Announcements: Next Week

- Tuesday’s lecture: discuss/review the Clark paper
  - Will assume we’re mostly doing Q&A so come with your questions!
  - Will use any time leftover to review material for the finals
- Thursday’s lecture: Balaji Sundaravel (Taara)
  - Rural connectivity and other “atypical” networks

Cellular Networks

CS168, Spring 2024
Sylvia Ratnasamy

Outline

- History and the ecosystem
- Central challenge: mobility
- What does a cellular network look like?
  - Infrastructure components
  - End-to-end operation
- Cellular networks within the Internet

Cellular Networks

- **The de facto access technology for <u>mobile</u> users/devices**
  - 5B+ users
  - Over 50% of web traffic originates from a cellular device!
- **Little doubt that mobile wireless access is the future**

Cellular Networks

- **The de facto access technology for mobile users/devices**
  - 5B+ users
  - Over 50% of web traffic originates from a cellular device!

- **Little doubt that mobile wireless access is the future**

- **Cellular operators now facing severe scaling challenges!**
  - New bandwidth-intensive mobile apps: AR/VR, self-driving cars, IoT, *etc.*
  - Deploying towers and buying spectrum is an expensive undertaking
  - Traditional telcos (at&t, verizon) don't have a reputation for rapid innovation
  - General consensus that this is an area ripe for disruption

History

- **Derived from the old telephone network**
  - NTT in 1979 supports voice calls for users in Tokyo (1G)
  - In 1983, Motorola sells first mobile phone in US for ~$4k

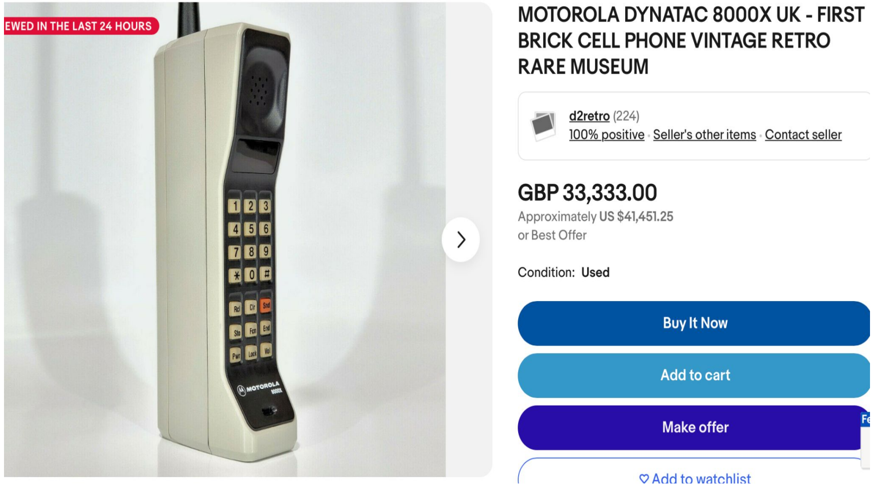

History

- **Derived from the old telephone network**

- These roots lead to architectural choices that differ from the Internet
  - Resource reservations, per-user state in the network, emphasis on accountability, *etc.*

- Today: can think of cellular networks as L2 networks within the Internet

Standards

- The 3GPP (3rd Generation Partnership Project) consortium oversees standardization efforts
- 3GPP standard ratified by the ITU (International Telecom Union), part of the United Nations!
- Typically, a new technology generation (“G”) introduced every 10 years

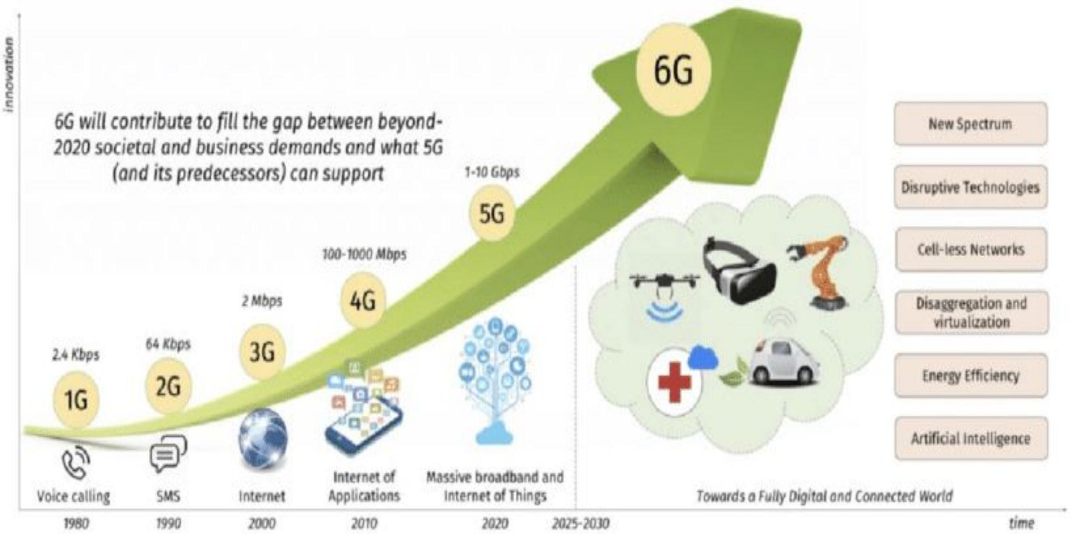

Marketing view

Standards

- The 3GPP (3rd Generation Partnership Project) consortium oversees standardization efforts
- 3GPP standard ratified by the ITU (International Telecom Union), part of the United Nations!
- Typically, a new technology generation ("G") introduced every 10 years

Enhancement of key capabilities from IMT-Advanced to IMT-2020

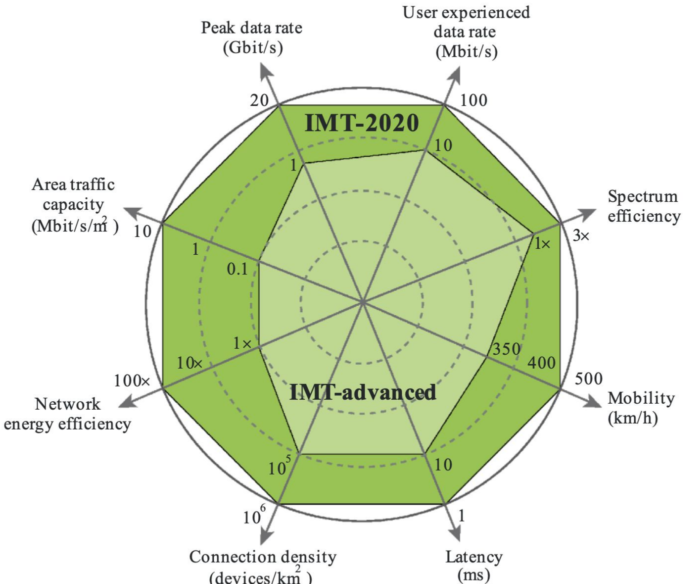

[ITU view]

Generations (*G) in cellular access

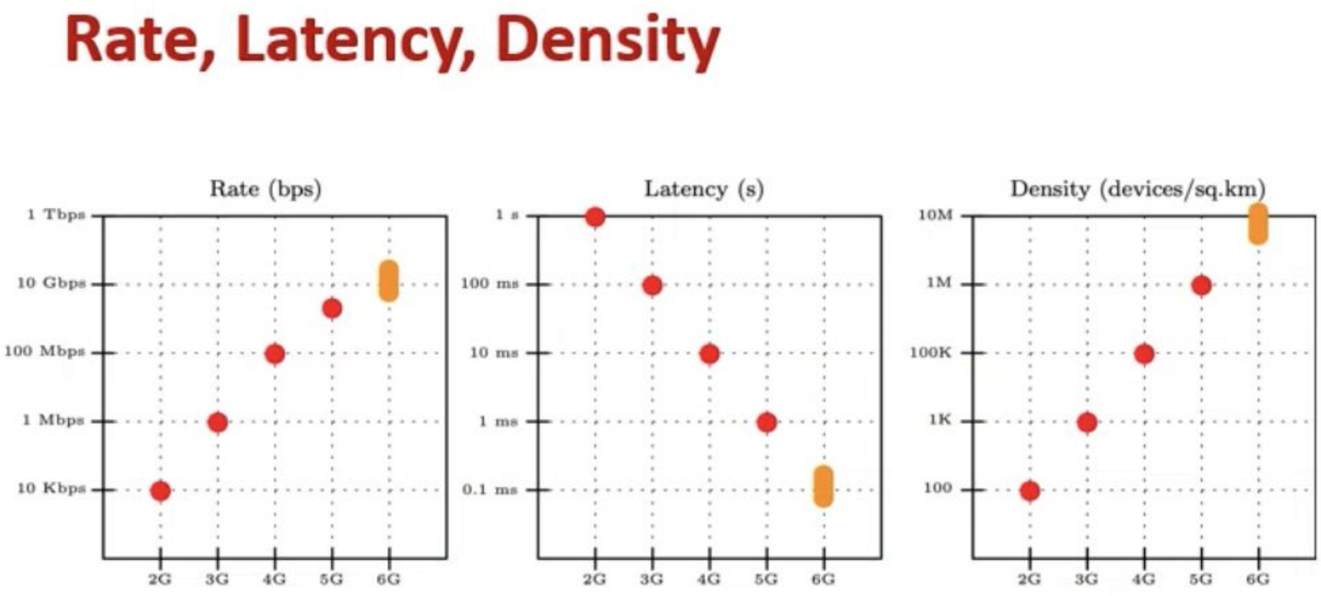

Standards

- The 3GPP (3rd Generation Partnership Project) consortium oversees standardization efforts
- 3GPP standard ratified by the ITU (International Telecom Union), part of the United Nations!
- Typically, a new technology generation ("G") introduced every 10 years

- Each generation achieves better performance and efficiency

Standards

- The 3GPP (3rd Generation Partnership Project) consortium oversees standardization efforts
- 3GPP standard ratified by the ITU (International Telecom Union), part of the United Nations!
- Typically, a new technology generation ("G") introduced every 10 years

