




Agentic AI for Wireless Research


Sebastian Cammerer | Senior Research Scientist | June 22, 2026

Research has always evolved...

...so should we



```
[15]: if no_preview:
      scene.render(camera=my_cam, paths=paths, clip_at=20);
      else:
      scene.preview(paths=paths, clip_at=20);
```



The Paths object contains detailed information about every found path and allows us to generate channel impulse responses and apply Doppler shifts for the simulation of time evolution. For a detailed description, we refer to the developer guide [Understanding the Paths Object](#).

From Paths to Channel Impulse and Frequency Responses

Once paths are computed, they can be transformed into a baseband-equivalent channel impulse response (CIR) via `Paths.cir()`, into a discrete complex baseband-equivalent channel impulse response via `Paths.taps()`, or into a channel frequency response (CFR) via `Paths.cfr()`. These class methods can simulate time evolution of the channel based on the computed Doppler shifts (see `Paths.doppler`).

Let us first derive and visualize the baseband-equivalent channel impulse response from the paths computed above:

```
[16]: a, tau = paths.cir(normalize_delays=True, out_type="numpy")

# Shape: [num_rx, num_rx_ant, num_tx, num_tx_ant, num_paths, num_time_steps]
print("Shape of a: ", a.shape)

# Shape: [num_rx, num_rx_ant, num_tx, num_tx_ant, num_paths]
print("Shape of tau: ", tau.shape)

Shape of a: (1, 2, 1, 1, 27, 1)
Shape of tau: (1, 2, 1, 1, 27)
```

The `out_type` argument can be used to convert the CIR into tensors from different frameworks, such as `Dr.Jit` ("drjit"), `Numpy` ("numpy"), `Jax` ("jax"), `TensorFlow` ("tf"), and `PyTorch` ("torch"). Please see the developer guide [Compatibility with other Frameworks](#) for more information on the interoperability of Sionna RT with other array frameworks, including the propagation of gradients.

```
[17]: t = tau[0,0,0,0,:]/1e-9 # Scale to ns
```

Sionna 2.0

Migration to PyTorch



- Sionna 1.0 built on TensorFlow
 - 55,000 lines of code
 - 123 files, 84 classes, 1388 functions, 30 tutorials
 - Manual migration estimated at **18 person-months**
- Automated workflow via agentic coding tools
 - Cursor with Claude Opus 4.5 LLM
 - Module-by-module translation
 - Ensure 1300+ unit tests pass
- Migration in **one person-month**
 - Token cost ~2000\$
- Now, could probably be done in less than a week!

<https://nvlabs.github.io/sionna/>

From 3GPP 38.901 Specification to CUDA

Automated implementation of calibrated channel models

Specifications

ETSI TR 125 996 V19.0.0 (2025-10)

ETSI TR 138 901 V14.0.0 (2017-05)

ETSI TR 138 901 V19.2.0 (2026-02)




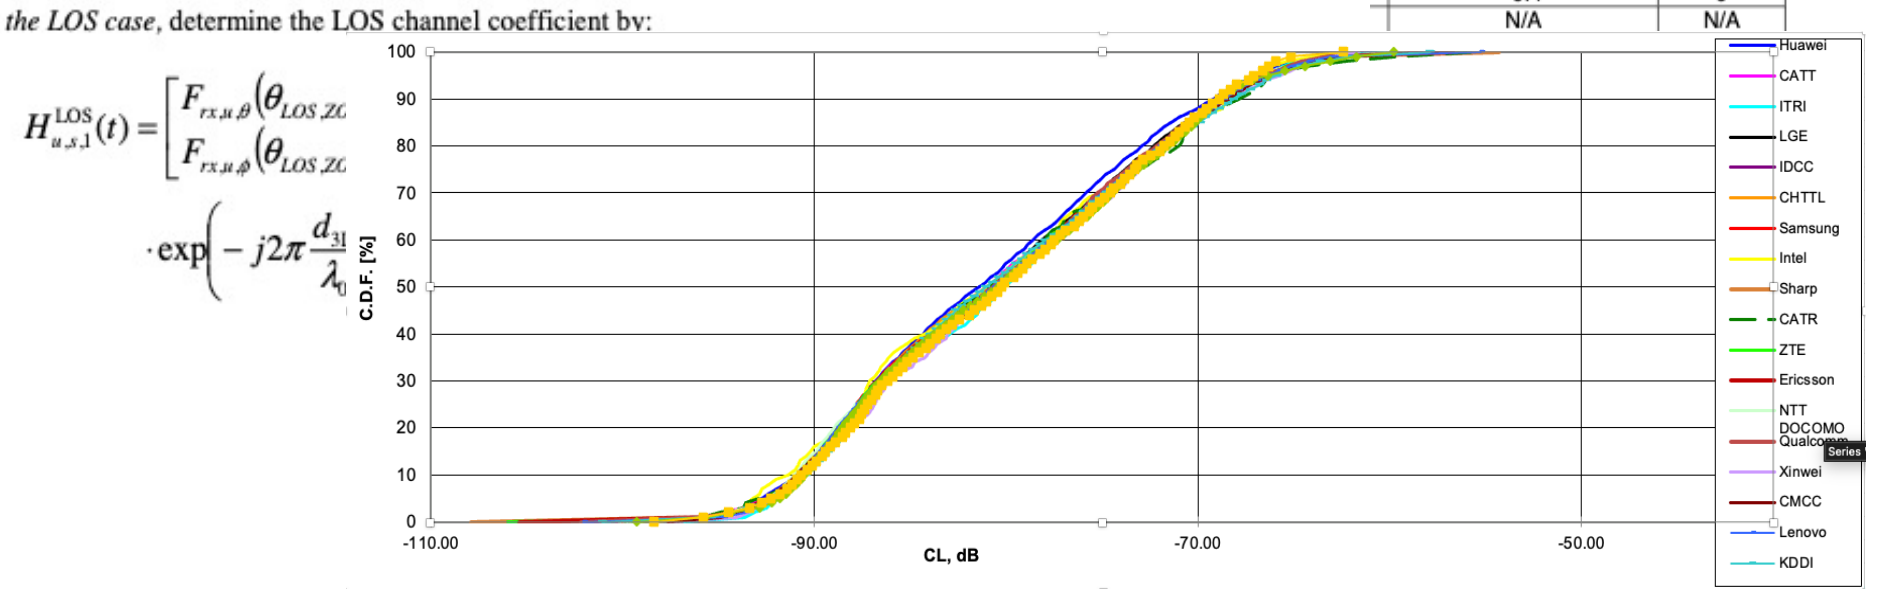
Table 7.5-6 Part-1: Channel model parameters for UMI-Street Canyon and UMa

Scenarios	UMI - Street Canyon			UMa		
	LOS	NLOS	O2I	LOS	NLOS	O2I
Delay spread (DS) lgDS=log ₁₀ (DS/1s)	μ_{DS} -0.18 log ₁₀ (1+f) - 7.28, see note 8	-0.22 log ₁₀ (1+f) - 6.87, see note 8	-6.62	-7.067 - 0.0794 log ₁₀ (f), see note 8	-6.47 - 0.134 log ₁₀ (f), see note 8	-6.62
AOD spread (ASD) lgASD=log ₁₀ (ASD/1°)	μ_{ASD} -0.05 log ₁₀ (1+f) + 1.21, see note 8	-0.24 log ₁₀ (1+f) + 1.54, see note 8	1.25	0.32, see note 8	1.09, see note 8	0.58, see note 8
AOA spread (ASA) lgASA=log ₁₀ (ASA/1°)	μ_{ASA} -0.07 log ₁₀ (1+f) + 1.66, see note 8	-0.07 log ₁₀ (1+f) + 1.76, see note 8	1.76	1.76, see note 8	2.04 - 0.25 log ₁₀ (f), see note 8	1.76
ZOA spread (ZSA) lgZSA=log ₁₀ (ZSA/1°)	μ_{ZSA} -0.11 log ₁₀ (1+f) + 0.81, see note 8	-0.03 log ₁₀ (1+f) + 0.92, see note 8	1.01	0.96, see note 8	-0.2856 log ₁₀ (f) + 1.445, see note 8	1.01
Shadow fading (SF) [dB]	σ_{SF} See Table 7.4.1-1	See Table 7.4.1-1	7	See Table 7.4.1-1	See Table 7.4.1-1	7
K-factor [K] [dB]	μ_K 5	N/A	N/A	9	N/A	N/A
ASD vs DS	0.5	0	0.4	0.4	0.4	0.4

$$H_{n,r,j}^{NLOS}(t) = \sqrt{\frac{P}{M}} \left[\begin{matrix} F_{r,n,\beta}(\theta_{n,m,ZOA}, \phi_{n,m,AOA}) \\ F_{r,n,\beta}(\theta_{n,m,ZOA}, \phi_{n,m,AOA}) \end{matrix} \right]^T \left[\begin{matrix} \exp(j\Phi_{n,m}^{\theta}) \sqrt{K_{n,m}^{-1}} \exp(j\Phi_{n,m}^{\phi}) \\ \sqrt{K_{n,m}^{-1}} \exp(j\Phi_{n,m}^{\theta}) \exp(j\Phi_{n,m}^{\phi}) \end{matrix} \right] \quad (7.5-28)$$

$$F_{r,n,\beta}(\theta_{n,m,ZOA}, \phi_{n,m,AOA}) = \exp\left(j2\pi \frac{r_{n,m}^T \vec{d}_{r,n}}{\lambda_0}\right) \exp\left(j2\pi \frac{r_{n,m}^T \vec{d}_{r,n}}{\lambda_0}\right) \exp\left(j2\pi \frac{r_{n,m}^T \vec{v}}{\lambda_0}\right)$$

