



Limitations

Table of contents

Known Limitations

Known Issues

Known Limitations

- The cuPHY library and binaries are intended for the Linux environment on the qualified platforms only.
- The supported configurations are limited to those listed above. Other configurations are not supported and may not perform well.
- Only homogeneous configurations supported for multiple cells.
- The configurable YAML parameters `enable_h2d_copy_thread`, `h2d_copy_thread_cpu_affinity`, and `h2d_copy_thread_sched_priority` are optional in the cuphycontroller YAML file. If these parameters are not present, the code uses the default values and throws the exception "YAML invalid key:" on the cuphycontroller console. This exception message has no impact on the functionality and can be disregarded.
- GPU Initiated Comms for DL (`gpu_init_comms_dl` flag in the cuphycontroller config yaml) is required to be enabled by default from 22-2.4 release onwards. The flag enables the feature within Aerial L1 to engage GPU kernels to prepare and send U-Plane packets on the DL as opposed to CPU Initiated Comms (`gpu_init_comms_dl=0`) which exercises CPU code/consumes CPU cycles to prepare/send U-plane packets on the DL.
- No simultaneous DL and UL scheduling in S-slot. However, DL-only s-slot is supported in E2E test with O-RU.
- When the FAPI messages for a given cell are sent via nvipc, L1 expects an explicit notify (once per cell) via nvipc. In the case of multiple cells, multiple explicit notify APIs be called from L2. When a cell doesn't have any messages for a given slot, L1 expects dummy DL_TTI and/or UL_TTI.request, that is (`nPDU = 0`), to be sent "per cell". If the Slot Response feature is enabled by compiling Aerial with `-DENABLE_L2_SLT_RSP=ON`, this step is optional.
- For multi cells operation, L2 can signal the L2Adapter in 2 ways:
 - Single event per slot: which contains SCF FAPI messages for all cells. The single event is raised by calling nvipc notify(1) once per slot after the messages for all the cells are sent.

- Single event per cell: which is signaled by L2 after all FAPI messages for a given cell are sent. It is expected that multiple `nvipc_notify(1)` are called for multiple cells. The number of times that `notify` is being called must be the same as the number of active cells. A cell is marked active after `START.req` is received from L2. In this case, L1 expects dummy `DL_TTI` and `UL_TTI` described above. This is the default behavior.

To select the operation mode, set the `ipc_sync_mode` in yaml:

```
# Option 1: Sync per slot ipc_sync_mode: 0 # Option 2: Sync per active cell
ipc_sync_mode: 1
```

If Slot Response feature is enabled by compiling Aerial with `-DENABLE_L2_SLT_RSP=ON`, this setting is a no-op as L1 does not expect any event from L2.

- Cell life cycle management:
 - All cells have to be configured before any cell start.
 - No In-service configuration update.
 - `CONFIG.request` received in CONFIGURED (Out-of-Service) state can be used to change PCI and the supported PRACH parameters specified in dynamic PRACH section in cuBB quickstart guide only. PHY ignores any other TLVs received in `CONFIG.request`. If `CONFIG.response` indicates success, then only PCI and supported PRACH parameters are changed. All other parameters remain as in the initial `CONFIG.request` received for the cell.
 - PHY reconfiguration of a cell in CONFIGURED (Out-of-Service) state can take upto 40ms to complete (details below). Another `CONFIG.request` for any cell during this time (around 20ms) that occurs before receiving a `CONFIG.response` returns a `CONFIG.response` with the error code "MSG_INVALID_STATE". The `ERROR.indication` will NOT be sent for this error. L2 needs to wait to receive a `CONFIG.response` before sending a `CONFIG.request` for another cell in CONFIGURED state.
 - If Aerial is configured for 4 cells and 3 cells are In-service with data running, reconfiguration of 1 cell (Out-of-Service) can take around 40ms to complete

