



HAL
open science

Data Consistency in the 5G Specification

Jonathan Sid-Otmane, Sofiane Imadali, Frederic Martelli, Marc Shapiro

► **To cite this version:**

Jonathan Sid-Otmane, Sofiane Imadali, Frederic Martelli, Marc Shapiro. Data Consistency in the 5G Specification. ICIN 2020 - 23rd Conference on Innovation in Clouds, Internet and Networks and Workshops, Feb 2020, Paris, France. pp.110-117, 10.1109/ICIN48450.2020.9059408 . hal-02943802

HAL Id: hal-02943802

<https://hal.science/hal-02943802v1>

Submitted on 21 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data Consistency in the 5G Specification

Jonathan Sid-Otmane
Orange Labs Network
Sorbonne-Université—LIP6
Paris, France
firstname.lastname@lip6.fr

Sofiane Imadali, Frédéric Martelli
Orange Labs Network
Paris, France
firstname.lastname@orange.com

Marc Shapiro
Sorbonne-Université—LIP6 & Inria Paris
Paris, France
firstname.lastname@lip6.fr

Abstract—Meeting the goals of 5G networks – high bandwidth, low latency, massive connectivity, and resiliency demands improvements to the infrastructure that hosts the network components. Mobile Network Operators will rely on a geographically distributed and highly scalable infrastructure that must handle and replicate user data consistently. This paper explores the management of user data with regards to data consistency in the first 5G specification. In particular we will focus on how the 5G system procedures handle and update data, and discuss failure scenarios where the correctness properties of the user data may be violated. In this work we present the necessary properties that an underlying data store must deliver in order to maintain correctness in the presence of failures.

Index Terms—5G, correctness, consistency, mobile networks, user data

I. INTRODUCTION

The first 5G specification (i.e. 3GPP release 16 [1], [2]) makes consistency assumptions on the database layer semantics that are in direct violation with well known impossibility results (Brewer’s CAP theorem). Indeed, to better serve their clients, Telcos focus on preventing network partitions in their virtualized infrastructure, involving their 5G User Data store: The Unified Data Management (UDM). Despite their best efforts, it is inevitable that partitions will occur.

The UDM is a distributed, transactional database with full ACID (Atomicity, Consistency, Isolation, and Durability) guarantees that *will not* be resilient to partitions: a transactional system that offers these guarantees cannot function correctly during a partition [3]. Many levels of consistency guarantees can be found in the literature that could prove to be better suited to the 5G network while still providing sufficient correctness guarantees, even in the event of failure [4].

In order to explore this spectrum of consistency guarantees, it is necessary to review the 5G Core Network architecture as it currently stands, and the data associated with a specific user device (User Equipment (UE) Context).

Along with the mobile data users needs, operators are also expected to deploy domain specific networks to deal with the demands of the industry, the IoT and the automotive world. These new functionalities will be provided by leveraging concepts that are novel to the Telco world: Software Defined Networks, Network Function Virtualization, and Network Slicing. Decentralization will be another key transformation of the infrastructure to keep up with these new scales, hence the need for distributed databases.

This paper provides an analysis of the main 5G procedures that ensure connectivity: registration, request, and handover. We explore the interactions of the different network components presented earlier as well as how they access and modify the UE Context. We link each variable of the UE Context to its specific procedures as well as the network components that access/modify it. Thus, we define a set of properties that are sufficient to ensure that even in the presence of failures it will not adversely affect the procedures it takes part in, preserving the correctness of the whole system.

We will discuss the properties of the UE Context variables that must be respected, and give an overview of different consistency guarantees that are sufficient for this purpose. We will demonstrate inconsistencies that the system is vulnerable to when failures or message delays occur, as well as consistency mechanisms to remedy them.

Section II of this paper presents the data consistency semantics, and implications of distributed databases for 5G. Section III presents the architecture of the 5G Core Network and defines its component. Section IV analyses the procedures in detail and the variables that they handle. Finally, section V discusses a few failure scenarios as well as consistency mechanisms to maintain correctness.

II. BACKGROUND

The Unified Data Management (UDM) acts as a distributed database for the 5G Network as we describe in section IV. Along with the UDM we will introduce Brewer’s CAP theorem and explain the need to circumvent it. Finally, this section reviews the work of Ojala et al.[5] on managing the state of a 4G network in a distributed database, also subject to the CAP theorem.

The UDM acts as a logical repository to store and access user data. Within the UDM, the data will be kept in UDR (User Data Repository). The UDM will help disseminate this data throughout the system by allowing all network functions to query and to subscribe to it. The semantics of the UDR are defined in the 5G specification as a distributed transactional system that offers ACID properties. These properties are known as Atomicity, Consistency, Isolation and Durability and are highly desirable[6].