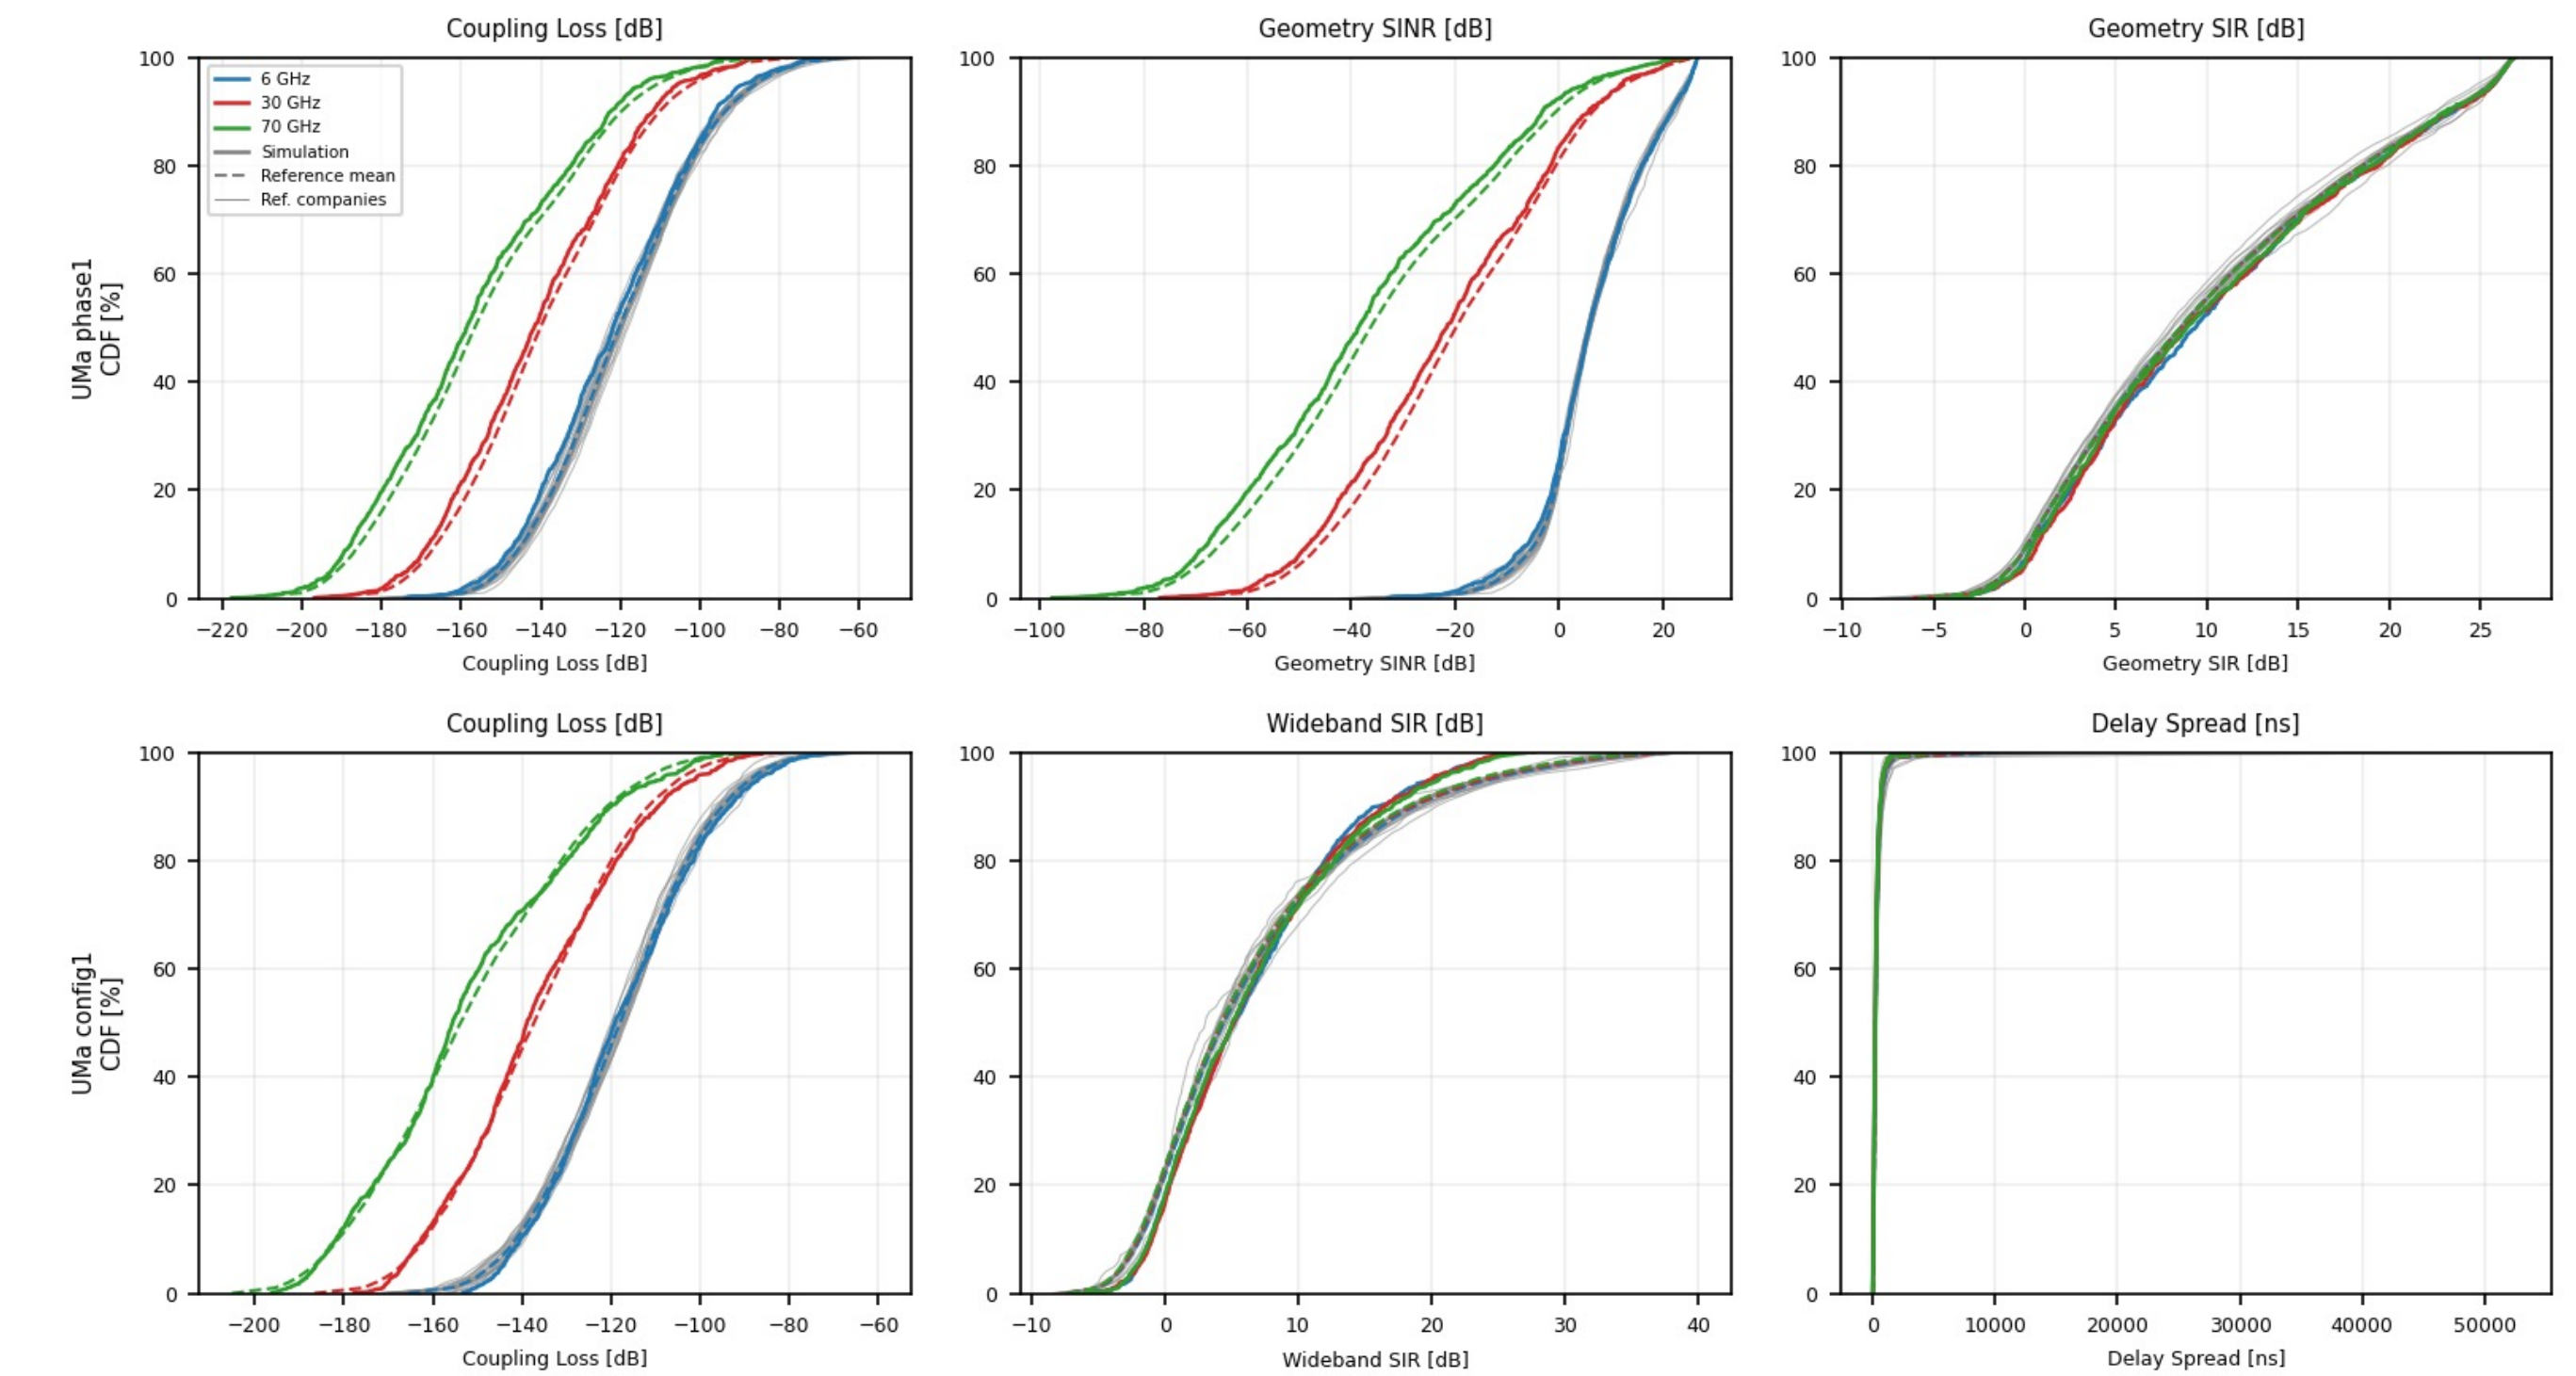
In the LOS case, determine the LOS channel coefficient by:

