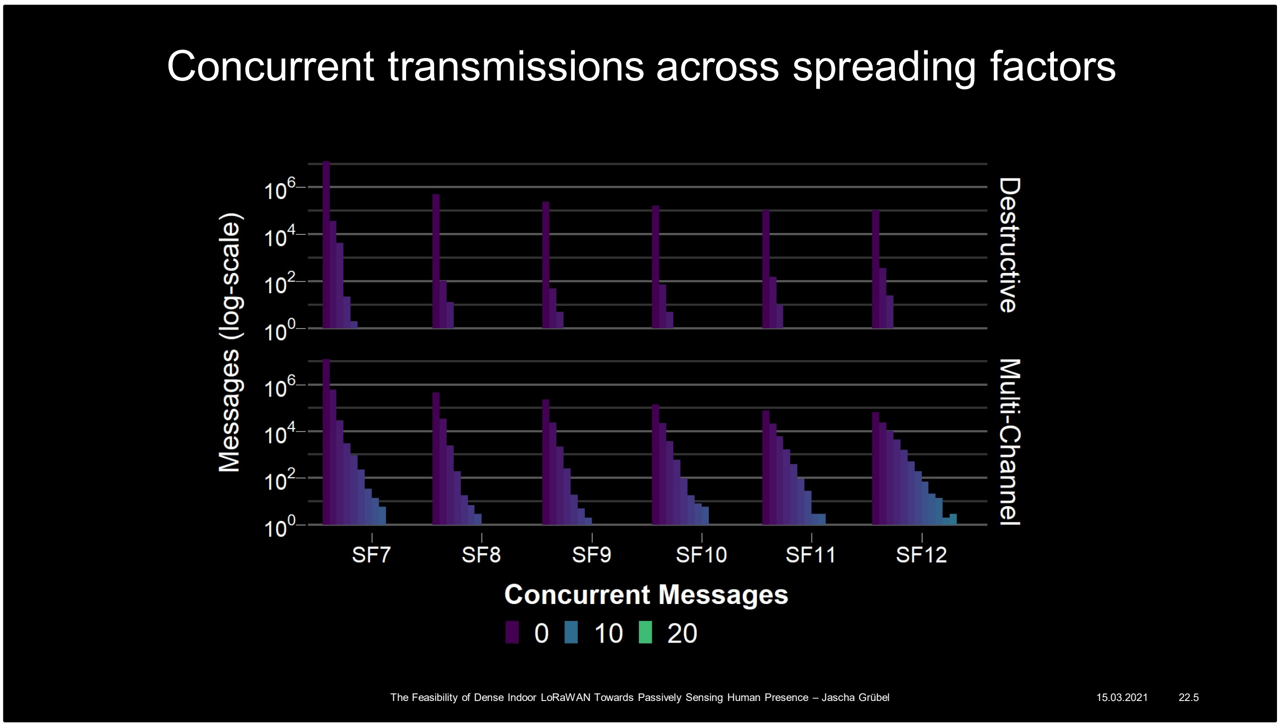

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

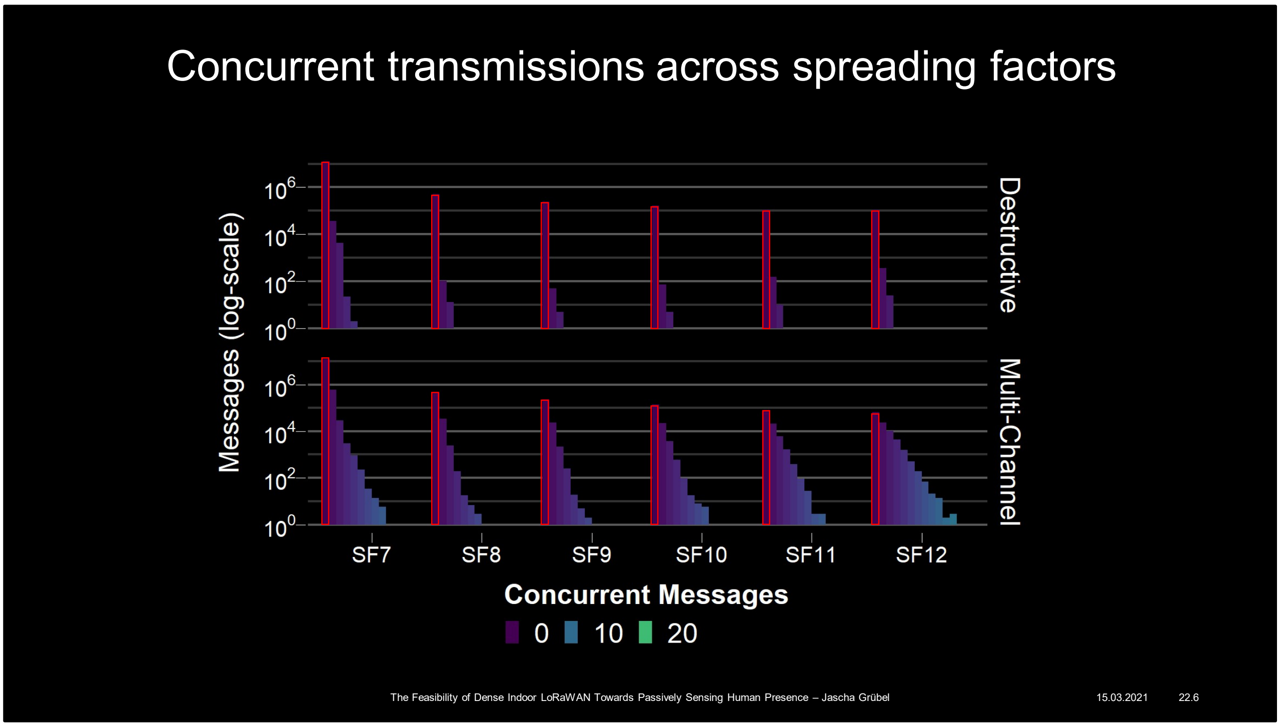

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

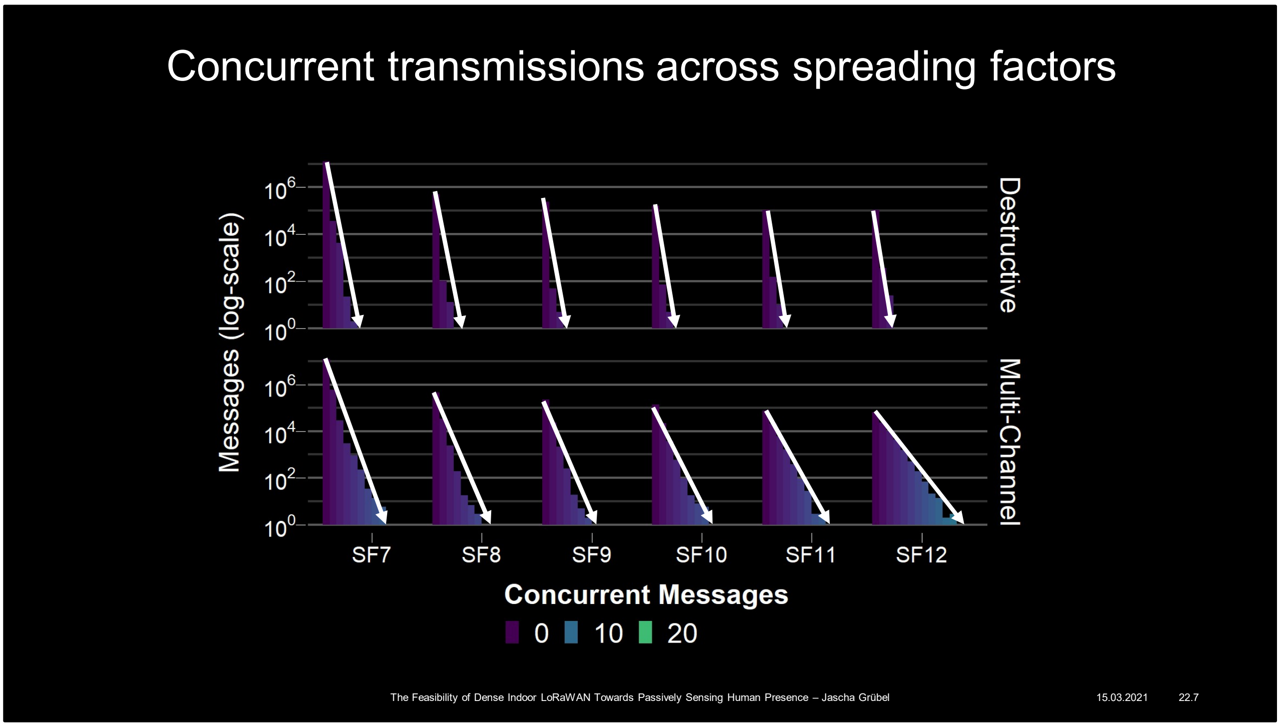

Theoretical discussions on LoRaWAN often discuss the maximal capacity and the risk of destructive concurrent transmissions limiting the system.
We would like to know whether this is the case in our system.
Here is a quick reminder from the definition of chirps:
Destructive concurrent transmission are send at the same time, on the same channel, and at the same SF.
Whereas multi-channel concurrent transmission are send at the same time and the same SF but on a different channel.
[C] The bars represent the count of how many concurrent messages where transmitted at the same time as the transmission and are
[C] facetted by the SF and whether they ae destructive or multi-channel.
[C] We can see that most transmissions have 0 concurrent transmissions.
[C] The occurrence of larger concurrency drops of exponentially – note the log scale – especially for destructive messages.
But there is also an increase of concurrency at higher SF due to the longer transmission time, especially notable in multi-channel transmission.
Despite our setup being order of magnitudes larger than previous work, we conclude that in our setup concurrency is still not an issue with 0.29% destructive and 5.5% multi-chancel concurrent transmissions of all

transmissions.
However, a trend can be observed that will play a more crucial role in an even denser system.