Brewer’s theorem states that it is impossible for a distributed system to offer the level of consistency necessary to ensure ACID properties and be resilient to network partitions at the

same time [3]. ACID requires that all replicas of a distributed database *see and apply* the same changes at the same time. All the replicas have to be involved. However during a partition, some replicas cannot be reached. This is why distributed database management systems usually either make themselves unavailable in the presence of faults or offer a weaker level of consistency guarantees[4], [7].

Consistency mechanisms regulate reads and writes on data so that all operations result in the data respecting constraints (invariants) [8]. The choice of these invariants depends on the needs of the application that uses the data, and results in a consistency level. Consistency mechanisms are classified according to the invariants they guarantee and the anomalies that they allow[9].

The database that serves 5G users (the UDM) must be distributed to serve its geographically distributed clients with the myriad of 5G use cases, involving slicing, for instance. The UDM must also be resilient to all sorts of faults and failures. This means that it falls squarely in the impossibility conundrum of Brewer’s theorem, and an appropriate consistency level must be guaranteed. Cloud based databases have faced this same challenge[7]. They must make the choice of either remaining correct during network partitions by limiting availability[6], or targeting a weaker level of consistency[10].

The UDM, that will serve the UE Context to components of the network must make the same trade-offs. Previous work by Ojala et. al. [5] has discussed a distributed Core Network that relies on a distributed NoSQL database to provide the UE Context. Their work highlights the impossibility of providing strong consistency for a geo-distributed database and makes the trade-off of a weakly consistent distributed database with high consistency for geographical zones that are centered around a given UE.

Weak consistency vs. strong consistency hybrid databases, while correct, could be inefficient for some of the data involved in the 5G features, such as Network Slicing. In particular, Ojala et al.[5] center high consistency zones around the geographical location of the UE, meaning that all information relevant to the user must be around him. But in the instance of network slicing a network slice treats a group of users as a single entity with geodistributed data. Consistency mechanisms cannot rely on the benefits of proximity.

III. 5G ARCHITECTURE AND UE CONTEXT

The purpose of the 5G Core network is to manage user equipments (devices), allow them to connect to their target Packet Data Network (PDN) according to the operator’s policies, and to charge its users. We have chosen 3 procedures that are at the core of these functionalities of the network involving its main components as shown in Figure 1. The procedures are:

- (i) *Registration*: When a device first connects to the network and its UE Context is queried from the UDM
- (ii) *Packet Data Unit (PDU) Session Establishment*: When a device initiates a data connection

- (iii) *Handover*: When a device has moved and must be recognized by the network at its new location and have it’s PDU session migrated when it tries to use them.




Fig. 1. 5G System Architecture [1]

The line dividing Figure 1 in two parts separates the control plane, above the line, from the data plane below. The Network Functions in the data plane are tasked with carrying data between the UE and DN, using the radio technologies of the Radio Access Network (RAN). The control plane is tasked with the setup of these connection, and will be the main focus of our work. This data plane as well as the RAN in charge of the radio technologies are beyond the scope of this paper.

A. 5G Network Functions

Figure 1 illustrates the components of the 5G network. Defined as potentially distinct physical appliances in previous generations, 5G will rely on Virtualized Network Function (VNF)[1]. To build an elastic infrastructure: the number of VNFs of each type that are responsible for any user or geographical region will be subject to change dynamically. Principles of the Service Based Software Architecture model are applied. VNFs will be more lightweight than their LTE counterparts, to allow them to be spun up and down more quickly.

As a significant upgrade to LTE, 5G leverages virtualization to introduce the concept of Network Slicing. Network slices are logical networks running on top of the common physical infrastructure of 5G. Their purpose is to serve new industries such as IoT or automotive whose connectivity needs are different from current mobile network uses.

1) Functions involved in the procedures:

- Access and Mobility Function (AMF)
The AMF is the central function that manages the UE. It receives requests directly from the UE and makes decisions accordingly or forwards them to different VNFs.
- Authentication Server Function (AUSF)
The AUSF is a server that authenticates the UE when it connects to the server. It keeps security related information in its own section of the UDM and installs a security context on the AMF and the UE so that all their communications are encrypted.
- Session Management Function (SMF)

The SMF is in charge of setting up and managing one or several PDU sessions. It is chosen during the session establishment procedure and will stay during the lifetime of the session.

- Unified Data Management (UDM)
The unified interface to 5G's data store. It allows 5G core network components to query, subscribe to, and modify values.
- User Equipment (UE)
The UE allows a subscriber to access the network.
- Policy Control Function (PCF)
The PCF ensures that all policies regarding a UE or a PDU Session are known and kept up to date.
- User Plane Function (UPF)
The UPF is the component that carries data packets between the UE and the PDN, either directly or through Intermediate UPF (I-UPF). The PSA (PDU Session Anchor) is the closest to the PDN and is chosen for the lifetime of the session.