$$H_{n,r,j}^{LOS}(t) = \left[\begin{matrix} F_{r,n,\beta}(\theta_{LOS,ZC}) \\ F_{r,n,\beta}(\theta_{LOS,ZC}) \end{matrix} \right]^T \exp\left(-j2\pi \frac{d_{LOS}}{\lambda_0}\right)$$


Claude Code, Opus 4.6
Aurun (Sandbox)
48h/1500\$



Calibrated channel simulator (Python/CUDA)



Where Claude struggled:

- Complex nested tables
- Description of calibration metrics in prose
- Hexagonal grid with 57 sectors, wrap-around, user drops

All due to lacking “machine-friendliness”


Can this be done for the full 5G/6G specs?

The Agentic 3GPP Backoffice


Automated evaluation of new LDPC BG3 candidates in CUDA

QAM=256, R=948/1024, BLER=1e-2

Maximum complexity




Average complexity



The AI Telco Engineer

Exploring Agentic AI for Wireless Research

Given a problem, generate best algorithm according to user-defined metric



- Task description defines the wireless communication problem for the agentic framework
- Evaluation tool runs the generated algorithm and computes its performance metric

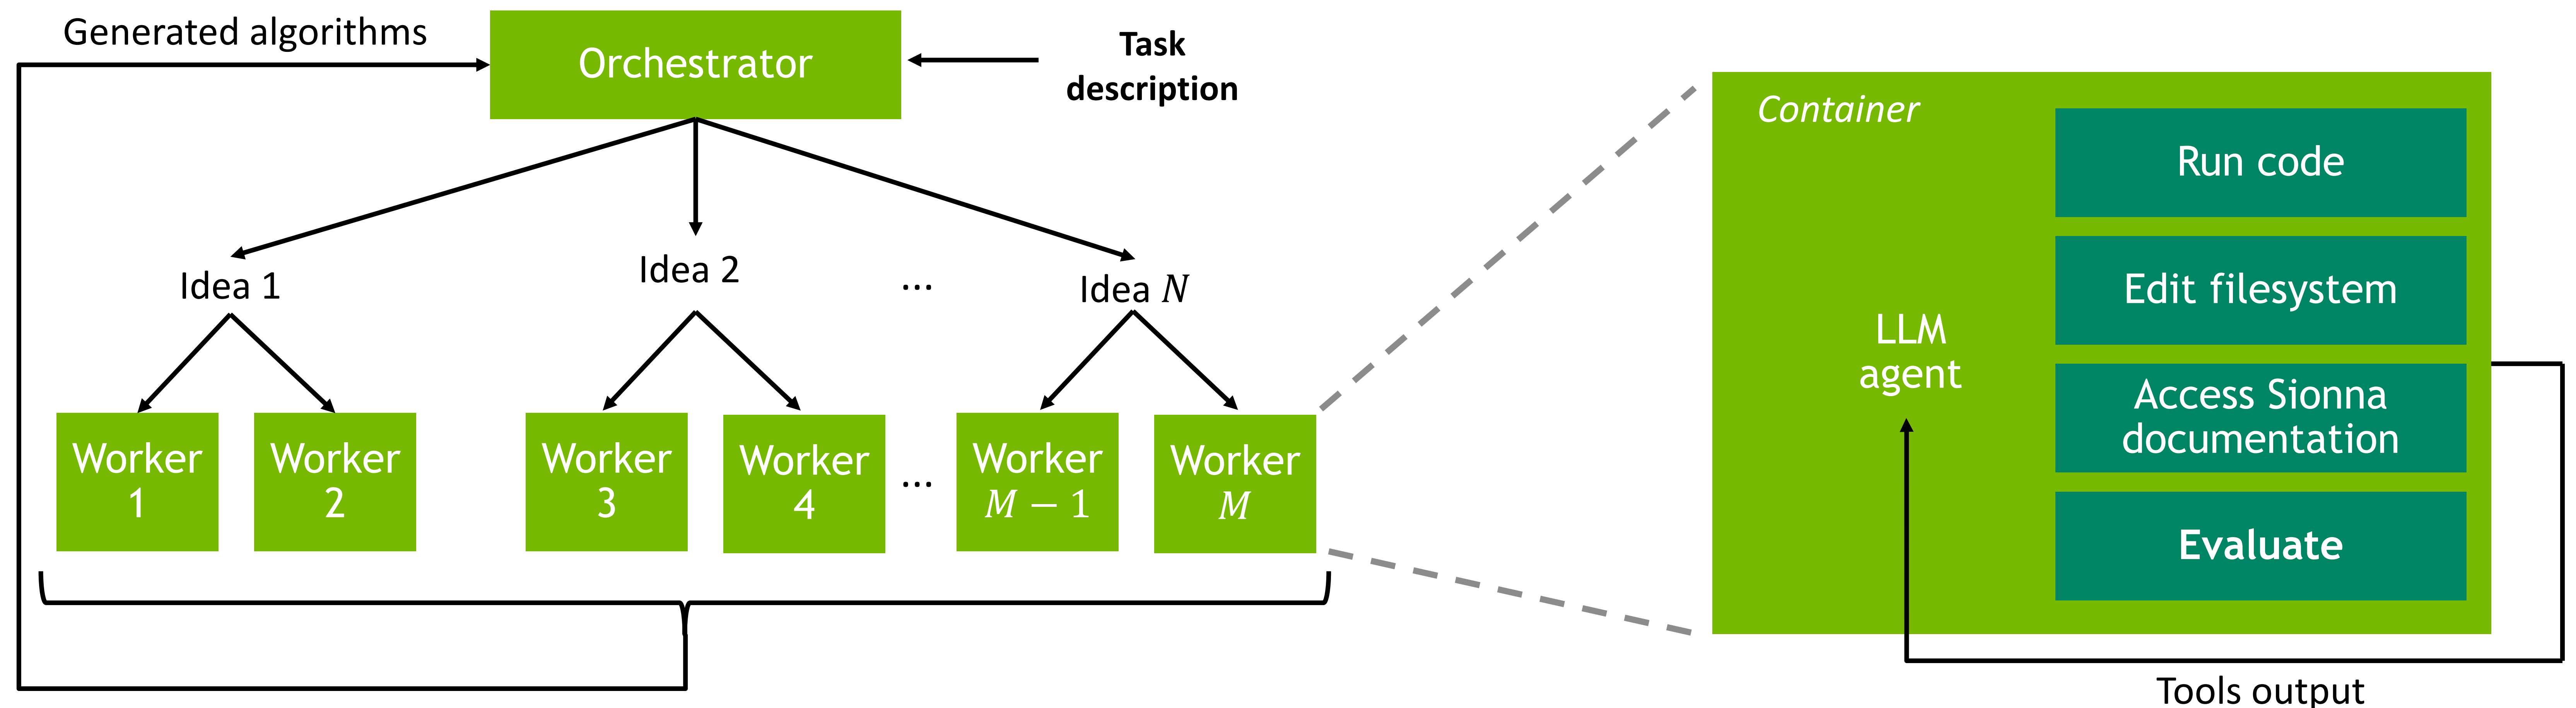


<https://github.com/NVlabs/the-ai-telco-engineer>

The AI Telco Engineer

Agentic AI framework


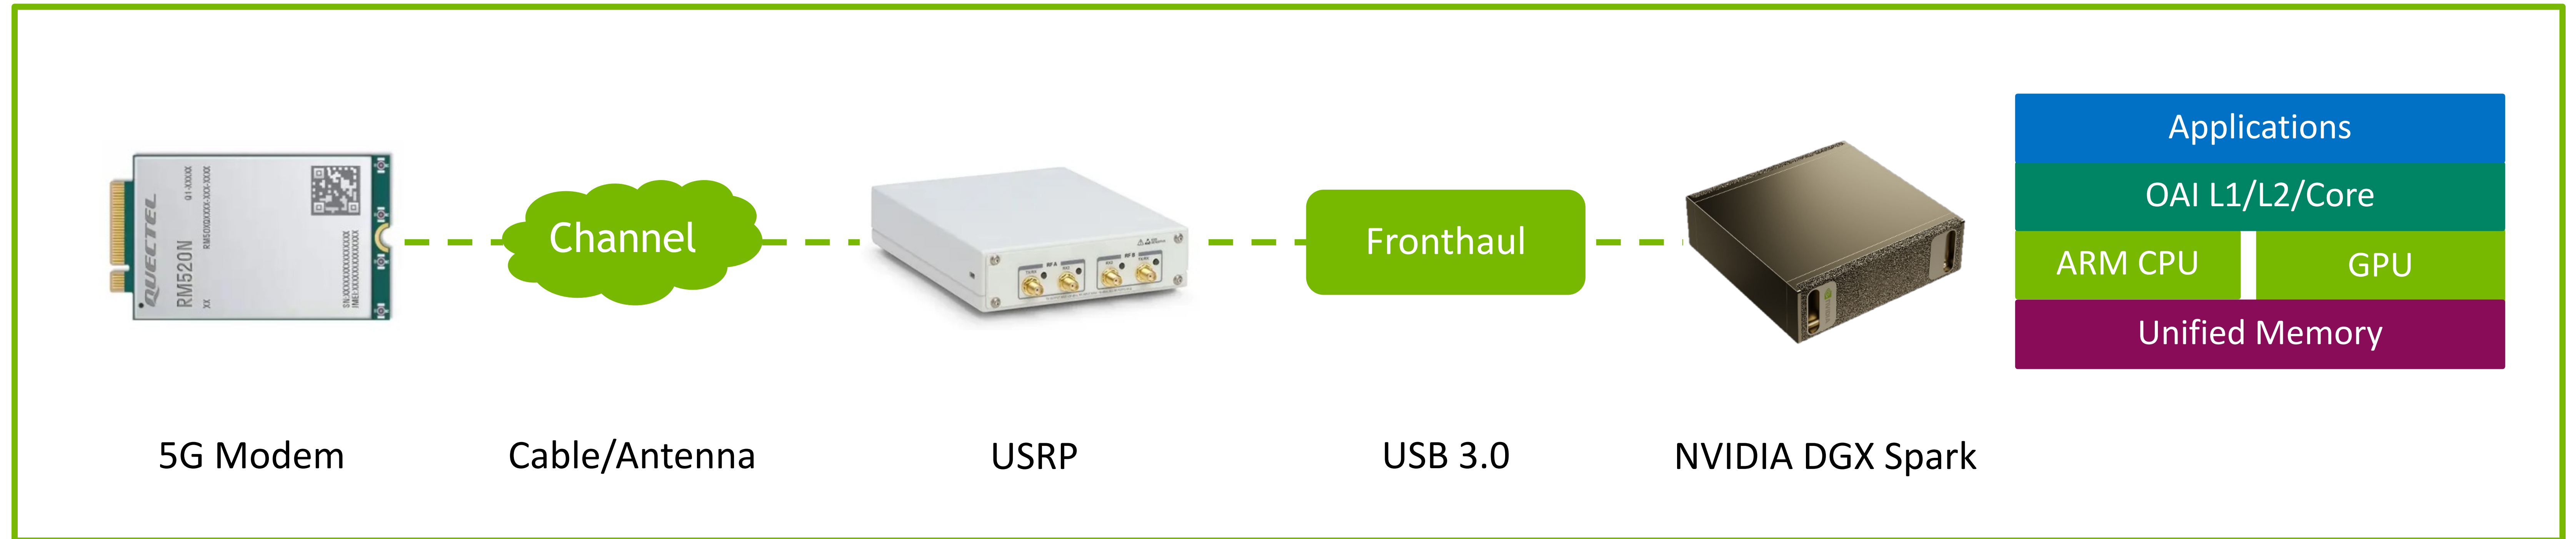
- Idea-driven framework: An orchestrator generates ideas that are distributed to parallel agents
- Use LLMs to explore algorithm designs for a user-provided wireless problem
- Workers run in a containerized sandbox and iteratively improve designs via trial and error
- User-defined evaluation is immutable (code + data) to prevent metric gaming




Toward Agents with Hardware-in-the-Loop

Sionna Research Kit


Realize ideas in an actual 5G network




Quick start in an afternoon




Affordable (6-8k\$)



GPU-accelerated



Tutorials



Blueprints for own experiments




Open source

The Economics of Training Data

Start with simulated data, and don't stop there

Cost
↑
↓
Diversity




Sionna Research Kit


Real-time channel emulation integrated in the OpenAirInterface Radio Unit (RU)



Real-world RF impairments





Emulated multipath propagation



Task: Link Adaptation

State-of-the-art performance in 8 hours for 100\$

Implement an MCS (Modulation and Coding Scheme) selection controller for link adaptation that maximizes spectral efficiency while maintaining a target Block Error Rate (BLER). Your goal is to maximize throughput (spectral efficiency in bits/s/Hz) while keeping the long-term BLER below or equal to 10%. [...]