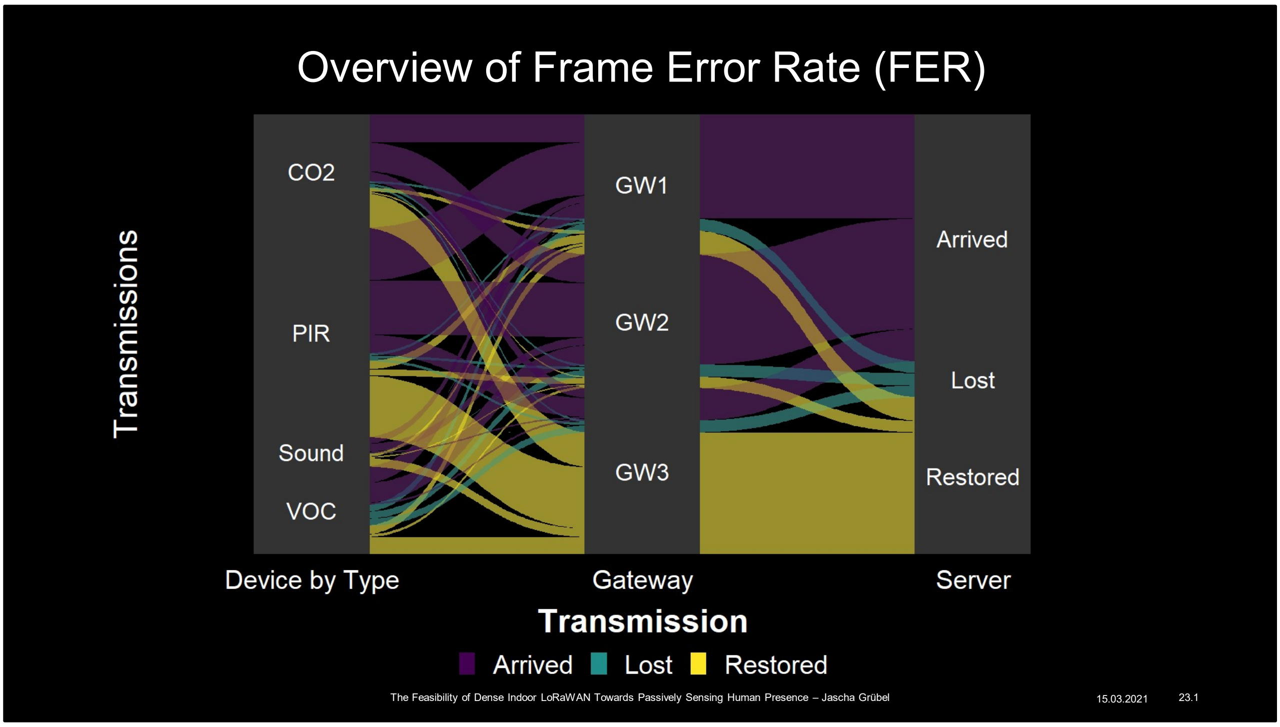

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

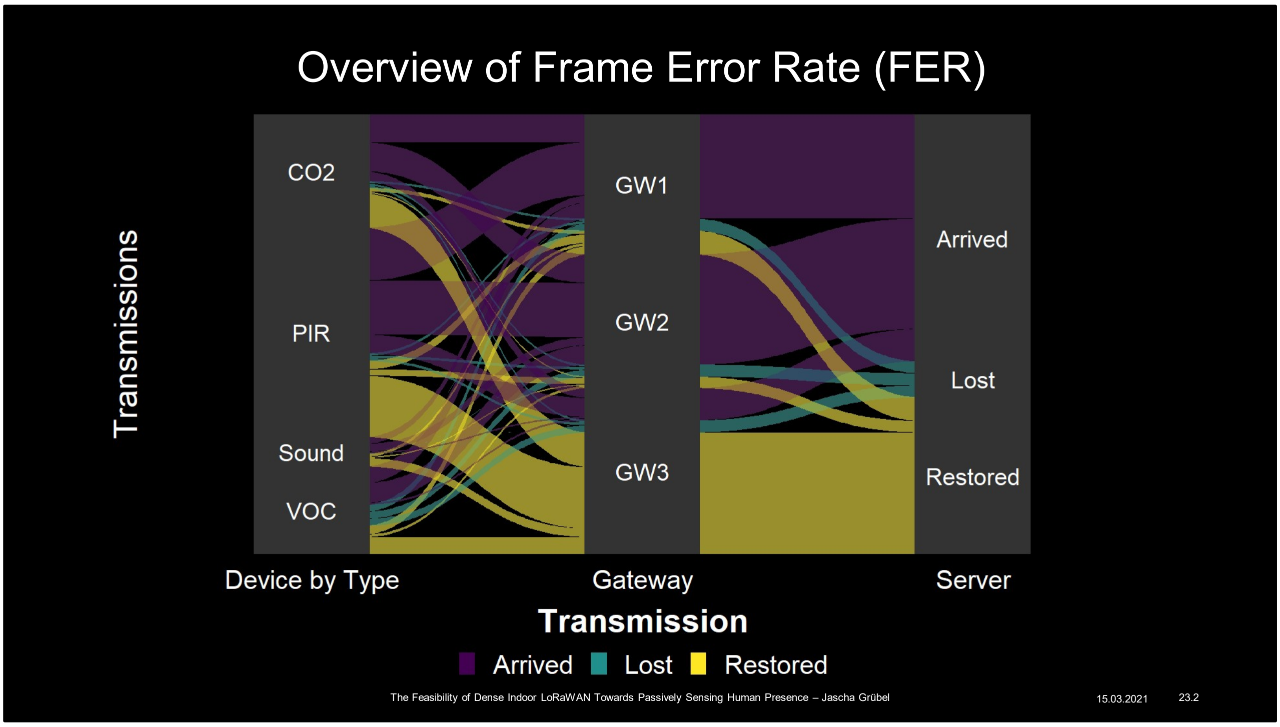

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

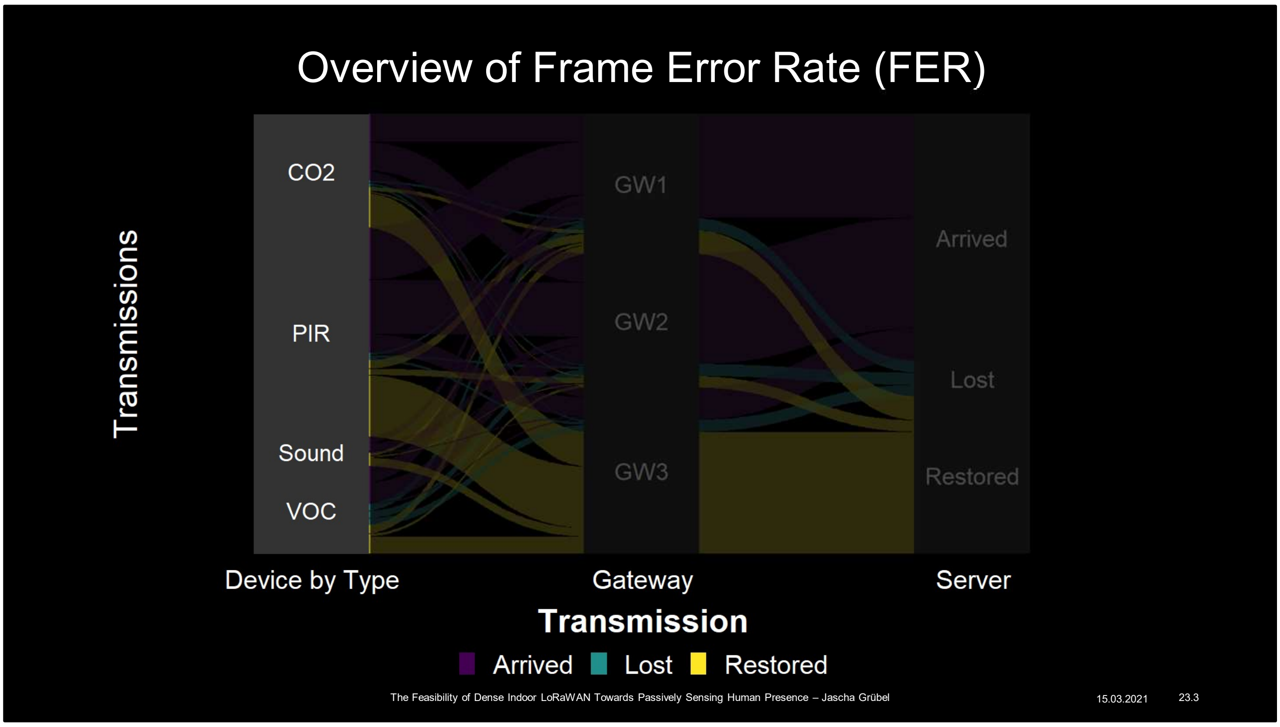

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

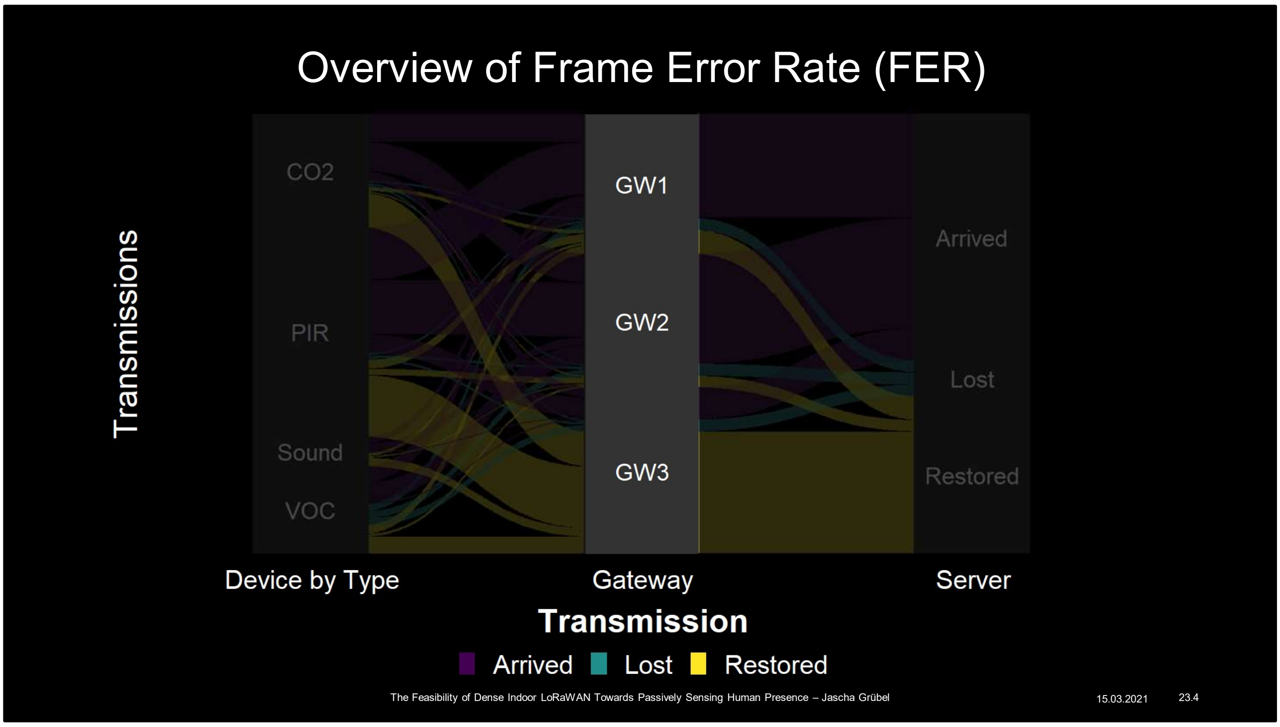

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

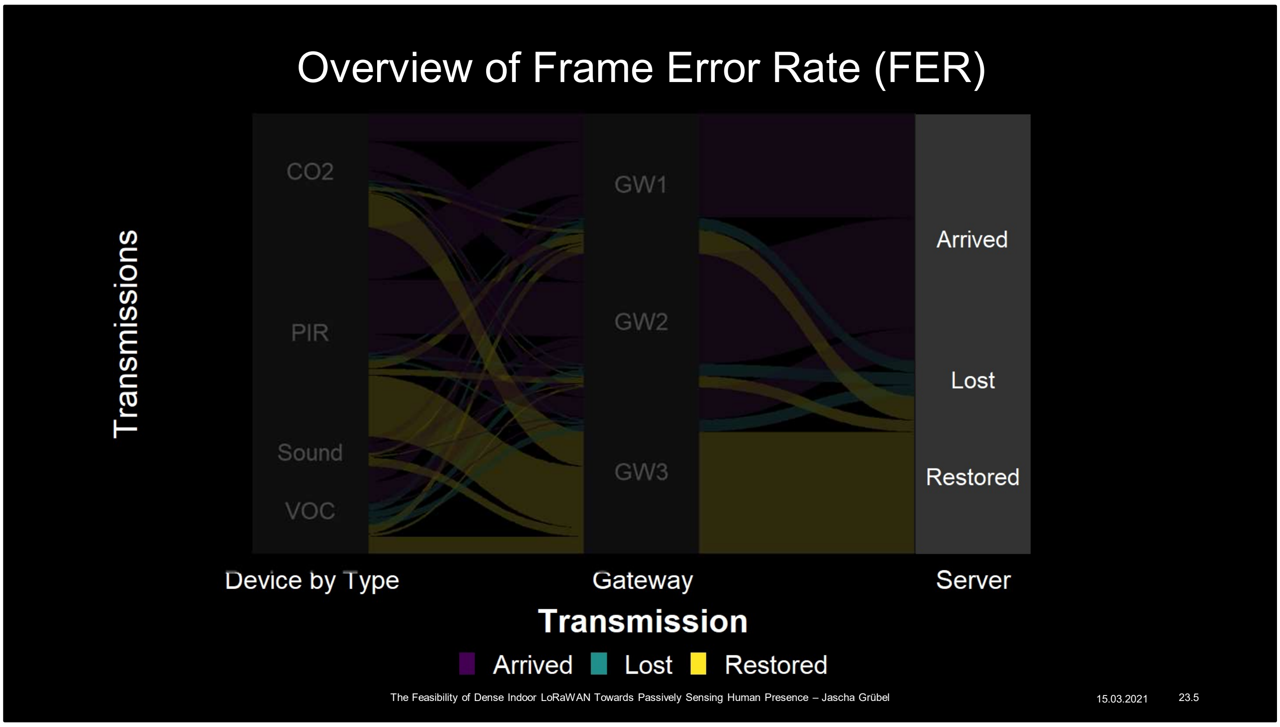

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

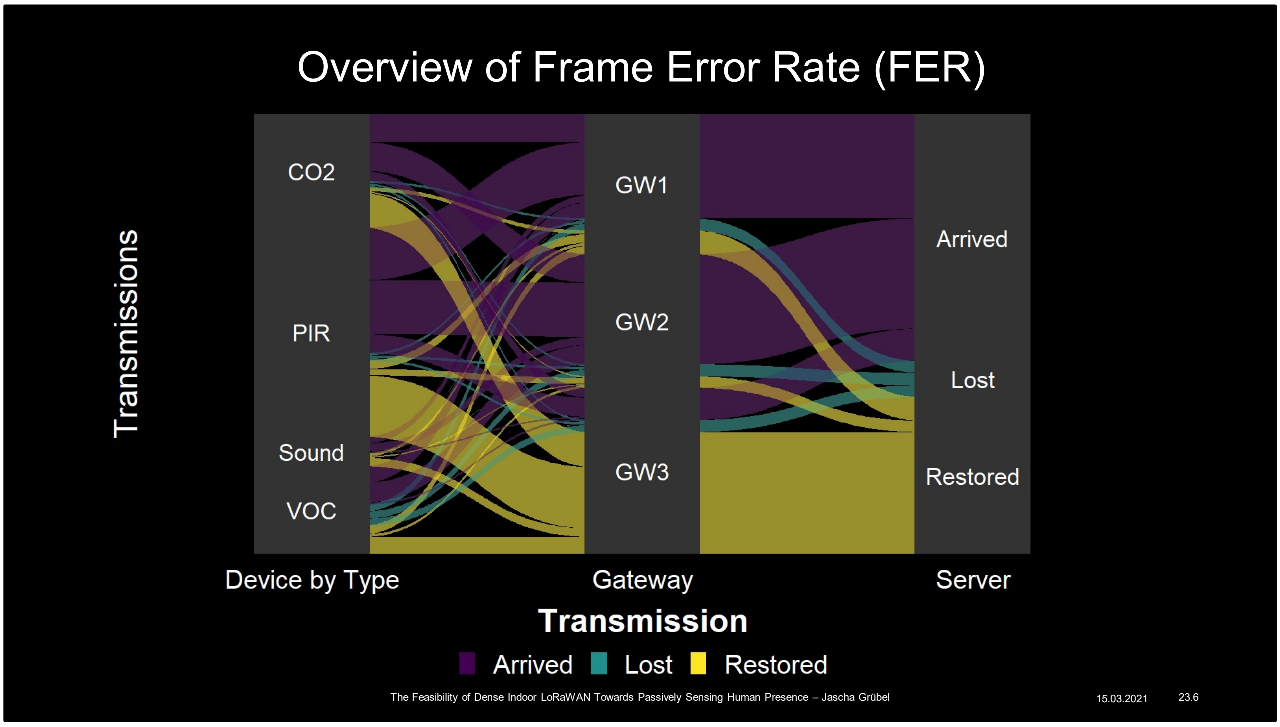

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

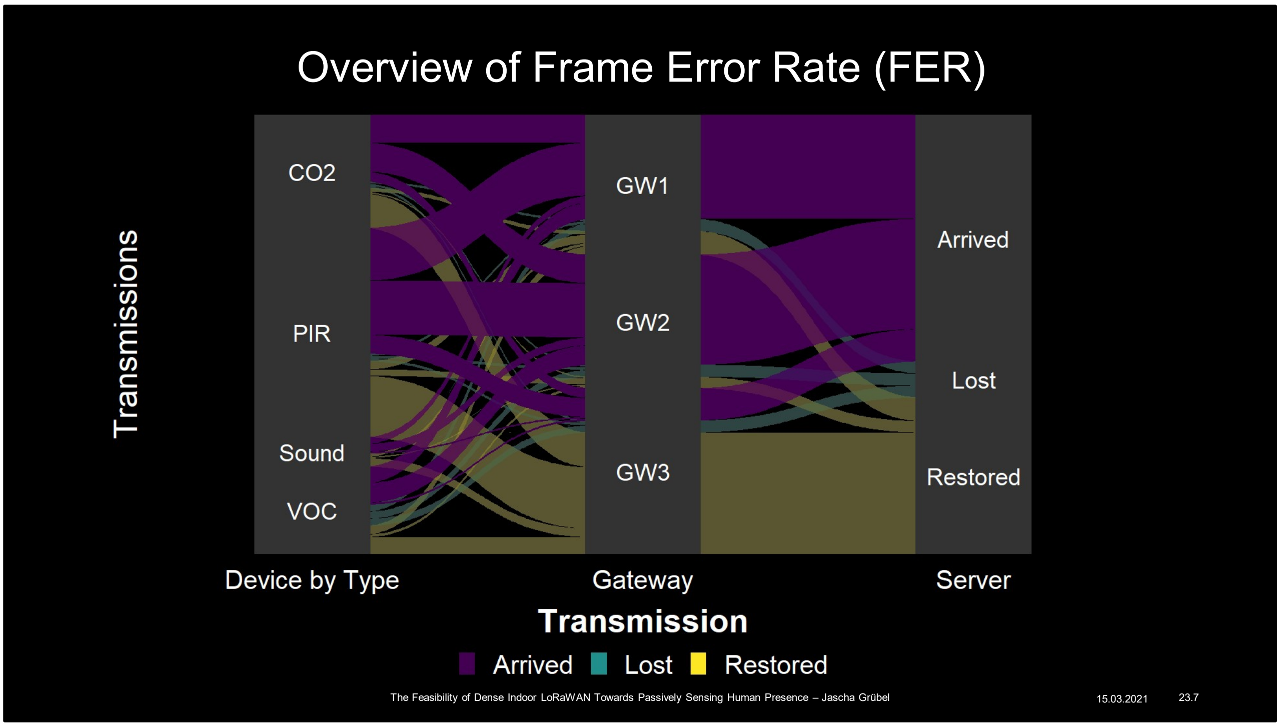

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

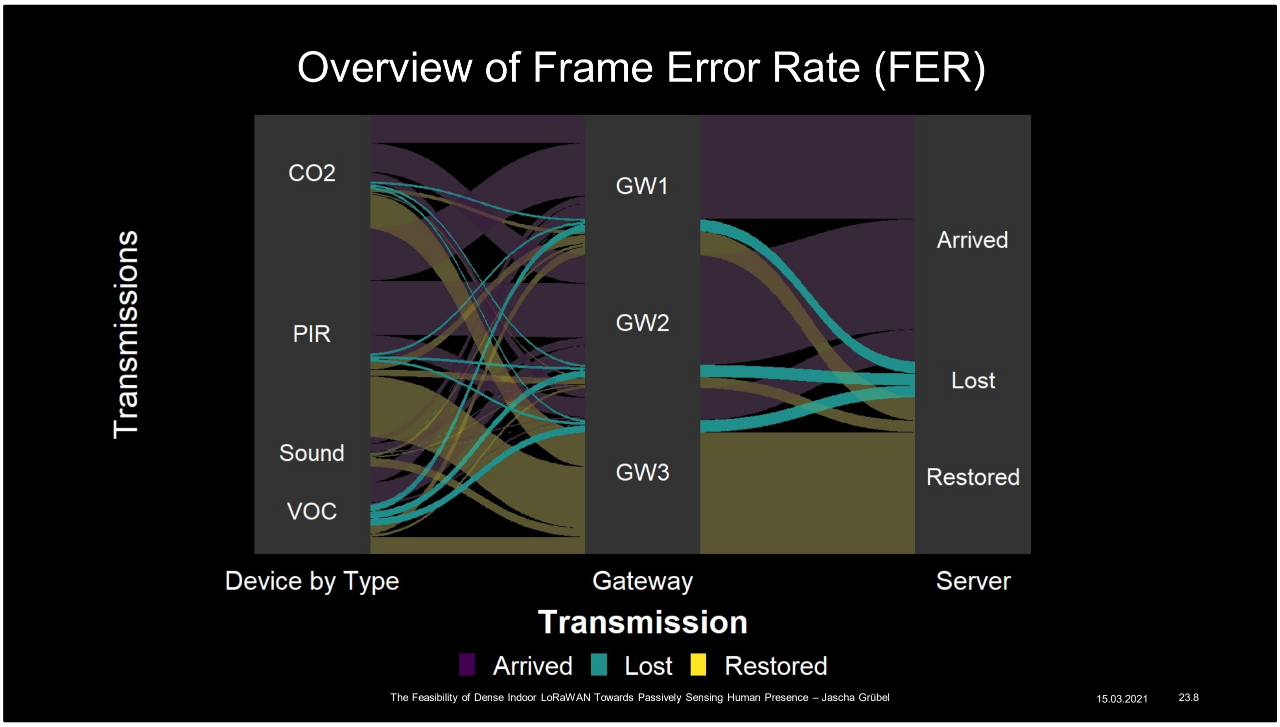

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

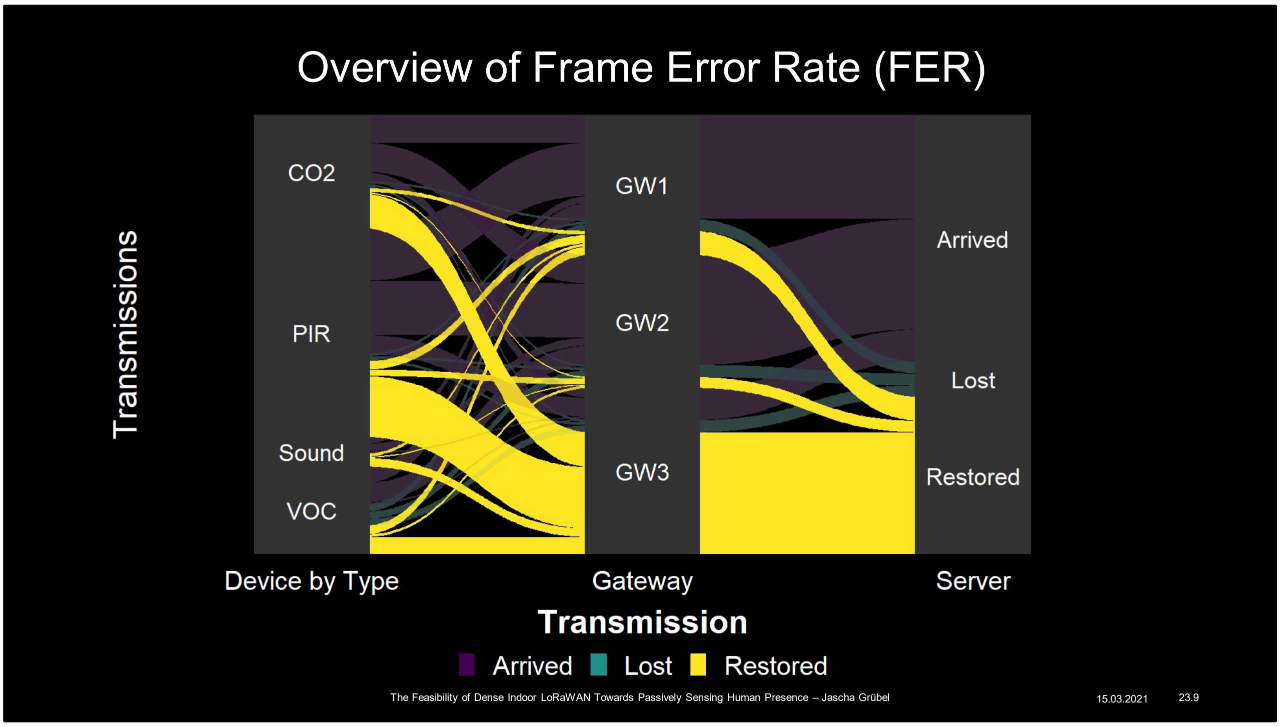

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

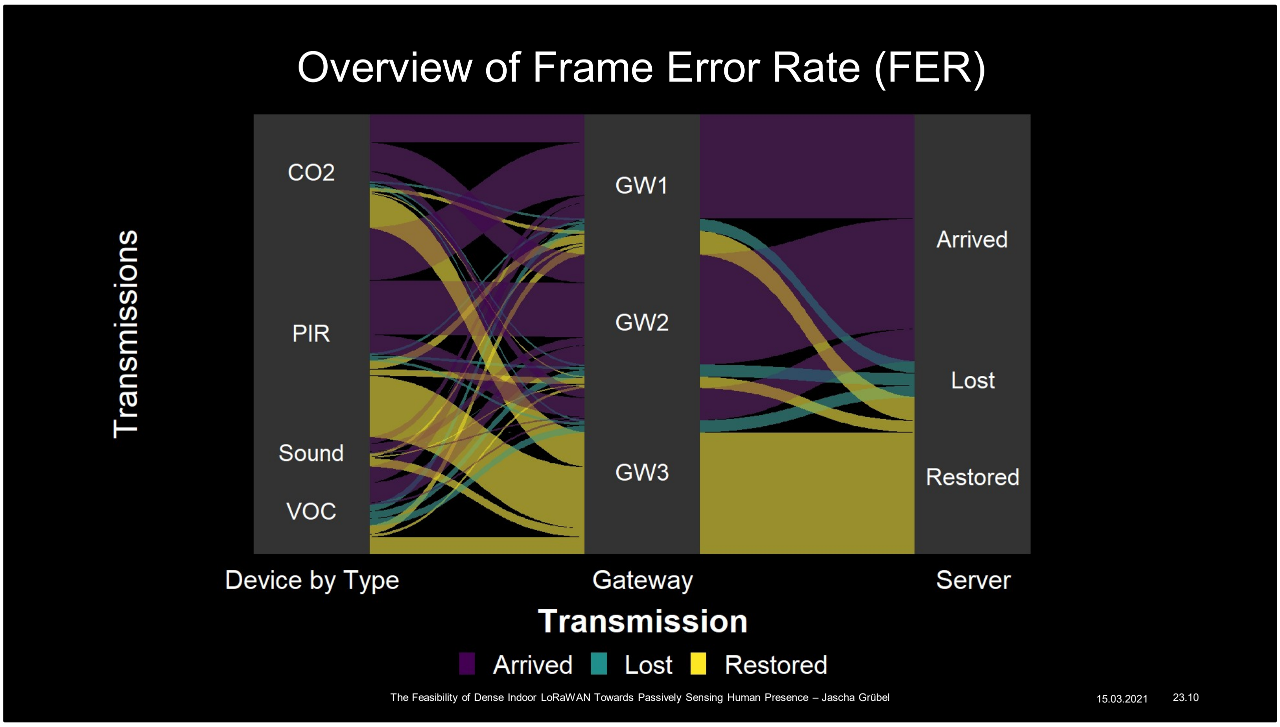

We would also like to know how many transmissions we did lose in our system.
The Frame Error Rate (or FER) can indicate how many transmissions were completely lost.
In this parallel set plot, each transmission is shown once per gateway.
[C] The first column consists of the 5 sensor types,
[C] the second column of the three gateways and
[C] the last column of the status of a transmission at the network server.
[C] The diagonals indicate the status of transmissions across the system.
[C] We can see that most transmissions arrived safely at the network server (55%).
[C] Some transmissions were completely lost (8%).
We use the frame count from the sensor devices to establish whether all transmissions are received in a sequence. A jump in the sequence indicates a frame that was not received by any gateway.
[C] Lastly, many transmissions could be restored because at least one gateway caught them (37%).
Note how every gateway would have lost transmissions if others had not received them.

Furthermore, pay special attention to gateway 3 which is on average more than 30m away from the sensors and required restoration of many transmissions.

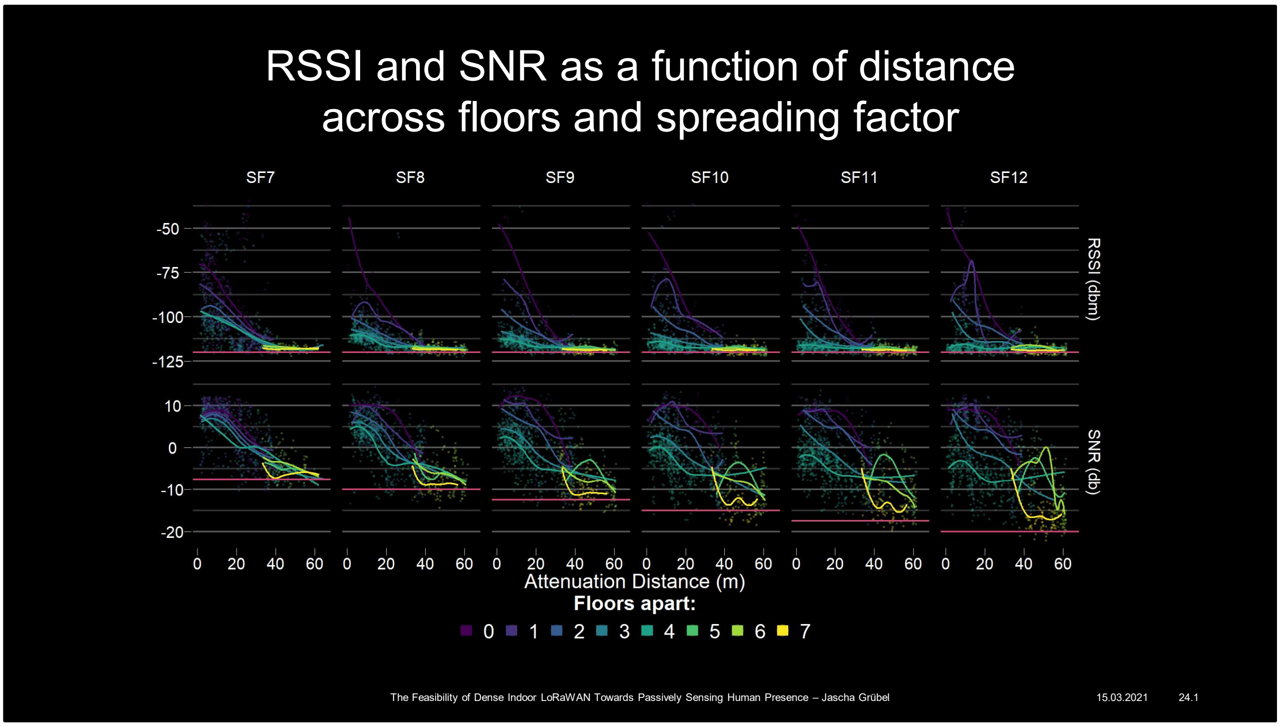

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

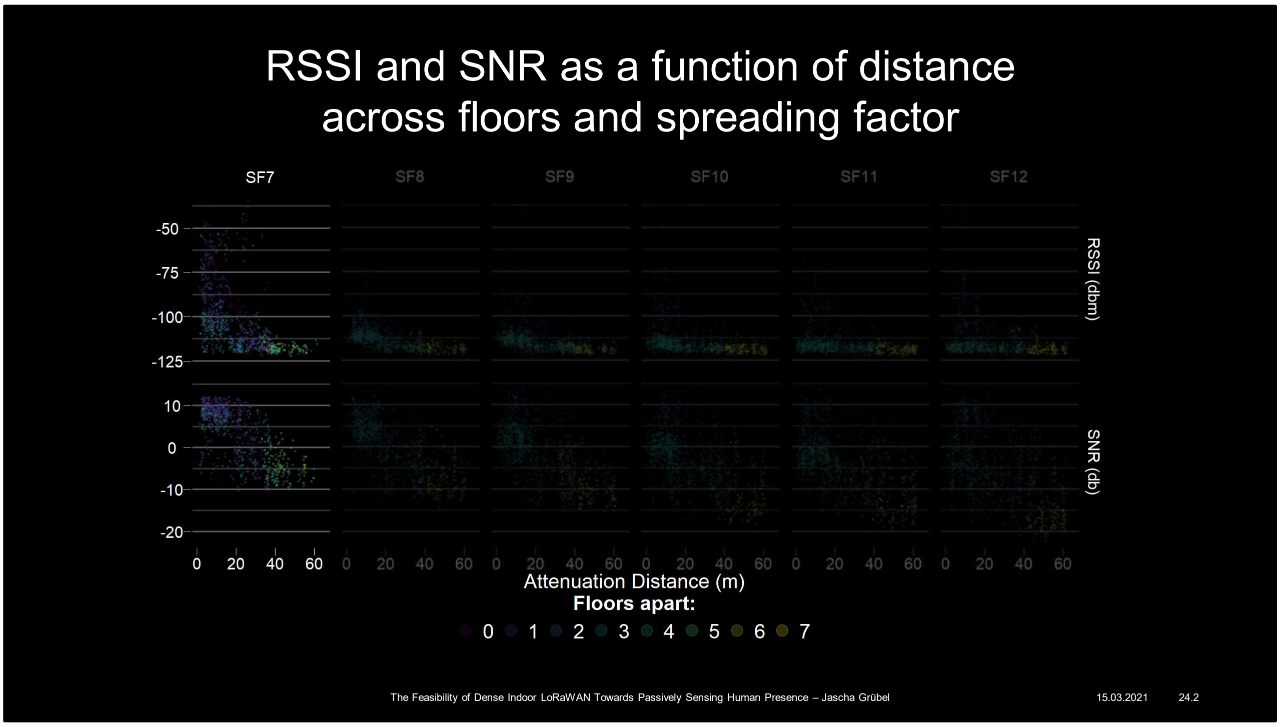

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

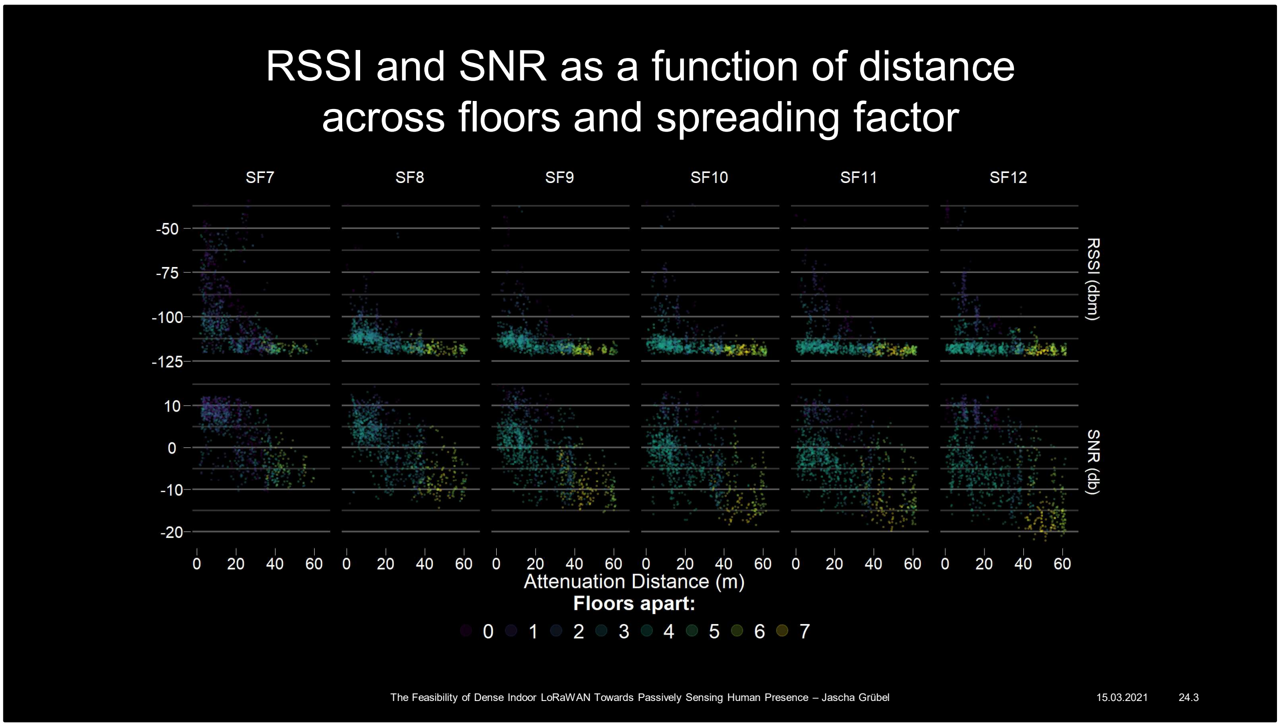

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.

[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.

[C] The approximate noise floor is indicated in purple.

[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.

