Themyscira Wireless Hera MSC Architecture

Mychaela N. Falconia

American 2G Cooperative

Version 1.24, last edited 2026/04/09

1. Introduction

American 2G Cooperative, operating under software brand Themyscira Wireless, seeks to implement a new software-based GSM MSC (Mobile services Switching Center) that shall function in the software universe of Osmocom+ThemWi. This new MSC, named Hera, is envisioned as an extension on top of OsmoMSC: the existing `osmo-msc` application, to be extended with some ThemWi-specific functions enabled via vty config, shall constitute just one internal component of the MSC as a whole. In addition to this suitably modified `osmo-msc`, the complete Hera MSC shall consist of the following components:

- A new transcoding MGW will need to be implemented, based on the existing `tw-border-mgw`, but speaking a dialect of MGCP: specifically based on the existing Osmocom dialect, but extended to support the required new functionality. This new MGW shall replace the MSC-associated `osmo-mgw` instance that comprises a part of the standard OsmoMSC setup: there will be no `osmo-mgw` in Hera MSC. OsmoMSC will need to be extended with a new vty config option that selects the type of MGCP-speaking MGW to be driven by the `osmo-msc` process: `osmo-mgw` (rtpbridge) or `hera-tc-mgw` (transcoder).
- There will be a collection of Hera-specific components, communicating via sockets, dedicated to call control functions. Multiple software components, each constituting its own Osmocom-style long-lived daemon process with vty, are needed in order to handle different types of calls: MO calls from GSM to SIP-I, MT calls from SIP-I to GSM, MO calls to special test numbers terminating inside the MSC, test MT calls made from the command line, MO CSD calls to special numbers that operate as IWF gateways to Telnet or to PPP, and possibly other types of calls not yet envisioned. All of these call handler components will connect to `hera-twcc` process, which will in turn connect to `osmo-msc` operating in the special vty-configured ThemWi mode. The interface between modified OsmoMSC and `hera-twcc` shall be TWCC (Themyscira Wireless Call Control), similar in principle to the existing Osmocom MNCC, but different in details.

1.1. Required changes to OsmoMSC

We currently anticipate needing the following changes to OsmoMSC in order for this pre-existing sw component to function as the building block needed for Hera MSC:

- A new vty config option needs to be added, selecting the type of MGCP-speaking MGW to be driven by the `osmo-msc` process: either **rtpbridge**, intended for the existing `osmo-mgw`, or **transcoder**, intended for `hera-tc-mgw`. The latter MGW will require some minor changes to the dialect of MGCP, and OsmoMSC needs to support these MGCP dialect modifications.
- In the case of standard voice calls that get transcoded to G.711 by `hera-tc-mgw`, DTMF insertion needs to happen within that MGW, controlled by our invented **XSIG** MGCP command. OsmoMSC needs to be taught to handle DTMF START and DTMF STOP requests from GSM CC by issuing this non-standard command to its slave MGW.
- CC handling code in OsmoMSC will need to be extended to support both the existing form of MNCC and TWCC. See §6.1 for more information regarding TWCC.

All patches to OsmoMSC that will be needed per the above summary will be developed on a feature branch. Our intent is to implement our changes in such way that they can eventually be merged into Osmocom mainline: make all behavioral and functional changes conditional on new vty config options, such that in the default configuration without ThemWi-specific vty settings, no behavioral or functional changes occur, relative to the current standard Osmocom way. However, if the desired merge into Osmocom mainline does not happen for whatever reason, American 2G Cooperative shall continue indefinitely with using our ThemWi branch version of OsmoMSC on our production network.

2. Scope of MSC functionality

Before embarking on the architectural design of a GSM MSC, it is important that all involved parties agree on the terminology: exactly what is an MSC? In other words, exactly what functionality constitutes a single MSC, versus what functionality falls outside the MSC? In the view of Themyscira Wireless, the functional scope of a single MSC for the purpose of voice and CSD services extends from the A interface to the local BSS (or multiple BSS served by the same MSC) to the back end interface that connects to outside PSTN.

2.1. Definition of PSTN

At Themyscira Wireless, we define the term PSTN (Public Switched Telephone Network) to mean the total set of all telephone destinations worldwide, *irrespective of technology*, that can be reached by dialing an E.164 number and paying any appropriate tolls. In contrast, many other participants in today's telecom industry object to this view and instead regard the term PSTN as derogatory, referring solely to analog or TDM or other specific technologies which they happen to not use. When we describe the problem of needing to build a GSM MSC that interfaces properly to PSTN at its back end, the latter group will often retort with statements like ‘‘we don't connect to PSTN, we connect to VoIP'', or ‘‘we connect to SIP'', or ‘‘we connect to other mobile networks'' and so forth, anything other than PSTN. However, those statements are disingenuous: if your GSM (or UMTS or IMS) subscribers can be reached by dialing an E.164 phone number, then your system is part of worldwide PSTN, even if you hate that term.

Our expansive definition of PSTN allows us to clearly describe the problem in need of solving, and we are sticking with it.

2.2. The need for transcoding

G.711 is the *lingua franca* of PSTN, the common language spoken by all interconnected parties on the global phone network. Each given end system may internally use GSM with its range of speech codecs that depend on handset capabilities, UMTS with its singular prescription of AMR as the sole codec for traditional narrowband telephony, VoIP with its own selection of compressed codecs (such as G.729), or whatever codecs are used in the world of IMS — yet for global interoperability on the worldwide PSTN, every system must transcode its internal speech codecs to G.711, or at least have the ability to do so when its more preferred TrFO paths fail.

The task of transcoding GSM or UMTS native speech codecs to G.711 is sometimes outsourced by Osmocom network operators to opaque third parties. If you disable all use of EFR and AMR, thereby forcing all GSM calls to the lowest common denominator of FRv1, and your SIP-based IP-PSTN connectivity provider happens to support ‘‘GSM'' VoIP codec (practically meaning that the provider's equipment transcodes from GSM 06.10 to G.711), you can choose to outsource the necessary transcoding function to that PSTN-via-SIP connectivity provider. A similar situation sometimes occurs with AMR, which some PSTN-via-SIP connectivity providers choose to offer as a supported codec in SIP, once again doing the transcoding to G.711 themselves.

However, this outsourcing approach is unattractive to American 2G Cooperative, a GSM network operator seeking to provide the same quality of service that was provided by legacy operators whose GSM networks already were or are in the process of being wrongfully shut down. The approach of restricting our interconnection choices only to those PSTN-via-SIP access providers who happen to support either GSM 06.10 or AMR would be at odds with economic interests (we would rather not pass up otherwise good providers who only offer standard G.711), and the approach of restricting GSM air interface speech codec selection to match the capabilities of the external connectivity provider would be at odds with the way traditional high-quality GSM networks operate, selecting the air interface codec of the highest voice quality supported by the MS. Therefore, we are committed to doing our own transcoding of all GSM calls to G.711 under our own control as our primary voice path.

2.3. Contrast with current Osmocom usage

Those current operators of Osmocom GSM (or UMTS) networks who run OsmoMSC (as opposed to parties who found themselves in the highly exotic circumstance of being allowed to connect their Osmocom BSS to some big operator's existing MSCs via AoIP) and who do their own transcoding, as opposed to outsourcing that task to their upstream SIP-to-PSTN connectivity provider, use some piece of non-GSM-specific, off-the-shelf PBX or SIP software (Asterisk, FreeSWITCH, Kamailio) to perform this critical function. As a result, that third-party SIP software, which was never meant to function as a GSM MSC or a component thereof, becomes a part of the complete MSC in reality!