- Each generation achieves better performance and efficiency

- Also significant architectural evolution, from a voice to data network
    - 1G: analog phones
    - 2G/3G: mostly circuit switched; focus still on voice traffic
    - LTE/4G onwards: packet switched; voice just another app

Note

- Reading a cellular specification is not for the faint of heart!
  - 100s of documents, 1000s of pages, obscure naming conventions, and endless acronyms
  - To make matters worse, components/protocols are renamed in every generation!
  - E.g., Base station → NodeB → evolved Node B (eNodeB) → next-gen Node B (gNB)

- In this class, we will exercise poetic license and invent our own terminology
  - Loosely based on the LTE architecture

- Conceptually correct but not a 1:1 match to textbooks, standards, *etc.*

Mobility

- **What fundamental new requirements does mobility introduce?**

Mobility

- **What fundamental new requirements does mobility introduce?**
    1. Discovery: what cell tower should a mobile device connect to?
    2. Authentication: should the tower provide service to this device?
    3. *Seamless* communication: no disruption to new/ongoing app sessions
    4. Accountability: enforcing resource limits based on the user's service plan

Outline

- History and the ecosystem
- Central challenge: mobility
- What does a cellular network look like?
  - Infrastructure components
  - End-to-end operation
- Cellular networks within the Internet

Infrastructure: Radio Towers

Implements the air interface to transmit/receive data

Radio Transceiver

Radio Controller

Controls the radio nodes: e.g., allocates radio resources

Radio Tower *(a.k.a Base Station, eNodeB, gNB, …)*

CableFree LTE Base Station Baseband Unit (BBU) for 4G & 5G RAN

CableFree 5G-NR Remote Radio Head (RRH)

Radio Resource Allocation

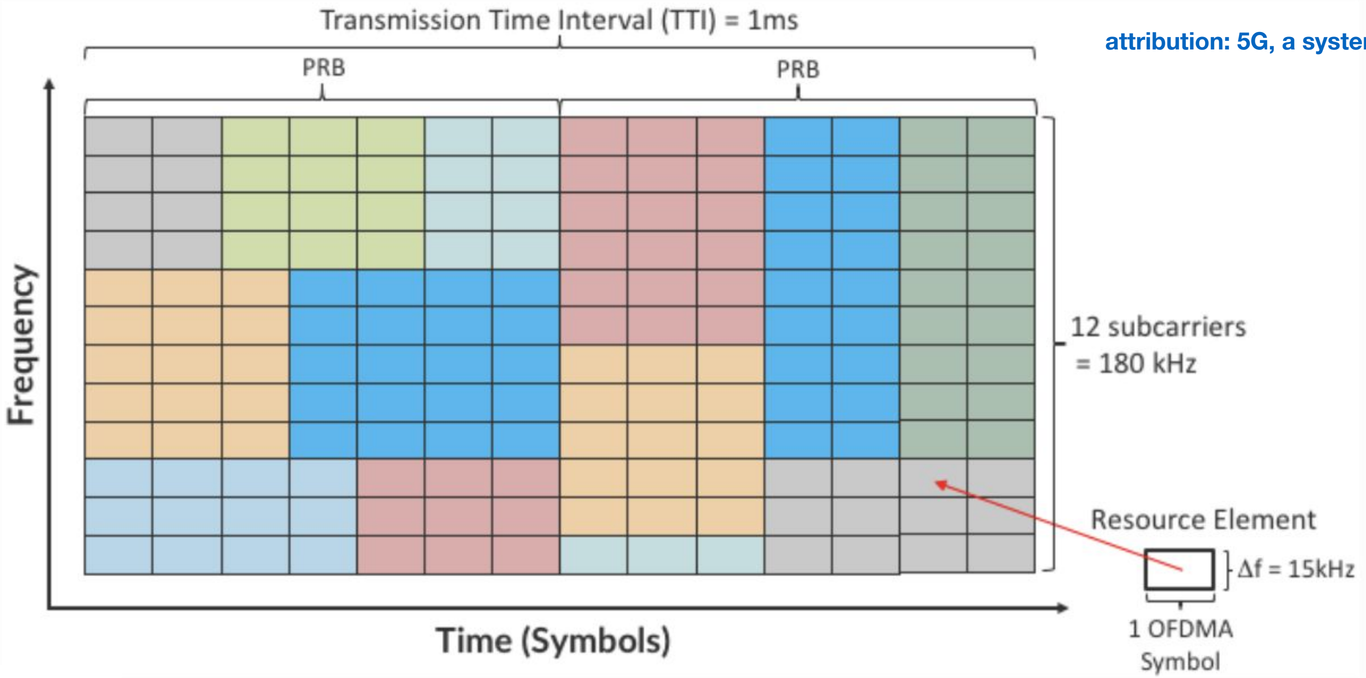

Spectrum abstractly represented by a 2-D grid of schedulable Resource Elements.

Simplified model: Radio controller determines who gets to transmit when and on what frequency

Infrastructure: Radio Access Network (RAN)

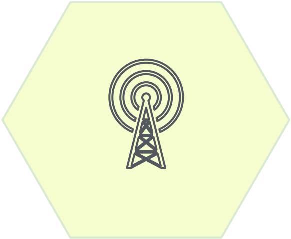

Infrastructure: Radio Access Network (RAN)

Cellular RAN

Infrastructure: Cellular Core

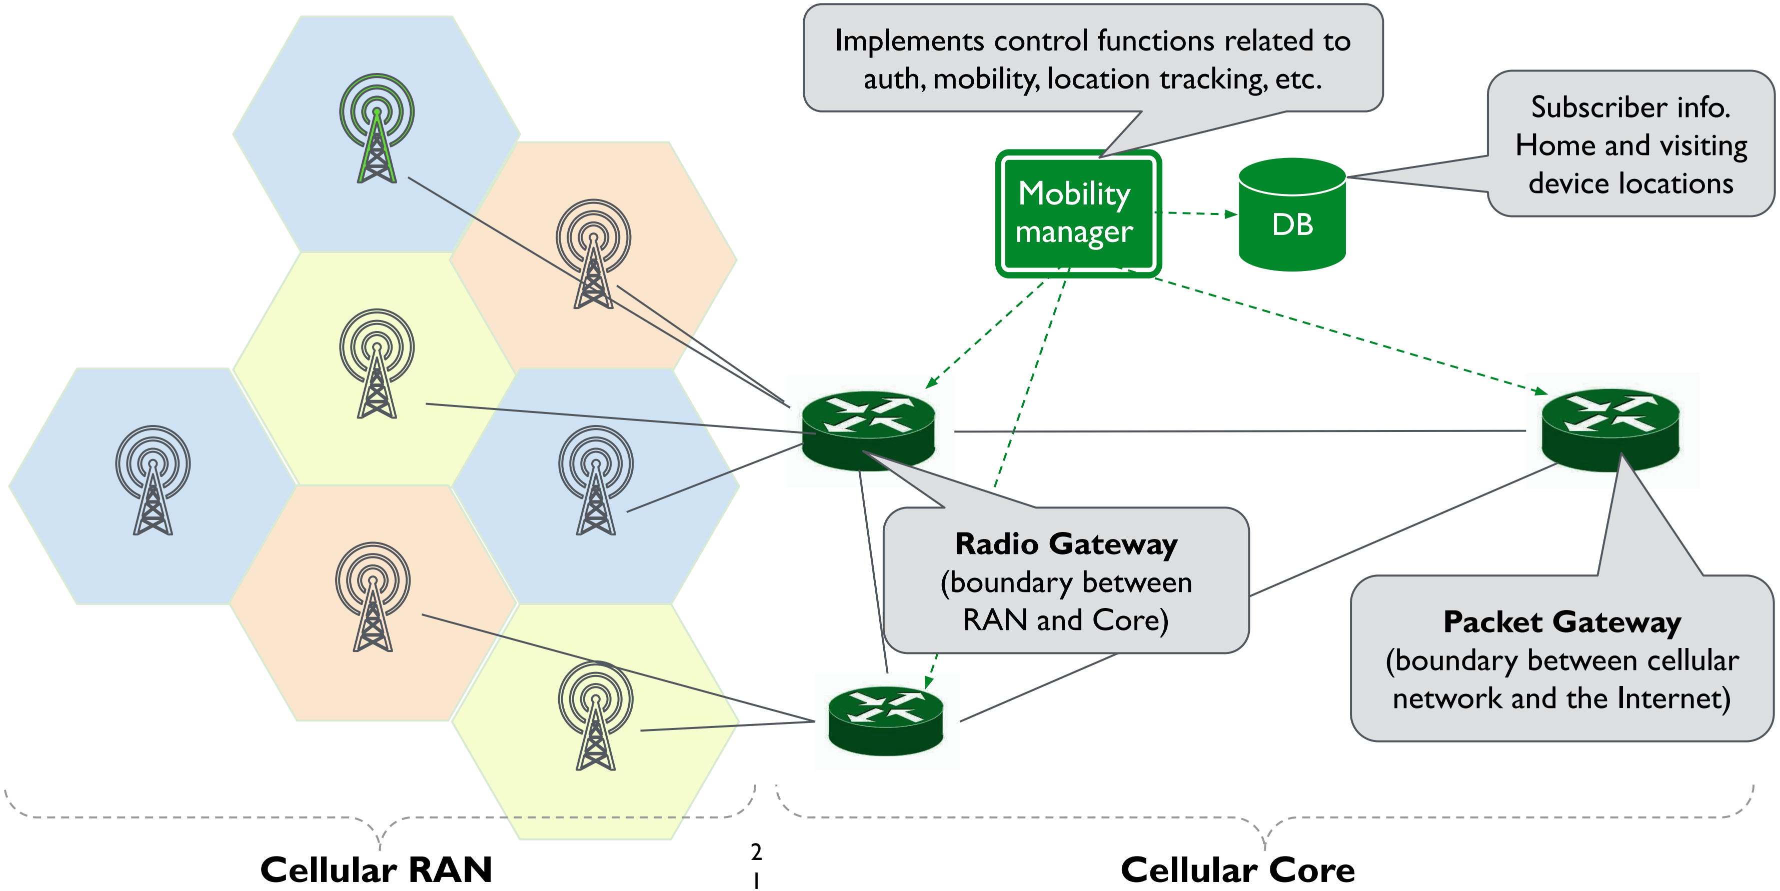

Recap: key components of cellular infrastructure

0. User device
1. Cell towers
2. Mobility manager
3. Cellular DB (to store subscriber information and device locations)
4. Radio gateways
5. Packet gateways

Questions?