[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

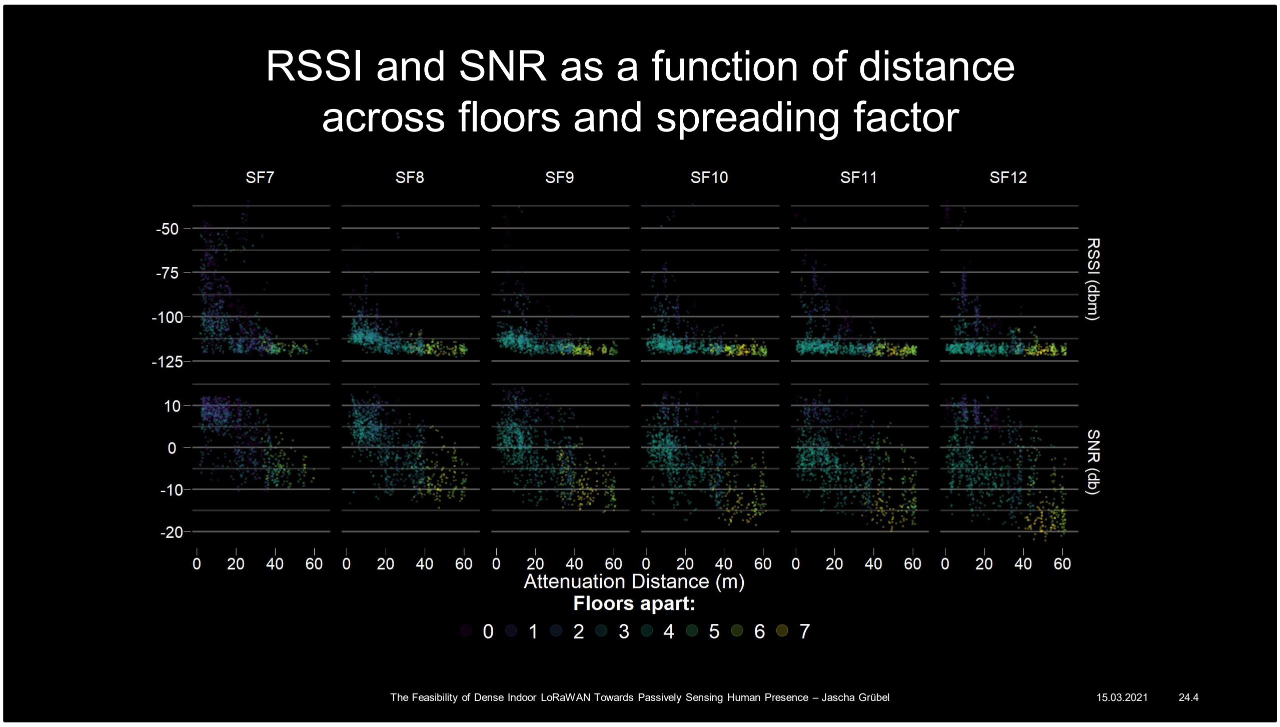

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

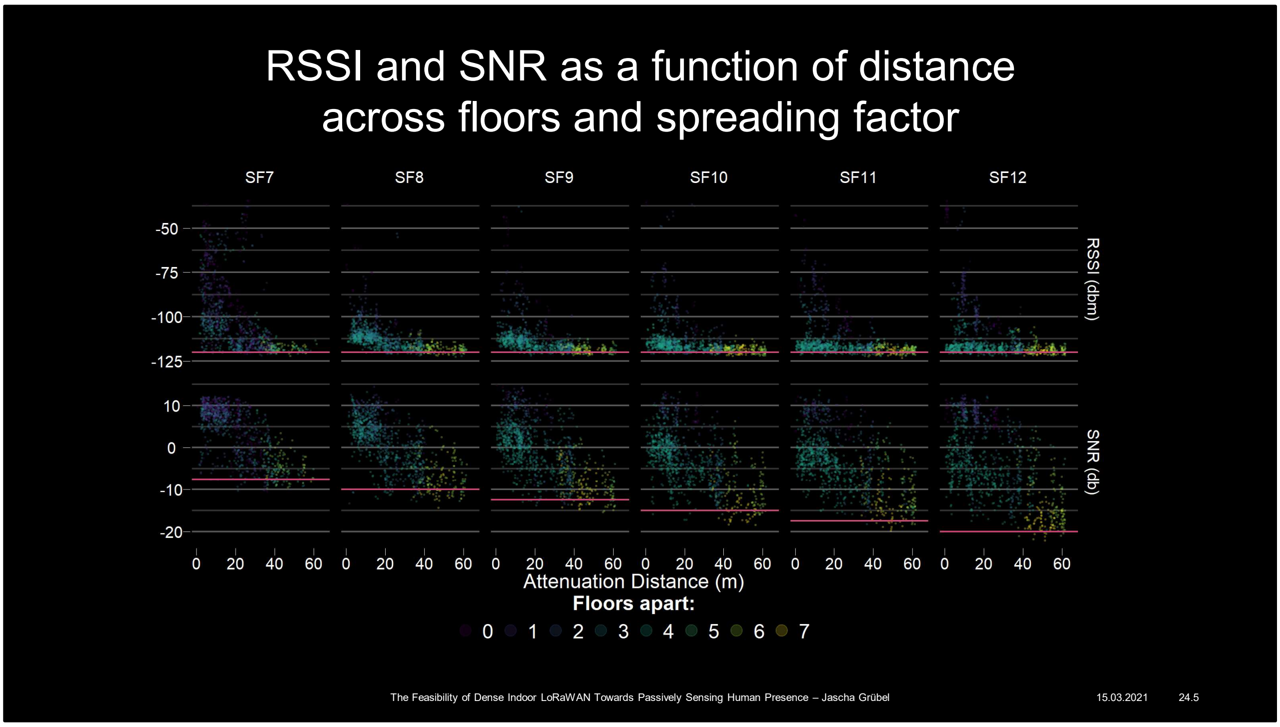

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

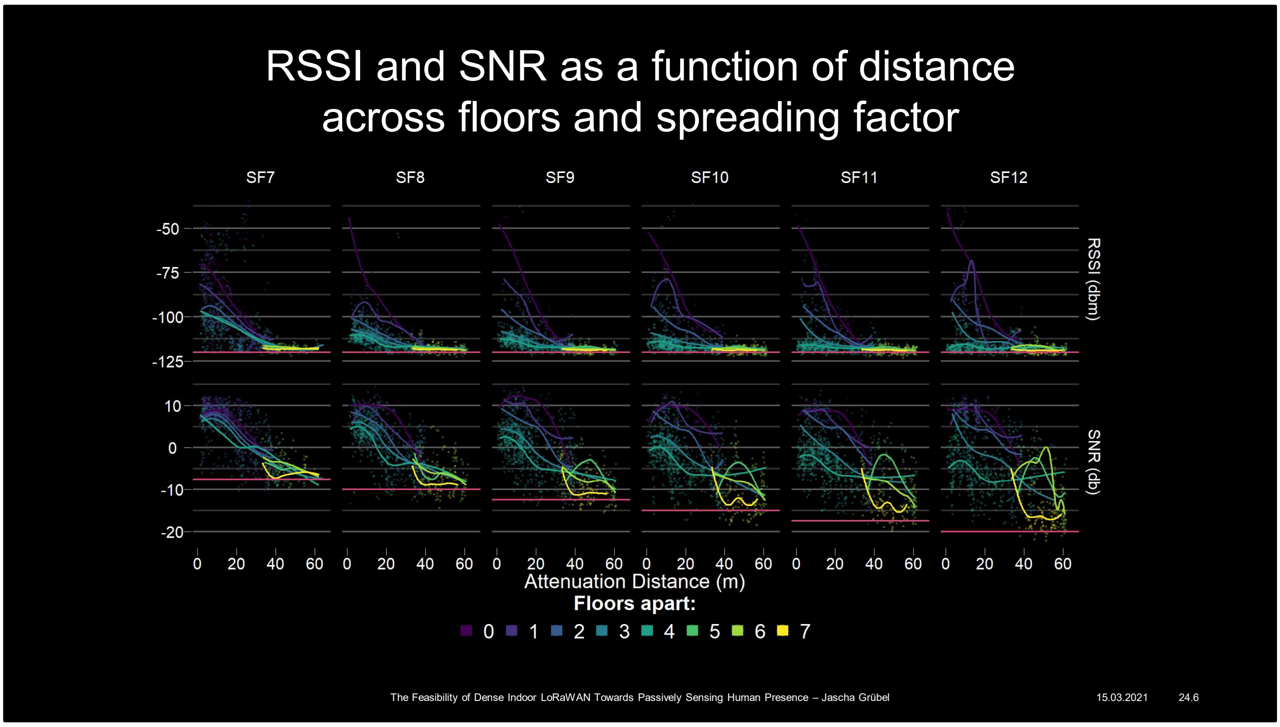

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

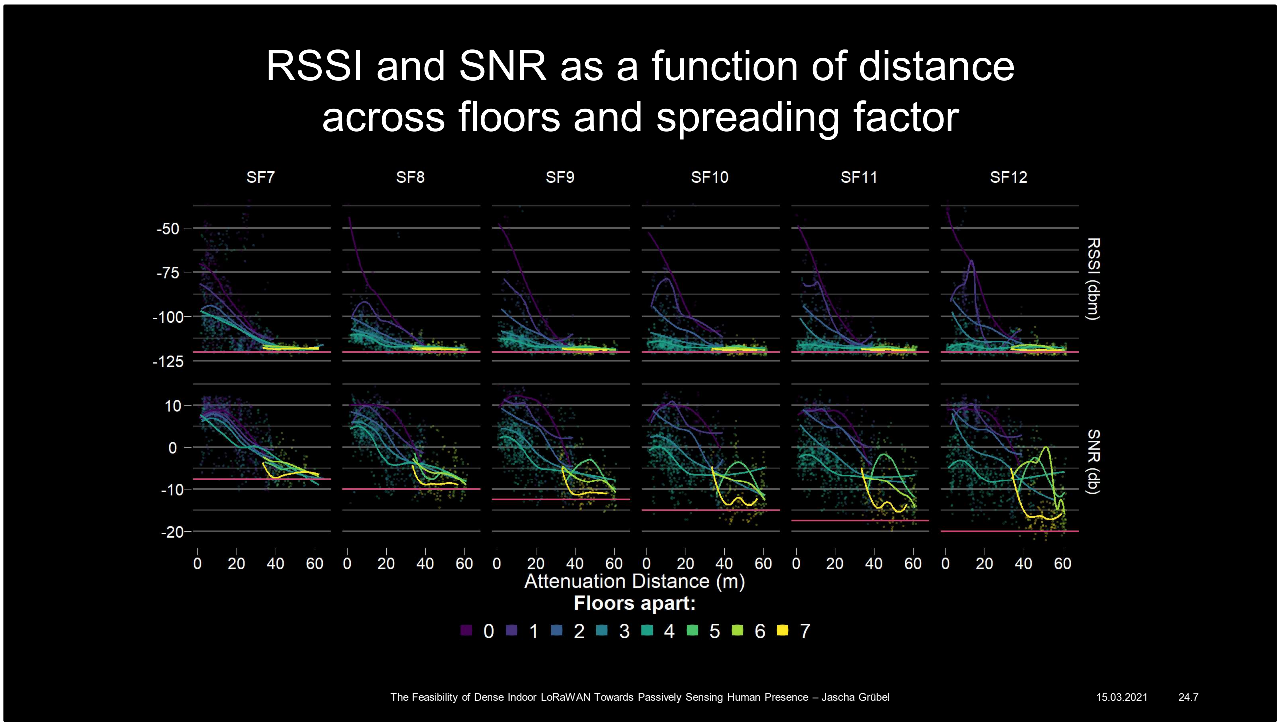

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

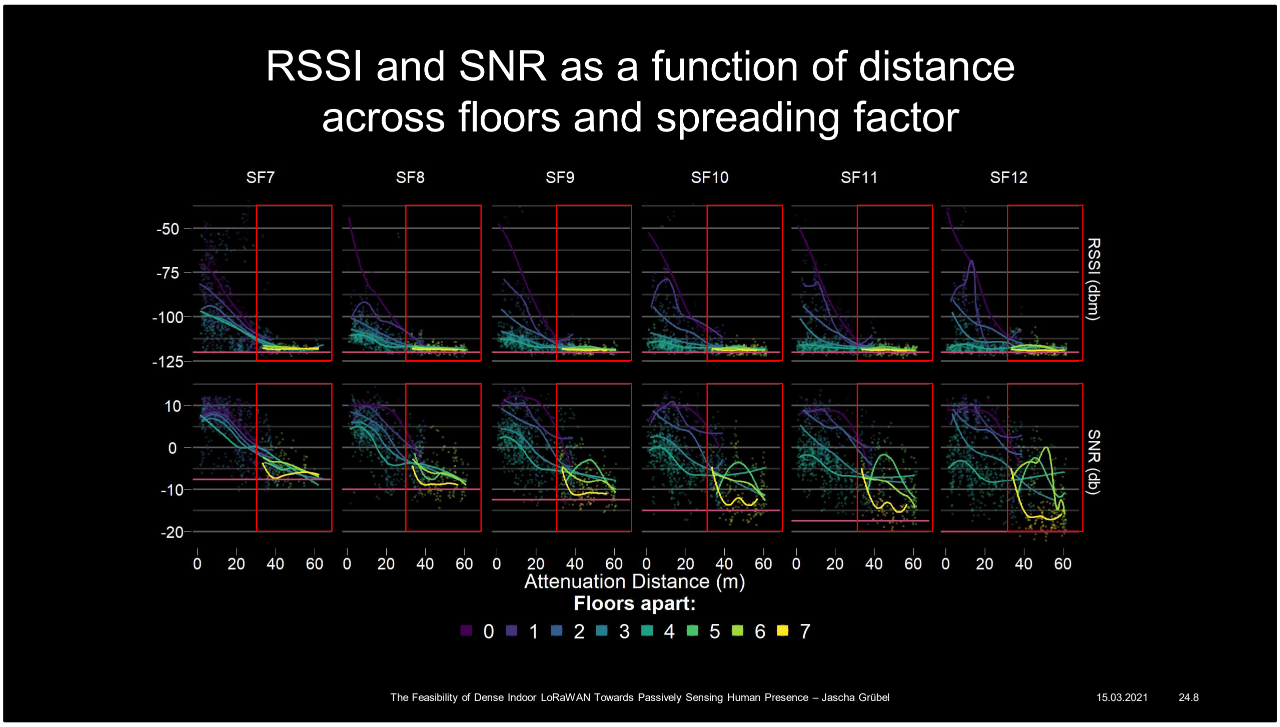

We have reviewed the transmission type, concurrency and the lost messages. Going beyond the overall reception of messages we would like to understand the maximal transmission distance to ensure that all sensor devices have good coverage and as few messages as possible are lost.
[C] Here, we show the RSSI and SNR of a random sample of transmissions by signal strength and distance, colored by floor level and facetted by SF.
[C] The approximate noise floor is indicated in purple.
[C] We use a LOESS interpolation of the average RSSI and SNR to indicate signal strength with increasing distance from the gateway.
[C] Note how beyond 30 meters, the signal quality strongly diminishes and quickly reaches the noise floor shown in red.

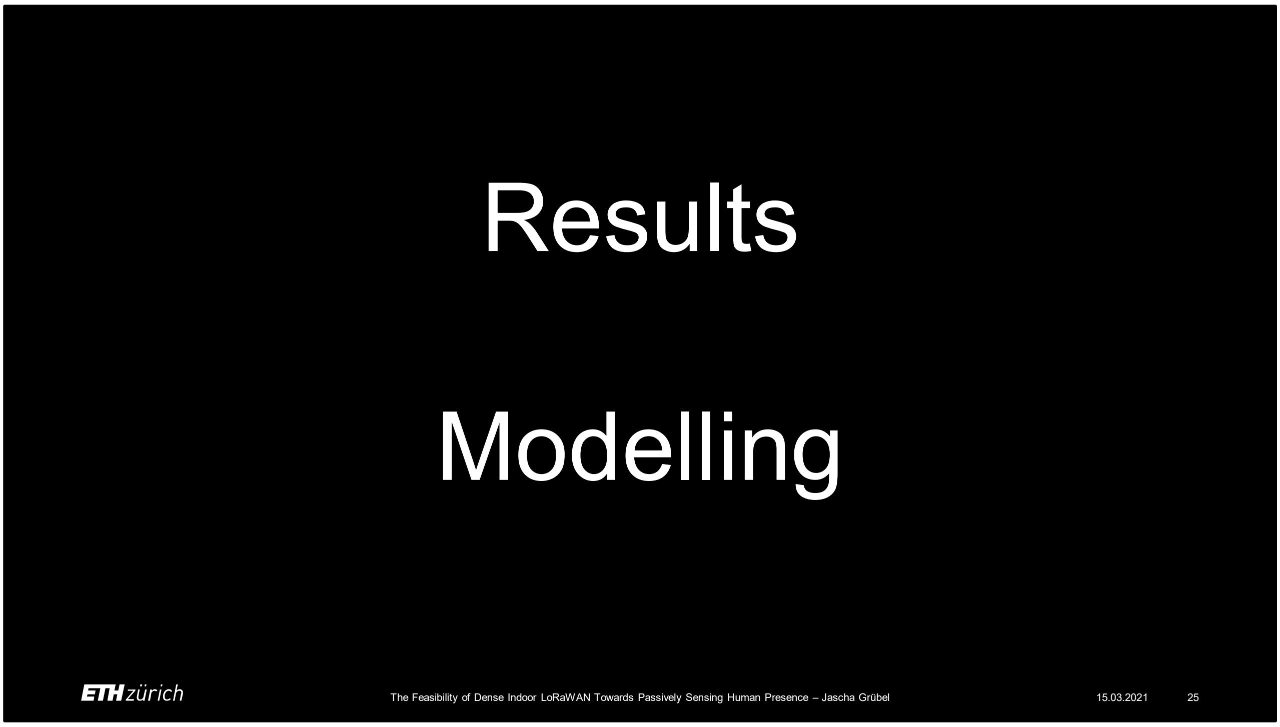

Beyond descriptives, we also constructed a model of the transmission indoors to make predictions on how to best place sensors.