Users/operators of Osmocom CNI software who find themselves in this situation will only think of OsmoMSC as the answer to the question of ''what MSC are you using?'', but in reality their complete MSC (as we define this term) includes not only `osmo-msc` process, but also `osmo-sip-connector` and whatever PBX or other software they are using to interface to IP-PSTN with transcoding! When we say that we seek to replace this status quo with our own superior MSC implementation, it is primarily the latter part which we seek to replace, rather than the basic core functions contained in `osmo-msc` process.

2.4. GSM 09.07 interworking for CSD

In order to be no worse than the legacy network of T-Mobile USA, the new GSM network of American 2G Cooperative must support CSD calls to and from regular ''voice'' phone numbers on outside PSTN, speaking waveforms on the external PSTN interface that emulate a V-series analog modem — functionality specified in GSM 09.07 §9.2 and §9.3. Similar functionality for fax also needs to be supported, as specified in GSM 03.45.

We also desire to support interworking of CSD calls to external calls in ISDN UDI format (GSM 09.07 §10.2), carried in CLEARMODE with SIP-I signaling. This facility will allow us to support mobile-to-mobile CSD calls (e.g., for end-to-end encrypted speech applications) in the same way how they worked in the original architecture in the glory days of ISDN and GSM, where the two connected mobile users may be on different MSCs or even on different MNOs in different countries, as opposed to stock Osmocom solution of raw RTP cross-connect within a single MSC. We also intend to find a way to connect these SIP-I calls in ISDN UDI format to OCTOI DIVF, thereby resurrecting the truest original form of data call interworking between GSM and ISDN.

2.5. Front and back sides of an MSC

For ease of visualization, we shall refer to the side of an MSC that connects to BSS (A interface) as the front, and the side that connects to PSTN or ISDN as the back. This orientation of front vs back is consistent with terminology used in modular RAN designs: the older term backhaul refers to transport between RAN nodes and the operator's CN, while more recently introduced ''front haul'' refers to the interface between baseband and RF parts of a single BTS. In the case of widely accepted RAN terminology, the front is closer to mobile users than the back — and our own terminology for front and back sides of an MSC does likewise.

3. Routing of the user plane inside an MSC

3.1. Traditional hardware-based MSCs for TDM-based GSM

In the original classic architecture, GSM networks were built on top of TDM and all network elements had to be implemented in hardware rather than software. In this architecture the Ater interface coming from physically distant BSC sites would go to the MSC-colocated bank of TRAUs first. Each A interface user plane channel identifiable by a given CIC (Circuit Identity Code) on BSSMAP would pass through its own dedicated TRAU instance inside the TRAU bank, expanding from a 16 kbit/s or 8 kbit/s subslot on Ater into a full 64 kbit/s channel going into the MSC proper. The MSC proper that follows was a repurposed ISDN switch, and as such it had a switch fabric that operated on 64 kbit/s channels, dynamically cross-connecting them as call traffic demands. The back side of this switch fabric connects to TDM-PSTN or ISDN trunks, producing a structure that may be visualized like this:

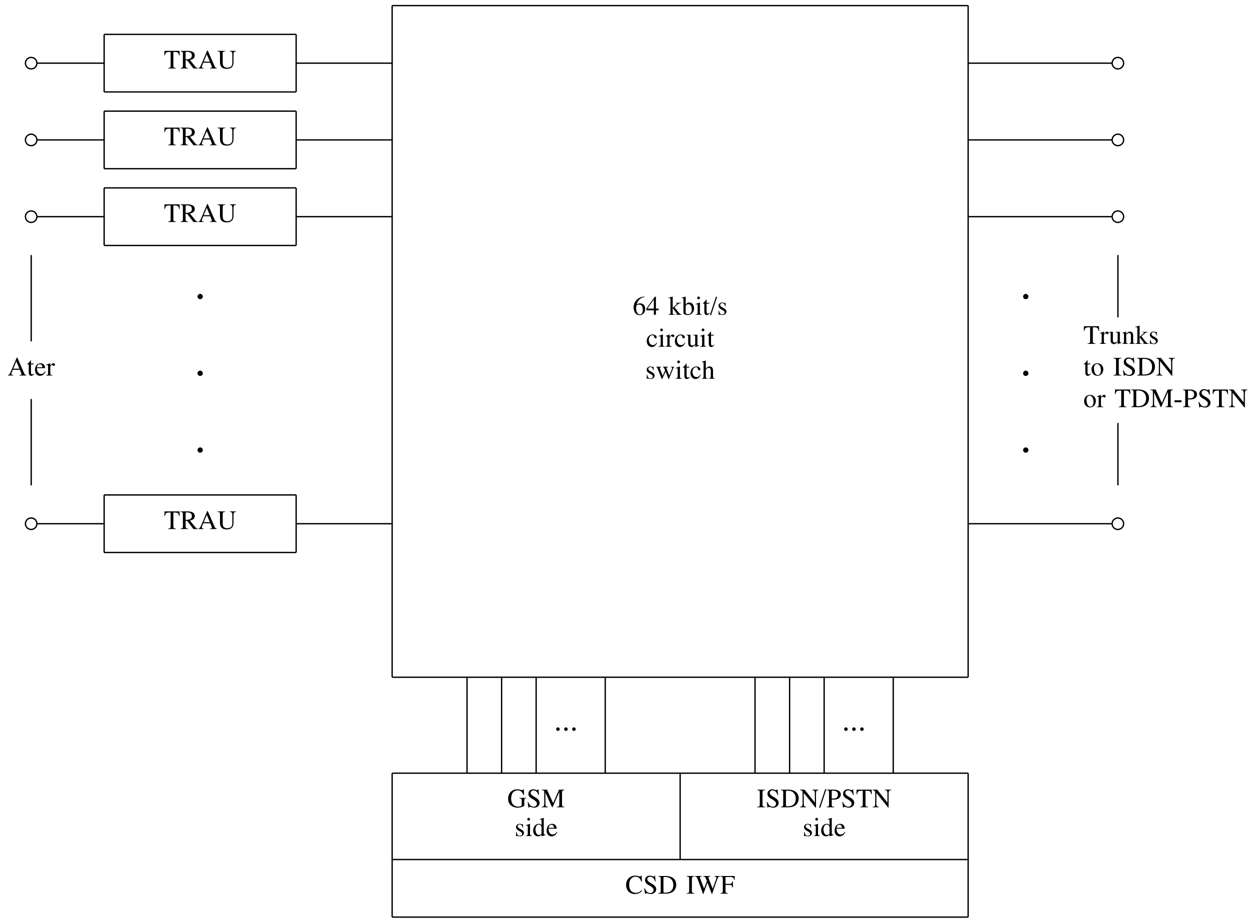

Figure 1: User plane routing in a traditional TDM-based MSC

The MSC depicted in Figure 1 supports both voice and CSD calls, which require different routing of the user plane:

- For voice calls, the switch fabric connects the 64 kbit/s A interface circuit coming out of the TRAU directly to an ISDN or TDM-PSTN trunk channel (DS0 time slot) on the back side of the MSC. Additional functions not shown in the figure are DTMF insertion (a function specific to GSM) and insertion of call progress tones, a function in common with traditional end office switches in land line telco world. But these tone generators connect to the 64 kbit/s voice path (A-law or µ-law) only as needed, not for the entire duration of the call.

  Some CSD calls (those that don't require an IWF) can be handled in the same way — but this manner of CSD handling is restricted to the case of Transparent connection element on GSM side, UDI mode on ISDN side, and user data rates of 9.6 kbit/s or less, excluding 14.4 kbit/s mode.

- All other types of CSD calls beyond the single special case detailed above require an IWF. In this MSC architecture the needed IWF is a separate hardware resource that has to be provided, and each MSC will have a finite number of these IWF instances. (This implementation cost aspect may have prompted some operators to charge extra for CSD calls compared to regular voice.) The 64 kbit/s channel from the TRAU (which can carry standard V.110, modified pseudo-V.110 frames for NT service, or A-TRAU frames of GSM 08.20 in the case of 14.4 kbit/s services) gets connected to the GSM side of the allocated IWF instance, while the other side of the same IWF connects to the outside world.