Outline

- History and the commercial ecosystem
- Central challenge: mobility
- What does a cellular network look like?
  - Infrastructure components
  - End-to-end operation
- Cellular networks within the Internet

10,000 ft. view of cellular operation

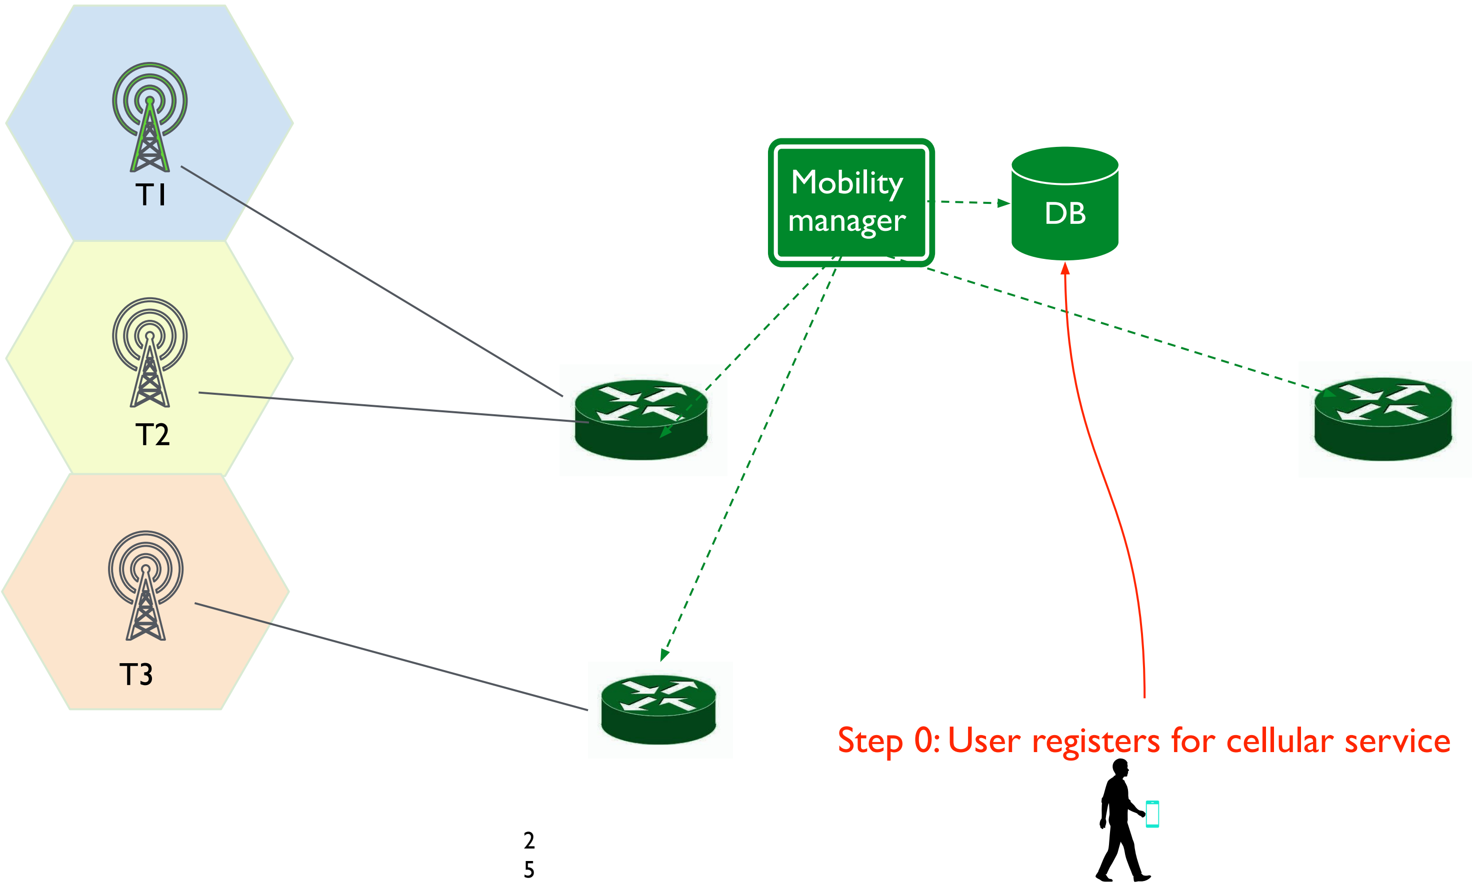

10,000 ft. view of cellular operation

User **discovers** available towers and picks one

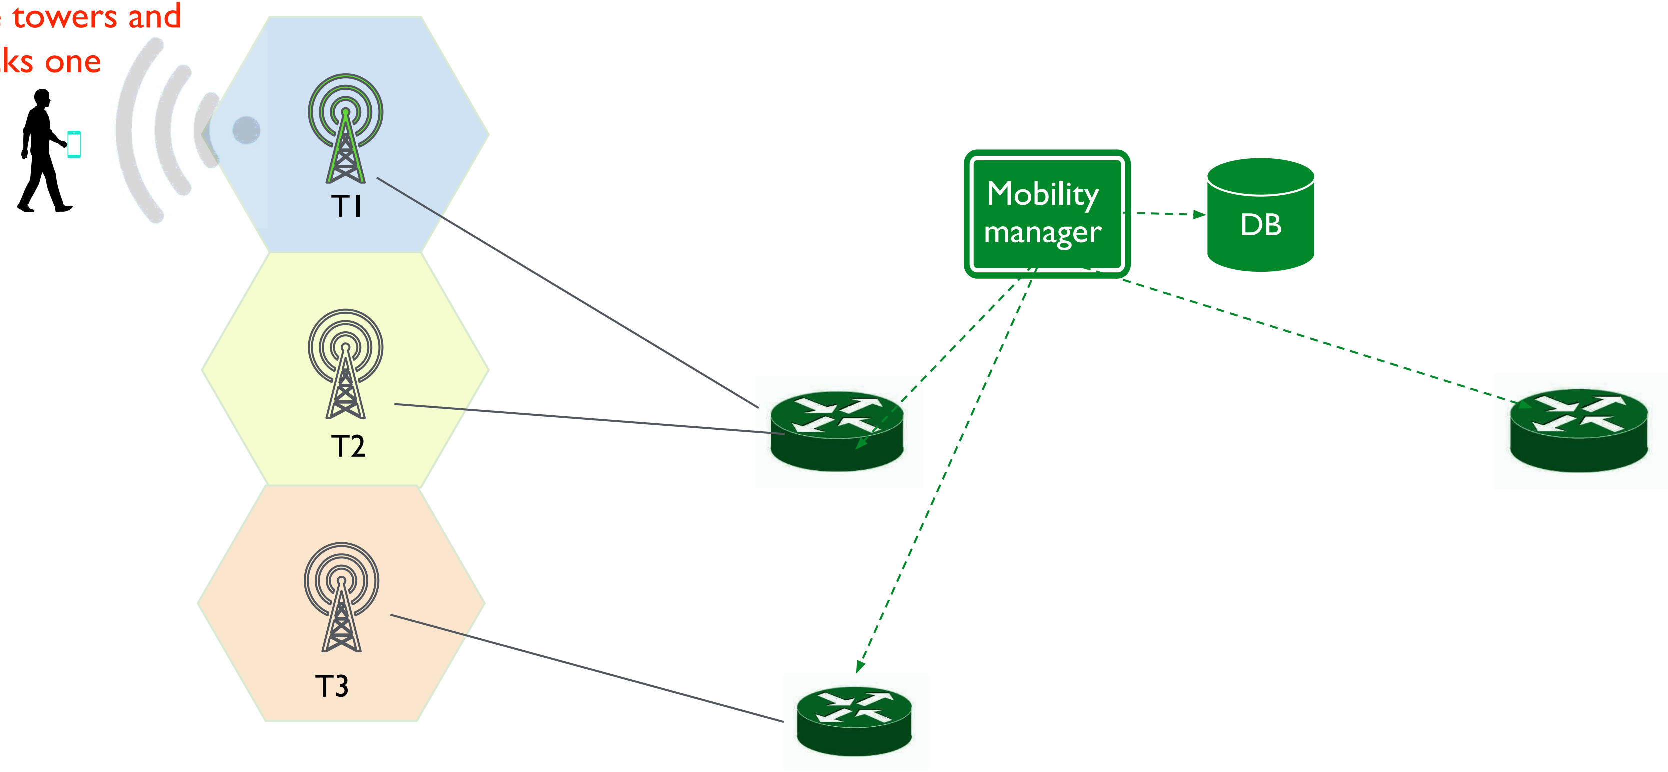

10,000 ft. view of cellular operation

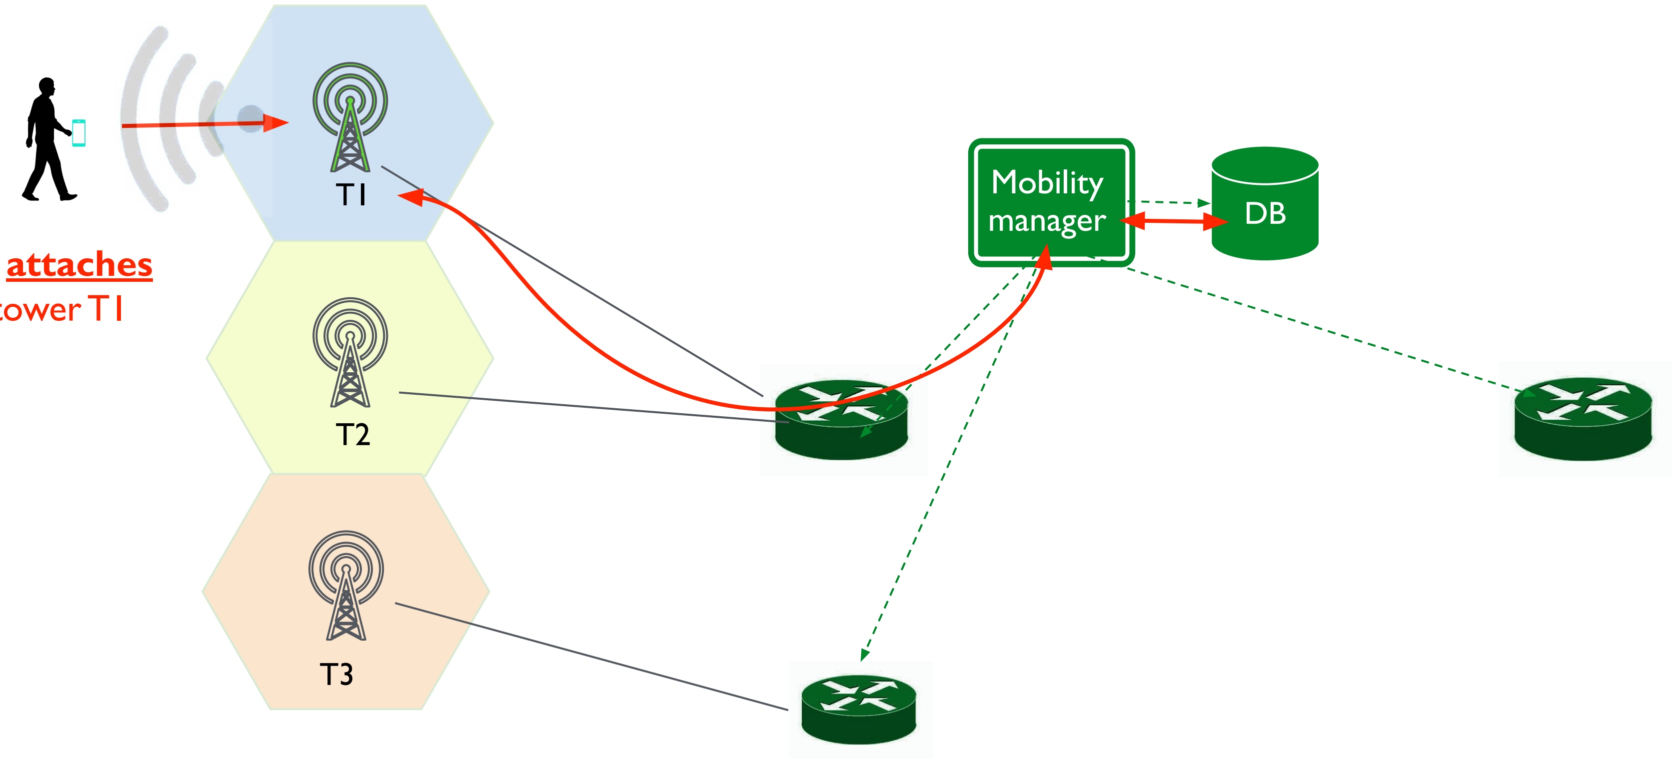

User **attaches** to tower T1

10,000 ft. view of cellular operation

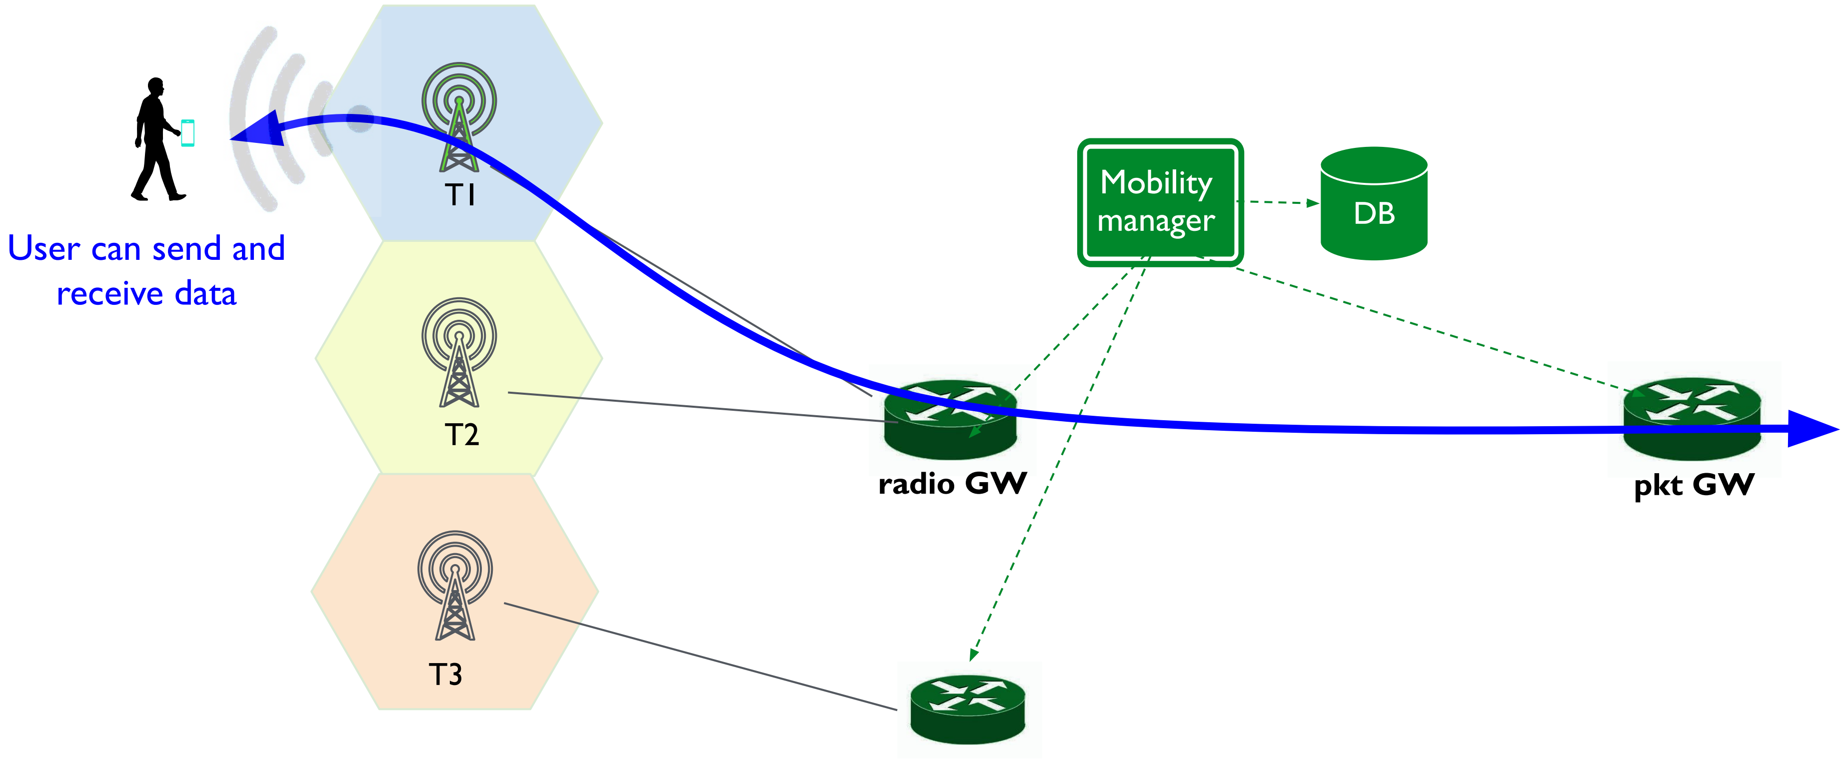

10,000 ft. view of cellular operation

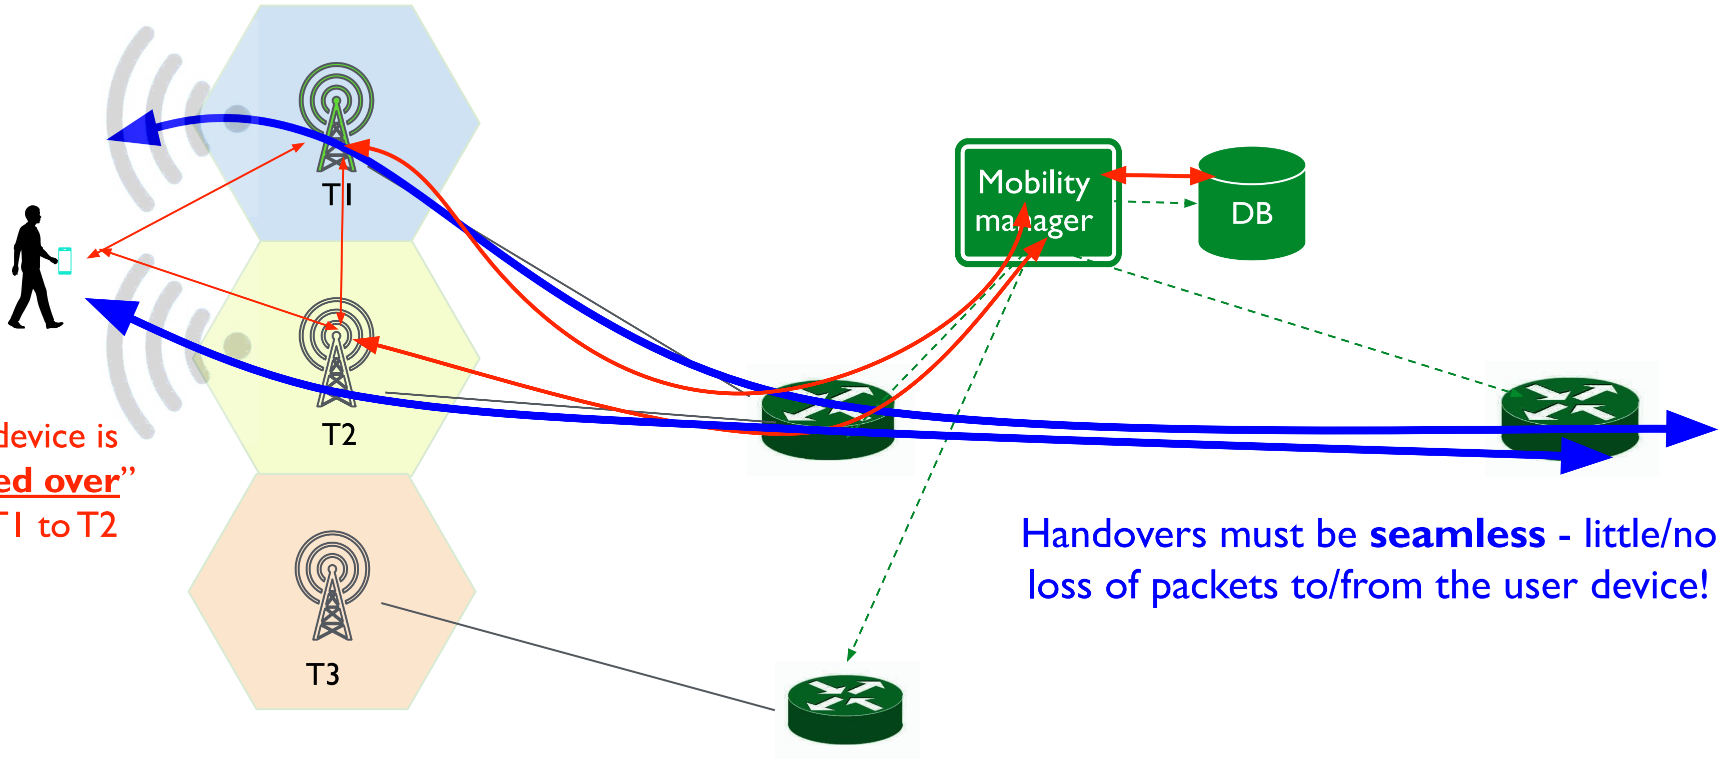