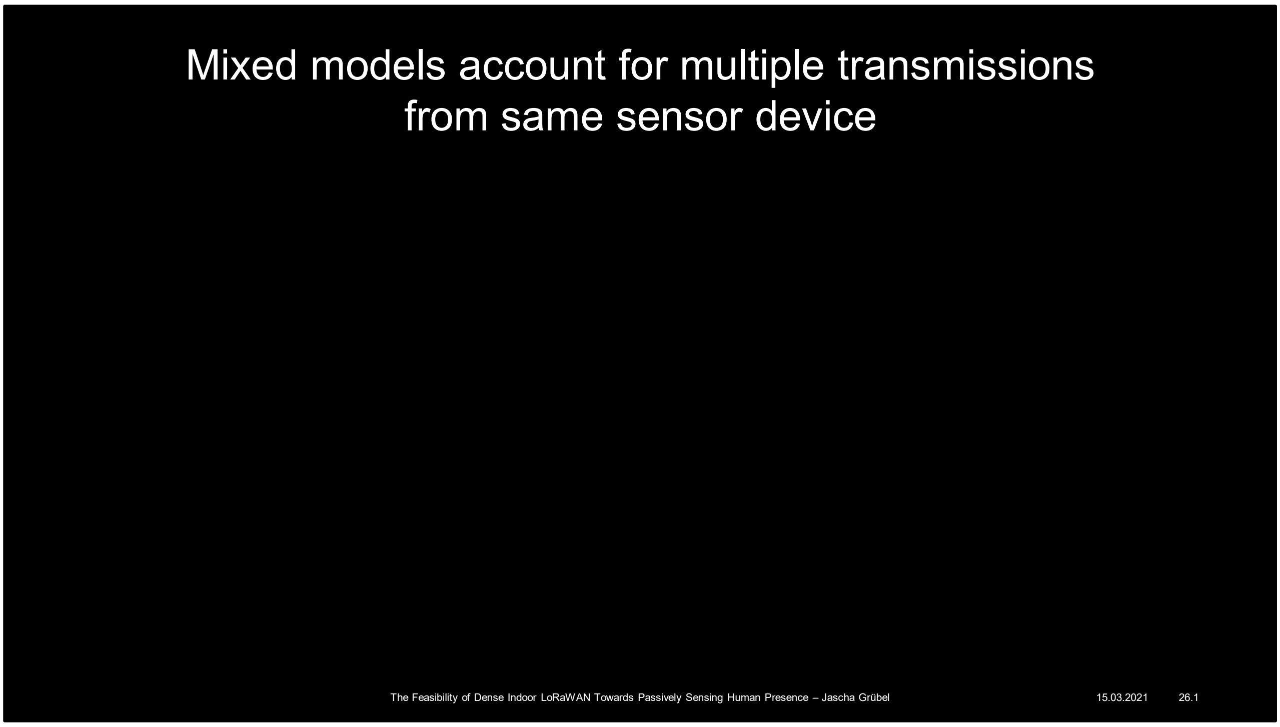

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

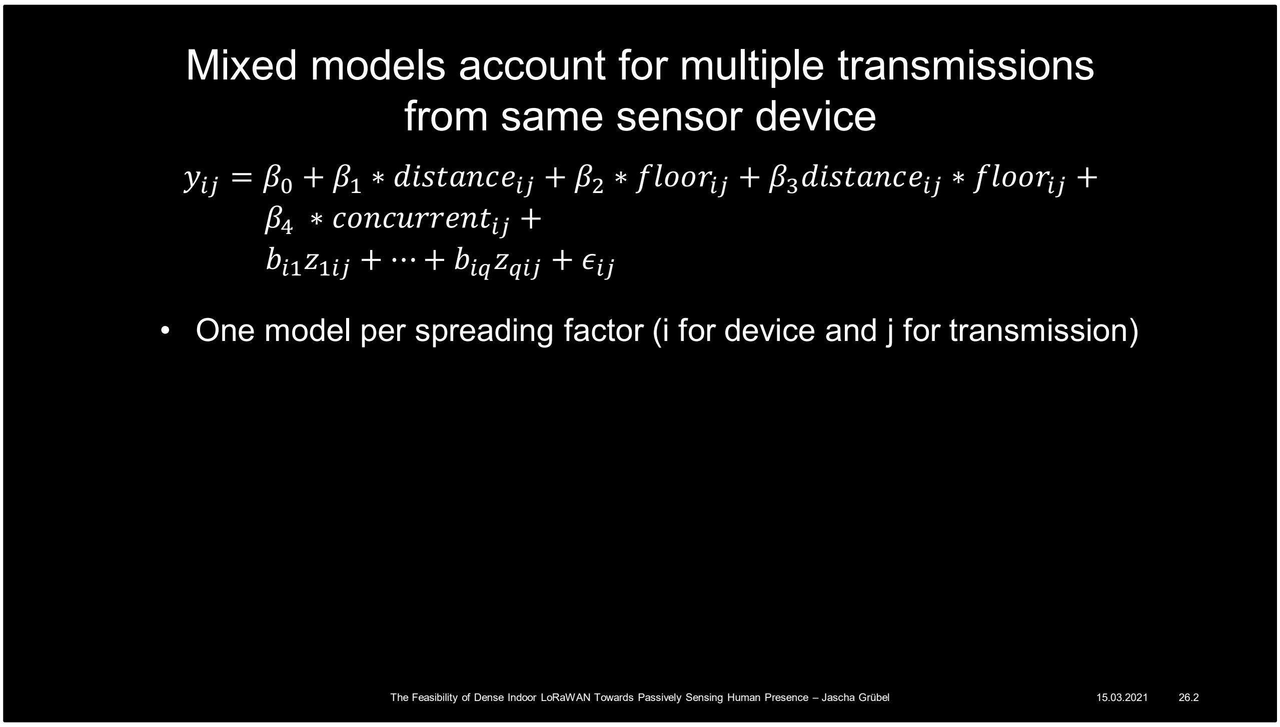

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

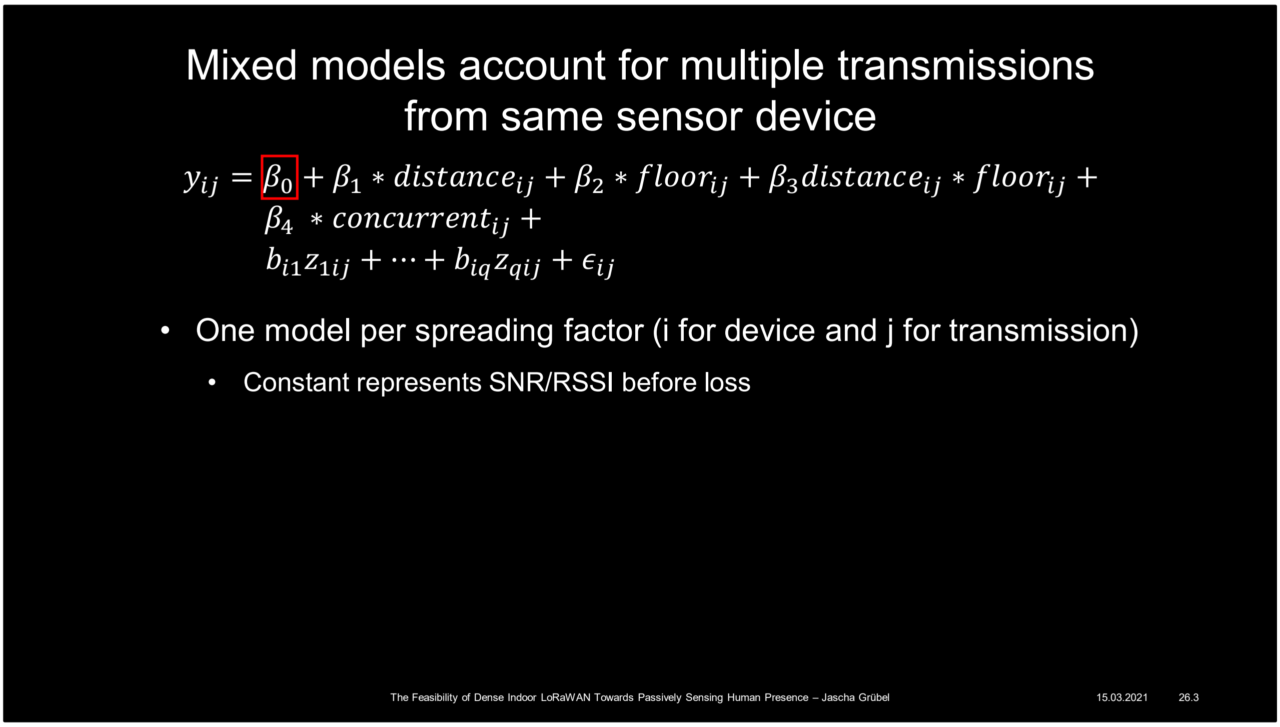

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

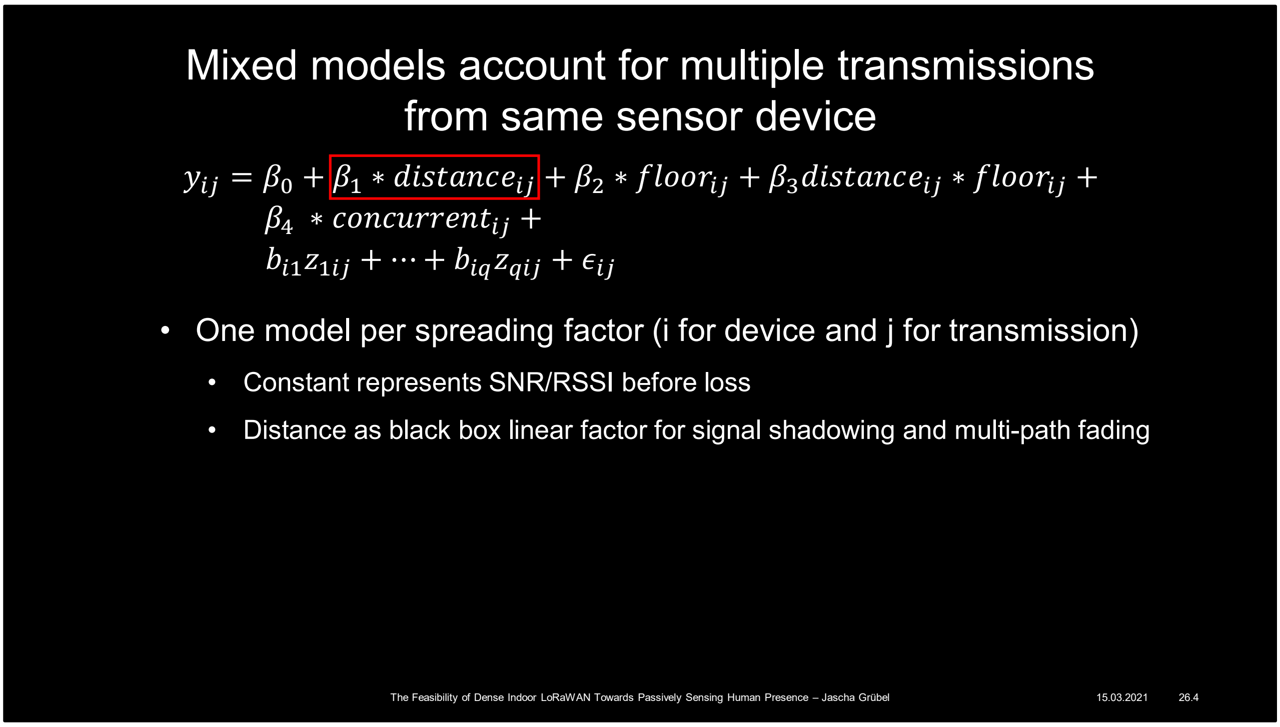

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

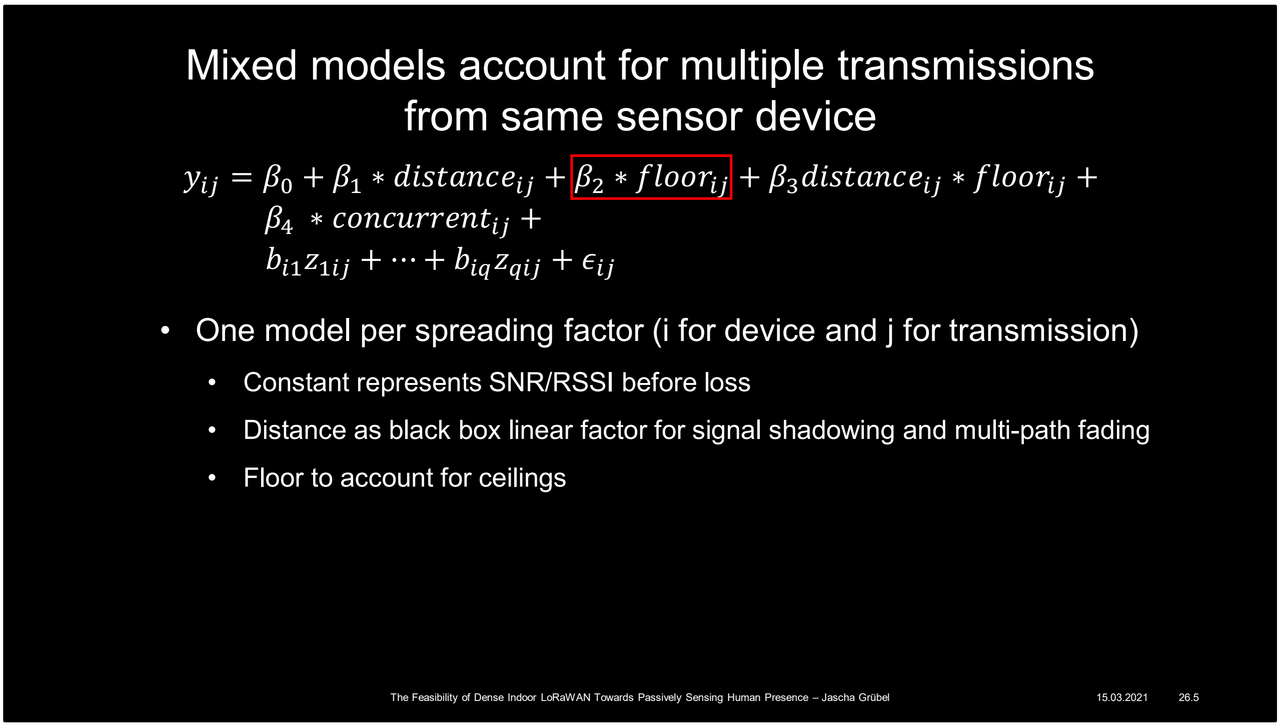

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

Mixed models account for multiple transmissions from same sensor device

$$y_{ij} = \beta_0 + \beta_1 * distance_{ij} + \beta_2 * floor_{ij} + \beta_3 distance_{ij} * floor_{ij} + \beta_4 * concurrent_{ij} + b_{i1}z_{1ij} + \cdots + b_{iq}z_{qij} + \epsilon_{ij}$$

- One model per spreading factor (i for device and j for transmission)
  - Constant represents SNR/RSSI before loss
  - Distance as black box linear factor for signal shadowing and multi-path fading
  - Floor to account for ceilings
  - Interaction to decompose floors into distance and non-distance component

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

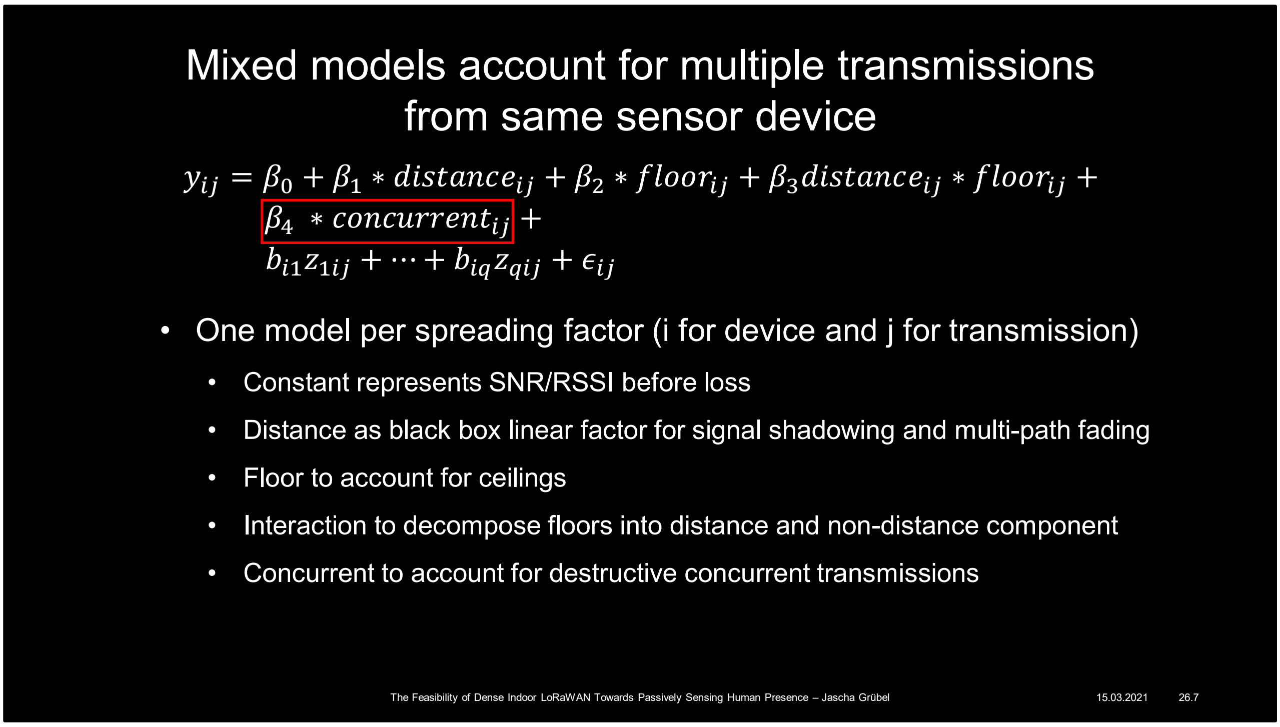

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

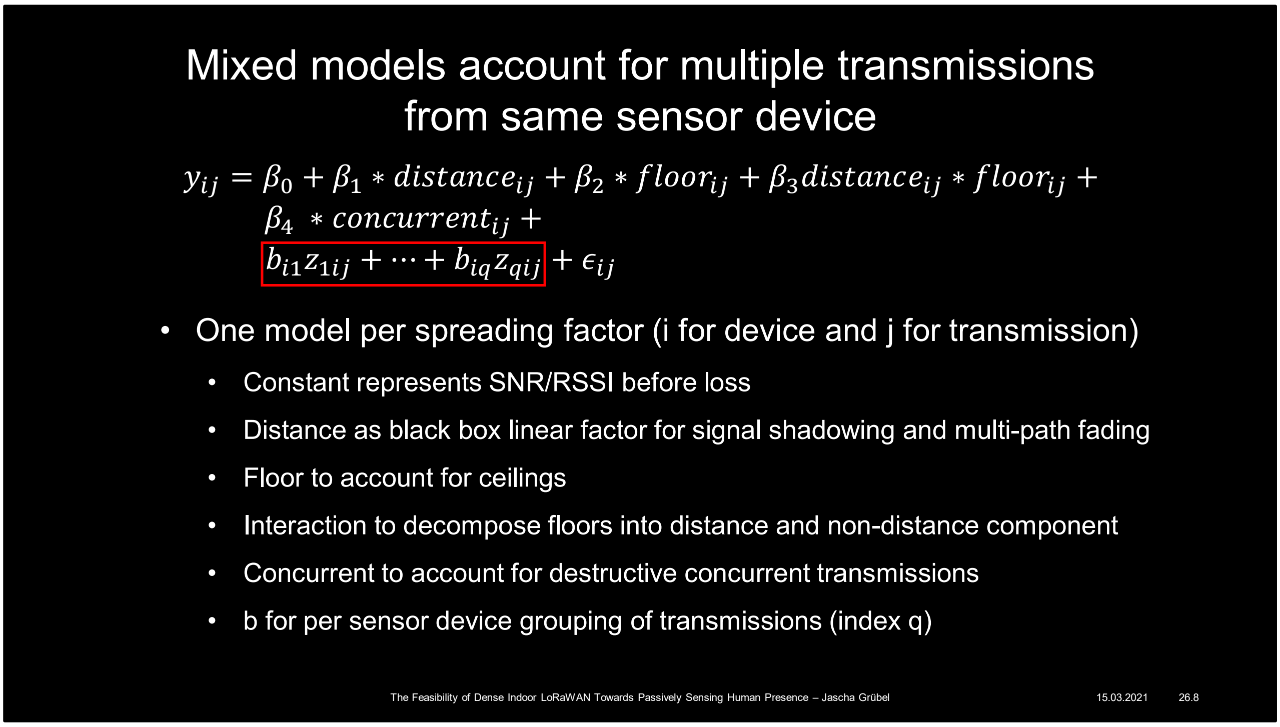

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

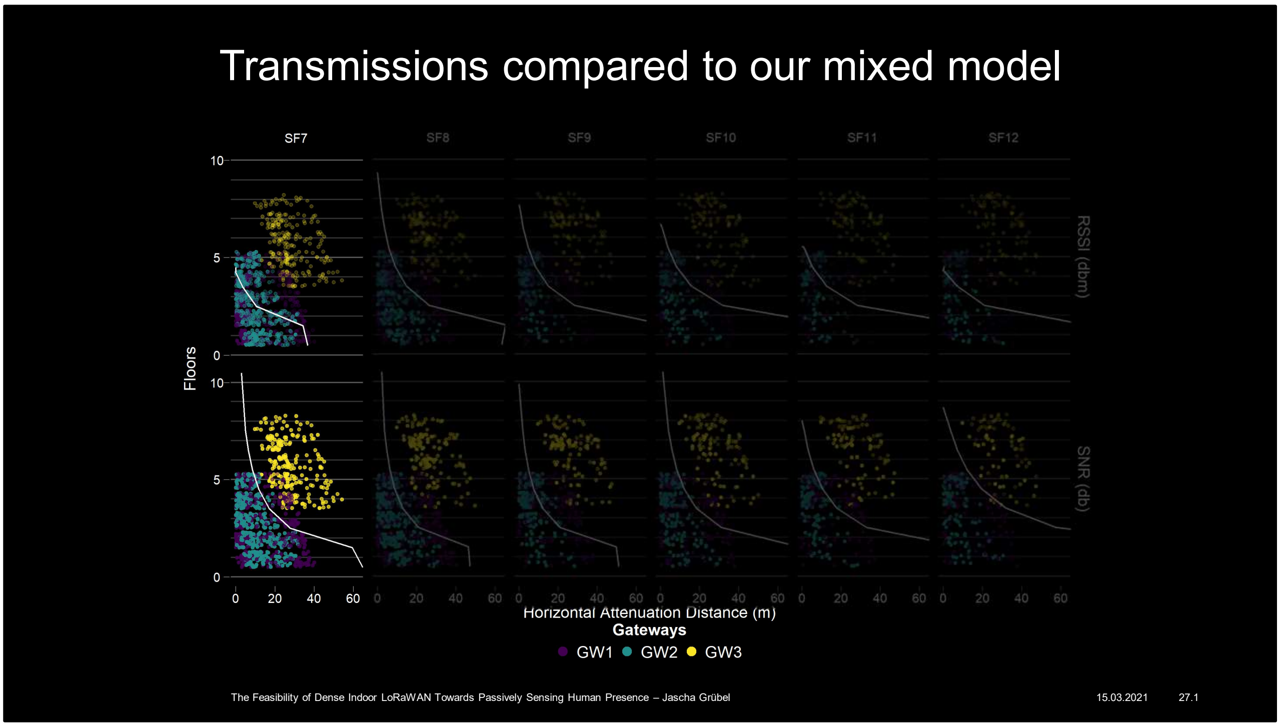

The points in this graph represent the sensor-gateway connections with the opacity representing the signal quality.
The y-axis is scaled by the floor and the x-axis by projected meters from the gateway located at (0,0).
The real sensor-gateway distance is represented by the diagonal but points are slightly jittered for visibility.
[C] They are facetted by SF, RSSI and SNR.
[C] Our mixed model predicts a conservative maximal transmission distance for each floor (left to the white line).
While several transmission were received beyond our threshold, they had a much lower sign quality.
Note how gateway 3 lost more transmissions and
[C] indeed is outside of our model prediction.
[C] A gateway every 30m and 5 floors ensures signal quality and redundancy – in short the 30x5 rule.