2) Additional Network Functions:

- AF Application Function
General functions interacting with the PCF and routing.
- DN Data Network.
The DN or Packet Data Network (PDN) is the network targeted by the UE, for data or voice services, the Internet for instance.
- NSSF Network Slice Selection Function
A repository of information on the slices.
- RAN Radio Access Network
Handles the radio access technology and connecting the UE to the network.

B. Contents of the UE Context

Figure 2 illustrates our procedures with the information stored and transmitted during the procedures depicted between the VNFs in section IV. In order to analyse relevant parts of the user data exchanged during these procedures, we annotate the exchange diagram with relevant information only. Most of this information is necessary to allow the Network to serve one specific UE and its connectivity sessions. This is why the variables below are part of the UE Context or the Session Management Context.

- Data Network Name (DNN)
- Local Area Data Network (LADN): a DNN that may only be accessed while in physical proximity. Every AMF knows all the nearby LADNs. (This is how the 5G network plans on implementing the Mobile Edge Computing (MEC) model[11])
- Subscription Concealed Identifier (SUCI), an identifier for the client that can be broadcast in clear to the AMF while still preserving their privacy.
- Multimedia Priority Service, Mission Critical Service (MPS/MCX): a set of flags that describe how the network should prioritize the request.
- Subscription Permanent Identifier (SUPI)

- Packet Data Network (PDN)
- Permanent Equipment Identifier (PEI, e.g. IMEI)
- Network Slice Selection Assistance Information (NS-SAI), Slice information and a list of *subscribed DNN* that comes with it.
- Single NSSAI (S-NSSAI) Designates a slice selected through the NSSAI.
- Packet Data Unit (PDU), a unit of information that is exchanged between the UE and any PDN that the user is subscribed to (i.e. IPv4, IPv6). PDU connectivity is the service that provides the exchange of PDUs between the UE and a PDN identified by its DNN by establishing a *PDU Session*. [1]
- PDU Session is a logical connection set up between the UE and the PDN.
- Slicing Change Indicator. Signals to the VNF that the subscription data for network slicing has been modified when its value changes and the UE configuration must be updated.
- Steering of Roaming, the information used by the UE to know which foreign networks it should connect to while abroad.
- Mobility Restrictions restricts mobility handling and service access for the UE (e.g. access closed to groups, forbidden areas)
- QoS Quality of Service.
- SM Context Session Management Context: all the information pertaining to current PDU sessions and the set up of future sessions.
- Subscriber Data Management (SDM): all the information pertaining to a Subscriber.
- Tunnel Info The information necessary to route the packets to the correct UPF, PDN and UE.
- Priority Flag a flag set in the header of the session establishment messages when appropriate according to the MPS/MCX priorities.

IV. 5G PROCEDURES ANALYSIS FOR USER DATA

The 5G specification [1], [2] introduces the connectivity procedures with exhaustive considerations to include all the edge cases. We have streamlined the procedures and considered a simpler user model that involves all the variables to study from the UE context, in a typical connectivity scenario.

Message exchange diagrams of Figures 2 and 3, display each VNF involved side by side with arrows to describe the exchange of messages or an operation. The vertical boxes show that a VNF is invoking a local procedure, as well as the information used during this local operation. The Figures 2 and 3 and the algorithms 1,2, and 3 describe the same procedures through a common numbering system.

Algorithms 2 and 3, show the procedures in a pseudo-code form. This allows us to decorate all messages and operations with the data that is sent over as an argument to better analyse the cases of data consistency violation. In particular, we highlight the lifecycle of the data, and what happens in failure scenarios.




Fig. 2. UE Registration Procedure[2]

A. Registration Procedure

The first *registration* procedure is used when a UE joins the network for the first time. A migration to another part of the network would be considered a *handover*, or *mobility registration*. The following steps are illustrated in Figure 2.

(1): The UE sends identification data to the AMF. The SUCI will be the only identifier that is sent without first being encrypted.

(2): The AMF uses the SUCI as an index to select the AUSF that is responsible for this Subscriber.

(3): The AUSF requests the authentication data concerning this subscriber from the UDM. It then authenticates the UE and, if successful, installs a security context on the AMF and the UE. This Security Context allows all further communications between the UE and the network to be encrypted. The AUSF also provides the Subscriber's SUPI to the AMF.

(4): Using this new identifier (i.e. SUPI), the AMF is able to select a UDM that is responsible for the UE context.

(5): The AMF contacts the UDM and requests all the information necessary to setup a local copy of the UE Context. Some of this information will be sent over to the UE in (8).

(6) The AMF subscribes to all future changes to the data it just obtained. Some of this data is likely to be updated and it is important for the system to ensure that all the VNFs see an up to date version. This is why step (5) returned a Slicing Change Indicator, that will allow the UDM to inform its subscribers and the UE that their version of the data is out of date.