User device is "**handed over**" from T1 to T2

Handovers must be **seamless** - little/no loss of packets to/from the user device!

Note: handovers may involve moving to a different Radio Gateway.

10,000 ft. view of cellular operation

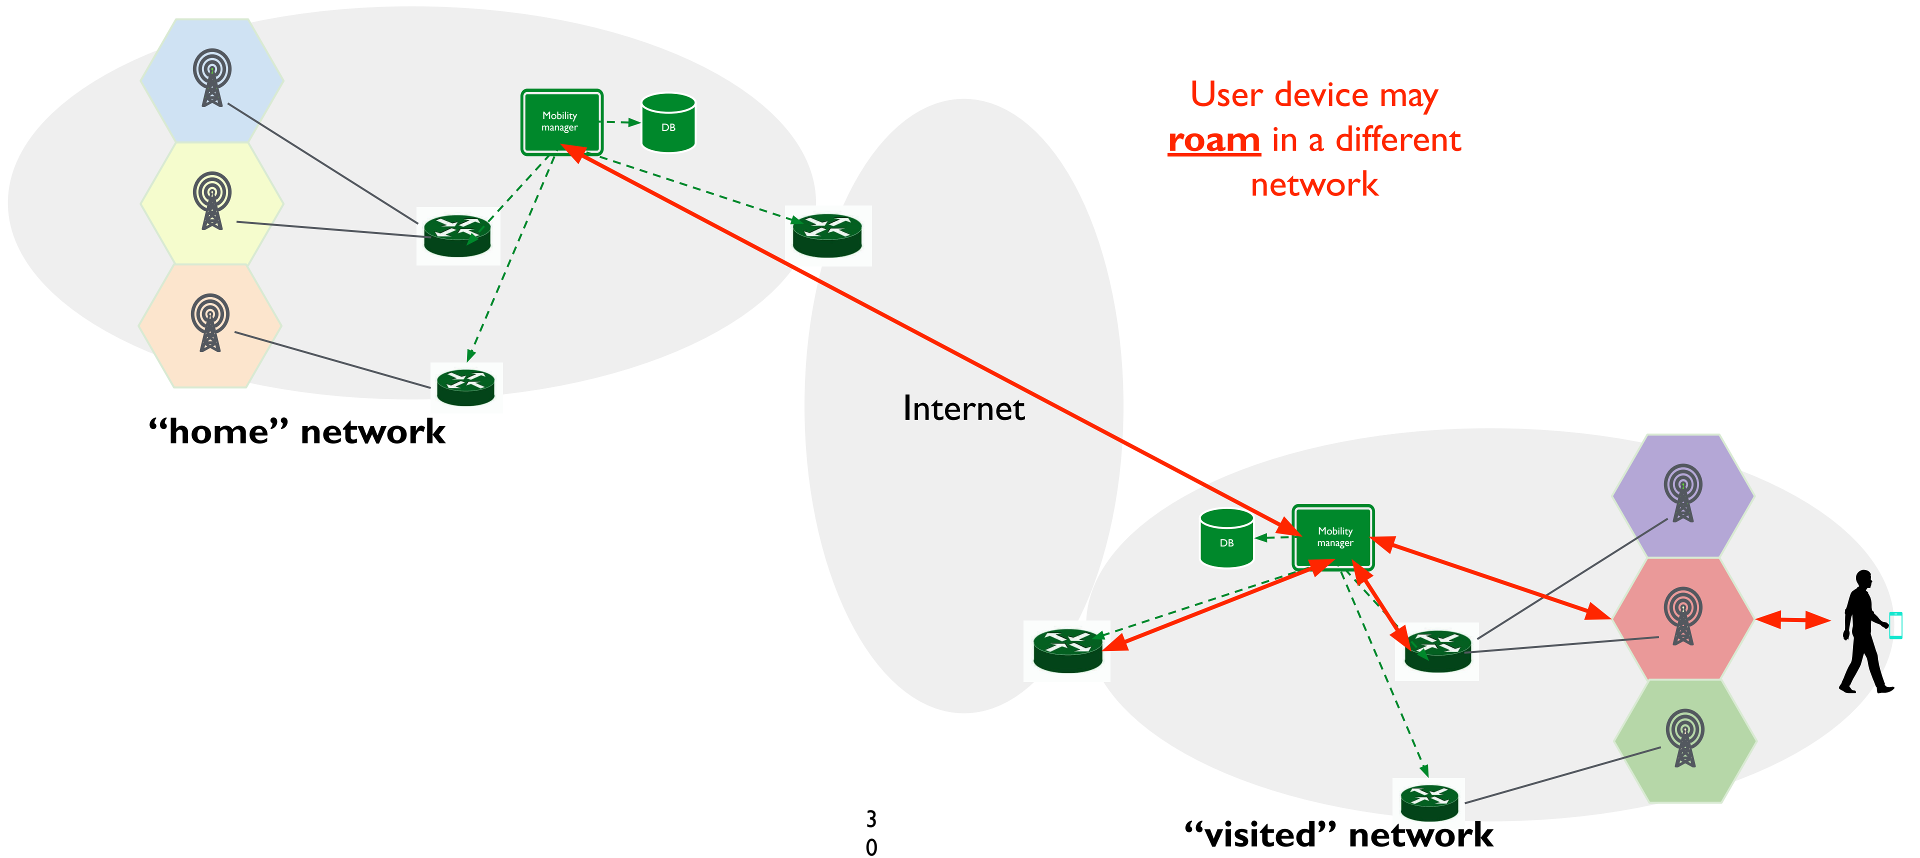

Recap: four key operations

1. Discovery
2. Attachment
3. Handover
4. Roaming

Let’s look at each in more detail.

Questions?

Identifiers associated with a user device

- **International Mobile Subscriber Identity (IMSI):** unique identifier associated with a user's subscription. Securely stored in hardware (SIM card)

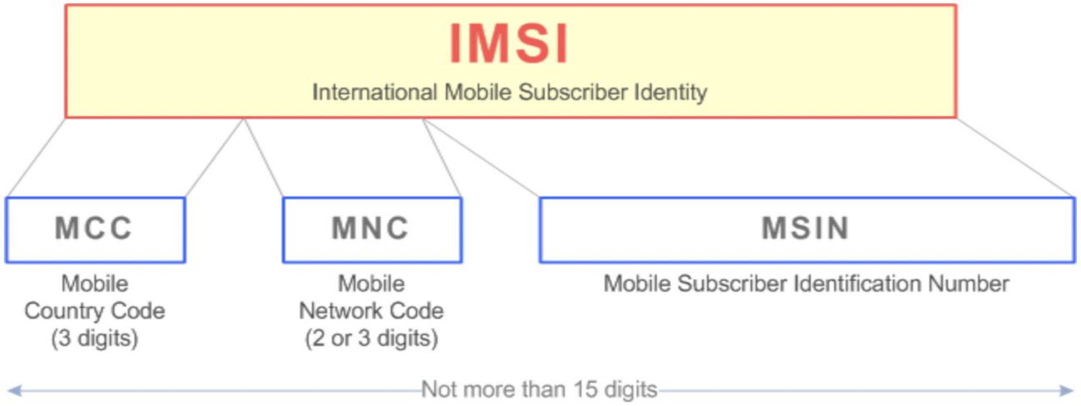

Identifiers associated with a user device

- **International Mobile Subscriber Identity (IMSI):** unique identifier associated with a user's subscription. Securely stored in hardware (SIM card)

- **IP address:** assigned to the user device on attachment; typically retained across handovers

Identifiers associated with a user device

- **International Mobile Subscriber Identity (IMSI):** unique identifier associated with a user's subscription. Securely stored in hardware (SIM card)

- **IP address:** assigned to the user device on attachment; typically retained across handovers

- **International Mobile Equipment Identity (IMEI):** unique identifier assigned to the physical device (vs. subscriber). Identifies the device manufacturer and model.