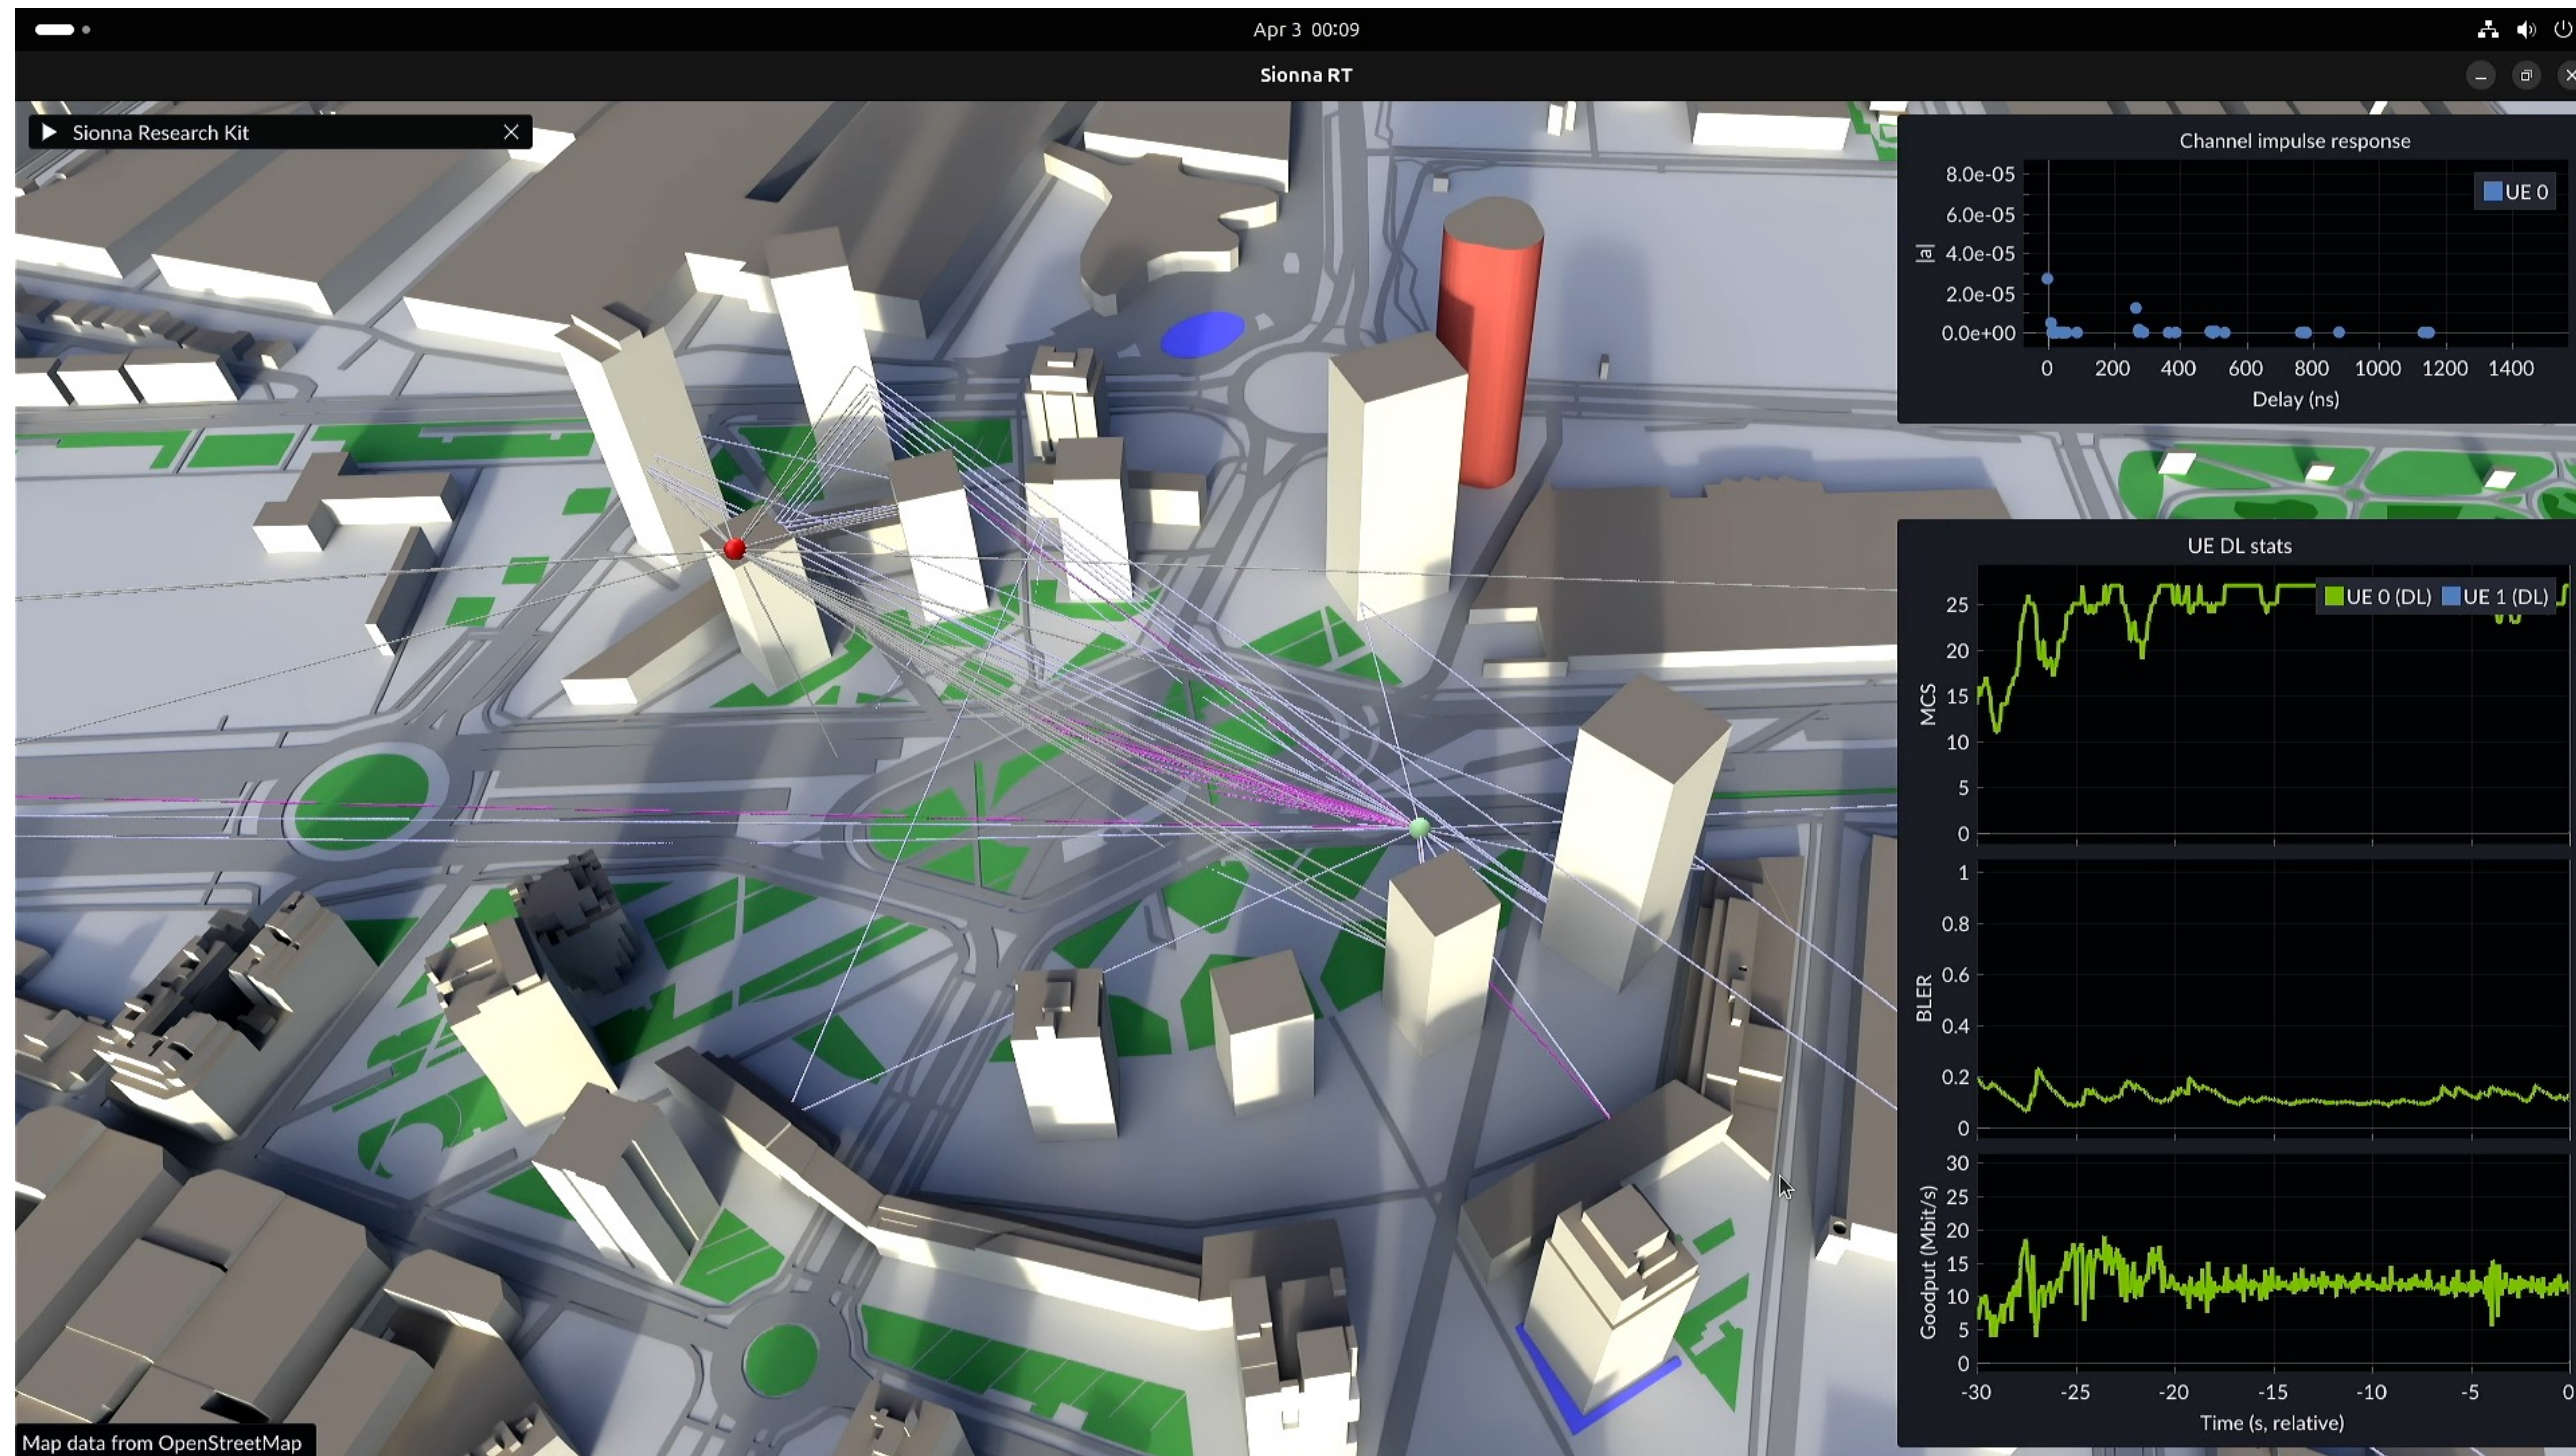


- Metric: Average spectral efficiency over 50 ray-traced channel sequences of 3000 time slots
- Algorithm has only access to ACK/NACK history