- If Aerial is configured for 4 cells and 3 cells are In-service with no data running, reconfiguration of 1 cell (O-RU) can take around 20 ms to complete
- If CONFIG.response is received with error code "MSG_INVALID_CONFIG", then reconfiguration was unsuccessful and the cell is still with the configuration received in initial CONFIG.request.
- No UE attach allowed in all cells during the reconfiguration time.
- Dynamic M-plane parameters:
 - When OAM sends gRPC message to change MAC address in M-plane, it must be a valid O-RU MAC address.
- Aerial supports only single section per packet in the UL.
- The nvlog_observer and nvlog_collect are deprecated in 23-1.
- F13 test cases are deprecated in 23-2.
- Early HARQ in UCI.indication:
 - This feature is supported only for single cell and for UCI-on-PUSCH only.
 - UCI.Indication with early HARQ will not have any measurement values.
 - If only HARQ is scheduled on PUSCH then with this feature enabled, no UCI.indication will be sent to L2 after full slot processing of PUSCH. Consequently no measurements for that slot will be reported to L2.
 - If CSI reports are also scheduled on PUSCH along with HARQ, then UCI.Indication with early HARQ will not have any measurement values. But the UCI.indication sent after full slot processing of PUSCH will have the measurements.
 - A constraint to enable early-HARQ is that these HARQ bits should be fully resident in OFDM symbols 0-3. So HARQ bits resident in OFDM symbols 0-3 will be in the 1st UCI.indication (that is, early-HARQ indication) and all other HARQ bits in the subsequent UCI.indication (that is, after full slot PUSCH processing completes).

- Multiple cell operation without issuing dummy config.req:
 - L2 should wait for at least 40msec between two CONFIG.request even at the initial stage, so that CONFIG.response is received by L2.
 - L2 can retry the failed CONFIG.request for a given cell after 1 sec.
- Multi-L2 with single cuphycontroller per GPU:
 - The total cell number of all L2 instances cannot exceed the cell_group_num configured in cuphycontroller yaml.
 - nvIPC only supports static cell allocation defined in the nvipc_multi_instances.yaml for multiple L2 instances. The number of cells and the cell mapping in each L2 instance cannot change after L1 is configured..
 - Support dynamic cell start/stop in each L2 instance. Do not support dynamic L2 restart. L2 instance needs to hold the nvipc instance after connecting to L1.
- 32T32R TDD single cell:
 - DL SU/MU-MIMO Number of Layers combinations
 - 1 UE in a Group: 1 or 2 or 4 layers
 - 2 UEs in a Group: 1+1 layers or 2+2 layers
 - 4 UEs in a Group: 1+1+1+1 layers
 - UL SU/MU-MIMO Number of Layers combinations
 - 1 UE in a Group: 1 or 2 layers
 - 2 UEs in a Group: 1+1 or 2+2 layers
 - PDSCH Resource Allocation Type 0 (RAT0) is not supported.

- Partial PRB allocation is not supported for the UE's in one group. Only full bandwidth can be allocated with PRG size 2.
- Currently PRG size of 2 is only supported for SRS & BFW.
- SRS channel is only supported in the special slot along with no other UL channels.
- SRS is not supported in the UL Slot, due to the presence of other UL channels.
- SRS reports related to antennaSwitching (FAPI 222.10.04, Table 3-133 - Channel SVD Representation) is not supported.
- The following test cases are not verified:
 - NZP-CSI-RS & TRS in combination with PDSCH.
 - Multiple UE Groups with UE's in each group containing same set of PRB's.
 - Multi-slot SRS + BFW + PDSCH/PUSCH + all other UL/DL channels.
- 64T64R TDD single cell:
 - PDSCH Resource Allocation Type 0 (RAT0) is not supported.
 - Partial PRB allocation is not supported for the UE's in one group. Only full bandwidth can be allocated with PRG size 2.
 - Currently PRG size of 2 is only supported for SRS & BFW.
 - SRS channel is only supported in the special slot along with no other UL channels.
 - SRS is not supported in the UL Slot, due to the presence of other UL channels.
 - SRS reports related to antennaSwitching (FAPI 222.10.04, Table 3-133 - Channel SVD Representation) is not supported.
 - The following test cases are not verified:

- NZP-CSI-RS & TRS in combination with PDSCH.
- Multiple UE Groups with UE's in each group containing same set of PRB's.
- Multi-slot SRS + BFW + PDSCH/PUSCH + all other UL/DL channels.