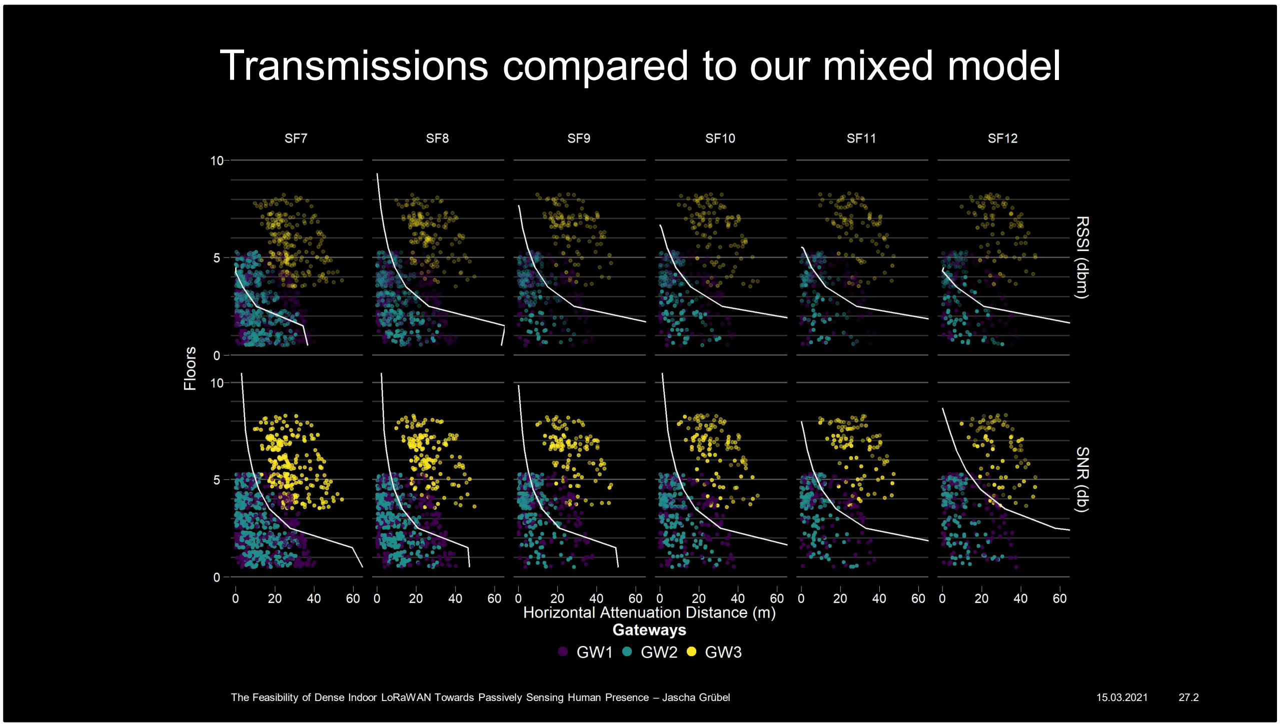

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

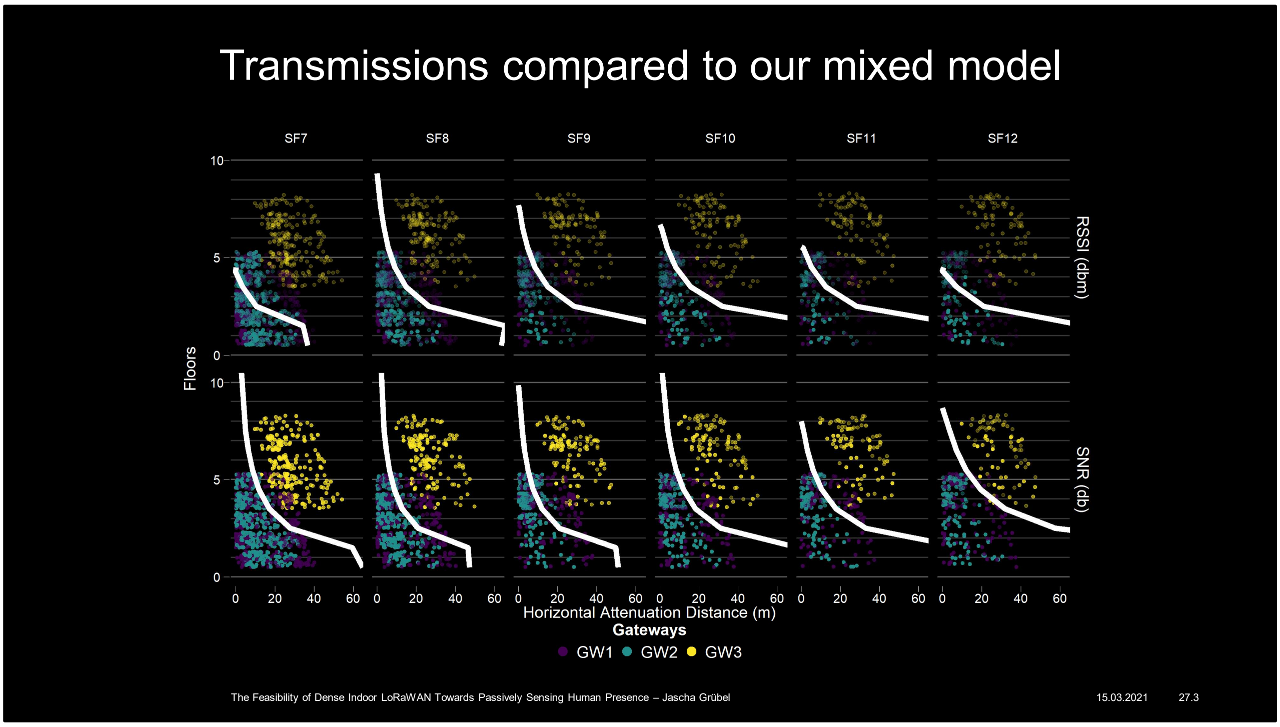

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

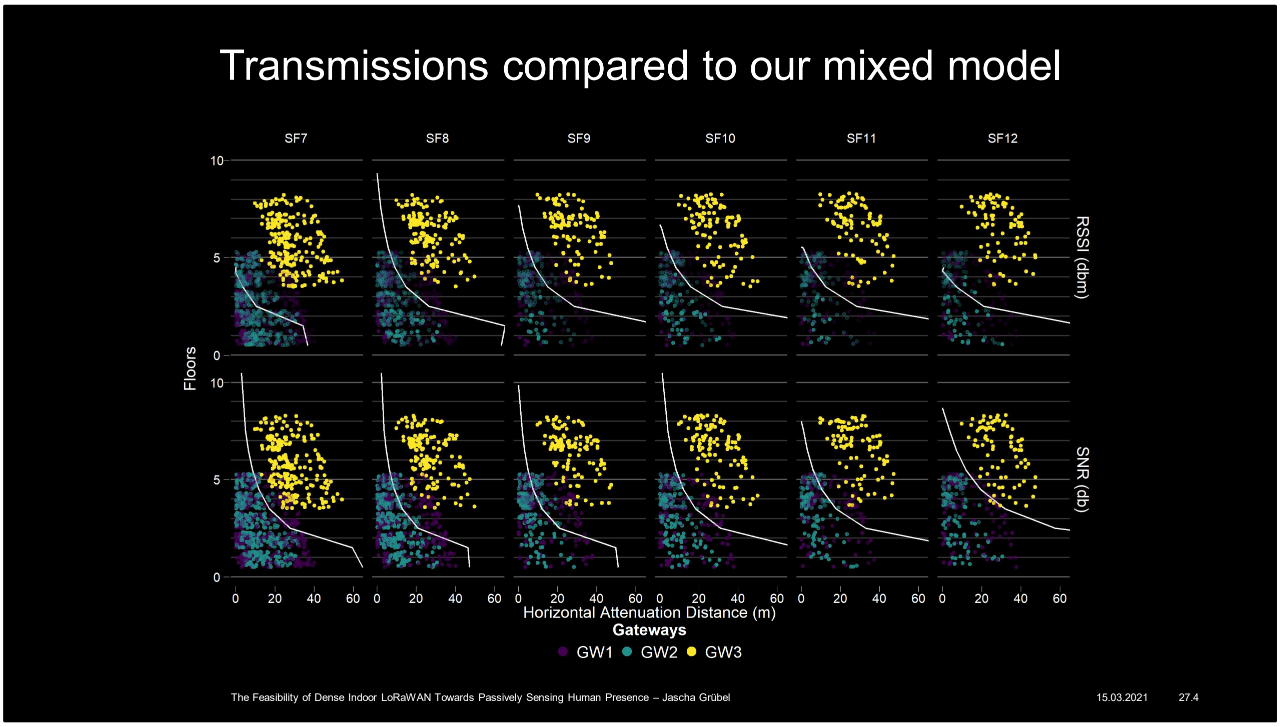

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

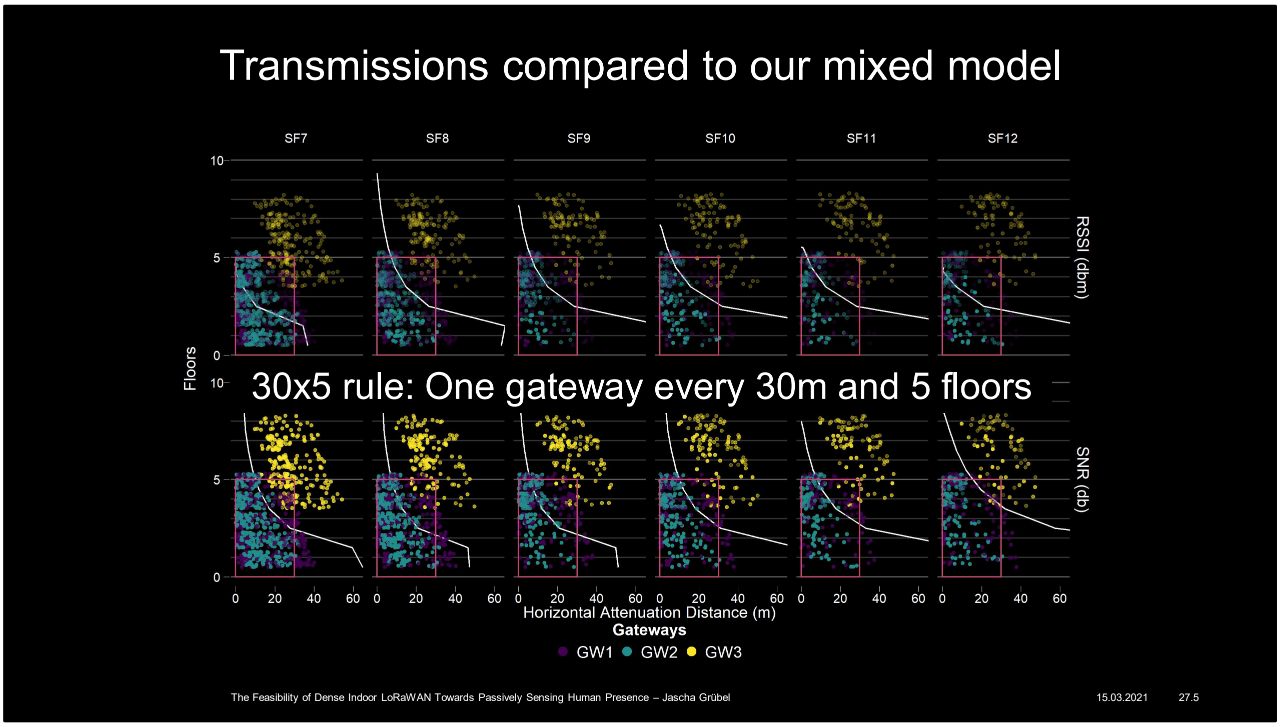

We use a mixed model to account for multiple transmissions from the same sensor.
[C]We also split the data by SF to account for the uneven number of transmissions per SF.
The formular represents a model with the indices i for the devices and j for the transmissions.

[C] The constant represents the signal quality at transmission at the device.
[C] The distance variable serves as a linear black box factor to account for signal shadowing and multi-path fading.
[C] The floor variable accounts for the ceilings and the interaction separates the floors from the distance.
[C] The concurrent variable accounts for destructive concurrent transmissions.
[C] The b variable groups all q transmissions from the same sensor device