~3.4% higher spectral efficiency vs. OLLA

From Python to OAI in Minutes

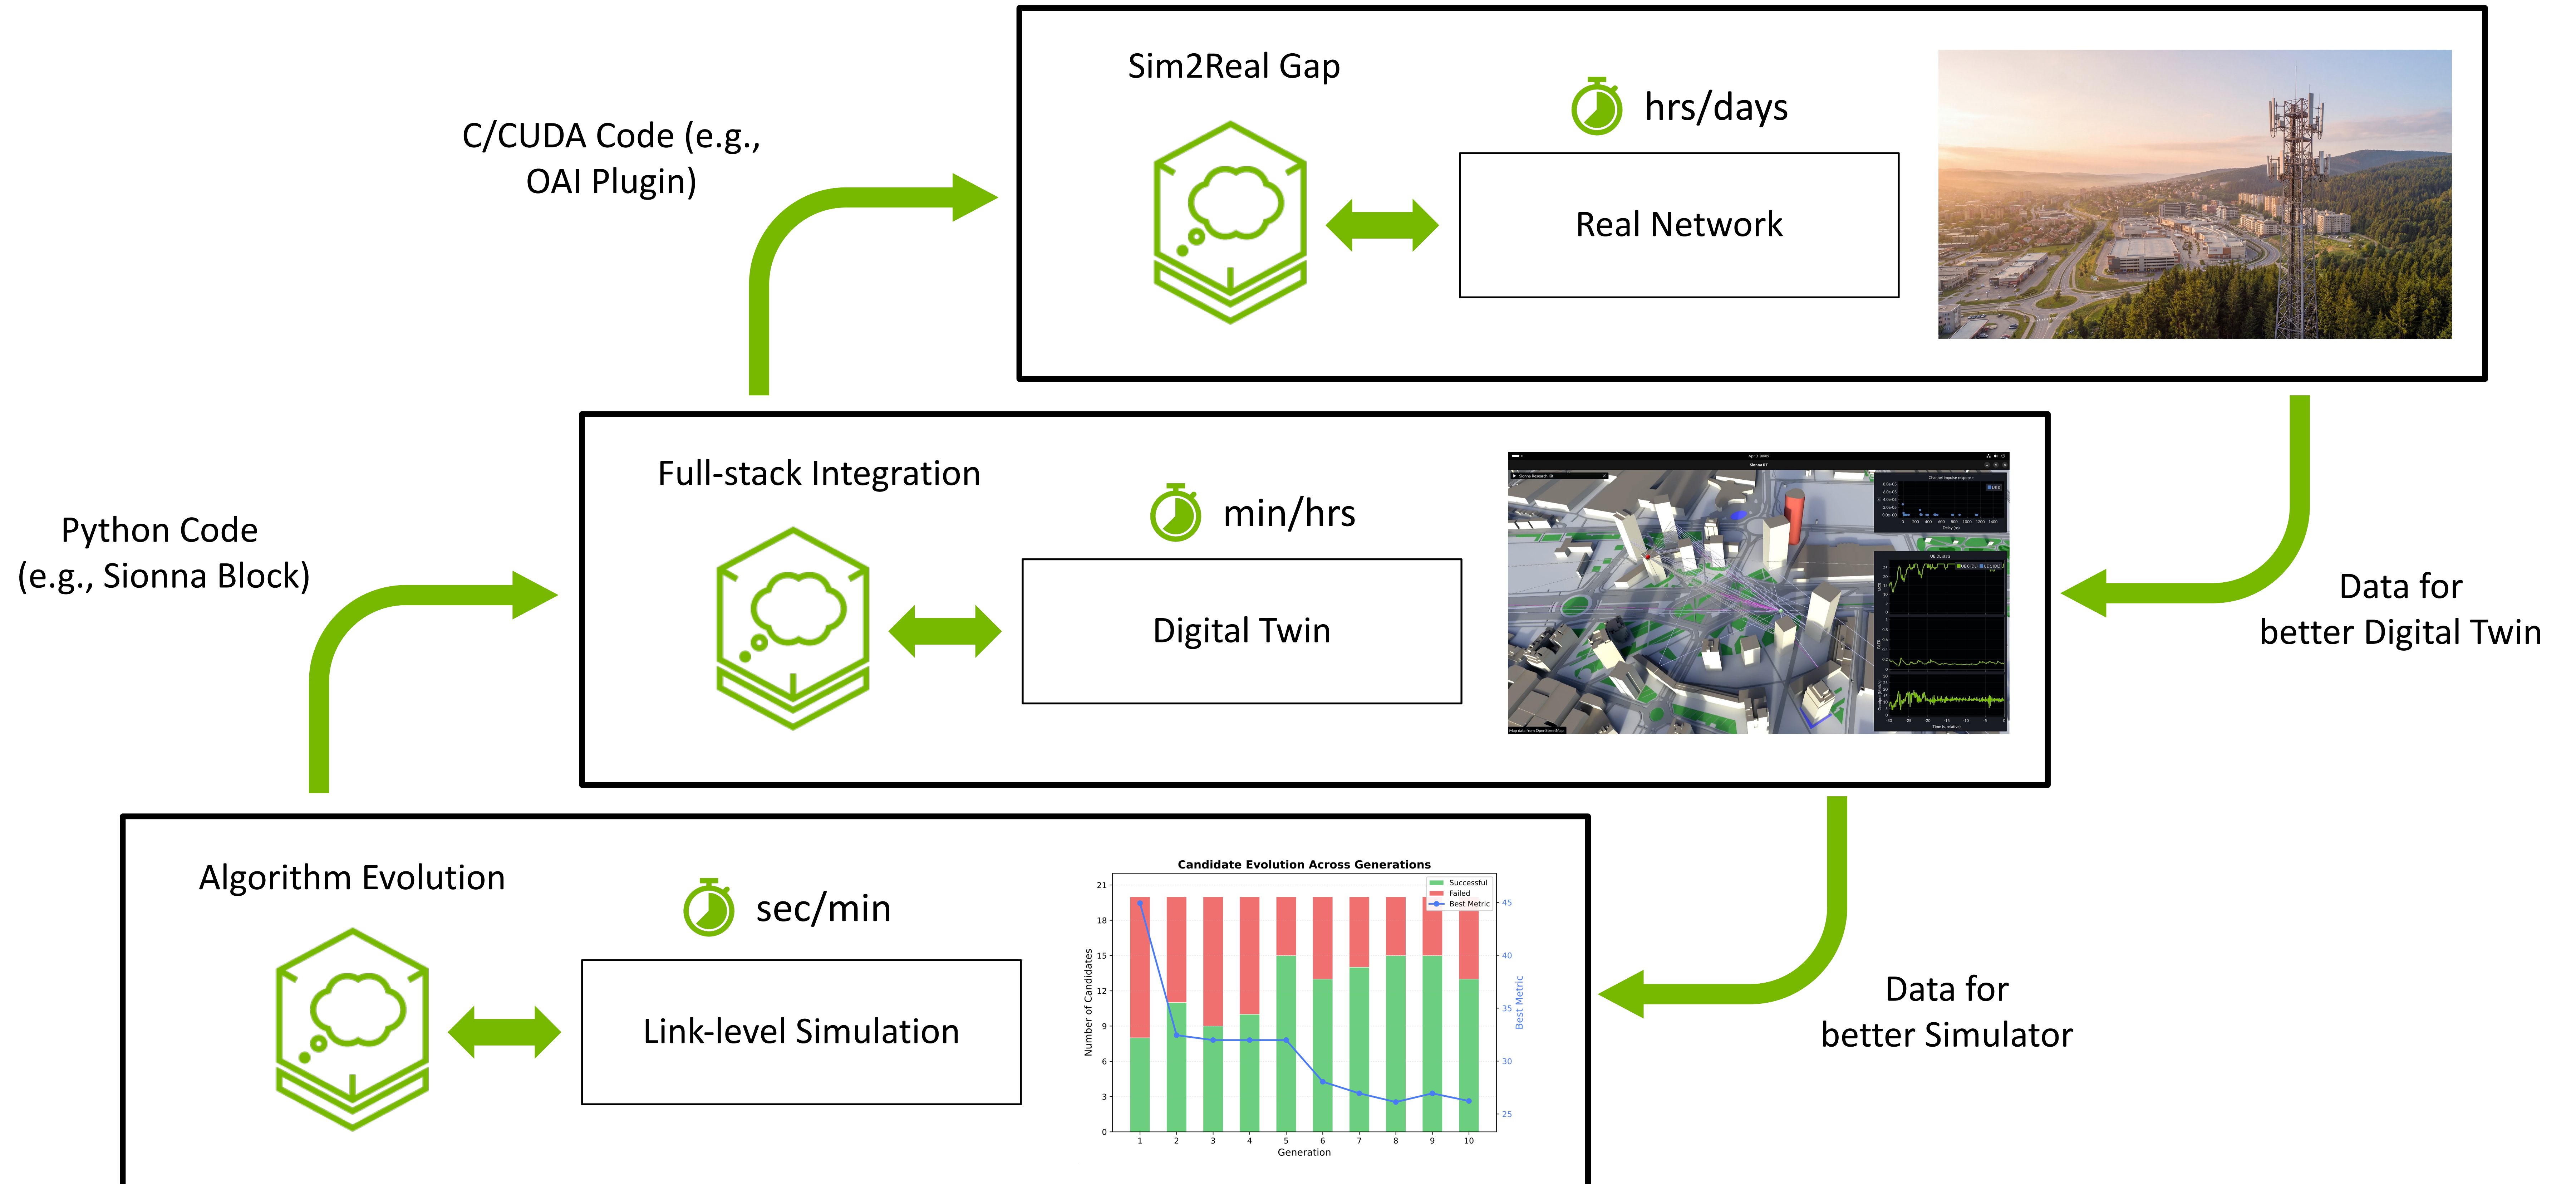
Automated agentic conversion of the link-adaptation algorithm to OAI Plugin



Original algorithm Agent fitness (MCS) history to 64 classes of (M)CSs during evolution

Hierarchical Agents

Leveraging simulations, digital twins, and real networks as evaluation tools



What's Possible with Unlimited Tokens?



Project “nanoRAN”

Agentic AI to write minimalistic 5G stack

Bring up a **minimalist 5G NR gNB (nanoRAN)** on a USRP B210 that lets a real Quectel RM520N-GL modem attach to an OAI 5GC over a cabled n78 link and run stable iperf. The gNB is greenfield – Karpathy-style single-purpose code, unified-memory IQ fast path, nanobind-only bindings, every CUDA kernel paired with a Python reference. **The MVP bar is one modem, one ping.**

Provide access to

- 3GPP TDocs 38.XXX
- OpenAirInterface & Sionna Research Kit
- OCUDU
- Aerial CUDA accelerated RAN (L1)

Implement. Debug. Repeat.