It also needs to be noted that the two DSP-intensive operations (speech transcoding and V-series modem emulation) happen in different places in this architecture. In the case of voice calls, TRAUs perform DSP-intensive transcoding, while IWF hardware resources are not needed at all. But in the case of CSD a very different situation arises: the work performed by the TRAU assigned to the CSD call reduces to almost nothing, while an entirely separate hardware resource gets allocated to perform the other DSP-intensive operation.

3.2. User plane in Hera MSC

Our Hera MSC will be implemented entirely in software, as opposed to hardware, with RTP streams taking the place of various physical circuit links in Figure 1. However, despite these significant differences in physical implementation, the logical or conceptual architecture remains very similar to this golden reference. Specifically, we intend to implement the following architecture in the user plane:

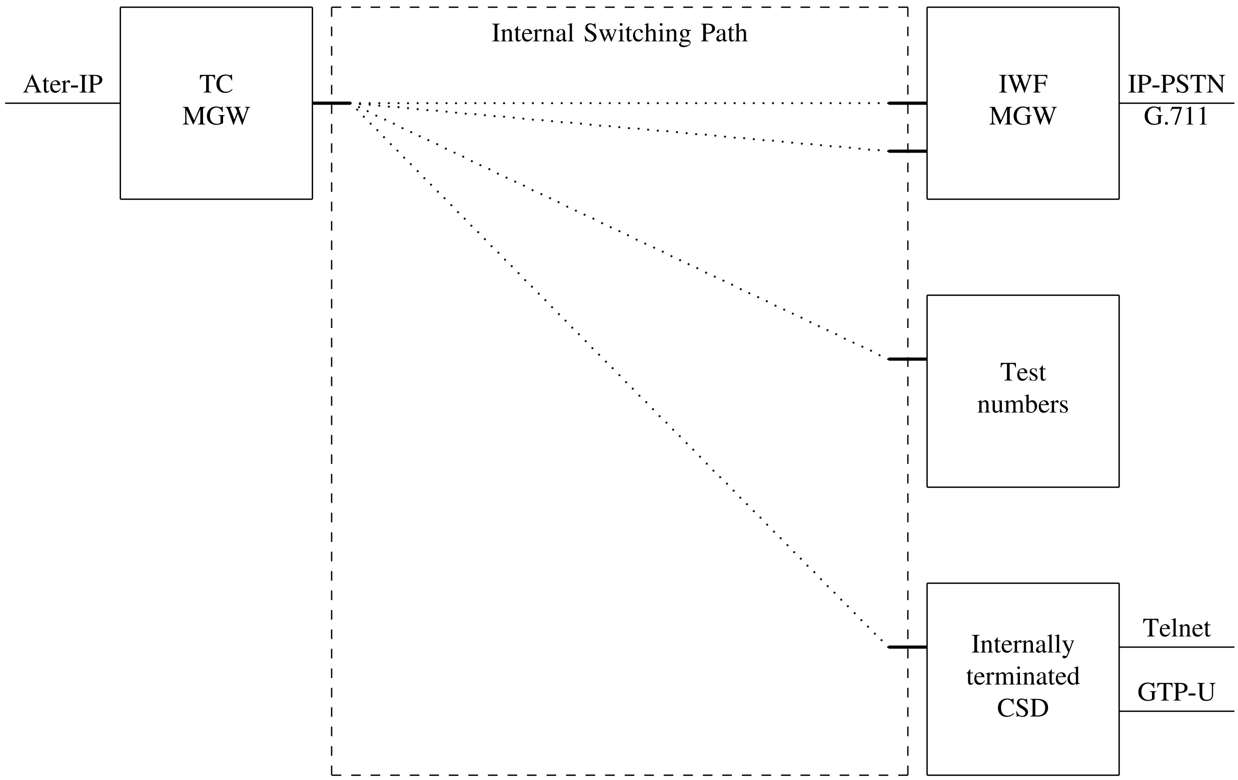

Figure 2: User plane routing in Hera MSC

One key aspect of Hera MSC user plane traffic path as depicted above is that every regular call that goes out to IP-PSTN passes through *two* custom MGWs: the transcoding (TC) MGW at the front of the MSC and the IWF MGW at the back of the MSC. This design stands in contrast to our previous one (ThemWi interim) in which there is only one `tw-border-mgw`. Components of Hera MSC architecture depicted in Figure 2 correspond to those of a traditional (hardware, TDM-based) MSC depicted in Figure 1 as follows:

- Hera TC MGW (an MGCP-speaking MGW to be controlled by `osmo-msc` process, interfacing to AoIP in the place of `osmo-mgw`) corresponds to an MSC-colocated TRAU bank in the traditional hw architecture. Voice calls get transcoded from GSM speech codecs to G.711 in this MGW, whereas for CSD calls it acts as a transparent (stateless) RTP forwarder with no added delay.
- The internal switching path (ISP) shown as a dashed box in Figure 2 corresponds to the 64 kbit/s circuit switching fabric of a traditional MSC. In physical realization, this ISP is an RTP leg that runs between Hera TC MGW on one end, and either Hera IWF MGW or one of the specialized call termination components shown in Figure 2 on the other end. The payload format carried on this RTP leg is G.711 PCMU or PCMA for voice calls, or TW-TS-007 for CSD.

  Note that TW-TS-007 is optional on the AoIP leg, with the legacy 160-octet format also supported for backward compatibility. However, because ISP is both a private interface internal to our MSC and an entirely new one, the use of TW-TS-007 for CSD is mandatory here: 64 kbit/s CLEARMODE won't be supported on ISP side by any of the components on the right side of Figure 2.

- Hera IWF MGW is conceptually modeled after the CSD-specific IWF box shown at the bottom of Figure 1, but it is not a direct correspondence. In Hera MSC architecture this IWF MGW is present in the user plane traffic path for all calls, both voice and CSD, but for voice calls it acts as a transparent (stateless) RTP forwarder with no added delay, whereas in the case of CSD it does real work. This IWF MGW also plays the announcement of ''this subscriber is not available, please try your call again later'' toward IP-PSTN (when it is not even known whether the subscriber would have accepted this call in voice or fax or modem emulation mode if they were reachable!), and for MO voice calls only, it plays the North American in-band ringing tone toward the local mobile user if the called party across SIP-I indicates ringing state, but does not provide in-band tones aka early media.

3.2.1. Ater-IP interface

In the original TDM-based GSM architecture, Ater is the interface between a BSC switch fabric (one that switches 16 kbit/s or 8 kbit/s channels between Abis and Ater) and a bank of TRAUs that sits between the BSC and the MSC, usually colocated with the latter. Traditional TDM-based Ater carries the same TRAU frame formats (and thus the same semantics) as better-known Abis, but is designed to be carried all the way to MSC sites, such that speech transcoding and bandwidth expansion from 16 kbit/s to 64 kbit/s per channel happen in the latter place. In Themyscira Wireless implementation of GSM, spanning from the MSC to the BSS, we introduce a conceptually similar interface which we call Ater-IP.

By design, Themyscira Wireless CN software works only with Osmocom BSS and no other. Thanks to this restriction (more precisely, exclusion of IP-based GSM RAN components from vendors other than Osmocom, while fully welcoming legacy E1-based components), we don't have to maintain compatibility with the crippled AoIP user plane format introduced in 3GPP Release 8 — instead we have the freedom to invent our own RTP payload formats that carry the semantics of traditional Abis or Ater over IP, as long as we can also implement them in Osmocom BSS.