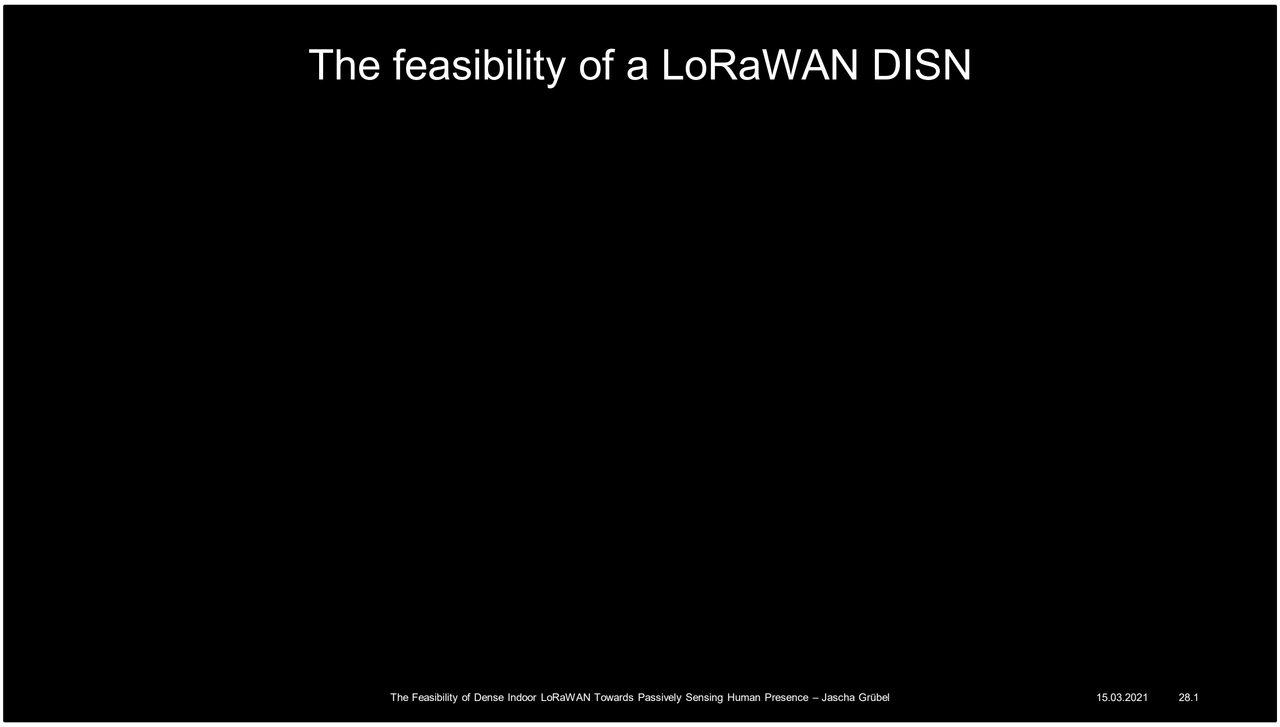

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

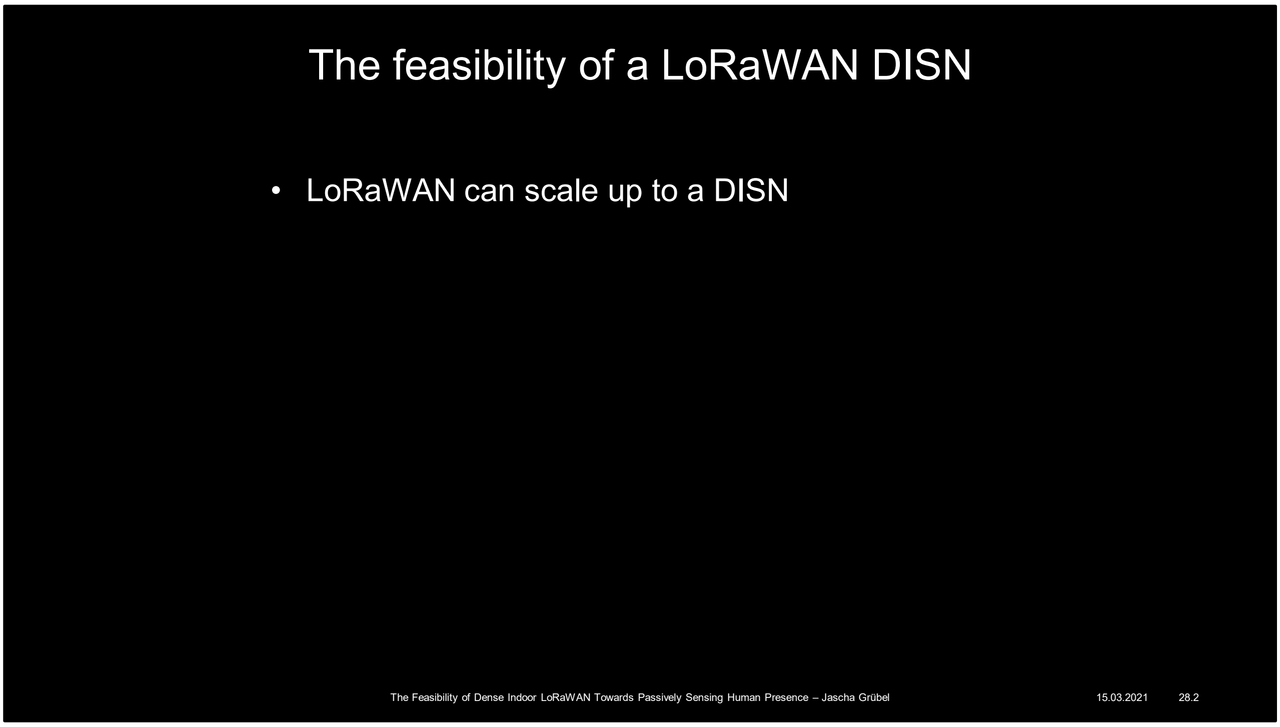

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

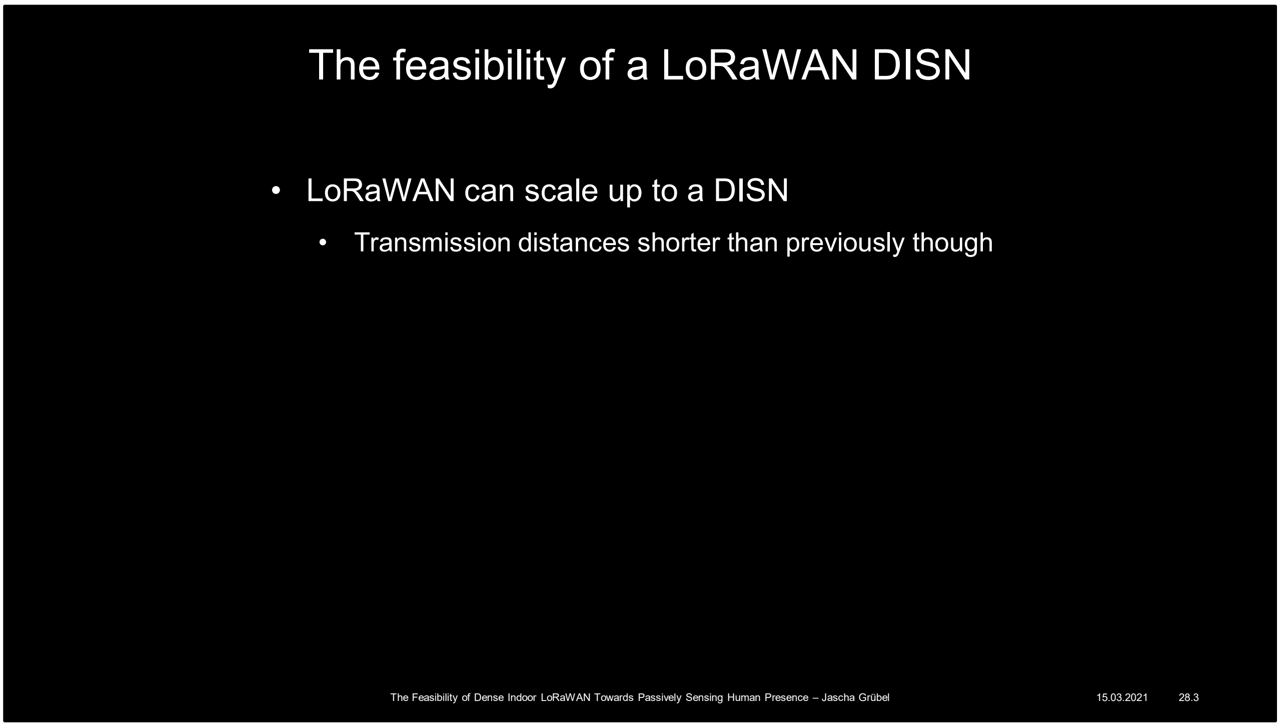

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

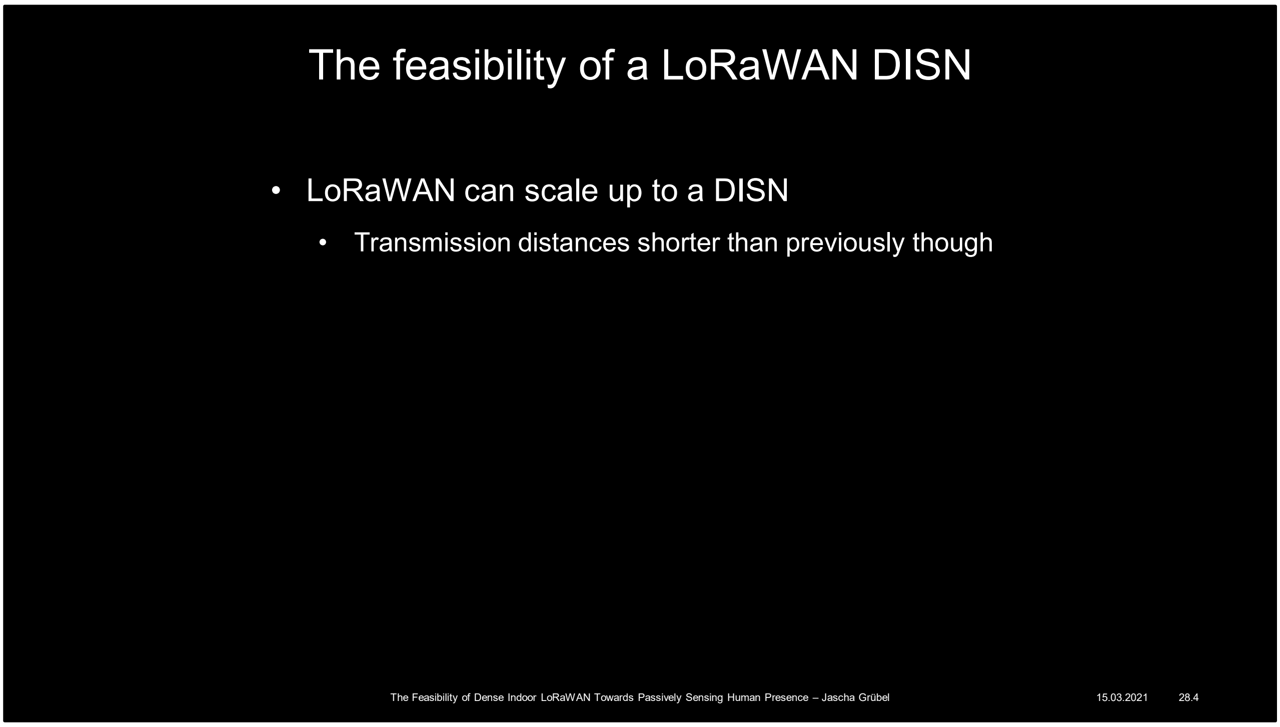

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

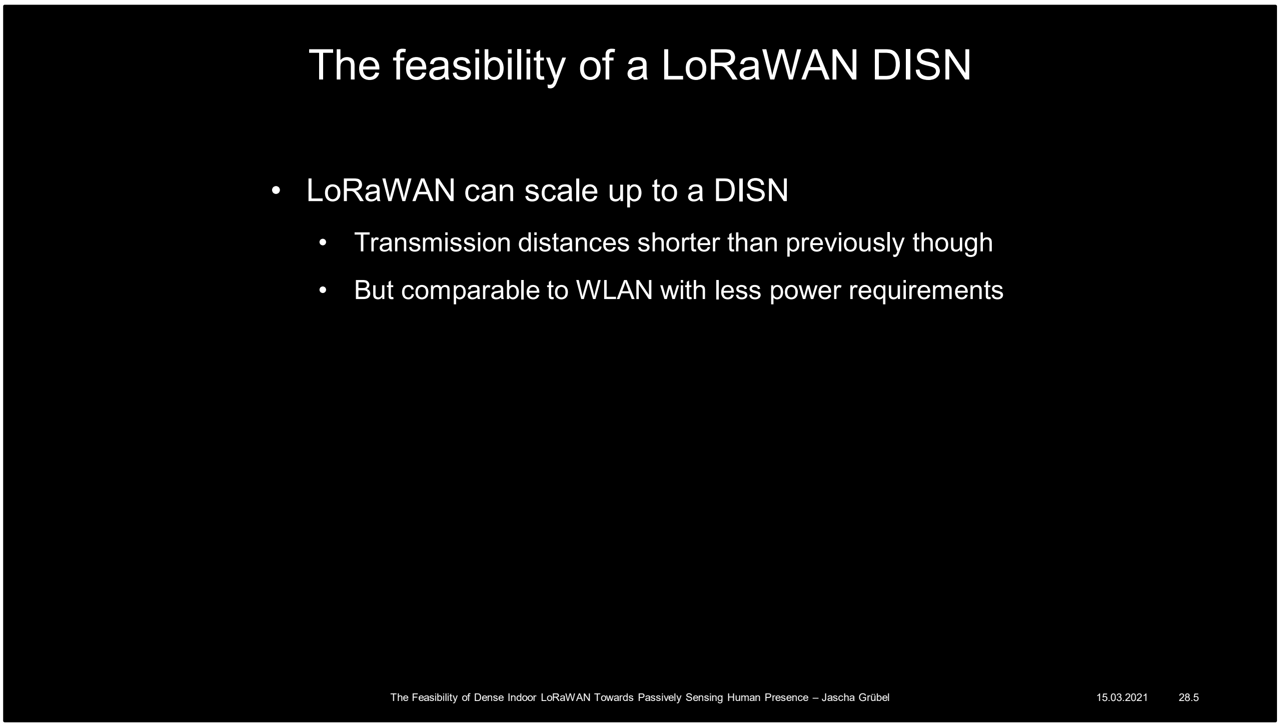

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

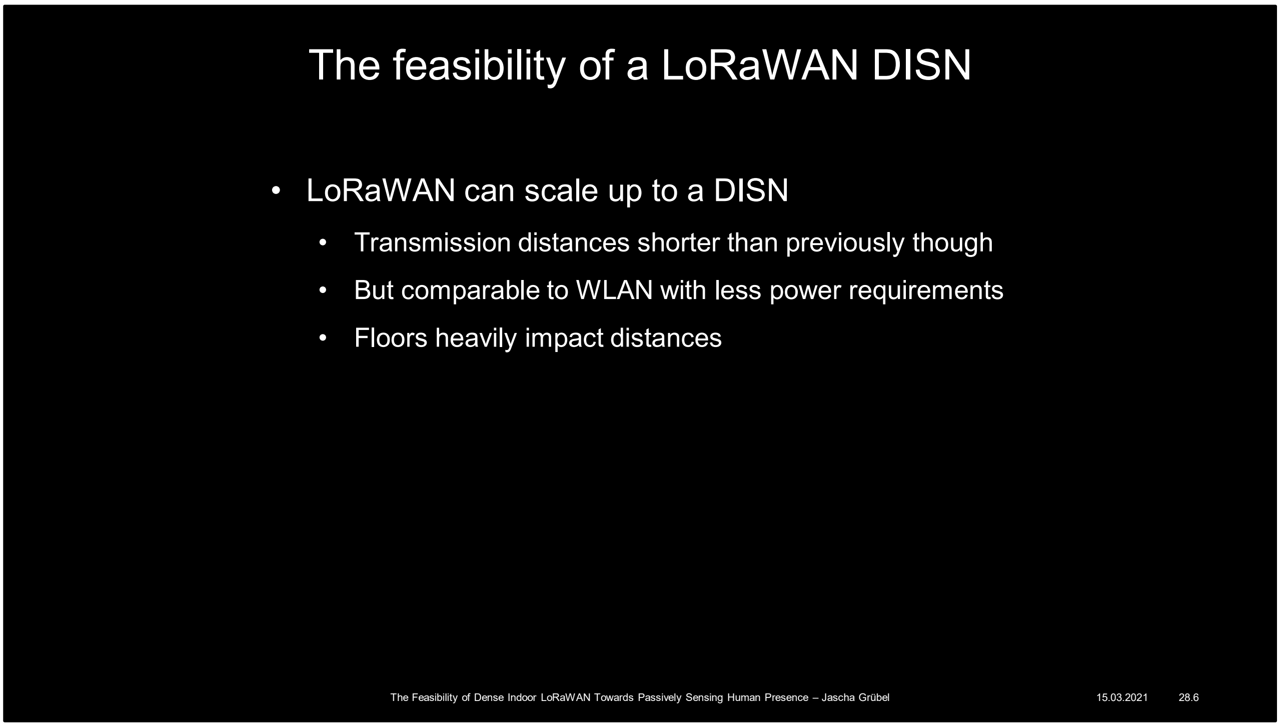

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

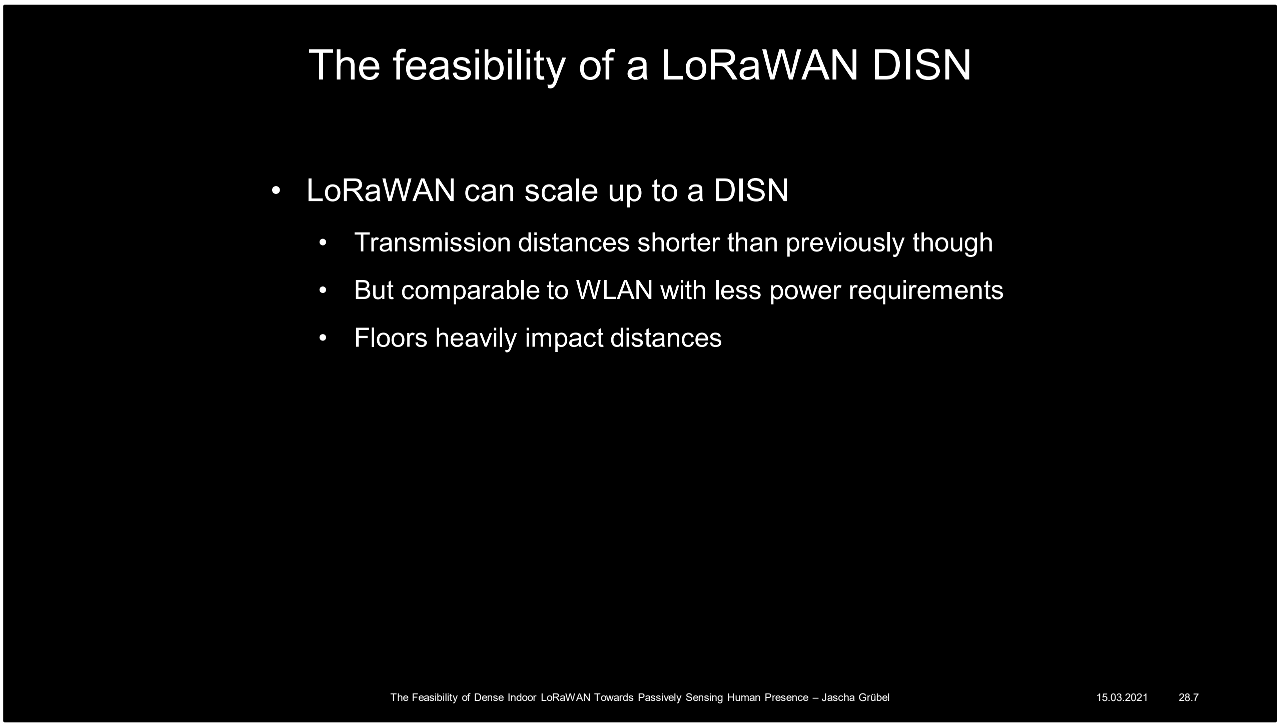

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

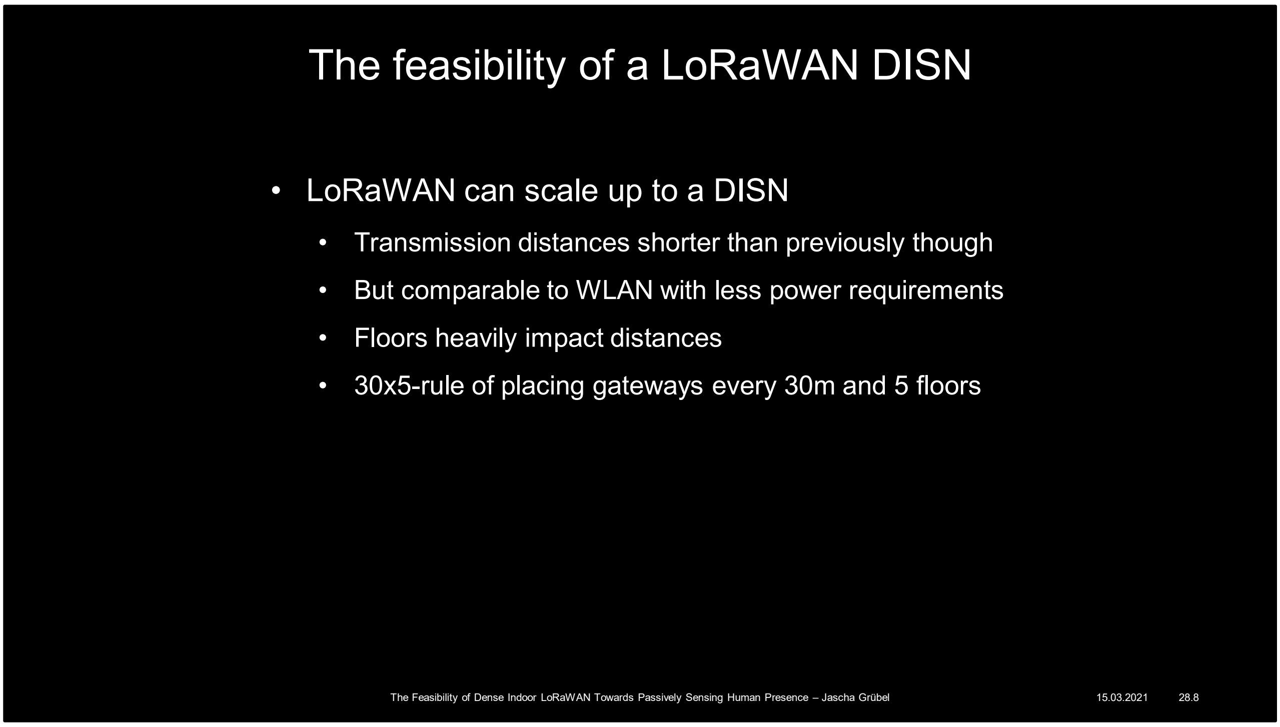

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

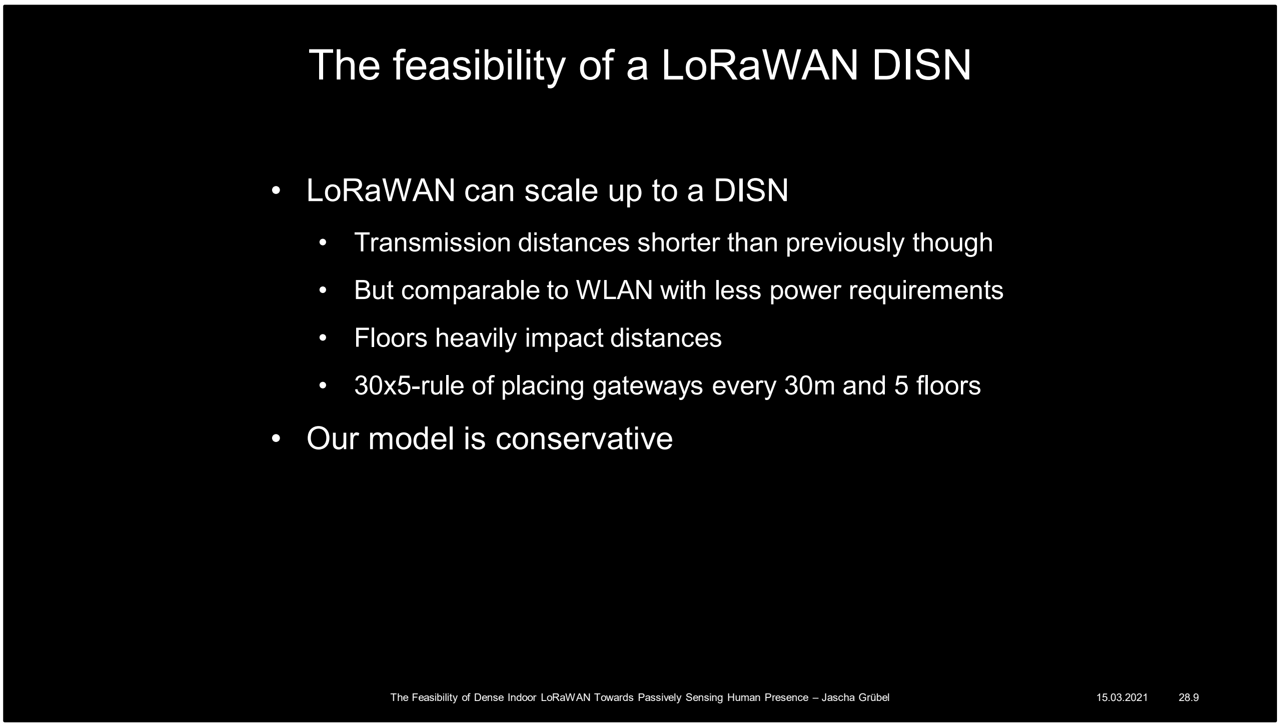

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

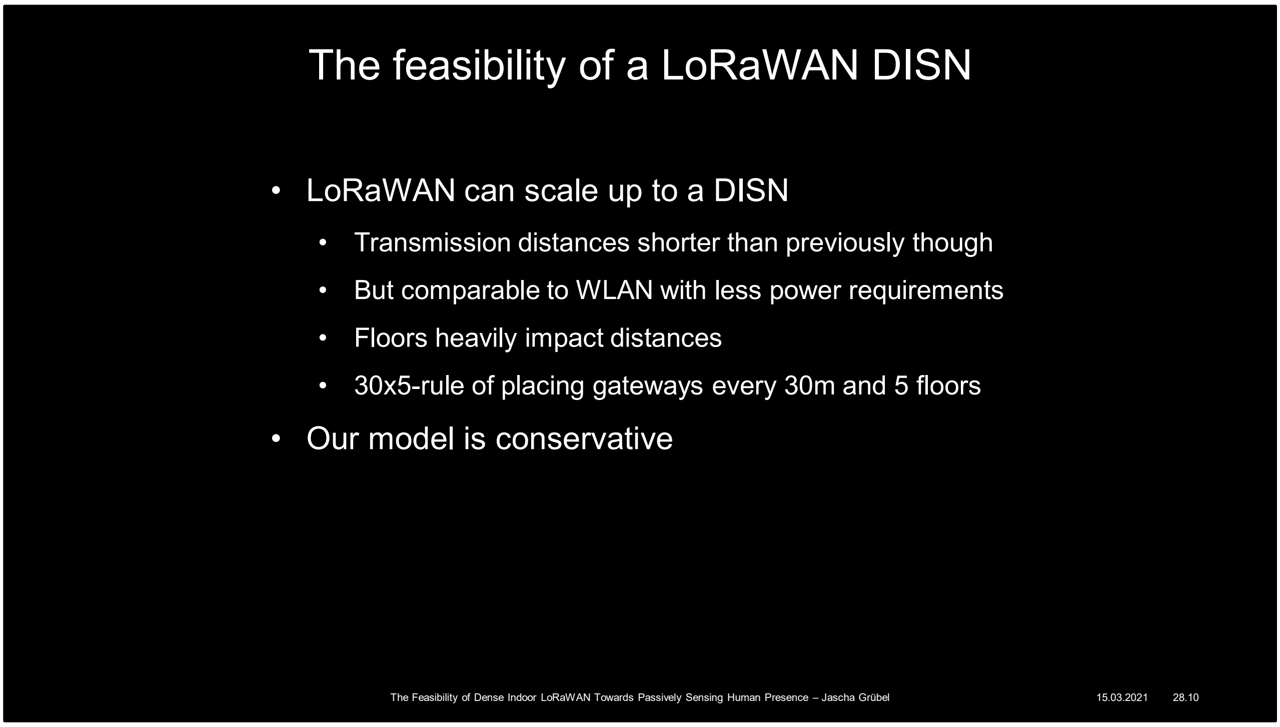

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

[C] We showed you how LoRaWAN can indeed scale up to the requirements of a dense indoor sensor network.
[C] We also showed that transmission distances indoors are much shorter than previous work claimed.
[C] However, it remains highly competitive with other technologies such as Wireless LAN that requires more power or Bluetooth that offers less coverage.
[C] In contrast to previous findings, the number of floors matter a lot.
[C] To sum up the results of our modelling, we recommend to use the 30x5-rule of placing one gateway every 30 meters and 5 floors.
[C] The model we used to arrive at these values is conservative because
[C] multipath-fading and signal shadowing are too complex to simulate or simplify.
[C] Therefore, in the real world, we actually achieve a better performance than our model.
But we want to ensure transmission quality and therefore stick with our conservative recommendation.
[C] We also investigated concurrent transmissions but in our setup they play no role yet.
Nonetheless, a trend is visible that may require action in even larger setups.

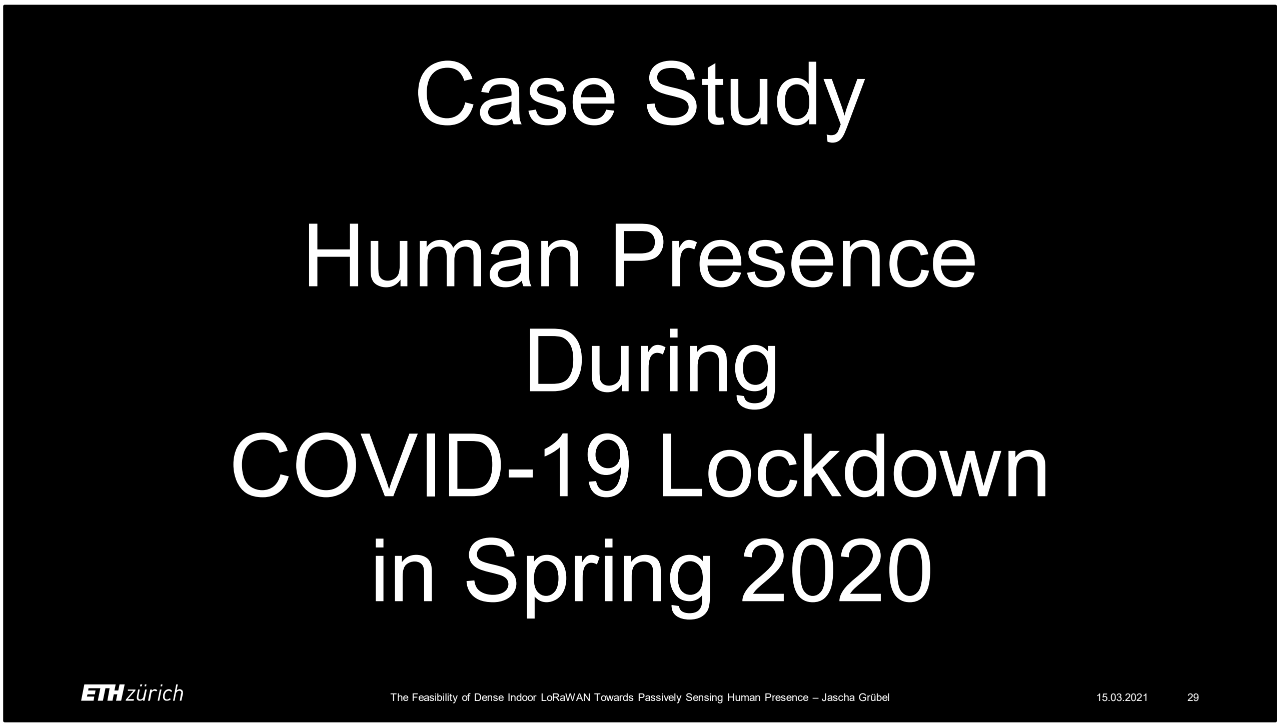

We established that LoRaWAN can be used to collect the data to study human presence.
However, we wanted to explore the data quality as well.
Since we do not have a ground truth for comparisons yet, we limit our selves to a small case study based on the Covid19-related lockdown in spring 2020.

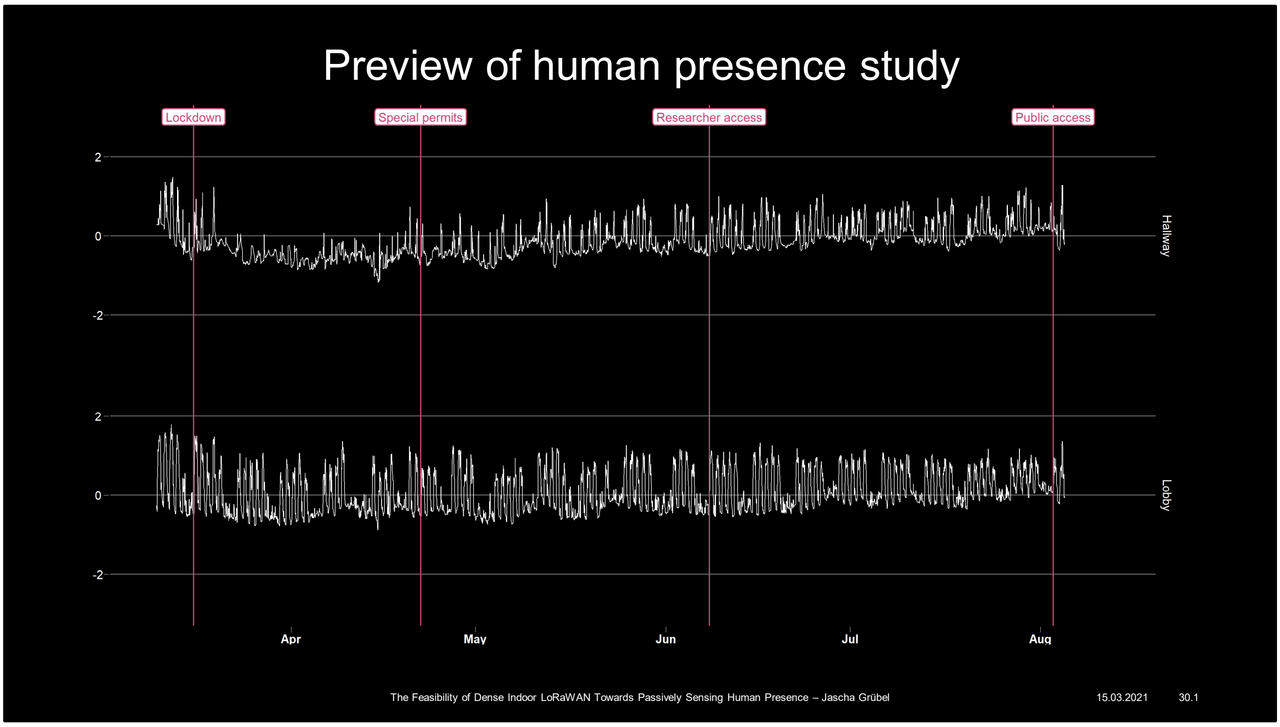

We take the COVID-19 induced lockdown as a natural experiment influencing a variance indicator that summarizes the 23 million data points by room and hour.
The lockdown removed nearly all human activity from the building.
A reduction in variance should thus be caused by the absence of humans.
In turn, the absent component is consequently caused by human presence and not by environmental factors.
We compare this variance in an office hallway with the permanently staffed lobby where essential personal resided.
[C] During the lockdown, variation in the hallway reduces strongly in contrast to the steady pattern in the lobby.
[C]Afterwards, the trend in the hallway slowly increases as lockdown measures ease.
In contrast, the lobby has nearly the same cyclical pattern throughout the whole observation period.
If you are interested on more details on the human presence data, we can discuss this in the Q and A session that follows.

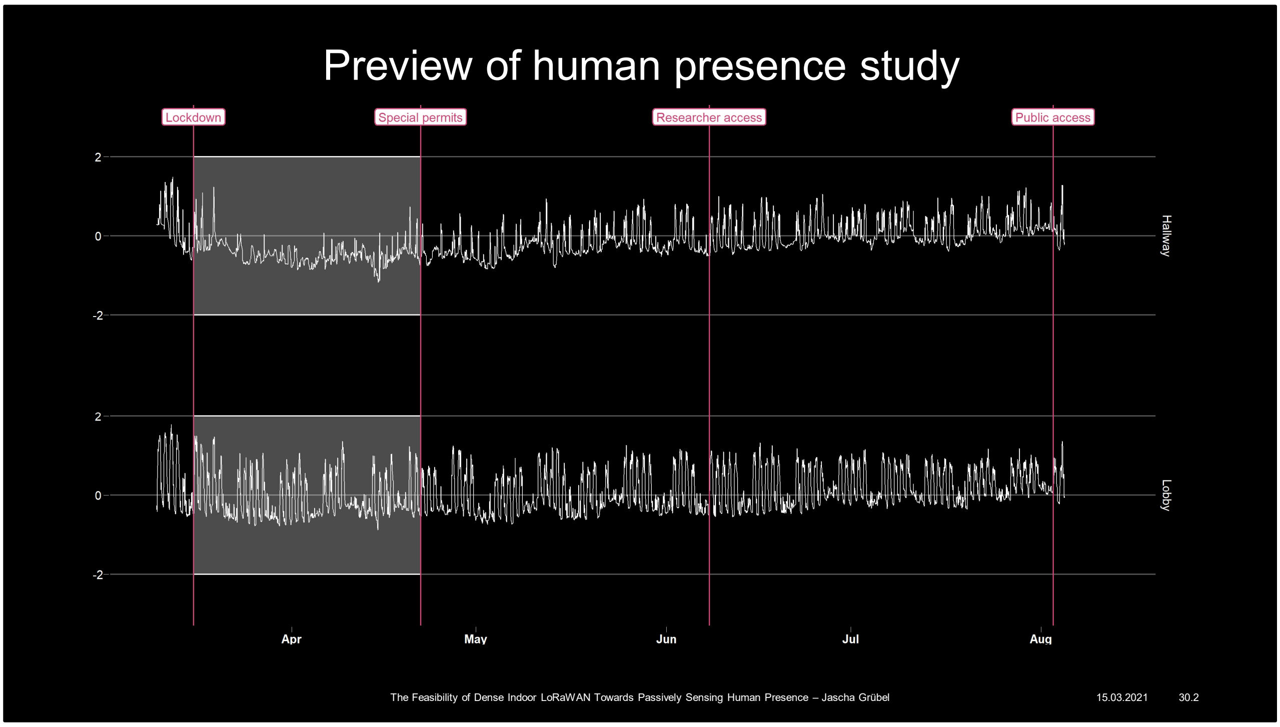

We take the COVID-19 induced lockdown as a natural experiment influencing a variance indicator that summarizes the 23 million data points by room and hour.
The lockdown removed nearly all human activity from the building.
A reduction in variance should thus be caused by the absence of humans.
In turn, the absent component is consequently caused by human presence and not by environmental factors.
We compare this variance in an office hallway with the permanently staffed lobby where essential personal resided.
[C] During the lockdown, variation in the hallway reduces strongly in contrast to the steady pattern in the lobby.
[C]Afterwards, the trend in the hallway slowly increases as lockdown measures ease.
In contrast, the lobby has nearly the same cyclical pattern throughout the whole observation period.
If you are interested on more details on the human presence data, we can discuss this in the Q and A session that follows.

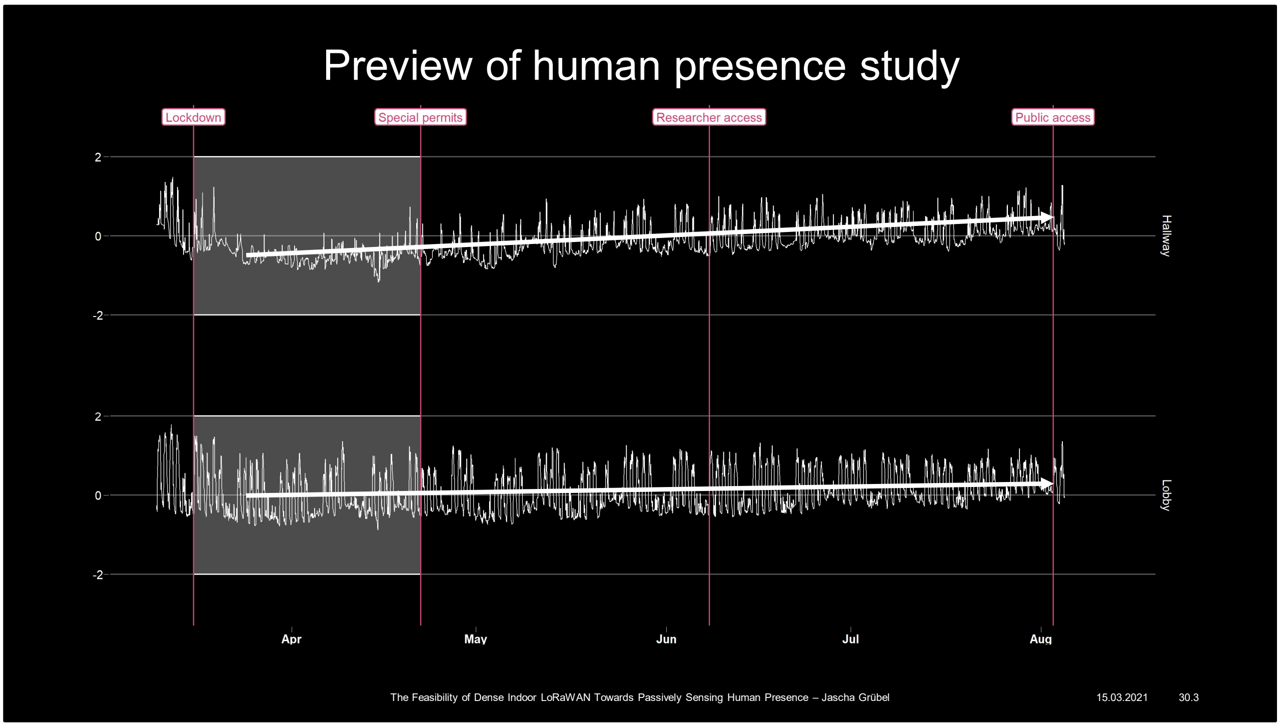

We take the COVID-19 induced lockdown as a natural experiment influencing a variance indicator that summarizes the 23 million data points by room and hour.
The lockdown removed nearly all human activity from the building.
A reduction in variance should thus be caused by the absence of humans.
In turn, the absent component is consequently caused by human presence and not by environmental factors.
We compare this variance in an office hallway with the permanently staffed lobby where essential personal resided.
[C] During the lockdown, variation in the hallway reduces strongly in contrast to the steady pattern in the lobby.
[C]Afterwards, the trend in the hallway slowly increases as lockdown measures ease.
In contrast, the lobby has nearly the same cyclical pattern throughout the whole observation period.
If you are interested on more details on the human presence data, we can discuss this in the Q and A session that follows.