- **Phone number (MSISDN - pronounced"miss den"):** phone number associated with an IMSI / SIM card.

Registration

- **User U** signs up for service with its home **cellular operator (O)**
  - O stores U's IMSI and associated plan info into its subscriber database
  - Establishes a shared secret key known to U's SIM card and O

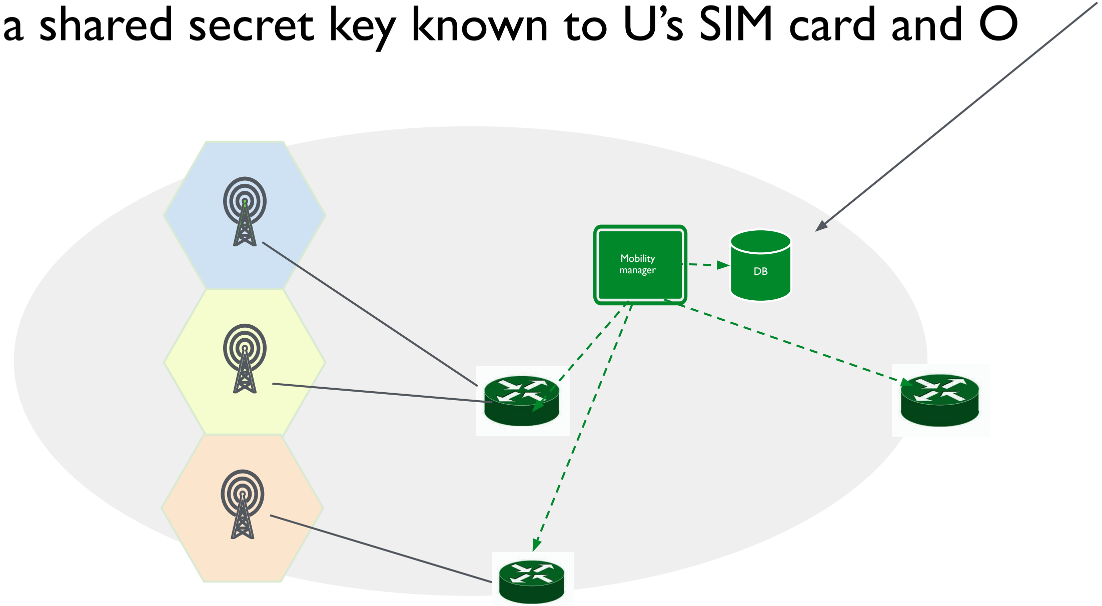

Discovery

Discovery

- Every tower transmits periodic "beacons" on a control channel associated with its assigned frequency range
  - Beacons identify the network operator (mobile network code in the IMSI)

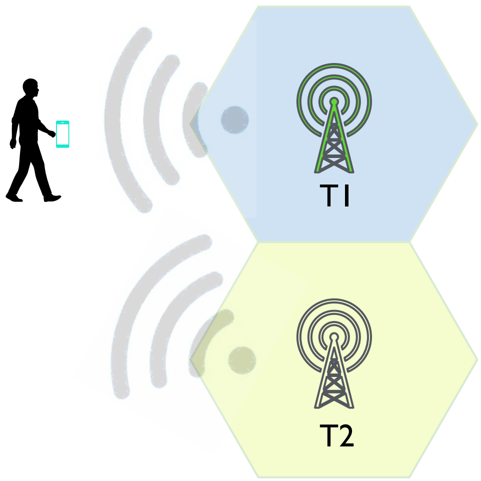

Discovery

- Every tower transmits periodic "beacons" on a control channel associated with its assigned frequency range
  - Beacons identify the network operator (mobile network code in the IMSI)
- User device U measures its signal strength to different towers and picks the tower (belonging to its operator) that has the best signal

Discovery

- Every tower transmits periodic "beacons" on a control channel associated with its assigned frequency range
  - Beacons identify the network operator (mobile network code in the IMSI)
- User device U measures its signal strength to different towers and picks the tower (belonging to its operator) that has the best signal
- How does U know which control channel to listen to?

Discovery

- Every tower transmits periodic "beacons" on a control channel associated with its assigned frequency range
  - Beacons identify the network operator (mobile network code in the IMSI)
- User device U measures its signal strength to different towers and picks the tower (belonging to its operator) that has the best signal
- How does U know which control channel to listen to?
  - Scan all: slow (10-100s of seconds) but sometimes necessary
  - Optimization: device is pre-configured with a set of likely frequency channels
  - Optimization: cache previously used channels
  - During handovers: tower T1 tells U what channel on T2 to use (10-100s of millisecs)

Attachment

1. U sends an "attach request" to tower T1

- U's request includes its IMSI

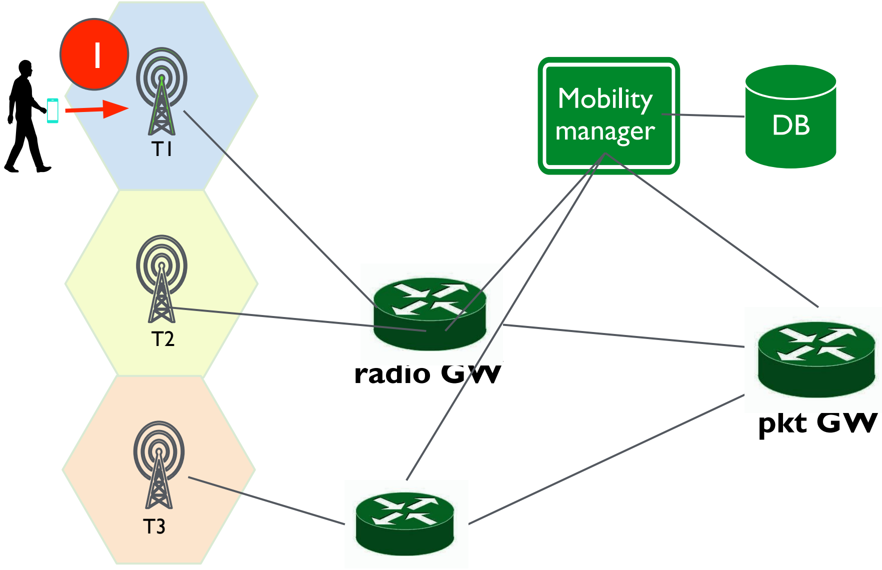

Attachment

1. U sends an "attach request" to tower T1

- U's request includes its IMSI

2. Request processed by mobility manager

- Looks up the IMSI in subscriber DB to determine the parameters of service based on U's plan
- Authentication based on pre-established shared key

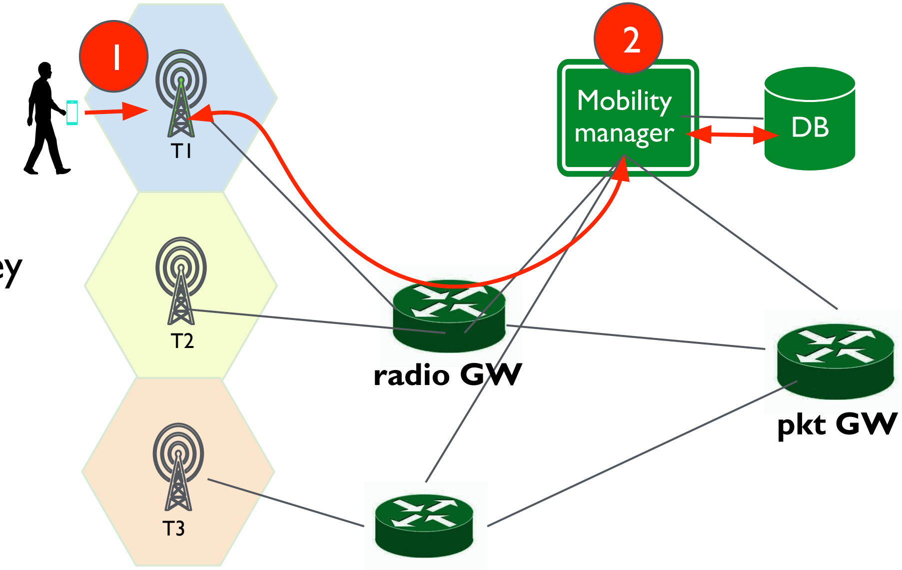

Attachment

1. **U sends an "attach request" to tower T1**
   - U's request includes its IMSI
2. **Request processed by mobility manager**
   - Looks up the IMSI in subscriber DB to determine the parameters of service based on U's plan
   - Authentication based on pre-established shared key
3. **Mobility manager configures data plane**
   - Assigns IP address to U (IP_u)
   - Configures T1 with appropriate radio parameters
   - Configures tunnel between T1 and Radio-GW
   - Configures tunnel between Radio-GW and Pkt-GW
   - Initializes counters, shapers to track/enforce U's QoS