(7) The AMF selects a PCF responsible for the UE. It will be called upon later. (See 9, 10)

(8) The AMF returns some of the UE Context it just pulled from the UDM to the UE along with an acceptance of its registration. At this point the registration is complete.

(9, 10) The AMF contacts the PCF chosen at (7) and registers the UE. If the PCF decides that the UE policies the UE has last seen are out of date, it sends an update.

B. PDU Session Establishment Procedure

The PDU session establishment is used when the UE wants to establish a data session (PDU) with any DN it is subscribed to. The first step is the establishment of the PDU session itself.

Algorithm 1 Registration Request

Network Functions: UE, AMF, AUSF, PCF, UDM

(1) Registration Request: UE \rightarrow AMF [SUCI, PEI, NSSAI]

(2) AMF **procedure** AUSF Selection (SUCI)

(3) Authentication: UE \rightleftharpoons AMF \rightleftharpoons AUSF \rightleftharpoons UDM [SUPI, Security Context]

(4) AMF **procedure** UDM Selection (SUPI)

(5) Get Subscriber Data: UDM \rightarrow AMF [NSSAI infos, Subscribed DNNs, Mobility Restrictions, Steering of Roaming, Subscribed LADNs, MPS priority, Slicing Change Indicator]

(6) SDM Subscribe: AMF \rightarrow UDM

(7) AMF **procedure** PCF Selection ()

(8) Registration Accept: AMF \rightarrow UE [Mobility Restrictions, PDU Status, Allowed NSSAIs, LADN Info, Slicing Change Indicator]

(9) UE Policy Association: AMF \rightarrow PCF []

(10) UE Policies Updates: PCF \rightarrow AMF \rightarrow UE

(1) The UE requests a PDU session to a data network identified over a specific slice. The UE chose this slice because of the guarantees it offers over the data connection. The UE also chooses a PDUid that will identify this PDU session later.

(2) According to the data network that must be reached and the slice chosen, the AMF selects an SMF to handle this session. The SMF will not change for the lifetime of the session.

(3) The SMF received a request to setup a data plane connection from the DNN to the UE Location. If the DNN is a LADN, the AMF also informs the SMF whether the UE is close enough to connect. The SMF will create a Session Context to store the information regarding this session and the AMF subscribes to modifications to it during this request.

(4, 5) The SMF requests a portion of the UE Context that is session management specific (not kept by the AMF) from the UDM. The information is the quality of service that is required when accessing the target data network over this slice, the PDU types that are allowed for this UE and, when it applies, whether this specific UE has a static IP address (or a static IP prefix) for this data network. Should this information be modified, the SMF will be notified.

(6) The SMF returns its internal context id for the data session context to the AMF.

(7,8,9) The SMF selects a PCF that will be responsible for this PDU Session and request the policy rules before the establishment of the session. (See 18-24 for future policy updates)

(10) The SMF selects one or several UPFs. The UPF will provide the user plane connectivity and form a data pipeline: The first one to be chosen will not change during the lifetime of the session (like the SMF). It will connect directly to the data network and act as a session anchor (see PSA): when an IP address is chosen, it will point to this UPF. Zero or more additional UPFs can be chosen, these Intermediate UPF (I-UPF) will form a tunnel between the PSA and the UE. The UPFs are chosen according to their location, the location of