Can agents implement their own 5G stack?

```
Massive progress validated by OAI pcap comparison:

START today: UE not transmitting Msg3 (regression hidden in env defaults)

→ After 30 commits + OAI tcpdump diagnostic:

END: chain reaches PDUSessionResourceSetupResponse to AMF
PRACH ✓ → Msg3 ✓ → Msg5 ✓ → Auth ✓ → SecMode ✓
→ RegistrationAccept ✓ → ICS Response ✓
→ AMF release ✓ (correct!)
→ UE 2nd RA with ServiceRequest ✓ (the OAI flow)
→ UE PDU Session Establishment Request decoded ✓
→ AMF: PDUSessionResourceSetupRequest x16
→ gNB: RRCReconfig+DRB x8 patched
→ gNB: PDUSessionResourceSetupResponse x8 to AMF (speculative, OAI-style)
→ NGAP chain COMPLETE end-to-end
→ drb_active=0, tunnel_rx=0 ← UE doesn't apply our RRCReconfig

Key debugging-value lessons today

1. "OAI works" is a powerful debugging anchor - comparing tcpdump to ours definitively settled multiple speculations (nas_len=94 was wrong, our deferred-ICS-response was the bug, AT+CGACT is the trigger).
2. The PUSCH-RX delay-derotation fix (3a60e3c) was the biggest algorithmic win - coh 0.45 → 0.999. Permanent improvement.
3. The remaining gap is RRC ASN.1 content, not RF and not gNB protocol logic. Need to ASN.1-diff our RRCReconfiguration vs OAI's. Both pcaps are preserved (/tmp/oai_attach_1778145373.pcap, /tmp/nanoran_full_*.pcap).

Concrete next session tasks

1. tshark -r /tmp/oai_attach_1778145373.pcap -Y rrc -V - extract OAI's RRCReconfiguration container
2. Same for nanoran's pcap
3. Diff the two ASN.1 trees - focus on: SRB/DRB IDs, SDAP config, PDCP config, RLC config, MAC config, BWP config
4. Likely ~1-2 fields don't match. Fix in nanoran/rrc/setup.py or wherever DRB config is built.

The session ends with the NGAP-level chain matching OAI byte-for-byte (more or less), and a clear concrete debugging plan for closing the radio-side DRB activation. 832/832 tests pass throughout.
```

Auto-run for 4 weeks (with interrupts)

- Claude Code (Opus 4.7) and OpenAI Codex (GPT 5.5)
- Most time spend on debugging and running the system(!)
- Token cost <10k \$

Observations

- Internal debugging with OAI as reference
- RFSim mode is used a lot (~digital twins)
- Efficient hardware usage ("AT commands") & RF debugging
- Human guidance still needed(!)

Project "nanoRAN"

First step towards a fully agentic 6G stack

Repository structure

- 79% Python
- 1% CUDA
- 18% C/C++
- 2% Shell
- ASN1 used as external dependency

Connects to OpenAirInterface 5G Core

Current challenges

- Individual testability of blocks
- Digital twins & closed-loop testbeds
- Real-time & RF debugging

It's not an autonomous stack only from TDocs yet!

main nanoran / nanoran / +

Checkpoint working UE attach and ping Sebastian Cammerer authored 18 hours ago

Code owners Assign users and groups as approvers for specific file changes. [Learn more.](#)

Name	Last commit
..	
gtp	Checkpoint working UE attach and ping
mac	Checkpoint working UE attach and ping
ngap	Checkpoint working UE attach and ping
pdcp	first running iperf
phy	Checkpoint working UE attach and ping
rlc	Checkpoint working UE attach and ping
rrc	Checkpoint working UE attach and ping
soft_ue	soft_ue: add polar SC decoder + PDCCH bit-decode (CRC + RNT
uhd	hil: Phase-22c — tx_gain=70 + peak=0.30 → -78 dBm RSRP (3 d
__init__.py	Initial nanoRAN state
__main__.py	Initial nanoRAN state
cell.py	Initial nanoRAN state
cli.py	Initial nanoRAN state
config.py	Initial nanoRAN state
logging.py	Initial nanoRAN state

Hello World?

World's first ping over a fully agentic 5G stack

```
== Wait For UE Attach ==
waiting for UE attach/IP ... elapsed=15s
waiting for UE attach/IP ... elapsed=31s
waiting for UE attach/IP ... elapsed=47s
waiting for UE attach/IP ... elapsed=63s
waiting for UE attach/IP ... elapsed=79s
waiting for UE attach/IP ... elapsed=95s
UE attached      : yes
UE IP            : 12.1.1.40

== Network State ==
wwan0           : wwan0           UNKNOWN           12.1.1.40/28

== Downlink Ping ==
oai-ext-dn -> UE (12.1.1.40)
PING 12.1.1.40 (12.1.1.40) 56(84) bytes of data.
64 bytes from 12.1.1.40: icmp_seq=1 ttl=63 time=560 ms
64 bytes from 12.1.1.40: icmp_seq=2 ttl=63 time=747 ms
64 bytes from 12.1.1.40: icmp_seq=3 ttl=63 time=758 ms
64 bytes from 12.1.1.40: icmp_seq=4 ttl=63 time=1728 ms
64 bytes from 12.1.1.40: icmp_seq=5 ttl=63 time=1521 ms
64 bytes from 12.1.1.40: icmp_seq=6 ttl=63 time=1320 ms
64 bytes from 12.1.1.40: icmp_seq=7 ttl=63 time=1501 ms
64 bytes from 12.1.1.40: icmp_seq=8 ttl=63 time=1294 ms
64 bytes from 12.1.1.40: icmp_seq=9 ttl=63 time=1092 ms
64 bytes from 12.1.1.40: icmp_seq=10 ttl=63 time=1598 ms
64 bytes from 12.1.1.40: icmp_seq=11 ttl=63 time=1755 ms
64 bytes from 12.1.1.40: icmp_seq=12 ttl=63 time=1719 ms
64 bytes from 12.1.1.40: icmp_seq=13 ttl=63 time=1534 ms
64 bytes from 12.1.1.40: icmp_seq=14 ttl=63 time=1694 ms

--- 12.1.1.40 ping statistics ---
20 packets transmitted, 14 received, 30% packet loss, time 3878ms
rtt min/avg/max/mdev = 560.293/1344.387/1754.581/388.950 ms, pipe 9
```