At the beginning of Hera MSC project, Osmocom BSS already includes full support for ThemWi enhanced RTP formats for GSM speech codecs (TW-TS-001 for FR and EFR, TW-TS-002 for HRv1 and TW-TS-006 for AMR), as well as TW-TS-007 compressed RTP format for CSD that also stipulates 20 ms alignment equivalent to that in TRAU frames. OsmoBSC accepts a BSSMAP extension IE (defined in TW-TS-003) that switches RTP payload formats from the standard 3GPP/IETF set to ThemWi alternatives, and OsmoMSC features a vty option that emits this BSSMAP extension IE and thereby enables ThemWi RTP formats. AoIP compressed speech interface that has been modified in this manner shall henceforth be called Ater-IP, and it is the only form of A interface supported by Hera MSC.

In speech transcoding mode, Hera TC MGW accepts only TW-TS-001 extended RTP format for FR and EFR, corresponding to TRAU-UL format of GSM 08.60. If and when support gets added for HRv1 or for AMR, TW-TS-002 and TW-TS-006 formats with radio Rx semantics will likewise be expected. However, a more permissive model is applied to CSD: UL traffic from BSS can arrive in either TW-TS-007 compressed format (preferred), or in the bloated CLEARMODE format with Osmocom-originating restriction of 20 ms alignment. This permissive model has been chosen for CSD because both formats carry the same information content and can be statelessly interconverted by calling already implemented library functions, and this permissiveness allows use of CSD on Osmocom BSS deployments with `osmo-bts-trx` and `osmo-bsc` operating with standard `osmo-mgw`, in addition to the primary target of Osmocom BSS with E1 BTS and `tw-e1abis-mgw`.

3.2.2. Transformation and retiming of RTP streams at MGWs

In the case of regular calls between the local GSM network and the outside world, the complete RTP path through the MSC, spanning between Ater-IP and IP-PSTN, always involves some transformation, usually very substantial:

- Speech calls are always transcoded between G.711 on IP-PSTN and GSM speech codecs on Ater-IP;
- CSD calls that interwork to analog fax or data modem emulation undergo very substantial content transformation inherent in that emulation;
- CSD calls that interwork to ISDN UDI also undergo a light transformation. CLEARMODE RTP streams on IP-PSTN side are presumed to potentially interwork to real ISDN B-channels at some point, and unlike the situation on Osmocom AoIP, there is no requirement that V.110 frames be perfectly aligned within

20 ms CLEARMODE RTP packets. Thus at the minimum, Hera IWF MGW must hunt for V.110 frame alignment and buffer portions of frames in the case of misalignment in the IP-PSTN to local GSM direction. Furthermore, in the case of 14.4 kbit/s transparent CSD services, the format on the leg between BSS and IWF is different from V.110 (see GSM 08.20), thus conversion is needed in both directions.

In addition to these content conversions, each RTP path through the MSC is also retimed: both the RTP stream that goes to BSS and the one that goes to IP-PSTN are timed by the MSC's local clock (more precisely, the local clock of one of the two Hera MSC MGWs, as will be explained shortly), as opposed to preserving the timing of the incoming RTP stream from the opposite side of the MSC. In the case of speech transcoding, the only correct way to perform this operation in the presence of unpredictable jitter or packet loss on RTP input to the TC is to pass that input through twjit, thereby retiming it to the transcoder's local time base, and then emit the output as an entirely new RTP stream with the TC's own SSRC. The same considerations apply to the complex IWF for analog modem emulation. In the case of CSD interworking to ISDN UDI, the need for retiming is less strict — but our current plan is to be consistent with other CSD IWF modes.

The problem of RTP retiming would have been simple and straightforward if we had only one MGW at our MSC, with one side connected to Ater-IP and the other side connected to IP-PSTN. Our previous `tw-border-mgw` design was just that — but as we progress toward a more complete MSC implementation fit for production deployment with more users than just the author, we have come to realize that the more complex architecture shown in Figure 2, modeled after the classic hardware MSC architecture of Figure 1, is necessary after all.

The design goal for Hera MSC is to have only one retiming for every RTP path through our complete MSC in either direction, despite having separate front and back MGWs connected by ISP. This goal is to be achieved as follows:

- For regular voice calls, or during speech portion of speech/data alternating calls to or from G.711 PSTN, Hera TC MGW does the retiming while Hera IWF MGW acts as a no-delay forwarder.
- For CSD calls to or from PSTN or ISDN, or during data portion of speech/data alternating calls, the opposite situation holds: Hera IWF MGW does the retiming while Hera TC MGW acts as a no-delay forwarder.

Each of the two MGWs thus must be capable of acting either as a retiming agent or as a no-delay RTP forwarder depending on current operating mode; but for every possible type of call, higher-level control logic ensures that one MGW retimes while the other forwards without delay. Whenever an MGW acts as a no-delay RTP forwarder, every received packet that is deemed valid (not rejected as a bogon) is passed through with RTP timestamp, sequence number and SSRC unchanged; the payload may also be unchanged, or it may be subject to simple transformations — but only stateless transformations are possible in this operating mode. OTOH, when an MGW acts as a retiming agent, its output is an entirely new RTP stream: the MGW emits its own SSRC and generates its own RTP timestamps and sequence numbers as if it were the ultimate source of this RTP stream.

The implications for parties seeing the overall behavior of the MSC on the defined interfaces (Ater-IP and IP-PSTN) are as follows:

- All BSS endpoints (`tw-e1abis-mgw` or IP-native OsmoBTS) always see timing generated by whichever MSC they are connected to, insulated from whatever packet loss or jitter occurs on IP-PSTN traffic arriving at that MSC.
- These BSS endpoints will see SSRC coming from either Hera TC MGW or Hera IWF MGW, depending on the type of call.
- When an alternating call switches between speech and data or fax, BSS endpoints will see an RTP stream discontinuity and a change of SSRC. However, because these switches between speech and CSD on GSM side are also Channel Mode Modify operations that propagate to the BTS and to the MS, no additional disruption occurs beyond what is inherent in the operation itself.
- All IP-PSTN peers always see RTP timing generated by the MSC. Whatever packet loss or jitter occurs on RAN transport links that carry traffic from cell sites to MSC sites is fully concealed from conversing parties on IP-PSTN. This concealment happens in all modes: regular voice, CSD to analog modem emulation and CSD to ISDN UDI.
- When an alternating call switches between speech and data or fax, the IP-PSTN peer will see a brief interruption in the G.711 RTP stream we emit, a change in SSRC (along with a switch to a different random

base for RTP timestamps and sequence numbers), and likely a small change in timing due to our TC and IWF MGWs having independent local clocks. Since a fundamental change in the nature of G.711 audio processing occurs at the same point, there is no possibility for human users to perceive any additional glitch or degradation — but if someone is spying on that G.711 IP-PSTN connection, they will notice an oddity.

Since RTP discontinuities and SSRC changes are generally allowed on IP-PSTN, subject only to unknown and unpredictable short-term effects on jitter buffers in various IP-PSTN middleboxes, the present blemish is deemed to be acceptable.

3.3. Considerations for call waiting

In order to be no worse than the many large commercial GSM networks that are being wrongfully shut down around the world, the new GSM network of American 2G Cooperative must support call waiting. Suppose that local GSM user Alice is on the phone with external party Bob, and in the middle of that conversation a new call arrives from Charlie, addressed to Alice's MSISDN. With all traditional GSM networks the following sequence ensues:

- The network sends a new MT SETUP CC message to Alice's MS with a new transaction ID, distinguishing the newly presented call from the one already in progress.
- The MS responds with CALL CONFIRM and starts mixing CW (call waiting) tones into its earpiece output, alerting Alice that she got another call. Charlie's number (the new calling party) from the new MT SETUP message will appear on the phone display.
- Alice can choose to ignore Charlie's call, reject it, or accept it. If she accepts Charlie's call, the existing conversation with Bob is placed on hold.
- At any time while Alice is talking to Charlie and has Bob on hold or vice-versa, she can press a ''switch calls'' button in her phone handset UI. Every time she does so, her MS executes CALL HOLD procedure on the currently active call, then CALL RETRIEVE on the other call.