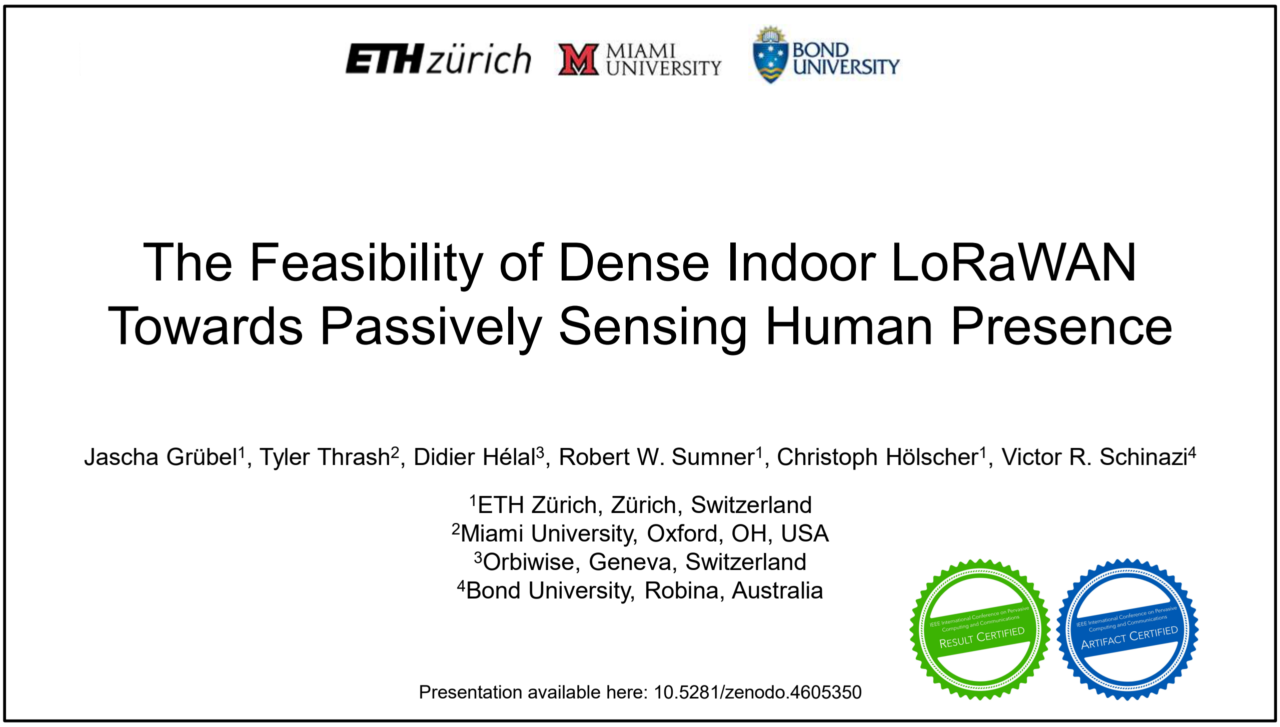

To conclude, our proof of concept goes beyond simulations to demonstrate the effective application of LoRaWAN indoors and offers us a way towards passively sensing human presence indoors.
We also introduce the 30x5 rule for how to place gateways indoors for dense indoor sensor networks to ensure signal quality and redundancy.

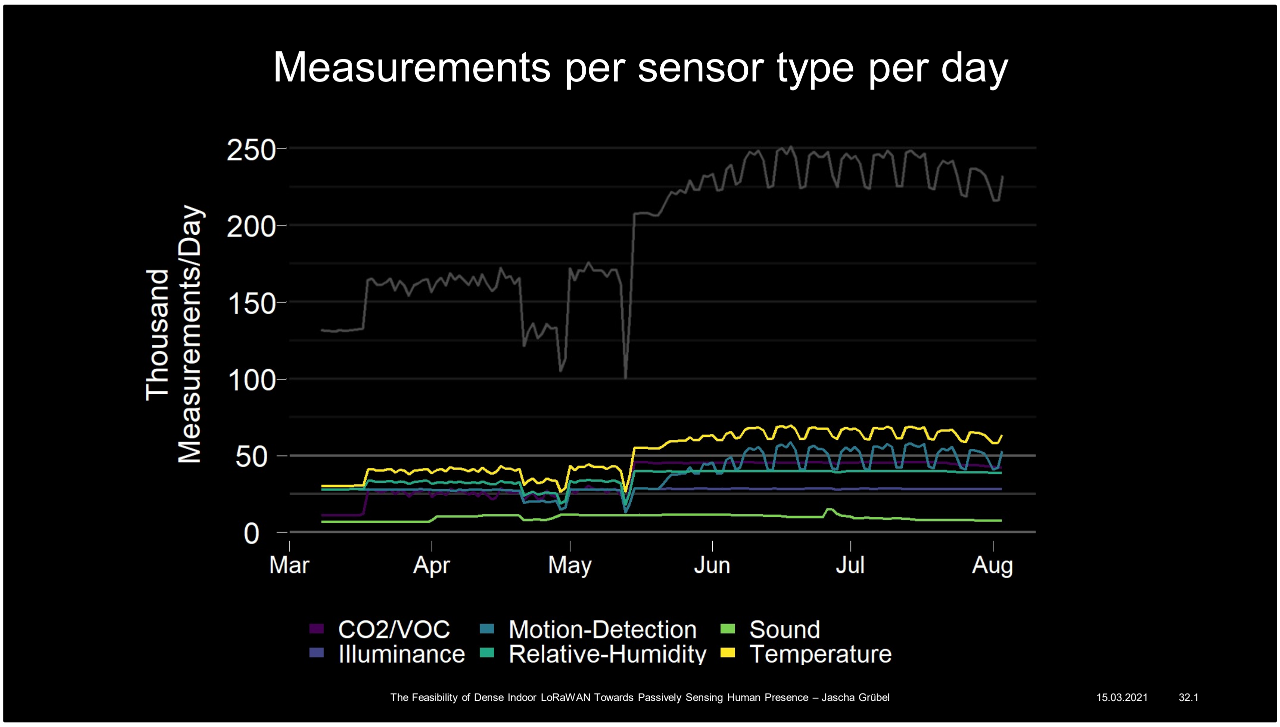

Let me start by giving you an overview of the number of measurements by sensor type and day during the 5 months of data collection for this paper.
Sensors are grouped and colored by 6 types listed at the bottom.
The sum of all measurements is shown in white
The ramp-up of the data collection was delayed as installation had to take place in stages due to the Covid10-related lockdown.
Currently, the system collect up to 250 thousand data points per day since March 2020 and probably at least until December 2021.

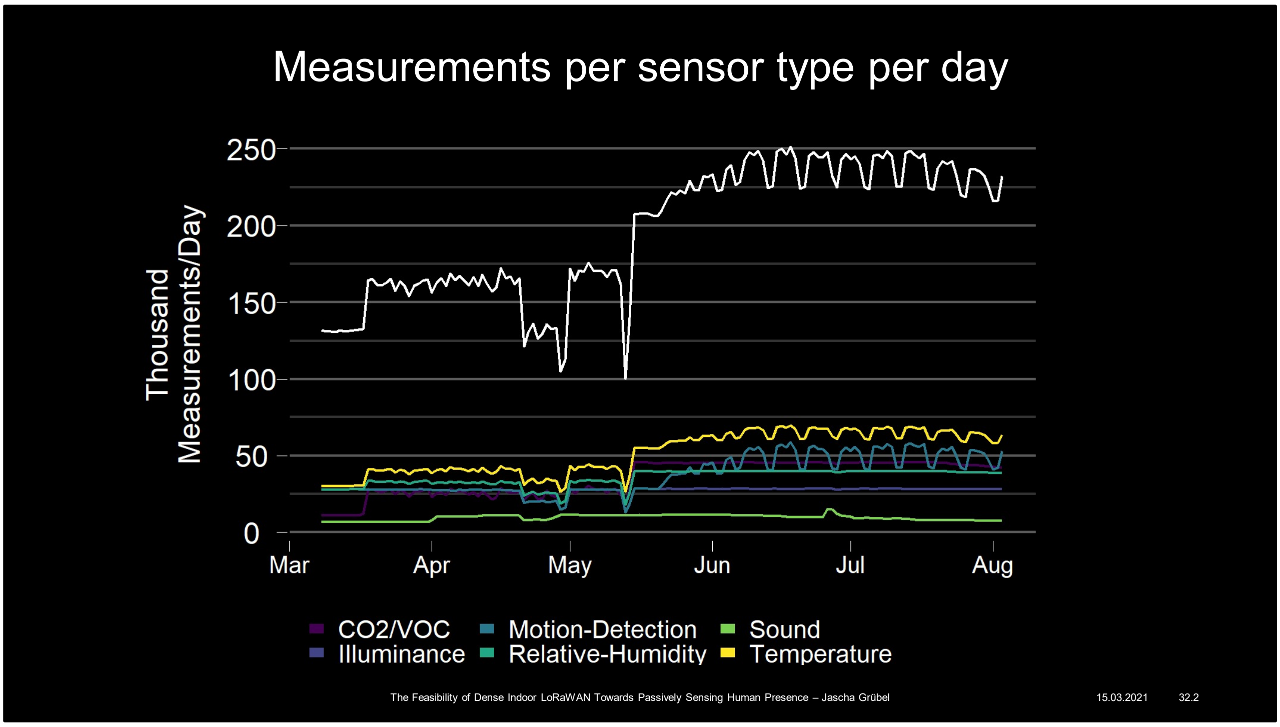

Let me start by giving you an overview of the number of measurements by sensor type and day during the 5 months of data collection for this paper.
Sensors are grouped and colored by 6 types listed at the bottom.
The sum of all measurements is shown in white
The ramp-up of the data collection was delayed as installation had to take place in stages due to the Covid10-related lockdown.
Currently, the system collect up to 250 thousand data points per day since March 2020 and probably at least until December 2021.

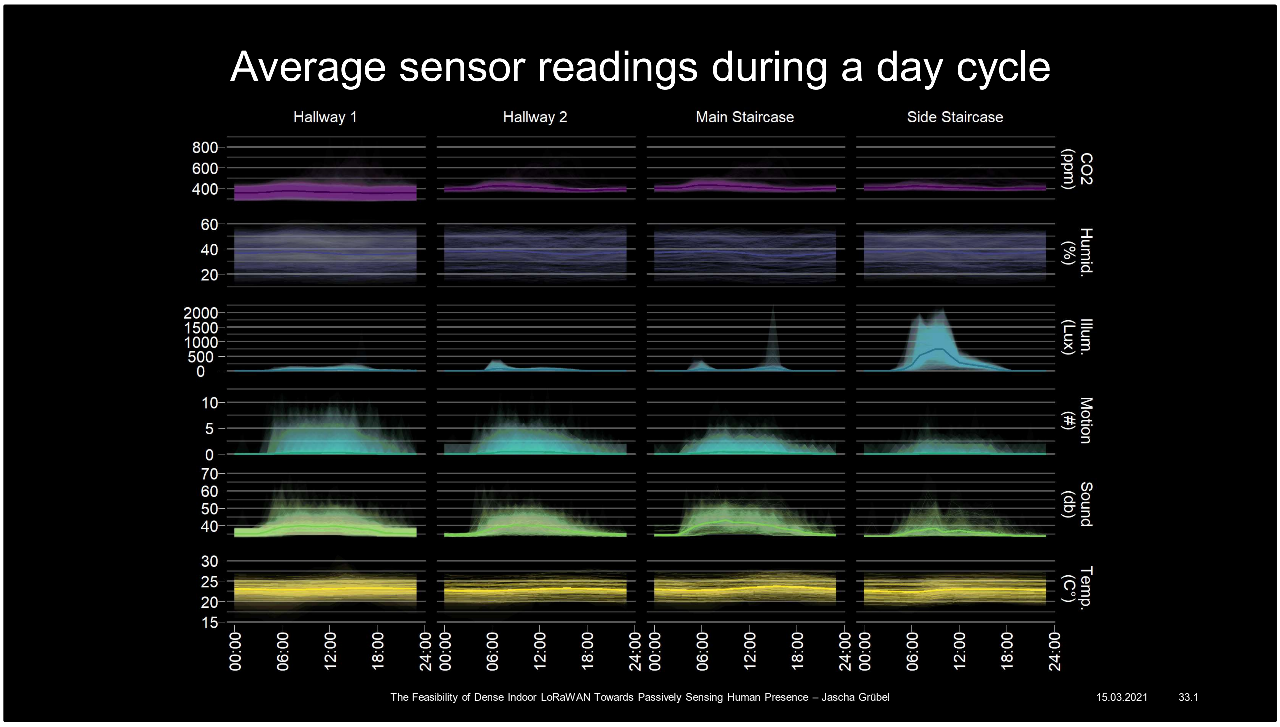

Before we look at the connection to human presence, let's peek at the data variation.
Here we show the average sensor readings for selected rooms and sensor types.
Each line represents a single day's sensor reading.
The thick line represents the average day.
The variation may be driven by both human activity and environmental factors.
For instance, the spike in illuminance in the side staircase is caused by the afternoon sun and not the hallway lamps.

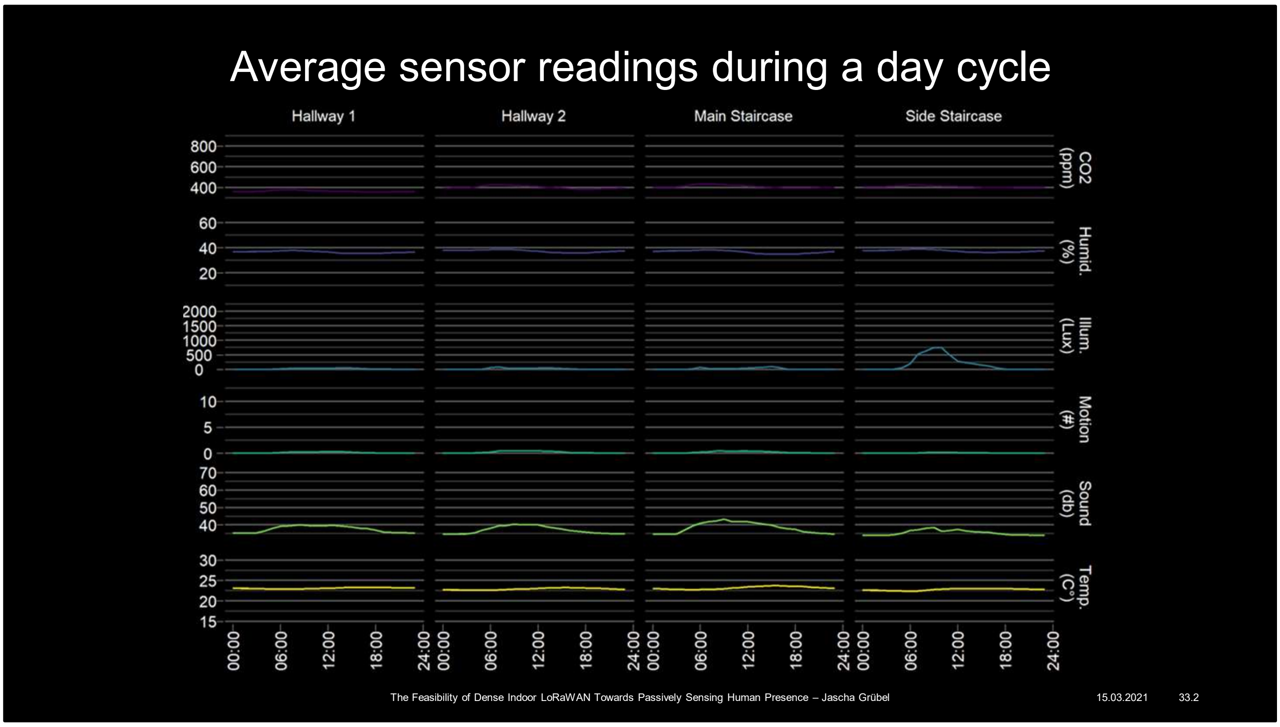

Before we look at the connection to human presence, let's peek at the data variation.
Here we show the average sensor readings for selected rooms and sensor types.
Each line represents a single day's sensor reading.
The thick line represents the average day.
The variation may be driven by both human activity and environmental factors.
For instance, the spike in illuminance in the side staircase is caused by the afternoon sun and not the hallway lamps.

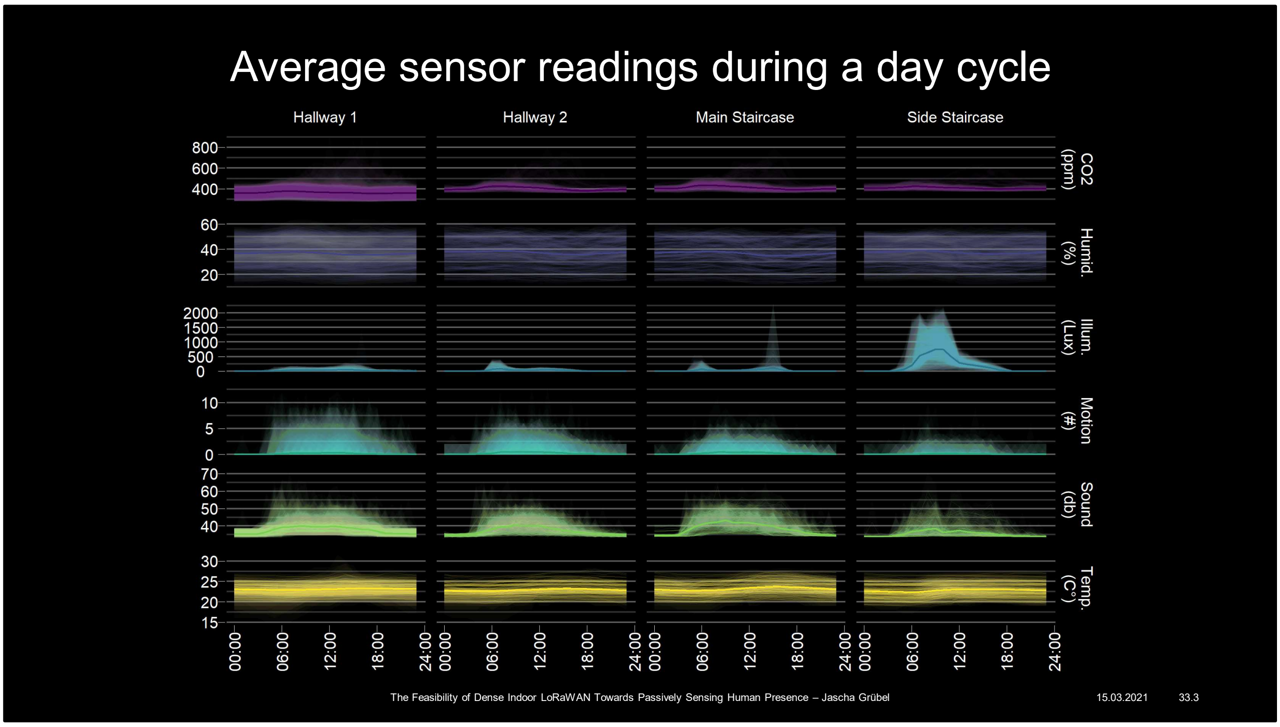

Before we look at the connection to human presence, let's peek at the data variation.
Here we show the average sensor readings for selected rooms and sensor types.
Each line represents a single day's sensor reading.
The thick line represents the average day.
The variation may be driven by both human activity and environmental factors.
For instance, the spike in illuminance in the side staircase is caused by the afternoon sun and not the hallway lamps.

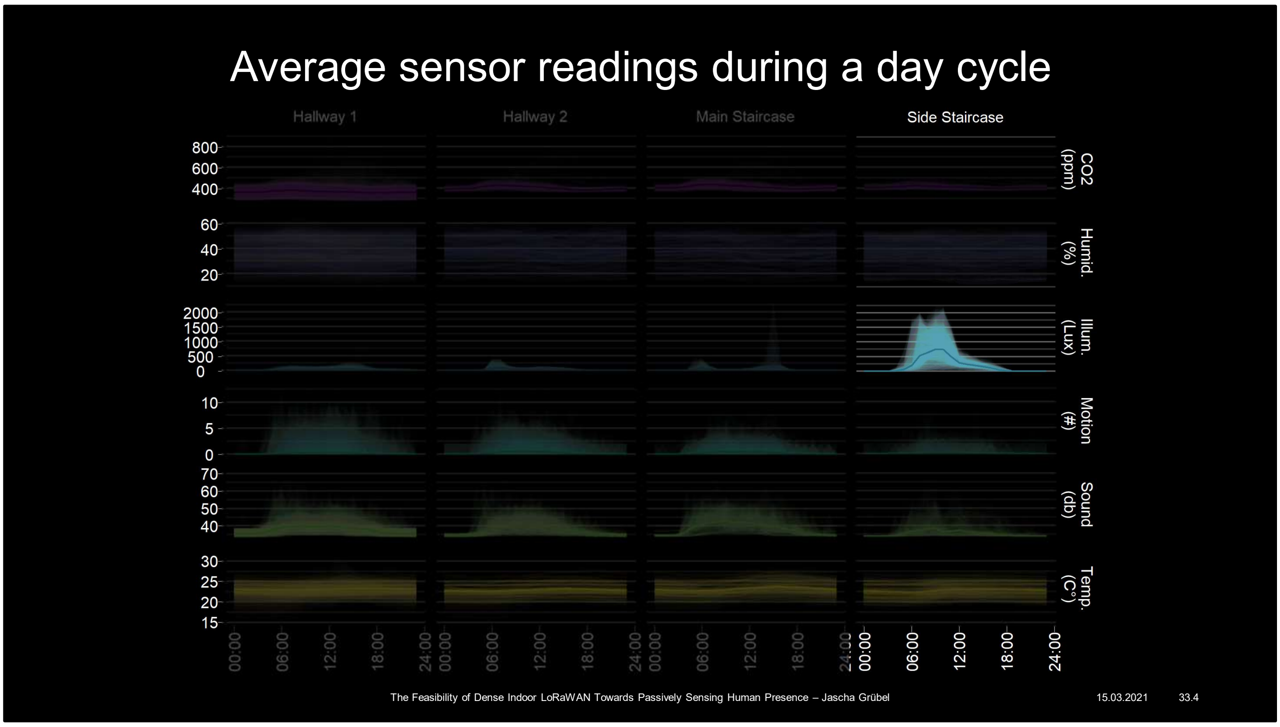

Before we look at the connection to human presence, let's peek at the data variation.
Here we show the average sensor readings for selected rooms and sensor types.
Each line represents a single day's sensor reading.
The thick line represents the average day.
The variation may be driven by both human activity and environmental factors.
For instance, the spike in illuminance in the side staircase is caused by the afternoon sun and not the hallway lamps.

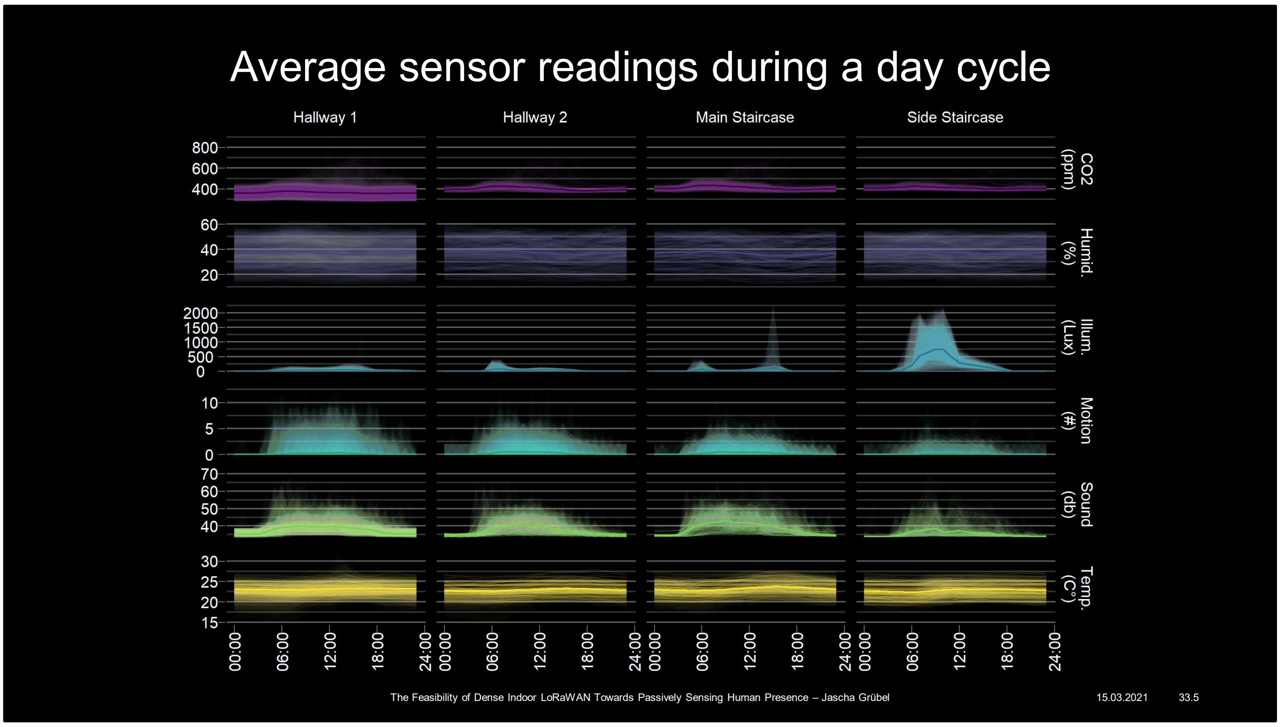

Before we look at the connection to human presence, let's peek at the data variation.
Here we show the average sensor readings for selected rooms and sensor types.
Each line represents a single day's sensor reading.
The thick line represents the average day.
The variation may be driven by both human activity and environmental factors.
For instance, the spike in illuminance in the side staircase is caused by the afternoon sun and not the hallway lamps.