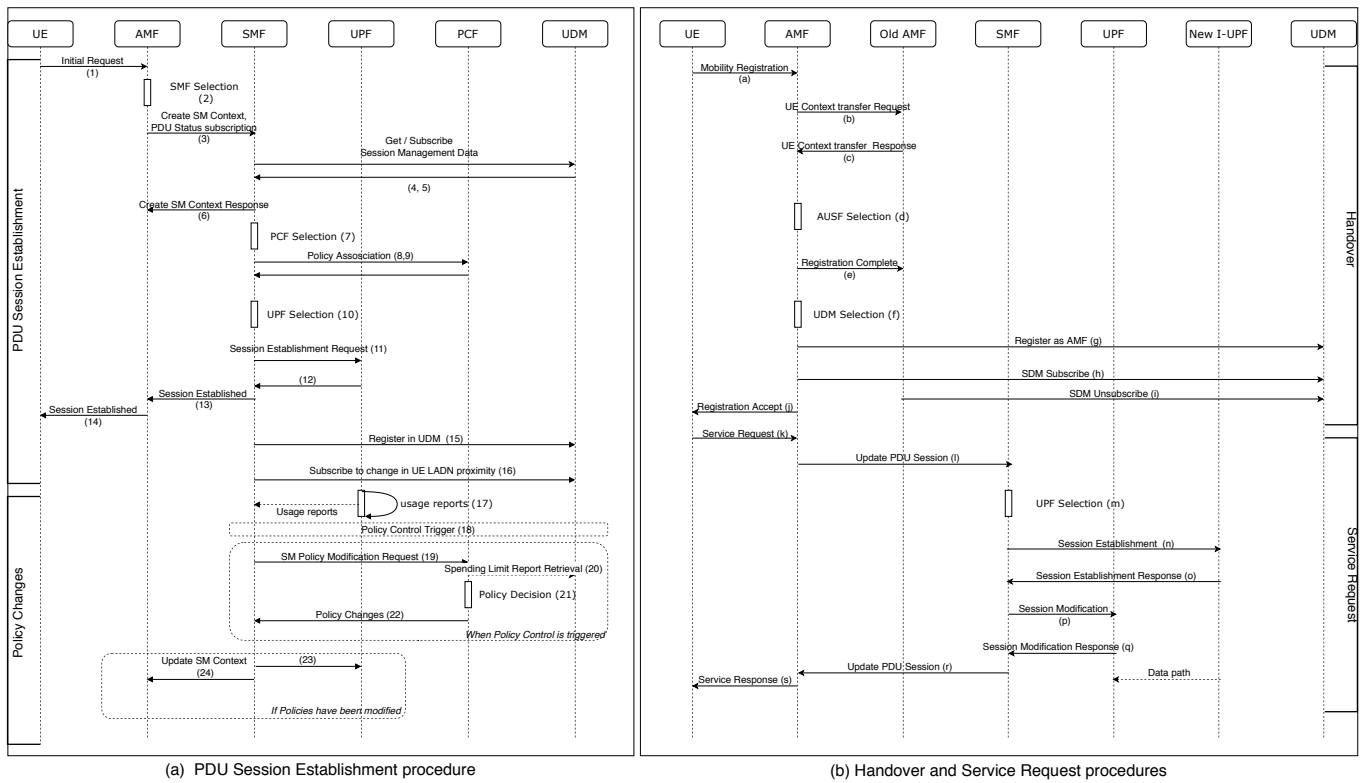


Fig. 3. Session Establishment and Service Request procedures[2]

the UE and their current workload.

(11,12) The SMF request the setup of a data tunnel from the UPF and gives them some rules on how and how often to report the PDU session data usage. They return information that will allow the UE to send data along this newly allocated tunnel.

(13,14) information is returned to the UE and the AMF on the tunnel address, the QoS profile of the newly established connection and the IP address of the connection in the PDN.

(15) The SMF registers itself in the SDM data of the UDM, along with the PDN name, the subscriber id and the session id.

(16) The SMF subscribes so that if the UE should be registered by the AMF in another location it will be notified.

After the establishment of the session, regular usage reports and policy checks are done for the PDU session.

(17) According to the reporting rules sent during step 11, every UPF sends regular reports to the SMF.

(18) Three mechanisms can trigger a policy control: (i) The SMF has received enough usage reports to trigger a check or it has received a notification from the UDM for a location change in the UE. (ii) An internal, unspecified event, in the PCF triggers the policy control. (iii) The UDM sends a notification to the SMF that triggers a policy control.

(19) When the triggers comes from the SMF, it sends a request to the PCF for an update.

(20, 21) The PCF retrieves a spending limit from the UDM and makes a policy decision accordingly.

(22) The PCF informs the SMF of any changes to the policy.

(23,24) If change must be made to the session to comply with the policy, the SMF makes the necessary changes to the data plane and informs the AMF of any notifications.

C. Handover Procedure

The handover procedure is executed when a UE that is already connected and authenticated in the network connects to a RAN that is the responsibility of a different AMF than the one the UE is currently served by. For instance, when a UE moves closer to a new antenna, it will connect to and register with this new antenna.

During the handover procedure, all the data that was held by the old AMF is sent over to the new AMF who populates its UE context with it. The AUSF re-authenticates the UE as needed.

Once the registration is complete, the two AMF modify the UE context information that is stored in the UDM by registering a new AMF id. The serving AMF must be kept up to date on the subscriber data, since other VNFs may modify it. This is why the AMF subscribes to be notified of all changes and the previous AMF must unsubscribe from the modifications it is no longer interested in. It also deletes the UE Context

(a) The UE requests to be registered with the new AMF, in the same fashion as during initial registration. But because it was previously connected to the system it has a local copy of some of the data in its UE Context.