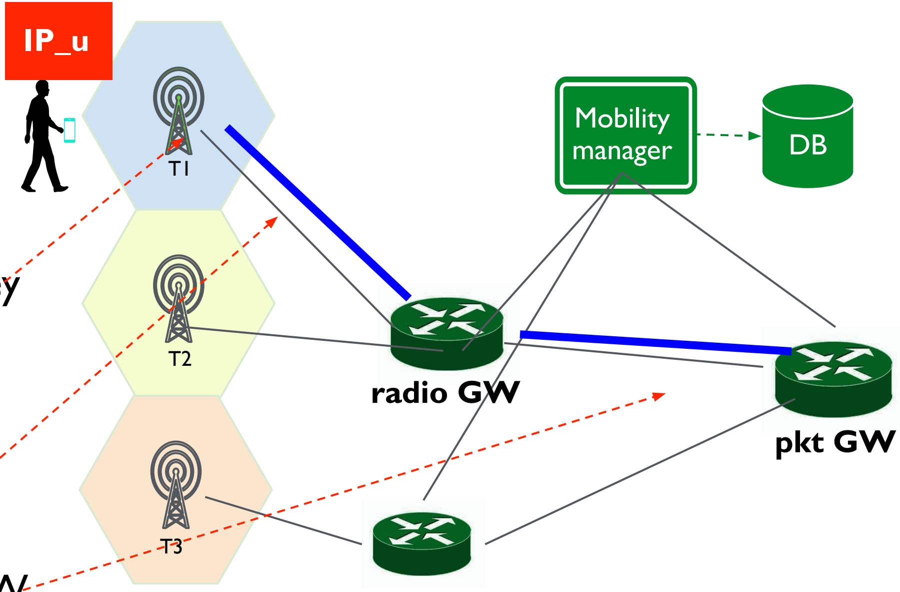

Attachment

1. U sends an "attach request" to tower T1

- U's request includes its IMSI

2. Request processed by mobility manager

- Looks up the IMSI in subscriber DB to determine the parameters of service based on U's plan
- Authentication based on pre-established shared key

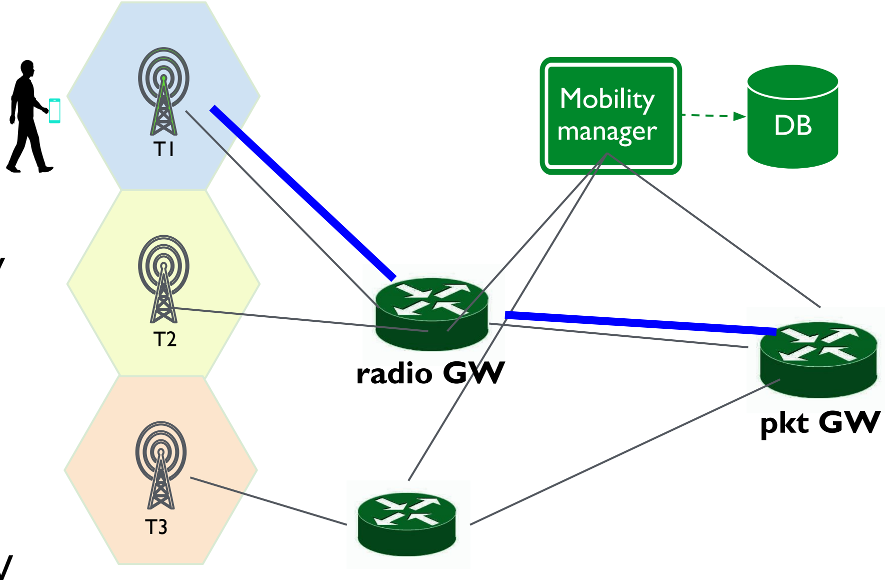

3. Mobility manager configures data plane

- Assigns IP address to U (IP_u)
- Configures T1 with appropriate radio parameters
- Configures tunnel between T1 and Radio-GW
- Configures tunnel between Radio-GW and Pkt-GW
- Initializes counters, shapers to track/enforce U's QoS

4. Mobility manager records U's location info in database (IMSI → T1, GWs, IP_u, etc.)

User can now send/receive data!

Note 1: up to this point, all communication with U takes place over control channels (no radio resources were assigned to U)

Note 2: packets to/from IP_u are tunneled / encapsulated; no direct forwarding on IP_u (Will return to this later)

Note 3: installing per-user state in the network's data and control plane (Different from traditional IP networks!)

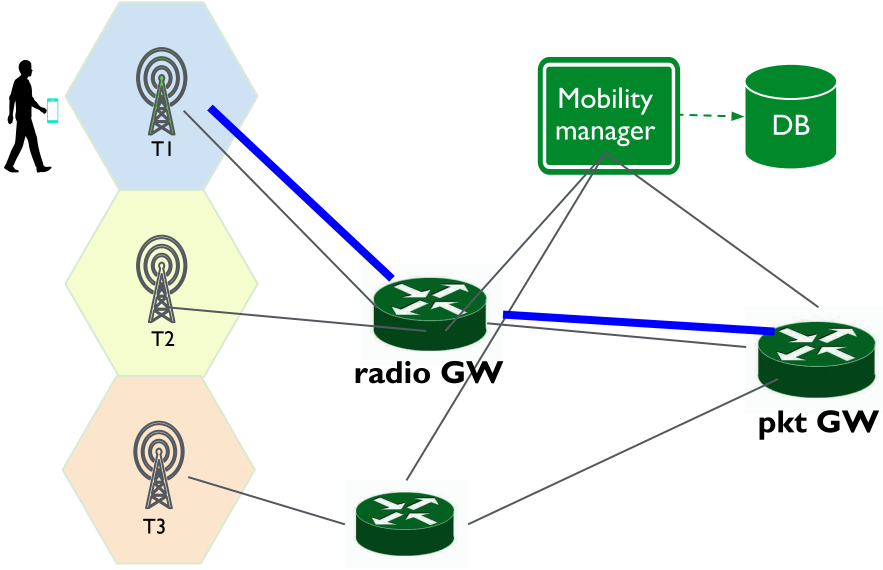

Handover: U moves from tower T1 to T2

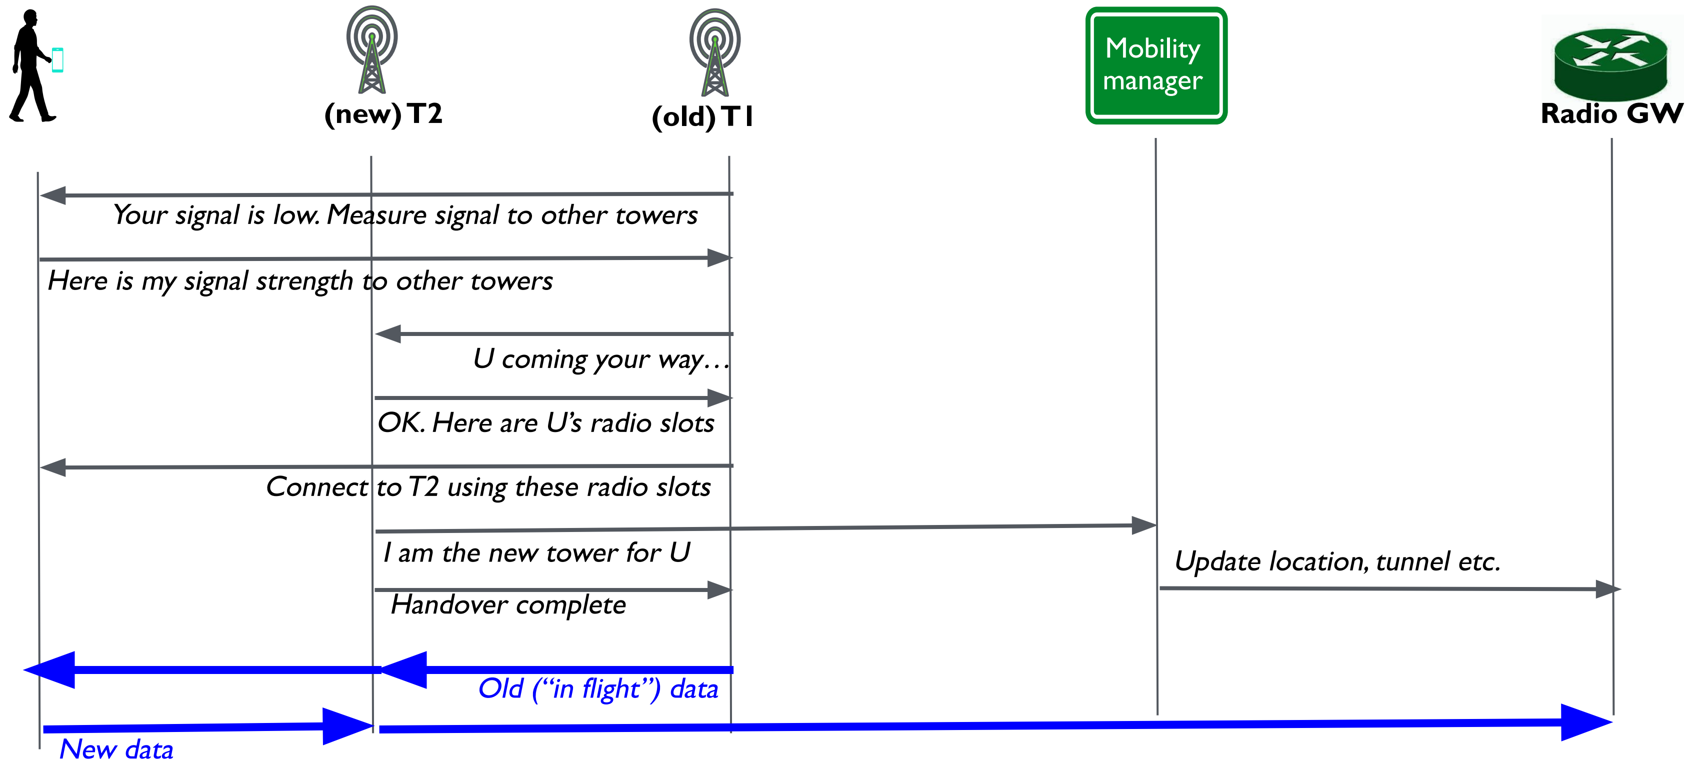

Acknowledgements: *adapted from Raj Jain's lectures*

Handover: U moves from tower T1 to T2

- Cooperative process between U, T1, T2, mobility manager, Radio-GW
  - More involved process when handover involves changing Radio or Packet GWs!
- Tower selection is network-driven: ultimate decision not made by user device U
  - tradeoffs?
- U's IP address remains unchanged! Instead, must reconfigure the necessary tunnels
  - Dynamically updating that per-user state!
- To avoid packet loss, T1 buffers/forwards packets that arrive during handovers

Thought exercise: how do you think this holds up as U starts to move really fast?

Roaming

- Very similar to what we've seen so far but now the mobility manager in the "visited" network contacts that in the "home" network for auth, etc.