What support is required in MSCs for this CW functionality to work as just described? Referring to Figure 1 (the block diagram of a traditional TDM-based MSC), because the MS and all radio resources remain the same, the TRAU and the A interface circuit (64 kbit/s) going into MSC switch fabric remain the same as Alice switches her active call between Bob and Charlie. However, on the back side of the same MSC, the call to or from Bob and the second call arriving from Charlie are two separate and independent calls, as far as PSTN call routing, call state and accounting (billing etc) are concerned. These two calls will most likely be carried on different trunks on the interface between the MSC and the fixed telephone network, and every time the MSC's control processor handles a CALL RETRIEVE request for a call that is different from the one that was just put on hold, a different path for the newly selected call has to be established through the switch fabric.

Hera MSC architecture with separate front and back MGWs within the same MSC allows us to replicate the same roles and semantics in our software implementation. The front (TC) MGW is akin to the bank of TRAUs in front of traditional hw MSCs, and for as long as the same MS on the same dedicated connection is involved in call switching, only one call instance exists in this TC MGW, akin to a single TRAU. OTOH, the back (IWF) MGW will have one instance allocated for each separate call that exists on IP-PSTN side, hence the call to or from Bob and the second call from Charlie will be handled by separate call instances in this back MGW. Figure 2 shows a single ISP endpoint in TC MGW being switched between two different ISP endpoints in IWF MGW — hence the internal switching path does exactly what its name indicates, switches call routing just like the switch fabric of a traditional hardware MSC.

3.4. Differences between TDM TRAUs and Hera TC MGW

While the architecture of Hera MSC mimics the traditional TDM-based GSM network design very closely, there are still some important differences that need to be noted. One aspect of the original GSM architecture with TRAUs that cannot be replicated sensibly in our Osmocom+ThemWi environment is logical ownership and control of those TRAUs. In classic GSM architecture all TRAUs logically belong to the BSS: they may be physically located at the MSC site, thus achieving the desired savings in transmission cost between sites, but they are still logically invisible to the MSC, which sees only G.711 in one fixed choice of A-law or μ-law. In contrast, Hera TC MGW is a part of Hera MSC both physically and logically, controlled directly by `osmo-msc` process operating in ThemWi mode.

In ordinary operation a TRAU or a functional equivalent thereof does not need any control plane aspects: it is controlled solely in-band by the UL frame stream from the BTS, whether that stream is carried in the original TDM format or in our novel Ater-IP. In this case the difference in logical ownership between classic TRAUs and our Hera TC MGW is moot. However, if it is desired to implement codec mismatch resolution for TFO (see §4.2.2), then a need arises for control plane communication between the transcoder (TRAU or otherwise) and whichever entity has the power to switch codecs (or AMR codec configurations) on the local radio leg. Here the architectural difference becomes apparent: in the original GSM architecture assumed in the TFO spec, that control entity is the BSC, but in our architecture it is the MSC.

In the present design phase, no detailed planning has been done yet for TFO codec mismatch resolution. The first implementation will support only FR and EFR codecs (not AMR), and TFO only on already-matching codecs, without mismatch resolution. Logic for TFO-AMR and for codec mismatch resolution will be implemented in later phases of the project, and the necessary control plane complications will be addressed then.

Another implication of the difference in logical placement of transcoders occurs during inter-BSC handovers. When TRAUs logically belong to specific BSS, an inter-BSC handover involves migration to a different TRAU, and the call moves to a different circuit entering MSC switch fabric on the left side of Figure 1. However, a transcoder instance owned and controlled by the MSC remains the same through inter-BSC handovers, with Ater-IP routing being switched to a different BSS IP:port address. Therefore, **all** handovers in Hera MSC architecture logically function in the way which the TFO spec regards as intra-BSC handovers, keeping the same TRAU.

3.5. Transcoder bypass mode for special test calls

Up until now it has been stated that the internal switching path (ISP) in Hera MSC can carry only two types of traffic: either voice in G.711 format or CSD in TW-TS-007 format. There is, however, a third option, supported only for special test calls: transcoder bypass for GSM speech calls. In this mode Hera TC MGW passes RTP packets transparently between Ater-IP and ISP without transcoding or retiming, thus allowing GSM compressed speech formats to propagate to ISP. This mode will be used only for special test calls, meaning either MO calls to special test numbers implemented inside the MSC itself or test MT calls made from the command line. TC bypass operation will be needed in the following test call scenarios:

- It is desirable to have some echo test numbers that echo call UL back to DL with a certain delay. This echo facility can be implemented at G.711 level, in which case UL speech will be decoded and then re-encoded by the test number before being sent back to DL. However, it is also desirable to implement an echo facility that operates without this double transcoding effect; such ‘‘tandem-free echo’’ can be achieved only via transcoder bypass mode, and it will have to be restricted to FR, HR and EFR codecs, not AMR.
- It is useful to be able to make a test call (MT from the command line to an MS is the most convenient setup in this case) and play a certain sequence of codec frames into the downlink, with the expectation that this sequence will be passed unaltered to the BTS. Transcoder bypass mode is naturally required for this operation.

The actual mode of operation for Hera TC MGW will always be selected by the first SDP codec (PCMU, PCMA, CLEARMODE, TW-TS-007 or one of GSM speech codecs) listed in the CRCX or MDCX command to the ISP side of this MGW, which is CN side in OsmoMSC view. OsmoMSC operating in ThemWi mode will in turn derive this information from the Assign & Connect (AssCon) command received on TWCC interface for a call (see §6.1.1).

4. Calls between two GSM subscribers

Classic GSM specs do not prescribe any special handling to be applied in the case when one GSM subscriber calls another GSM subscriber on the same MSC or on the same MNO. Instead the model implicitly assumed in those classic specs is that every MO call passes through standard steps of speech transcoding, any appropriate IWF for CSD and signaling conversion to ISUP, and only then the fixed telephone network past the MSC routes this call based on the dialed number. If this dialed number happens to belong to another GSM subscriber, whether on the same MNO or a different one, a GMSC node of the called party’s number-owning operator figures out which MSC that subscriber is roaming on at the moment (HLR and VLR queries are involved), and routes the ISDN call to that MSC. If it just so happens that both Alice and Bob are served by the same MSC at the moment when Alice calls Bob, absolutely no special handling is applied to this case: the MSC will handle MO call leg A and MT call leg B as two unrelated calls without any awareness of their connection. This approach is highly inefficient when a single MSC serves a semi-isolated community with a lot of internal calls, but works wonders when the two GSM subscribers may be served by different MSCs, roaming on different visited networks, or customers of different MNOs.

In designing Themyscira Wireless MSC software, the present author faces a dilemma: should we implement the classic approach just described, or should we implement optimized handling for Alice-to-Bob calls that stay within the same MSC, or within the multi-MSC network of the same MNO? The answer to this question depends on the expected call traffic pattern: will there be a significant number of internal calls (local subscriber A to local subscriber B) that meet the criteria for this optimized handling, or will such calls occur so rarely that the minuscule performance gain cannot justify the effort of implementing and supporting an entirely separate path for this call configuration?

In the case of American 2G Cooperative and the network we seek to build, we expect most A2GC members to have family and other frequent personal contacts who use mainstream phones served by mainstream carriers. Therefore, calls between one A2GC member on GSM and outside phone numbers belonging to major carriers (mostly VoLTE) are expected to be much more frequent than calls from one A2GC member to another, with both parties using A2GC GSM network at the same time. Therefore, at the present time there is no justification for implementing and supporting a separate optimized path (without transcoding or signaling conversion to SIP-I) for calls within the same MSC, or from one A2GC MSC to another. The classic approach described at the beginning of this section will be implemented instead.