Algorithm 2 PDU Session Establishment

Network Functions: UE, AMF, SMF, UPF, PCF, UDM

1 Initial Request: UE \rightarrow AMF [S-NSSAI, DNN, new PDU Session id]
2 AMF **procedure** SMF Selection (S-NSSAI, DNN)
3 Create Session Management Context and PDU Status Subscription: AMF \rightarrow SMF [SUPI,DNN,S-NSSAI,PDU Session id, UE location, LADN proximity flag]
4 Get and Subscribe Session Management Data: SMF \rightarrow UDM
5 Session Management Data: UDM \rightarrow SMF : [Allowed PDU types, QoS for DNN and S-NSSAI, Static UE IP]
6 Create SM Context Response: SMF \rightarrow AMF : [SM Context id]
7 SMF **procedure** PCF Selection()
8 Policy Association : SMF \rightarrow PCF
9 Policy Association Response : A \rightarrow B : [Policy rules for PDU Session]
10 SMF **procedure** UPF Selection (UE Location, UPF Loads, target DNN, PDU type, S-NSSAI)
11 Session Establishment Request: SMF \rightarrow UPF : [Usage reporting rules]
12 Session Establishment Response: UPF \rightarrow SMF : [Tunnel info]
13 Session Established: SMF \rightarrow AMF : [Tunnel info, IP addr, QoS profile]
14 Session Established: AMF \rightarrow UE : [PDU Session id, Tunnel Info, IP addr, QoS profile]
15 Register in UDM: SMF \rightarrow UDM : [SUPI, DNN, PDU id, SMF id]
16 Subscribe to changes in UE LADN proximity: SMF \rightarrow UDM
17 UPF **procedure loop** send usage reports to SMF
18 **if** Policy Control Triggered **then**:
19 SM Policy Modification Request: SMF \rightarrow PCF
20 Spending Limit Report Retrieval: PCF \rightarrow UDM
21 PCF **procedure** Policy Decision()
22 Policy Changes: PCF \rightarrow SMF : [UE policies update]
endif
if Polices have been modified **then**
23 Session Modification: SMF \rightarrow UPF
24 Update SM Context: SMF \rightarrow AMF : [PDU Session id, Tunnel info, QoS Flag]
endif

- (b) The AMF uses this information to identify the previous AMF and sends a request for the UE Context.
(c) The information that the old AMF had without sharing it with the UE, is sent at this time. The old AMF now knows that a UE it was previously responsible for, has moved away.
(d) The AMF uses the SUCI to identify the AUSF responsible for the UE. If necessary, the UE may be re-authenticated by the AUSF at this point.
(e) The new AMF confirms to the old AMF that the registration is complete. At this point the old AMF is no longer

responsible for the UE.

(f,g,h) The AMF identifies the UDM where the UE Context resides. It modifies the value of the responsible AMF within the UDM to itself and registers for future updates.

(i) The old AMF unregisters from future updates. At this point, it no longer has to store the UE Context.

(j) The registration is complete. The AMF sends the UE all the information it needs, as it did during the first registration request.

The second step of the handover procedure occurs when the UE requests access to any previously configured data session. The SMF is notified and the user plane is setup in order to bring connectivity to the UE.

(k) The UE sends a request with the id of an already setup PDU Session.

(l) The AMF contacts the SMF that is associated with the PDU Session in the UE Context.

(m) According to the new UE location, the SMF selects one or several UPF that will bring connectivity to the UE. If the new UE is too far or if the LADN proximity flag is no longer true, the SMF may simply close the session.

(n,o) One or several new Intermediate UPF (I-UPF) connect to UPF that are already setup. They create a new data tunnel in the data plane.

(p,q) Some of the old UPF are no longer needed and are freed, the others are given tunneling information to connect to the new I-UPF. Even though all other may be removed from the PDU Session, the PSA is not freed for the lifetime of the session. This is because it's directly connected to the PDN that the UE is trying to reach.

(r,s) The SMF returns new Tunnel information that may be used by the UE to send and receive data over the PDU Session. The Session is now setup correctly for the UE's new location.

V. FAILURE SCENARIOS AND INVARIANTS

This section discusses how to provide consistency and face likely failure scenarios of the 5G infrastructure.

As section II described, the system cannot provide strong consistency over all operations and all variables. We must instead deduce from the procedures analysed in Section IV, where stronger consistency is *required* for the system to remain correct. Instead of considering the consistency semantics of the system globally, we intend to focus on the specific requirements of the user data to ensure that it remains consistent according to its use by the system.

Variables that are *constant* during the procedures do not require consistency constraints. They are never modified and always remain correct. We must consider the variables that are modified during the procedures we have shown. They come in three kinds.

- Variables that are always used by operations in a specific geographical zone. They never leave their geographical location. For instance, *data plane variables*: the tunneling information. Only UPFs need them to connect to each other, and the UE to make use of the tunnel.