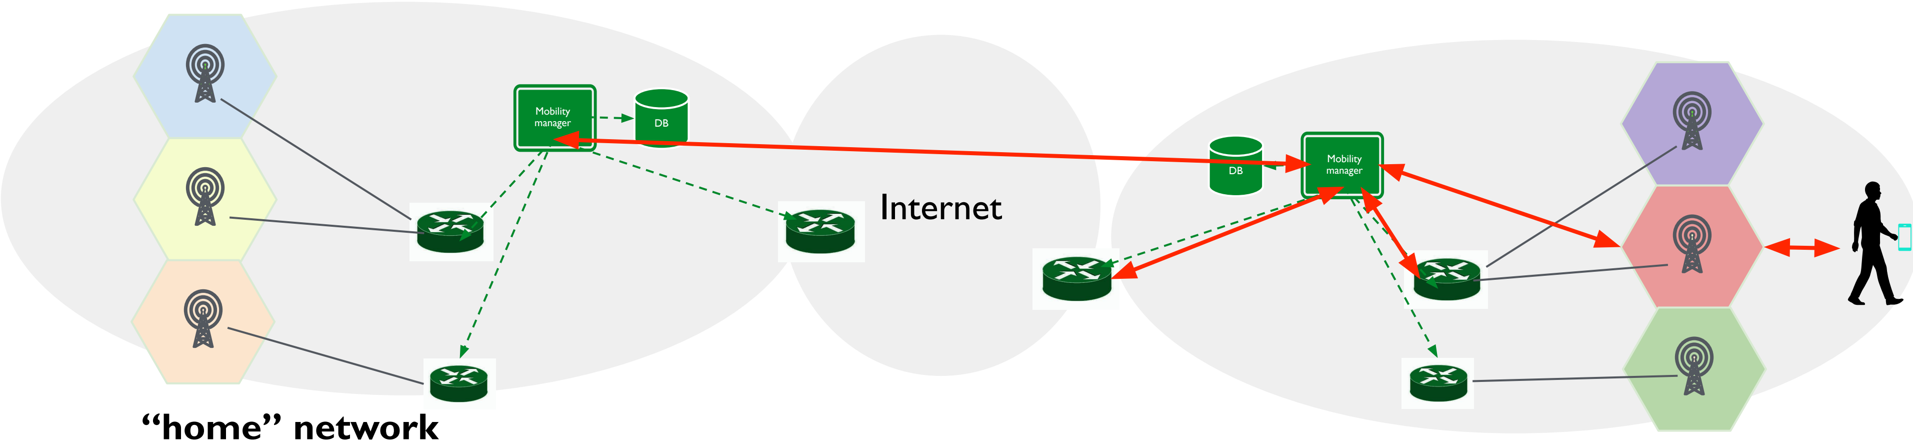

Roaming

- Very similar to what we've seen so far but now the mobility manager in the "visited" network contacts that in the "home" network for auth, etc.
  - How does the visited network know what home network to contact?

- Two common options for how data plane is set up
  - Home routing: tunnel traffic through the home network's packet GW

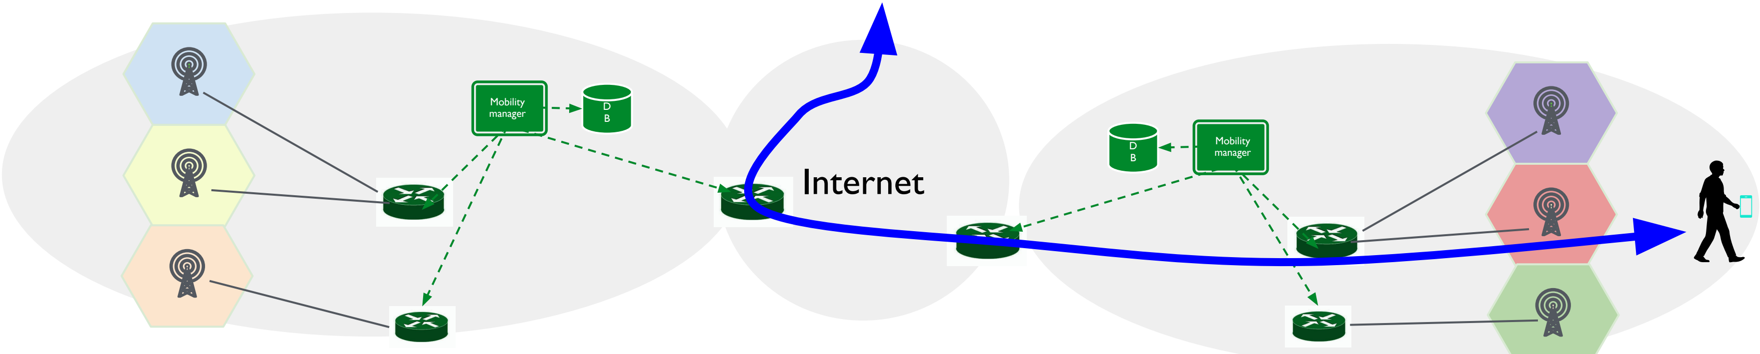

Roaming

- Very similar to what we've seen so far but now the mobility manager in the "visited" network contacts that in the "home" network for auth, etc.
  - How does the visited network know what home network to contact?

- Two common options for how data plane is set up
  - Home routing: tunnel traffic through the home network's packet GW
  - Local breakout: traffic directly exits from visited network's packet GW

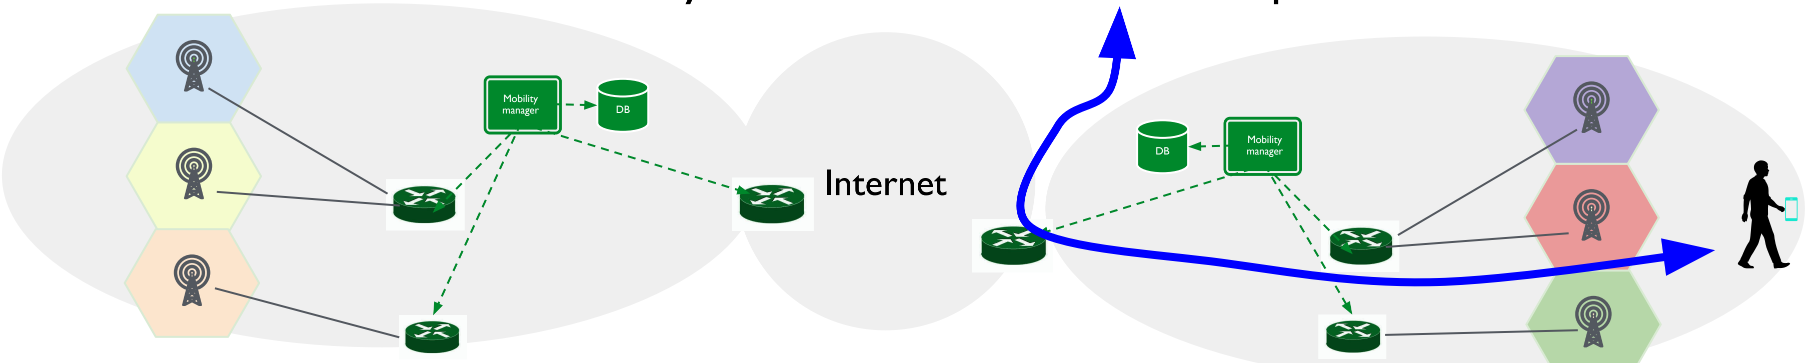

Roaming

- Very similar to what we've seen so far but now the mobility manager in the "visited" network contacts that in the "home" network for auth, etc.
  - How does the visited network know what home network to contact?

- Two common options for how data plane is set up
  - Home routing: tunnel traffic through the home network's packet GW
  - Local breakout: traffic directly exits from visited network's packet GW

- Home network's mobility manager records U's current visited location in its DB
  - Why is this important?

Thought exercise: why the home vs. visited distinction?
Why can't U treat any available operator as its home network?

Recap: four key operations

1. Discovery
2. Attachment
3. Handover
4. Roaming
5. Additional operations exist
   a. Lawful intercept
   b. Paging
   c. Stolen phone registries
   d. And more …

Taking Stock: Cellular networks

- Compared to the Internet architecture we've been studying
  - Authentication and accounting are central goals
  - Allocation of radio bandwidth is based on reservations
  - Lots of in-network state that is dynamic and per-user
- Stateful networks are complex and challenging!
  - Reconfiguring tunnels per mobility event limits the rate of mobility and scalability
- Thought experiment: why did we need all these tunnels?
  - User's IP address is not directly routable since its location in topology keeps changing!
  - Why can't we just let the user's IP address change on ~each handover? TCP breaks!
  - What if we had an alternative to TCP that is OK with changing IPs? Check out QUIC!

Summary

- **Cellular networks have evolved from a standalone voice network to being an integral part of the Internet**

- **Based on a very different design philosophy**
  - Authentication and accountability are primary goals
  - Generality was not an early goal
  - Mobility is the central challenge

- **Yet, we've been able to seamlessly integrate cellular networks into the Internet!**
  - And the cellular architecture continues to evolve towards that of the Internet

Questions?

Goal: Relay packets between devices and Internet

- **Cellular infrastructure consists of two components:**
  - Radio Access Network (RAN): implements the air interface between *mobile devices* and *cell towers*, tunneling packets from devices to the cellular core
  - Cellular Core (Core): get packets from the tunnels and send them over to the Internet and vice-versa.
- **Design goals / requirements**
  1. Discovery: what cell tower should a mobile device connect to?
  2. Authentication: should the tower provide service to this device?
  3. *Seamless* communication: no disruption to new/ongoing app sessions
  4. Accountability: enforcing resource limits based on the user's service plan
- **Main challenge: mobility**! I.e. device switches towers. New design questions:
  - How to correctly route device packets to the Internet despite the tower switching?
  - How to mitigate the effects of switching, e.g., apps can see the same source IP?
  - How to allow the device switches to the towers of a different network provider? Etc.