4.1. Support for other community GSM networks

Because USA is not the only country in which legacy GSM networks are being wrongfully shut down by indifferent operators ignoring the cries of Vintage Mobile Phone users and communities, we (A2GC) sincerely hope that our GSM network won't be the only one in the world. Therefore, instead of focusing on internal call paths that provide richer features or improved performance or quality when calling other numbers served by our own system (and thus giving everyone else a lower tier of service), we shall strive to replicate the full original paradigm that treats GSM as wireless ISDN, such that GSM to GSM calls are effectively GSM→ISDN→GSM as originally envisioned. In this paradigm there is no discrimination between our own GSM subscribers versus those of any other system anywhere in the global E.164 number space.

4.2. In-band TFO instead of TrFO

If a call between two GSM subscribers passes through G.711 in the middle, with each GSM end doing full transcoding, the resulting effect is double transcoding or two lossy speech codecs operating in tandem. This arrangement produces worse voice quality for the two communicating GSM users compared to a single lossy codec, with encoding in Alice's MS and decoding in Bob's MS and vice-versa. The two solutions to this problem that have been developed by the mainstream telecom industry over the years are TFO (tandem-free operation) and TrFO (transcoder-free operation). TFO is in-band: the call still passes through a transit network that only supports G.711 and a transcoder is still present in the call path at each end, but the two TCs communicate via a special in-band protocol, dynamically discover each other's existence and negotiate to enter a special pass-through mode in which the 2 least significant bits of every G.711 PCM sample are replaced with a modified version of TRAU-UL frames. In contrast, TrFO eliminates speech transcoders from both GSM ends, requires that these two ends agree on the same GSM speech codec out of band, and connects the two ends via an IP-native RTP path that carries the common GSM speech codec instead of G.711.

Those mainstream carriers and equipment vendors who embraced the all-IP model of 3GPP Release 8 have effectively discontinued support for in-band TFO (3GPP were too cowardly to say it outright, but AoIP compressed speech mode introduced in that release breaks TFO by stipulating RTP payload formats from IETF that throw away information content of TRAU-UL frames) and implement only TrFO. However, TrFO has a fundamental flaw that can be seen as a moral defect: it works only when Alice and Bob are served by the same MNO, or when Alice's MNO recognizes Bob's number as belonging to a different MNO with whom they have a direct interconnection path that bypasses the regular G.711-speaking PSTN. If Alice's operator does not recognize the E.164 number she just dialed and sends the call out to a PSTN common carrier, the call gets transcoded to G.711, and if the transcoder does not support in-band TFO, the possibility to have a tandem-free call with Bob on another GSM network not known to Alice's MNO is completely foregone.

This TrFO approach is thus at odds with the moral-philosophical framework set out in §4.1, and on the basis of this consideration, we (Themyscira Wireless) choose to implement and deploy in-band TFO instead. Every call which we send out to IP-PSTN in G.711 format will include TFO capability, preserving the possibility, however slim, that the other end may be another community GSM network of which we have no *a priori* knowledge. For approximately the first 5 s of the call the least significant bit of every 16th PCM sample in the G.711 stream emanating from our TC will be replaced with in-band TFO_REQ messages that announce our TFO capability and the GSM speech codec in use. If the far end responds with similar TFO_ACK messages, further TFO protocol will ensue, hopefully leading to active TFO and elimination of double transcoding for this call. In the far more common case of no TFO_ACK response, these TFO_REQ messages will cease after approximately 5 s and a regular G.711 transcoding session will ensue for the remainder of the call.

4.2.1. TFO for FR/HR/EFR codecs vs AMR

TFO first appeared in GSM Release 98, the same release that introduced AMR — however, this first definition of TFO supported only FR, HR and EFR codecs, and not the new AMR codec introduced in the same release. TFO for AMR was specified only much later in 3GPP Release 4 — thus for two GSM/3GPP release cycles (R98 and R99), a peculiar situation existed in that an operator could reap the benefits of TFO *or* those of AMR, but not both at the same time. Nokia TCSM2 is a perfect example of transcoder equipment that represents this stage of GSM evolution: it supports FR, HR and EFR codecs with TFO, and it also supports AMR without TFO — but not the combination of AMR with TFO.

In the present time when TFO specifications exist for both FR/HR/EFR and AMR, with both specs equally available for reading and implementation, the situation remains that TFO for FR/HR/EFR is much simpler and easier to implement than TFO-AMR. For this reason, the current plan for Hera MSC is to implement only FR and EFR codecs in our TC MGW initially, supporting TFO for both of them, and then address AMR in a later phase of the project.

4.2.2. Codec mismatch resolution

What happens if a digitally transparent G.711 transport path connects two TFO-capable transcoders (the necessary condition for TFO establishment), but the two GSM radio legs operate with different and thus incompatible codecs? In a network that supports FR and EFR, this situation will arise if one subscriber decides to exercise a super-old phone that supports only plain FR while all other phones are EFR-capable; in a network that additionally supports AMR, this situation will be even more likely, as phones that support EFR but not AMR are much easier to find than those that only support FRv1.

In some of these codec mismatch situations it is desirable to be able to switch the call leg using the ''higher'' codec to the ''lower'' codec used by the other call leg, while in other situations it may be preferable to keep the tandem connection with double transcoding. In a call between an AMR user and a non-AMR-capable phone, it may be inadvisable to switch the AMR user to FR or EFR, in case the AMR user is at a great distance from the cell site or otherwise experiences poor radio conditions in which AMR provides clear benefits. OTOH, if a call is set up between an EFR user and an FRv1 user, there is absolutely no benefit to keeping the tandem connection: an end to end FRv1 connection (which can be achieved by switching the EFR user to FRv1) will always have a higher quality than an FRv1 leg and an EFR leg operating in tandem.

Switching the local radio leg to an alternative GSM codec in response to a dynamically discovered TFO mismatch situation requires control plane communication between the transcoder and a suitable entity in the standard control plane for GSM. In the original TDM-based GSM architecture assumed in the TFO spec, the latter entity is the BSC — but in our Osmocom+ThemWi environment the MSC has to take on this role instead.

Devising the necessary extensions to the MGCP dialect used between Hera TC MGW on one end and OsmoMSC modified for ThemWi on the other end, providing a way for the MSC core to tell the MGW what alternative codecs (or AMR codec configurations) are available for TFO, and then an event notification mechanism by which the MGW asks the MSC core to actually switch the locally used codec, will be a very major task. Therefore, we are deferring it to later phases of our GSM revival project; the initial version will only support

TFO establishment on already-matching codecs. Restricting our initial operation to FR and EFR codecs (no AMR) will significantly reduce the occurrence of codec mismatch situations in this initial phase.

5. Circuit-switched data services

5.1. Possible CSD call destinations

A mobile-originated (MO) CSD call is initiated just like voice calls, by dialing a phone number for the called party; the only difference between voice and CSD calls in MO direction is the information transfer capability (ITC) indicated in the bearer capability information element (BC-IE) in the MO SETUP CC message, along with some other fields being different in the same BC-IE. What is the set of destination phone numbers to which CSD calls may be addressed?

Classic GSM architecture generally assumes that CSD calls are addressed to the same destinations as regular phone calls: just as regular voice calls can be made to a number on analog PSTN, a number on ISDN or another GSM subscriber, CSD calls can be made to all of the same entities. A CSD call between two GSM users or between a GSM user and a fixed-line ISDN user is a straightforward digital data connection, but a CSD call between GSM and analog PSTN involves V-series modem emulation performed by an IWF at the GSM MSC, as covered in GSM 09.07 §9.2.