Algorithm 3 UE Handover

Network Functions: UE, AMF, Old AMF, SMF, UPF, new I-UPF, UDM

Handover:

(a) Mobility Registration: $UE \rightarrow AMF$: [SUCI, PEI, NSSAI, PDU Status, LADN info, AMF id]

(b) UE Context transfer request: $AMF \rightarrow OldAMF$: [SUCI, PEI, NSSAI, PDU Status, LADN info, AMF id]

(c) UE Context transfer response: $OldAMF \rightarrow AMF$: [SUPI, PDU Sessions, SMF list, NSSAI, DNN, Steering of Roaming, Slicing Change Indicator]

(d) AMF **procedure** AUSF Selection(SUCI)

(e) Registration Complete: $AMF \rightarrow OldAMF$

(f) AMF **procedure** UDM Selection(SUPI)

(g) Register as AMF: $AMF \rightarrow UDM$: [AMF id]

(h) SDM subscribe: $AMF \rightarrow UDM$

(i) SDM unsubscribe: $OldAMF \rightarrow UDM$

(j) Registration Accept: $AMF \rightarrow UE$: [Mobility restrictions, PDU Status, Allowed NSSAIs, LADN Info, Slicing Change Indicator]

Service Request:

(k) Service Request: $UE \rightarrow AMF$: [PDU Status, PDU id to activate]

(l) Update PDU Session: $AMF \rightarrow SMF$: [Priority flag, UE Location, LADN proximity flag, PDU id]

(m) SMF **procedure** I-UPF Selection (UE Location, UPF loads, target DNN, PDU type, S-NSSAI)

(n) Data Consistency in the 5G Specification Editor mode.

Session Establishment: $SMF \rightarrow I-UPF$: [Tunnel info, old UPF id, reporting rules]

(o) Session Establishment Response: $I-UPF \rightarrow SMF$: [Tunnel info for I-UPF]

(p) Session Modification: $SMF \rightarrow UPF$: [Tunnel info for I-UPF, reporting rules]

(q) Session Modification Response: $UPF \rightarrow SMF$: [Tunnel info]

(r) Update PDU Session: $SMF \rightarrow AMF$: [PDU id, Tunnel info]

(s) Service Response: $AMF \rightarrow UE$: [PDU id, Tunnel info, Mobility Restrictions, S-NSSAI]

- Variables that belong in the UE Context but are modified by operations from a limited number of VNFs. They are stored and spread by the UDM, but with a limited number of actors that modify them to a central version. e.g. The AMF id: modified during handover (step g). The new AMF overwrites the previous value in the UDM.
- Variables that are inherently more distributed: A greater number of VNFs modify their own version independently. And the value of these variables must reflect the most recent writes from each of the actors. For instance: after a Policy Control trigger the PCF retrieves the current Spending Limit from the UDM (PDU session establishment step 20). The value of this spending limit is the result of the spending report received by every SMF for

every PDU session associated with the subscriber.

Despite being modified, the data plane variables are not intended to be distributed over the network. We will not consider them. To better reason on the consistency constraints for distributed variables, we present a few failure scenarios. We will delay and reorder messages, to create a worst case scenario.

A. Registration procedure failures

When some central data is modified from several actors at the same time, such as during the handover operation (step g), when the AMF id in the UDM is modified by the AMF, two issues arise:

(i) The system could temporarily be in a state where no AMF is designated for a specific UE. This could happen when the old AMF unsubscribes from the UDM before the new AMF has successfully completed the registration.

(ii) The system could also be in a state where too many AMF are trying to register. The handover occurs at the request of the UE. It starts to register to a new AMF A, when the UE is moving quickly it tries to register to AMF B before completing AMF A registration. In this situation, AMF B completes registration with the UDM. And then due to a network delay, AMF A registers with the UDM and replaces the AMF id value for the subscriber.

It is important to ensure that the AMF id remains correct, because the serving AMF is responsible for keeping the UE Context. So it should be held through a continuous chain of serving AMF, and the responsibility passed on from older AMF to newer AMF. There has to be a consensus on which AMF is the current serving AMF, so that the current one will not delete its state.

The UE can leave at any moment and the AMF are always changing, so it should be the responsibility of a *central database* to decide on this. This consensus cannot be naively reached at the edge. To ensure correctness, the handover should be handled through a transactional operation and the UDM should serve as a transaction manager with Strict Serializability properties.

B. Partial update failures

Some data within the UE Context must always be kept up to date, such as the Slicing related data. The UE Context contains a Slicing Change Indicator that reflects the current version of the data. It cannot be effective if variables in the system are accessed and updated independently. If the UDM shows one version of the Slicing Change Indicator, every other key within the same operation must be accessed at the same version.

To tackle this issue, we propose that the Slicing related data and its indicator must be accessed and modified through *Highly Available Transactions*[7]. In this case, the UDM will return the result of the last successful transaction. The modifications of currently ongoing transactions will be not be visible.

C. Usage report ordering anomalies

(i) During data plane setup, the SMF requires usage reports from the UPFs. It receives them during the lifetime of the PDU Session. (PDU session establishment, steps 11, 17) The value of the PDU Session usage depends on the last value it received from each UPF. When one of the usage reports is delayed, and at least one more recently sent report arrives first, the PDU Session usage will seem to decrease. Depending on the specific mechanisms of computing the data usage according to the reports, this may result in an incorrect view of the system state.