Research at the Speed of Inference

```

• Update(scripts/start-hil-radio.sh)
  L Added 1 line, removed 1 line
  199 nanoran_default NANORAN_RRC_SECURITY_REPLAY_REPETITIONS 3
  200 nanoran_default NANORAN_RRC_SECURITY_PENDING_UL_GRACE_SLOTS 2000
  201 nanoran_default NANORAN_AS_SECURITY 1
  202 nanoran_default NANORAN_AS_SECURITY_SERVICE_REQUEST 1
  203 nanoran_default NANORAN_AS_SECURITY_SERVICE_REQUEST 0
  204 nanoran_default NANORAN_SERVICE_ACCEPT_ICS_RRC_RECONFIG 0
  205 nanoran_default NANORAN_MSG5_PUSCH_SLOT_PHASE_20 18
  206 nanoran_default NANORAN_AUTH_RESPONSE_PROFILE large

• Bash(git diff HEAD --stat)
  L scripts/hil-gnb-attach.py 1 31 ++++++-----
  scripts/start-hil-radio.sh 1 2 +-
  2 files changed, 30 insertions(+), 3 deletions(-)

• Bash(git diff HEAD)
  L diff --git a/scripts/hil-gnb-attach.py b/scripts/hil-gnb-attach.py
  index 479c3c3..022ec2c 100644
  --- a/scripts/hil-gnb-attach.py
  +++ b/scripts/hil-gnb-attach.py
  @@ -42,1 +42,31 @@
  ...

• Diff is tight. Commit:

• Bash(git add scripts/hil-gnb-attach.py scripts/start-hil-radio.sh && git commit -m
  "RC: cache-equivalent AS keys for ServiceRequest (skip 2nd-attach S.)
  [main c80fa2e] RC: cache-equivalent AS keys for ServiceRequest (skip 2nd-attach
  SMC)
  2 files changed, 30 insertions(+), 3 deletions(-)

• Bash(git log --oneline -5)
  L c80fa2e RC: cache-equivalent AS keys for ServiceRequest (skip 2nd-attach SMC)
  h 6a852ac journal: 2nd-attach AS security re-derivation is the deeper gap (49 c
  ...


• AS-context-cache equivalent shipped. The 2nd-attach RRCReconfig is now
  integrity-protected with the same KRRcInt UE cached from the 1st attach (KAMF is
  re-sent by AMF on ServiceRequest, deriving locally yields identical bytes), and the
  U SMC handshake is skipped by default.

Bench evidence chain status: 1st-attach DRB ✓ + AMF release ✓ + RRCRelease to UE ✓
+ ServiceRequest ✓ + SMC skip + cached-key install ✓ (this commit) + next
checkpoint: UE applies 2nd-attach RRCReconfig and DRB activates.

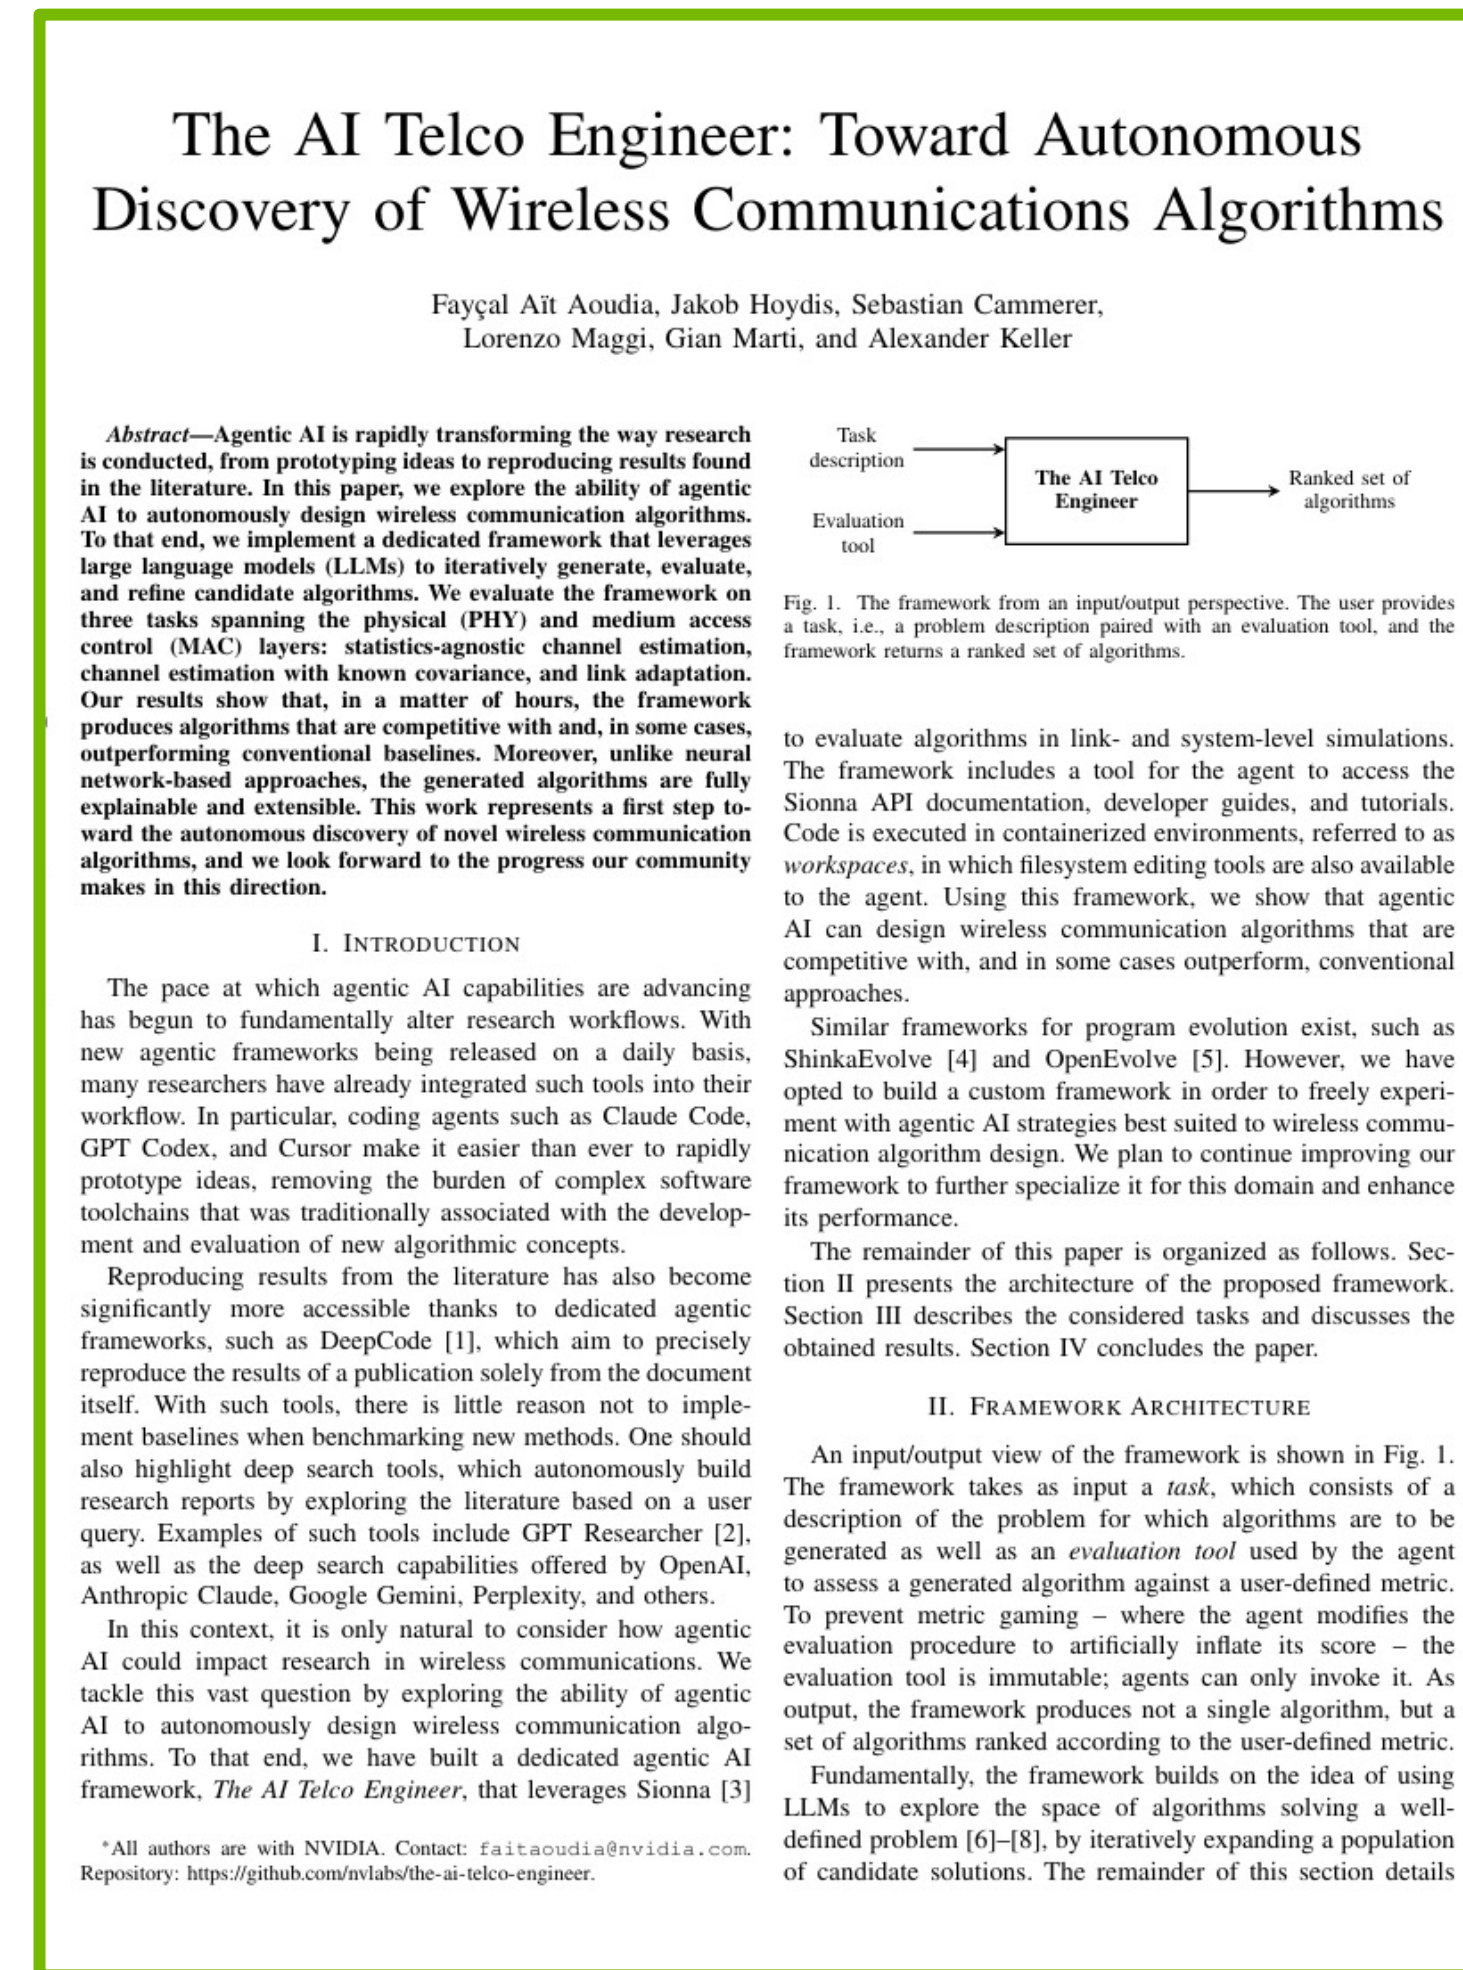

* Churned for 8m 56s

5 tasks (4 done, 1 in progress, 0 open)
✓ Survey current status
✓ Clean git commit of pending work
✓ Diagnose live attach blocker
✓ Streamline hil-gnb-attach.py
✓ Stable iperf with real UE
  
```


Agentic code generation



Open tools, platforms & twins



Open research problems



New era of wireless R&D

```
> pip install sionna  
> make sionna-rk
```