Known Issues

- The support for CPU Initiated Comms (gpu_init_comms_dl=0) mode is no longer available after the 22-2.4 release and it is recommended that this mode not be enabled for testing purposes.
- Support up to 8 DMRS ports, if the allocations are contiguous in PDSCH and PUSCH.
- Some DOCA error messages are not real errors. For example, the following messages contain debug info:

```
E [FH.QUEUE] Doca RxQ created! ... [DOCA][ERR][DOCA_GPU::mlx5:1229] ...
```

- Changing shm_log_level to 6 or 7 in nvlog_config.yaml causes a crash in the msg_processing thread.
- SCHED_FIFO + 100% CPU poll thread causes the system to hang on the 5.4.0-65-lowlatency kernel. The solution is one of the following:
 - Configure the kernel option CONFIG_RCU_NOCB_CPU=y, recompile, and install the kernel.
 - Upgrade the host system to 5.15.0-71-lowlatency or later.
- The cuphycontroller prints an error and aborts while running F08_8C_59_BFP14 over 2 100Gbps ports. The workaround is to downgrade the CX6-DX NIC FW to 22.35.1012 on the RU emulator server.
- 32T32R TDD single cell:
 - SRS + UL_BFW + PUSCH is failing due to Order Kernel Timeout.

- GPU/CPU Stats (DRV.MAP_DL & DRV.MAP_UL) not coming in phy logs for Aggr DL_BFW & Aggr UL_BFW.
- Partial bandwidth allocation for 100Mhz is failing i.e. launch_pattern 91104 and 91105.
- CUDA application on Grace Hopper:
 - CUDA applications on the Grace Hopper platform require ATS support. Currently, ATS is not enabled on the arm64 platform when IOMMU passthrough is enabled.
- NIC string conversion issue on Grace Hopper:
 - While working on dynamic CPU core assignments in K8s pod, we need to parse and dump the cuphycontroller config yaml file. On the Grace Hopper, the *nic: 0000:01:00.0* will be converted to *nic: 60.0*. This is because the PCIe address might be interpreted as a 60 based integer according to ['https://yaml.org/type/int.html'](https://yaml.org/type/int.html). The fix is to explicitly tell yaml parser to interpret the PCIe address as a string by putting single quotation marks around or *!!str* before the pcie address, e.g., *nic: '0000:01:00.0'* or *nic: !!str 0000:01:00.0*.

```
From sed -i "s/nic:*/nic: 0000:01:00.0/" ${cuBB_SDK}/cuPHY-CP/cuphycontroller/config/cuphycontroller_P5G_FXN.yaml to sed -i "s/nic:*/nic: '0000:01:00.0'/" ${cuBB_SDK}/cuPHY-CP/cuphycontroller/config/cuphycontroller_P5G_FXN.yaml (or sed -i "s/nic:*/nic: \!\!str 0000:01:00.0/" ${cuBB_SDK}/cuPHY-CP/cuphycontroller/config/cuphycontroller_P5G_FXN.yaml)
```

- The following test cases are not passing. They could be functionality issues or test framework issues:

Channel	Test Cases	Feature
PDSCH	3853	32TR
PUSCH	7374, 7476	RKHS ChEst
	7599, 7600	Multiple CSIP2

SRS	8037-8044	PRG size
mSlot_mCell	90027-90029, 90061-90067	DSUUU
	90068	harq 4-rx
	90070-90086	32TR
	90090-90109	64TR

© Copyright 2024, NVIDIA.. PDF Generated on 06/06/2024