The primary goal of American 2G Cooperative is to build a new GSM network that will be no worse than the one which T-Mobile USA are itching to shut down. Since the legacy GSM network of T-Mobile supports CSD calls to and from PSTN numbers in analog mode, with V-series modem emulation, we must support this mode of operation as well. We also wish to support mobile-to-mobile transparent CSD calls in ISDN UDI mode (supported by stock Osmocom but not by T-Mobile network in its present degraded state), as well as CSD calls between GSM and real ISDN, where the latter will most likely have to be OCTOI. Finally, we seek to implement a novel service where special numbers that exist only within our GSM network (implemented inside each MSC) act as gateways from CSD to Telnet or from CSD to PPP to mobile Internet. Subsequent sections of this chapter cover each of the several types of CSD service we wish to implement.

5.2. Emulation of V-series analog modems

The feat called for in this mode of operation requires a software library that implements V-series modulations, interfacing to data bits on one end and PCM samples at 8000 samples/s rate on the other end. The most mature library of this type that is currently available is SpanDSP by Steve Underwood — hence we plan on using SpanDSP as our foundation in this subproject. However, the highest V-series modulation supported by current SpanDSP is V.22bis, which maxes out at 2.4 kbit/s. Support for V.32 (4.8 kbit/s and 9.6 kbit/s analog services) and V.32bis (14.4 kbit/s analog service) is currently missing, hence we have a gap in terms of analog modem services which **should** be able to interwork with GSM CSD, but for which we lack the necessary library support.

Our current plan is to implement V.22 (1.2 kbit/s) and V.22bis (2.4 kbit/s) services initially, and then see if Mr. Underwood can be convinced to implement the missing V.32 support, or if anyone else in the community might be able to do so.

5.3. CSD calls from one mobile to another

The cleanest, most native way to establish a transparent CSD call from one GSM MS to another (e.g., for end-to-end encrypted speech) in the classic GSM architecture is to request ITC (information transfer capability) of UDI (unrestricted digital information), as opposed to ‘‘3.1 kHz audio external to PLMN’’. An example configuration is **AT+CBST=71,0,0**. The alternative request of ‘‘3.1 kHz audio ex PLMN’’ (e.g., **AT+CBST=7,0,0**) is also possible, but would be highly inefficient and wasteful: each GSM end would perform V-series analog modem emulation, and the two call legs would connect at the level of these emulated waveforms traveling over the internal G.711 connection.

In Themyscira Wireless implementation of GSM, using Hera MSC described in this document, the ability to do mobile-to-mobile CSD calls ‘‘the analog way’’, with each end applying V-series modem IWF, will still be there: we are not going to divert GSM-to-GSM calls to a special path different from other calls (local GSM to outside world and outside world to local GSM), and we don’t believe in artificial blocking of harmless fun.

However, this path will be subject to SpanDSP limitations covered in the previous section, hence it may be a long time before the full range of CSD user data rates becomes supported.

For this reason along with others, it is important to implement CSD in ‘‘ISDN mode’’ (ITC of UDI) in addition to ‘‘analog’’ mode. Furthermore, we wish to stick with our philosophy of treating calls from one GSM subscriber to another subscriber of our own no different from calls to users of other community GSM networks or users of real ISDN, if any such users still exist anywhere in the world. Therefore, we need to convert CSD calls to and from real ISDN semantics (ISUP signaling and true V.110 in the user plane, including 14.4 kbit/s mode) in exactly the same way how this feat was done by classic MSCs in the glory days of ISDN and GSM, and we need a way to transport these ISDN calls across what is otherwise IP-PSTN carrying G.711 traffic.

Our currently planned solution is to use SIP-I, which combines SIP and ISUP. For regular PSTN-compatible calls (voice, analog modem emulation or fax) we fully expect many IP-PSTN middleboxes to strip away ISUP and keep only plain SIP, and we may have to configure our outbound gateway (GSM MO calls to SIP-I) to send plain SIP rather than SIP-I to non-permissive destinations. But in the case of MO CSD calls interworking to ISDN UDI, the outgoing SIP INVITE will always be SIP-I (with ISUP IAM included), with SDP offer indicating CLEARMODE instead of {PCMU, PCMA} choice, and successful completion of such calls will require a route to the destination number that allows CLEARMODE and preserves SIP-I. Such routes will exist between all of our MSCs, thereby allowing CSD calls from one own subscriber to another even when they are served by different MSCs, and we will gladly establish such routes with any other community networks (GSM or ISDN) who would like to have a peering interconnection with us.

5.4. GSM fax service

GSM 03.45 spec defines the teleservice of fax over GSM as interworking between a special fax-carrying type of CSD call and actual analog fax (T.30) on external PSTN. In the case of fax transmission from one GSM terminal to another, no special exception is made: each GSM end interworks per GSM 03.45 to a conceptual analog PSTN interface (G.711 in practice) just as if it were communicating with an analog fax machine on traditional PSTN, and the two ends connect in PSTN-emulating G.711 land.

In Themyscira Wireless implementation of GSM, using Hera MSC described in this document, we plan to support fax teleservice exactly per this canonical architecture: interwork every fax-type CSD call to G.711 on IP-PSTN, with our IWF MGW implementing the Fax Adaptor of GSM 03.45. We plan on using SpanDSP library to do so; fortunately the state of this library is much more mature for fax services than for data modems, hence the problem described in §5.2 is expected to not apply to fax.

5.5. Alternating speech/data and speech/fax calls

Glory-days versions of GSM specs include a provision whereby a single call that interworks to analog (or G.711) PSTN may be switched on the fly between speech mode (one state) and data or fax mode (the other state). Upon such ICM (in-call modification) events, which can be requested only from the MS, the GSM leg between the MS and the MSC switches between speech and CSD, but to the other party on the PSTN the single call remains unbroken. The MSC switches between simply through-connecting the call in speech phase or applying the appropriate IWF in data or fax phase.

We plan on implementing this support in Hera MSC, with user plane switching as described in §3.2.2.

5.6. Internally terminated CSD services

We plan to implement a novel service in Themyscira Wireless whereby some special CSD-callable numbers will be defined and implemented inside Hera MSC itself. These numbers will accept asynchronous (character-oriented) data calls only (no speech, fax or alternating calls), and these calls will be answered and terminated by a special software component within Hera MSC suite. There is some precedent in GSM specs for doing what we propose: see GSM 03.10 §5.4, figures 4c and 4d. The 3 ‘‘special resources’’ (in the terminology of those GSM 03.10 spec figures and related passages) which we plan to implement are covered in the following subsections. It should be noted that with all 3 of these ‘‘special resources’’, the distinction of whether the MS requested an ‘‘analog’’ or ‘‘ISDN’’ call (ITC set to UDI or ‘‘3.1 kHz audio ex PLMN’’) will be ignored except for diagnostic messages: no actual interworking of either type occurs with any of these within-MSC CSD destinations.

5.6.1. CSD test number

The simplest MSC-internal CSD service we plan to implement is a test number that allows exercising CSD functionality in GSM MS (AT commands) and the network path for CSD, passing through the BTS, the BSC-associated MGW and Hera TC MGW in CSD pass-through mode. Upon connection, the service at this test number will print some messages into the data call channel (telling the user which number she is calling from, the specific CSD mode in which she connected, etc), and then drop into character echo mode in which all characters typed by the user will be echoed back to her.

5.6.2. CSD to Telnet services

The most canonical use case for CSD in the glory days of GSM was for a mobile user to use her GSM phone as a modem, dial into the dial-up modem bank of her institution, and log into her mainframe or minicomputer account, as in big UNIX, VMS etc. In the spirit of Holy retrocomputing and retronetworking, we wish to recreate this user experience — but there are two problems:

- Almost no one operates dial-up modem services on PSTN any more;
- SpanDSP limitations described in §5.2 will get in the way of connecting to modem services at 9.6 kbit/s, which was the standard speed in the days which we seek to recreate.