(ii) When the data usage of the PDU Session reaches a given threshold, the SMF request the PCF to reevaluate the PDU policies. The PCF does so according to the usage report for the subscriber (PDU session establishment, steps 18 to 20). This global information is requested from the UDM, and is subject to updates from every SMF currently serving the subscriber. Each with their own PDU Session and their own reporting UPFs.

Network usage is a constantly increasing metric, we propose that the system compute the usage through a monotonously increasing function, e.g. the maximum out of all previously received values.

VI. CONCLUSION AND FUTURE WORK

There have been few research activities on matters of data consistency in relation with the mobile networks, and fewer still take into account the full spectrum of consistency guarantees that is available in the literature. Because of the complexity and the diversity of information that the 5G system must handle, no single consistency scheme will be sufficient.

We have presented different consistency mechanisms that help the system avoid an incorrect state, and these could be generalised to data that exhibits the same properties. There have been some works on hybrid consistency databases, where a different consistency scheme is applied to different parts of the database.

We propose to further our study through experimentation. We want to extend our study of failure scenarios and consistency guarantees through a mock-up of a 5G core network deployment. We will also use Slice related scenarios, because some of the data involved is much more distributed than what is seen in the UE Context.

Another extension that should be evaluated is storing monotonous data in convergent data types[12] and ensuring that thresholds are respected through Bounded Counters[13].

VII. ACKNOWLEDGEMENT

We would like to thank Sihem Cherrared and Louiza Yala for their early review of this article. We would like to acknowledge Sara Hamouda's contribution and thank her for lending her expertise for invariants analysis.

REFERENCES

- [1] "3GPP 23.501 system architecture for the 5g system, release 16.1.0," 2019.
- [2] "3GPP 23.502 procedures for the 5g system, release 16.1.1," 2019.
- [3] S. Gilbert and N. Lynch, "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services," vol. 33, no. 2, pp. 51–59, 2002, ISSN: 0163-5700. DOI: <http://doi.acm.org/10.1145/564585.564601>.
- [4] E. Brewer, "CAP twelve years later: How the "rules" have changed," *IEEE Computer*, vol. 45, no. 2, pp. 23–29, Feb. 2012. DOI: <http://dx.doi.org/10.1109/MC.2012.37>.
- [5] F. Ojala, A. Rao, H. Flinck, and S. Tarkoma, "NoSQL stores for coreless mobile networks," in *2017 IEEE Conference on Standards for Communications and Networking (CSCN)*, IEEE, 2017. DOI: 10.1109/cscn.2017.8088622. [Online]. Available: <https://doi.org/10.1109/cscn.2017.8088622>.
- [6] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, W. Hsieh, S. Kanthak, E. Kogan, H. Li, A. Lloyd, S. Melnik, D. Mwaura, D. Nagle, S. Quinlan, R. Rao, L. Rolig, Y. Saito, M. Szymaniak, C. Taylor, R. Wang, and D. Woodford, "Spanner: Google's globally-distributed database," Hollywood, CA, USA, Oct. 2012, pp. 251–264. [Online]. Available: <https://www.usenix.org/system/files/conference/osdi12/osdi12-final-16.pdf>.
- [7] P. Bailis, A. Davidson, A. Fekete, A. Ghodsi, J. M. Hellerstein, and I. Stoica, "Highly available transactions: Virtues and limitations," vol. 7, no. 3, pp. 181–192, Nov. 2013. DOI: 10.14778/2732232.2732237. [Online]. Available: <http://dx.doi.org/10.14778/2732232.2732237>.
- [8] H. N. S. Aldin, H. Deldari, M. H. Moattar, and M. R. Ghods, "Consistency models in distributed systems: A survey on definitions, disciplines, challenges and applications," *CoRR*, vol. abs/1902.03305, 2019. arXiv: 1902.03305. [Online]. Available: <http://arxiv.org/abs/1902.03305>.
- [9] H. Berenson, P. Bernstein, J. Gray, J. Melton, E. O'Neil, and P. O'Neil, "A critique of ansi sql isolation levels," in *ACM SIGMOD Record*, ACM, vol. 24, 1995, pp. 1–10.
- [10] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels, "Dynamo: Amazon's Highly Available Key-value Store," en, p. 16,
- [11] E. M. E. Computing, I Initiative, *et al.*, "Mobile-edge computing: Introductory technical white paper," *ETSI: Sophia Antipolis, France*, pp. 1–36, 2014.
- [12] M. Shapiro, N. Pregoça, C. Baquero, and M. Zawirski, "Conflict-free Replicated Data Types," Research Report RR-7687, Jul. 2011, p. 18. [Online]. Available: <https://hal.inria.fr/inria-00609399>.
- [13] V. Balegas and U. N. de Lisboa, "Bounded Counters: Maintaining numeric invariants with high availability," en, Tech. Rep., Oct. 2016, p. 2. [Online]. Available: <https://pages.lip6.fr/Marc.Shapiro/papers/Just-Right-Consistency-RR-9145-2018-01.pdf>.