Variance indicator

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

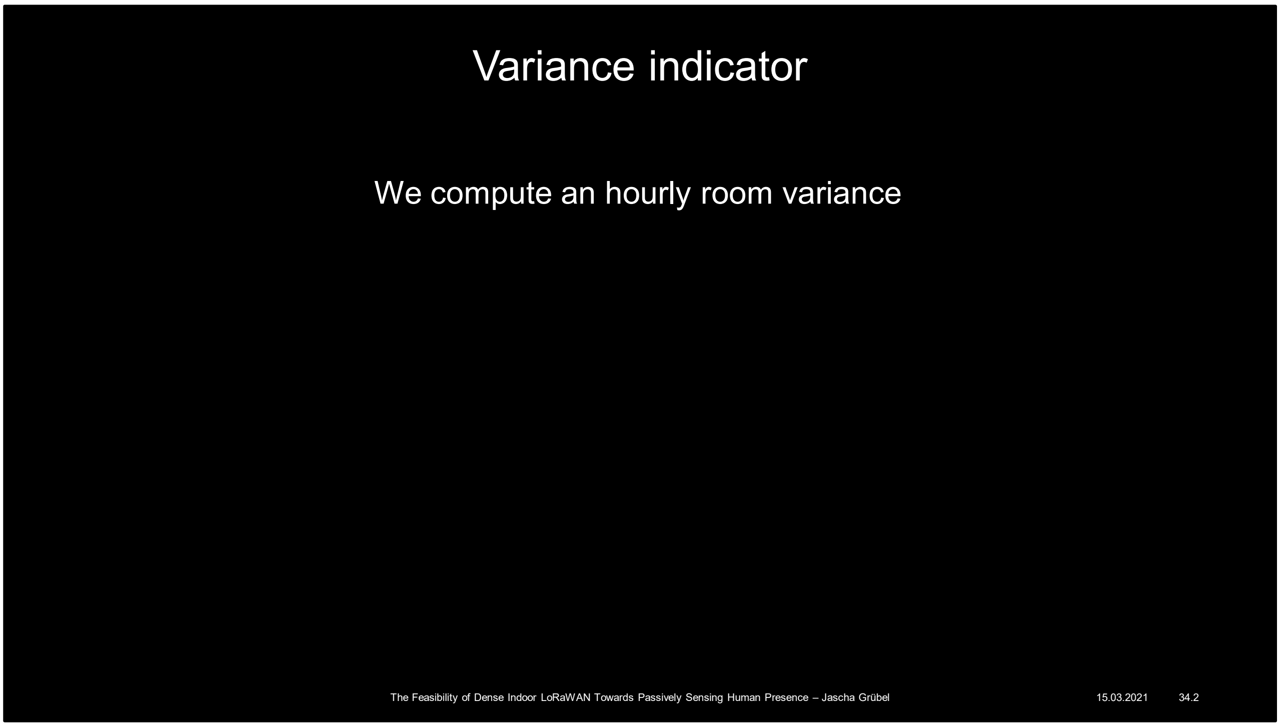

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

Variance indicator

We compute an hourly room variance

The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence – Jascha Grübel 15.03.2021 34.3

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

Variance indicator

We compute an hourly room variance

Per room (r) per hour (t) per sensor (i) z-score: $\widetilde{\mathbb{V}}_{r,i,t} = \frac{1}{n_{r,i}} \sum_{j,k} \frac{v_{r,i,j,t_k} - \overline{v}_i}{\sigma_{v_i}}$

The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence – Jascha Grübel 15.03.2021 34.4

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

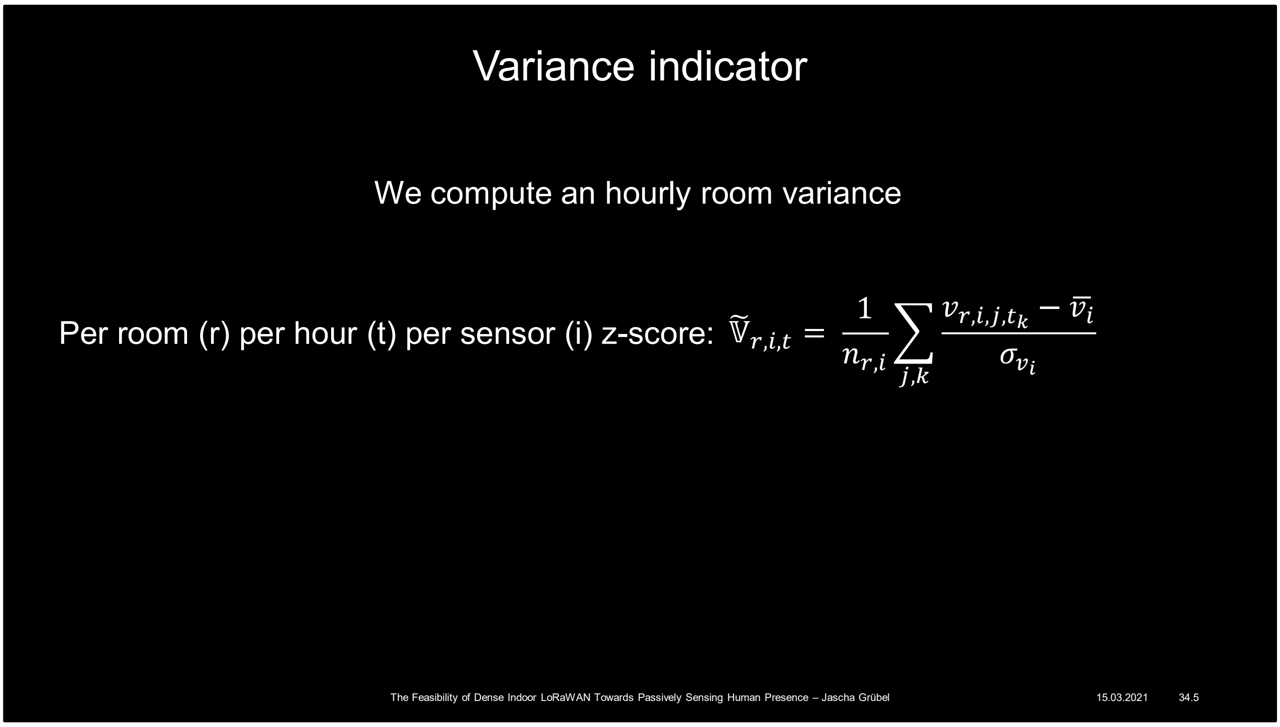

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

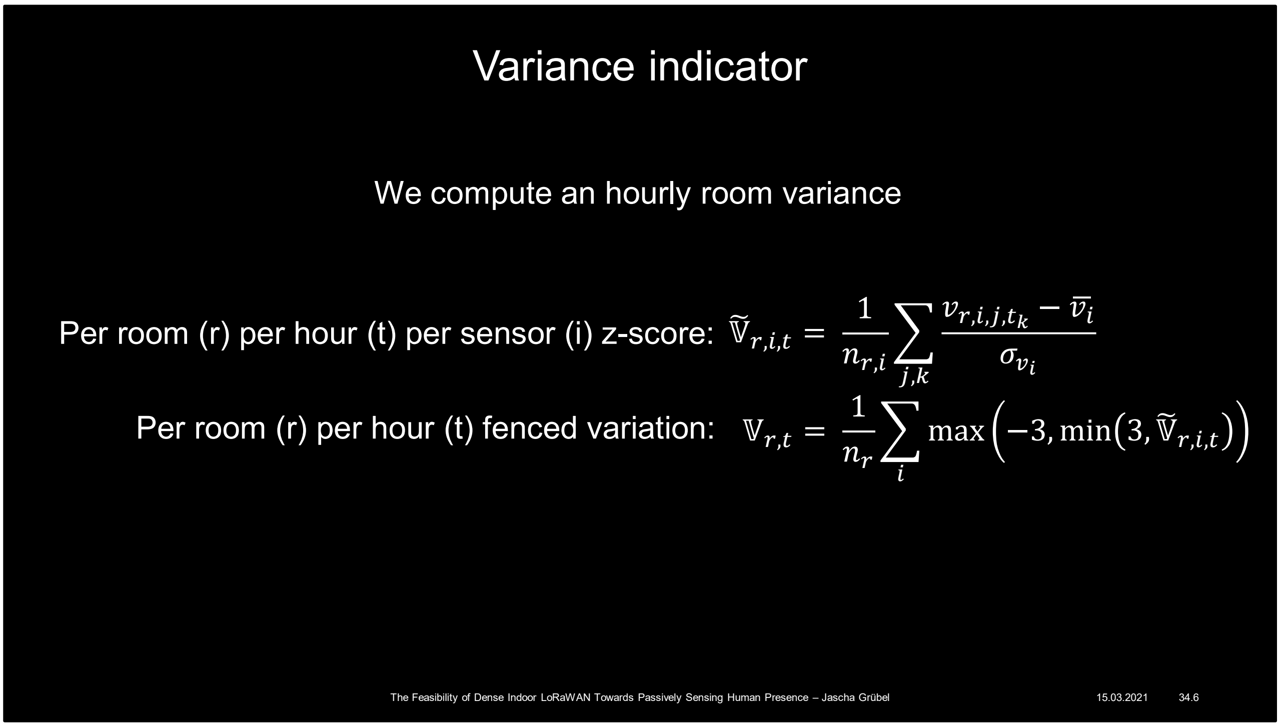

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

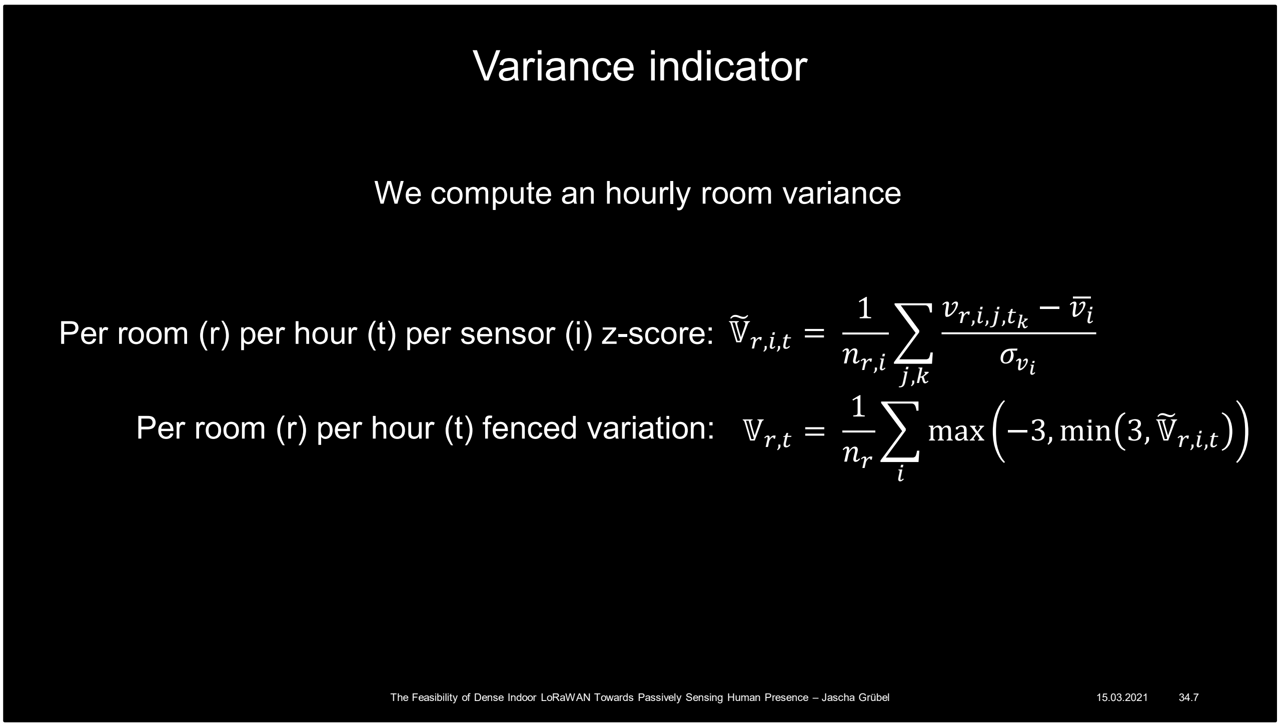

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

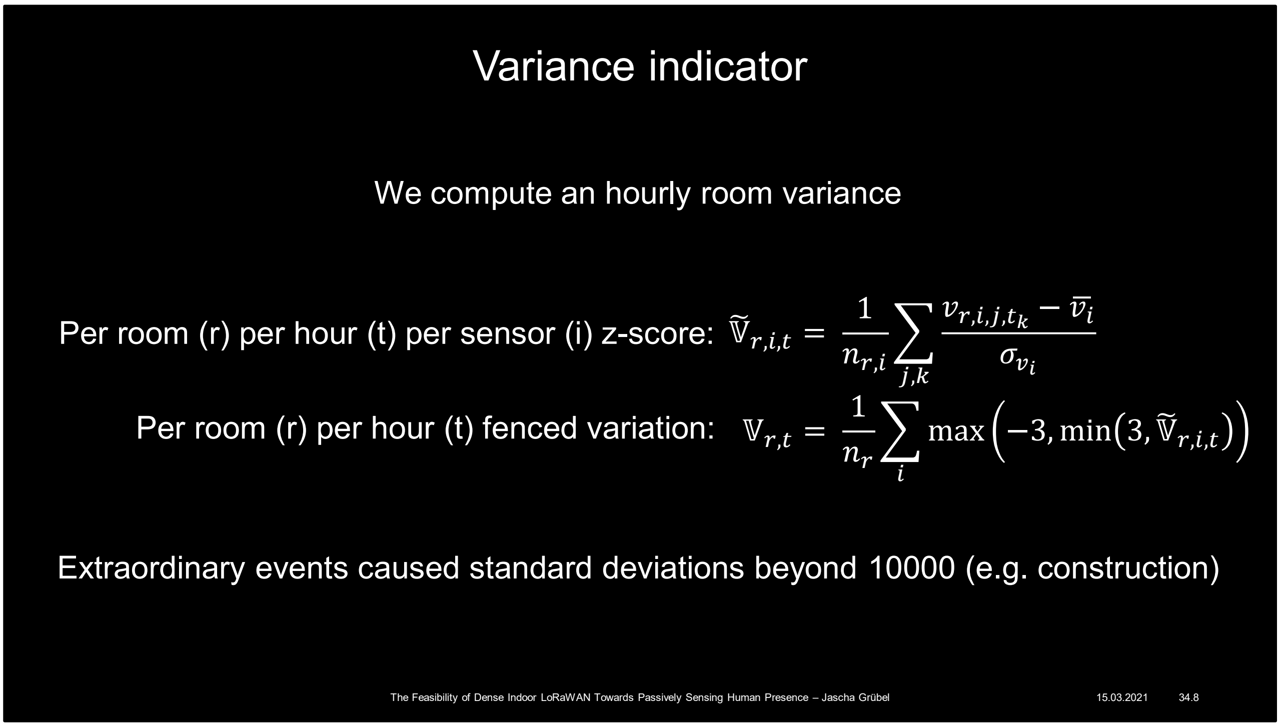

To understand the overall variation, we construct a variance indicator that summarized the sensor data every hour by room across all sensors.
We z-score the variance with a per sensor type mean.
We also fence the data to three standard deviations because extraordinary events like construction sometimes drove the standard deviation beyond 10000.

Mixed Model Fixed Effects coefficients

TABLE IV
MIXED MODELS TO DETERMINE SNR AND RSSI FOR EACH DATA RATE

	SF7		SF8		SF9		SF10		SF11		SF12	
	SNR	RSSI	SNR	RSSI	SNR	RSSI	SNR	RSSI	SNR	RSSI	SNR	RSSI
Average signal strength at node (constant)	10.03* (0.09)	−63.75* (0.46)	10.35* (0.23)	−94.75* (0.61)	10.00* (0.29)	−97.33* (0.64)	8.07* (0.35)	−87.07* (0.91)	10.45* (0.36)	−85.05* (0.90)	8.38* (0.43)	−60.35* (0.77)
Attenuation in m	−0.27* (0.00)	−1.61* (0.00)	−0.43* (0.00)	−0.49* (0.00)	−0.44* (0.00)	−0.45* (0.00)	−0.32* (0.00)	−0.55* (0.00)	−0.34* (0.01)	−0.61* (0.01)	−0.20* (0.01)	−1.09* (0.01)
One floor between	−0.26* (0.00)	−9.48* (0.01)	0.04 (0.03)	−1.92* (0.03)	−0.92* (0.04)	−2.89* (0.03)	−0.96* (0.05)	−5.29* (0.04)	−2.50* (0.07)	−7.10* (0.06)	−2.68* (0.08)	−14.86* (0.08)
Floor-attenuation interaction	−0.01* (0.00)	0.29* (0.00)	0.02* (0.00)	0.07* (0.00)	0.04* (0.00)	0.08* (0.00)	0.02* (0.00)	0.12* (0.00)	0.04* (0.00)	0.15* (0.00)	0.02* (0.00)	0.29* (0.00)
Colliding message	0.13* (0.01)	1.57* (0.05)	0.17 (0.28)	0.67 (0.28)	0.24 (0.44)	0.66 (0.34)	0.77 (0.43)	0.36 (0.36)	1.37* (0.29)	0.27 (0.25)	0.63 (0.21)	0.27 (0.21)
Observations	12,948,936	12,948,936	495,967	495,967	246,184	246,184	160,618	160,618	104,499	104,499	106,025	106,025

*p<0.001

The Feasibility of Dense Indoor LoRaWAN Towards Passively Sensing Human Presence – Jascha Grübel 15.03.2021 35

Most transmissions arrived, only few are lost as many could be restored from neighbouring gateways. Multiple gateways improve reception but have diminishing returns beyond 30m.

Mixed Model Maximal Transmission Distances

TABLE V
SNR AND RSSI LIMITS FOR TRANSMISSION IN OUR SYSTEM BASED ON MIXED MODELS

	SNR						RSSI					
Max. reachable	SF7	SF8	SF9	SF10	SF11	SF12	SF7	SF8	SF9	SF10	SF11	SF12
Distance 0 floors up	64.88 m	47.29 m	50.96 m	72.6 m	83.34 m	141.84 m	36.81 m	63.49 m	71.1 m	82.14 m	79.3 m	69.65 m
Distance 1 floor up	59.5 m	46.39 m	49.73 m	70.87 m	82.73 m	141.59 m	34.36 m	65.4 m	74.49 m	88.78 m	85.46 m	72.69 m
Distance 2 floors up	27.85 m	20.96 m	21.01 m	30.98 m	33 m	57.66 m	10.4 m	26.48 m	28.42 m	31.8 m	28.34 m	21 m
Distance 3 floors up	17.19 m	12.76 m	11.96 m	18.26 m	17.79 m	31.89 m	3.43 m	14.72 m	14.9 m	15.72 m	12.6 m	7.1 m
Distance 4 floors up	11.84 m	8.7 m	7.53 m	11.99 m	10.43 m	19.38 m	0.11 m	9.03 m	8.44 m	8.13 m	5.24 m	0.64 m
Distance 5 floors up	8.61 m	6.29 m	4.89 m	8.27 m	6.07 m	11.99 m	-*	5.69 m	4.66 m	3.71 m	0.96 m	-*
Floors	14.82	15.25	9.87	11.74	8	8.68	4.26	9.33	7.68	6.68	5.55	4.31

Note: * No meaningful output.

Most transmissions arrived, only few are lost as many could be restored from neighbouring gateways. Multiple gateways improve reception but have diminishing returns beyond 30m.

References

[1] Laput, Gierad, and Chris Harrison. "Exploring the efficacy of sparse, general-purpose sensor constellations for wide-area activity sensing." *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3.2 (2019): 1-19.

[2] Wang, Wei, Jiayu Chen, and Tianzhen Hong. "Occupancy prediction through machine learning and data fusion of environmental sensing and Wi-Fi sensing in buildings." *Automation in Construction* 94 (2018): 233-243.

[3] Brian Ray. “What Is LoRa? A Technical Breakdown“ *https://www.link-labs.com/blog/what-is-lora (2018)*

[4] The LoRa Alliance. “About LoRaWAN“ *https://lora-alliance.org/about-lorawan/ (2018)*

[5] Abramson, Norman. "The ALOHA system: Another alternative for computer communications." *Proceedings of the November 17-19, 1970, fall joint computer conference*. 1970.

[6] E. D. Ayele, C. Hakkenberg, J. P. Meijers, K. Zhang, N. Meratnia, and P. J. Havinga, “Performance analysis of lora radio for an indoor iot applications,” in *2017 Int. Conf. Internet of Things for the Global Community (IoTGC)*. New York: IEEE, 2017.

[7] M. M. Erbati, G. Schiele, and G. Batke, “Analysis of lorawan technology in an outdoor and an indoor scenario in duisburg-germany,” in *2018 3rd Int. Conf. Comput. and Commun. Syst. (ICCCS)*. New York: IEEE, 2018, pp. 273–277.

[8] L. Gregora, L. Vojtech, and M. Neruda, “Indoor signal propagation of lora technology,” in *17th Int. Conf. Mechatronics-Mechatronika*. New York: IEEE, 2016.

[9] J. Haxhibeqiri, A. Karaagac, F. Van den Abeele, W. Joseph, I. Moerman, and J. Hoebeke, “Lora indoor coverage and performance in an industrial environment: Case study,” in *22nd IEEE Int. Conf. Emerg. Technol. Factory Autom. (ETFA)*. New York: IEEE, 2017.

[10] M. I. Muzammir, H. Z. Abidin, S. A. C. Abdullah, and F. H. K. Zaman, “Performance analysis of lorawan for indoor application,” in *9th IEEE Symp. Comput. Appl. & Ind. Electron. (ISCAIE)*. New York: IEEE, 2019, pp. 156–159.

[11] P. Neumann, J. Montavont, and T. Noel, “Indoor deployment of low-power wide area networks (LPWAN): A LoRaWAN case study,” in *12th IEEE Int. Conf. Wireless and Mobile Comput., Netw. and Commun. (WiMob)*. New York: IEEE, 2016.

[12] J. Petaj ¨ aj ¨ arvi, K. Mikhaylov, R. Yasmin, M. H ¨ am ¨ al ¨ ainen, and J. Iinatti, ¨ “Evaluation of LoRa LPWAN technology for indoor remote health and wellbeing monitoring,” *Int. J. of Wireless Inf. Netw.*, vol. 24, no. 2, pp. 153–165, 2017.

[13] N. Vatcharatiansakul, P. Tuwanut, and C. Pornavalai, “Experimental performance evaluation of lorawan: A case study in bangkok,” in *14th Int. Joint Conf. Comput. Science and Softw. Eng. (JCSSE)*. New York: IEEE, 2017.

[14] T. Wendt, F. Volk, and E. Mackensen, “A benchmark survey of long range (lora) spread-spectrum-communication at 2.45 ghz for safety applications,” in *16th IEEE Annu. Wireless and Microw. Technol. Conf. (WAMICON)*. New York: IEEE, 2015

[15] C. Jiang, M. K. Masood, Y. C. Soh, and H. Li, “Indoor occupancyestimation from carbon dioxide concentration,”Energy and Buildings,vol. 131, pp. 132–141, 2016.

[16] Z. Yang and B. Becerik-Gerber, “Modeling personalized occupancyprofiles for representing long term patterns by using ambient context,”Building and Environment, vol. 78, pp. 23–35, 2014.

[17] P. Liu, S. Nguang, and A. Partridge, “Occupancy Inference UsingPyroelectric Infrared Sensors Through Hidden Markov Models,”IEEESensors J., vol. 16, no. 4, pp. 1062–1068, 2 2016.

Most transmissions arrived, only few are lost as many could be restored from neighbouring gateways. Multiple gateways improve reception but have diminishing returns beyond 30m.