One possible solution is to skip the analog modem leg and connect directly (by Telnet or ssh) to the same mainframe, minicomputer or BBS services to which the terminal server attached to the dial-up modem bank would connect anyway, initiating this Telnet (or ssh equivalent) connection from the special MSC component that provides this internal CSD service.

5.6.3. CSD to mobile Internet

Unfortunately a very large portion of today's Vintage Mobile Phone user community has no concept of using CSD the traditional way, by connecting the phone to a computer and using it as a modem, dialing into some BBS or minicomputer account in ASCII text terminal mode. Instead the only CSD they know of is the kind where some phones feature a WAP browser that can connect to mobile Internet via CSD-PPP in addition to GPRS, or solely via CSD-PPP in some early models. And since American 2G Cooperative needs to attract as many members as possible to gain strength in numbers, we need to support this Internet-ized flavor of CSD.

In the days when this CSD-PPP mode of connecting to mobile Internet was current, the actual service consisted of two components: the GSM MNO provided CSD connection to analog PSTN only (complete with V-series modem emulation at waveform level), and a separate unrelated company was the traditional dial-up ISP, connecting from analog dial-up lines to Internet. But the same problems apply as in the previous section: analog dial-up ISPs are mostly extinct, we have no SpanDSP support for the necessary higher modem speeds, and given the community we are serving, it would be best to provide a self-contained service that does not rely on third parties.

Because the segment of VMP user community that desires this flavor of CSD does not care one bit about the analog modem portion of the WAP service they are after, the same shortcut that was presented in the previous section applies even more readily to the present use case. Our current plan is to terminate PPP inside the same special MSC component that terminates the lower layers of V.110 or RLP transported in TW-TS-007 on our internal switching path, and then interwork to GTP, connecting to the same GGSN as will be used for GPRS services.

5.7. Transparent vs non-transparent CSD services

T and NT are two alternative connection elements for the leg that runs from the MS to the IWF in the MSC. In the case of a CSD call between two mobiles, the choice between T and NT is independent and private for each call leg; in the case of CSD calls to other entities, the non-GSM entity has no way of knowing whether the GSM leg is T or NT. For further understanding of T vs NT modes, please refer to GSM 03.10 §4.2.

The aspect that matters the most for Themyscira Wireless is that transparent mode is much easier to implement in MSC IWFs, but the most canonical application for CSD (interactive text mode login into a mainframe or minicomputer account, see §5.6.2) calls for NT mode, making that mode the most popular one — to the point that many GSM MS implementations only support NT mode!

The current plan for CSD IWFs in Hera MSC is to implement both T and NT modes in the special IWF component that handles internally terminated CSD services of §5.6 (all subsections), but implement only transparent mode in the IWFs of §5.2 and §5.3. More precisely, NT mode for those IWFs will be implemented only if someone really wants it, and will be given a lower priority compared to many other needed tasks.

It should also be noted that the Radio Link Protocol (RLP) of NT mode is intended to run only between one MS and an IWF in the network, not between two GSM MS! Mainline OsmoMSC implements the latter manner of operation, but it is not correct, and many GSM MS don't work with it because they get confused by seeing another MS across the protocol, instead of the expected IWF. We do **not** plan to support this bogus mode of operation in ThemWi Hera MSC.

6. MSC control plane and TWCC

6.1. Interface between OsmoMSC and ThemWi components

The existing MNCC interface implemented in OsmoMSC is not suitable for Themyscira Wireless: too much intelligent logic and transformation is applied between CC air interface messages and MNCC, and the external MNCC cannot exercise any real control over TCH assignment operations on BSSMAP, neither the initial assignment Channel Type nor subsequent reassignments for Channel Mode Modify. Instead we are defining our own control plane interface between internal components of a single MSC which we call TWCC (Themyscira Wireless Call Control), and we need to develop patches to OsmoMSC to support both the existing form of MNCC and TWCC.

The complete details of TWCC will be defined only when we reach the point of implementing it on a feature branch of `osmo-msc`, and the complete interface specification will be included with those implementation patches. However, the general principles will be as follows:

- Most CC air interface messages will be passed through directly between DTAP and TWCC components, without OsmoMSC inserting its own mind.
- OsmoMSC won't do any parsing of BC-IE, and it won't initiate traffic channel assignment on BSSMAP until commanded to do so by TWCC.

6.1.1. Assign & Connect operation

Assign & Connect (AssCon) command from a TWCC call handler to OsmoMSC core process indicates what Channel Type is needed for this call (or for the current phase of an alternating call), and communicates this call handler's RTP IP:port address for the internal switching path of Figure 2 — see §3.2. Upon receiving this command from TWCC, OsmoMSC core process shall perform the following steps:

1) If the call has not been assigned yet, perform initial assignment procedure on BSSMAP, using the Channel Type indicated in the AssCon command.
2) If the call has already been assigned to TCH previously, compare the newly requested Channel Type against the saved previous one. If a different Channel Type is required, perform a new assignment procedure on BSSMAP.
3) Irrespective of whether a new assignment or reassignment needed to be done on BSSMAP, command the MGCP-controlled MGW (which will be Hera TC MGW in this configuration), via CRCX or MDCX as appropriate, to connect its CN side to the ISP RTP IP:port address supplied in the AssCon command. The same AssCon command will also indicate **sendrecv** or **recvonly** mode and the codec type to be used on ISP (PCMU, PCMA, TW-TS-007 or raw pass-through of GSM speech codecs); this information shall also be passed through to the CRCX or MDCX command to Hera TC MGW.

OsmoMSC response to TWCC AssCon shall indicate success or failure of these operations and the ISP RTP IP:port address of the front MGW controlled by this OsmoMSC core process.

6.2. Breakdown of Hera control plane components

A dedicated `hera-twcc` process will maintain the main TWCC socket connection to `osmo-msc` core process: existing OsmoMSC supports only one MNCC socket connection, and there is no need to change this design. All other call handler components in Hera MSC suite will connect to `hera-twcc` via additional Hera-specific socket connections, and the latter process will act as a TWCC message switch. The following call handler components are currently envisioned:

- `hera-sip-in` will handle inbound (MT) calls from SIP-I to local GSM subscribers currently served by this MSC;
- `hera-sip-out` will handle MO calls that need to go out to SIP-I — all regular MO calls;
- `hera-testnum` will handle MO calls to special test numbers like ‘‘just sit there’’ (silence), echo test and playout of a predefined voice announcement followed by disconnection;
- `hera-intcsd` will implement internally terminated CSD services covered in §5.6;
- We also desire a command line tool (run by test operators as needed, **not** a long-lived daemon process like all others) that generates test MT calls to local GSM subscribers directly from the command line. These test calls will typically be silent in the sense of having no audio fed into the downlink, but the same command line utility will also have the ability to play TW-TS-005 hex files (sequences of RTP payloads) into DL path, allowing manual tests that exercise various aspects of the BTS or the MS.

In MT direction, `hera-twcc` will provide a socket to which other components can freely connect and send MT calls to the local GSM network via OsmoMSC; this mechanism will be used by `hera-sip-in` and the test MT call command line tool. In MO direction, `hera-twcc` will have knowledge of the telephone numbering plan used for the local deployment (real PSTN of the host country or a private network numbering plan) and some basic functionality for routing of calls based on the dialed number. It will have vty configuration for directing MO calls (based on the dialed number) to different secondary TWCC sockets; the latter sockets will connect to `hera-sip-out`, `hera-testnum` and `hera-intcsd`.