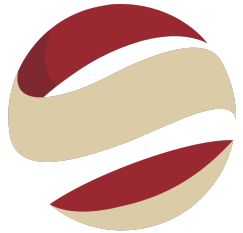

Preparando o SPRACE para a era do HL-LHC

Thiago Tomei, Rogerio Iope,
Jadir Marra, Marcio Costa

SPRACE

High-Luminosity LHC Schedule

SPRACE Computing Power Evolution

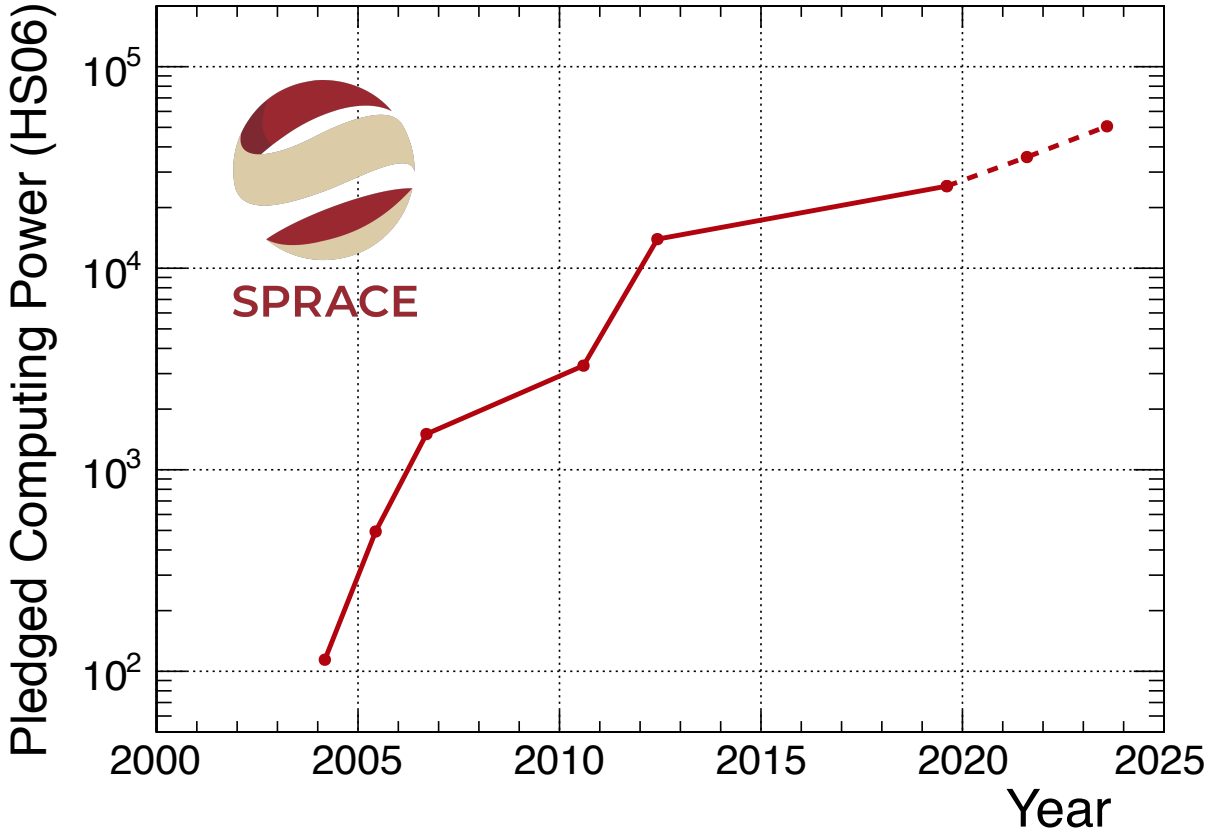

CMS HL-LHC Resource Projections

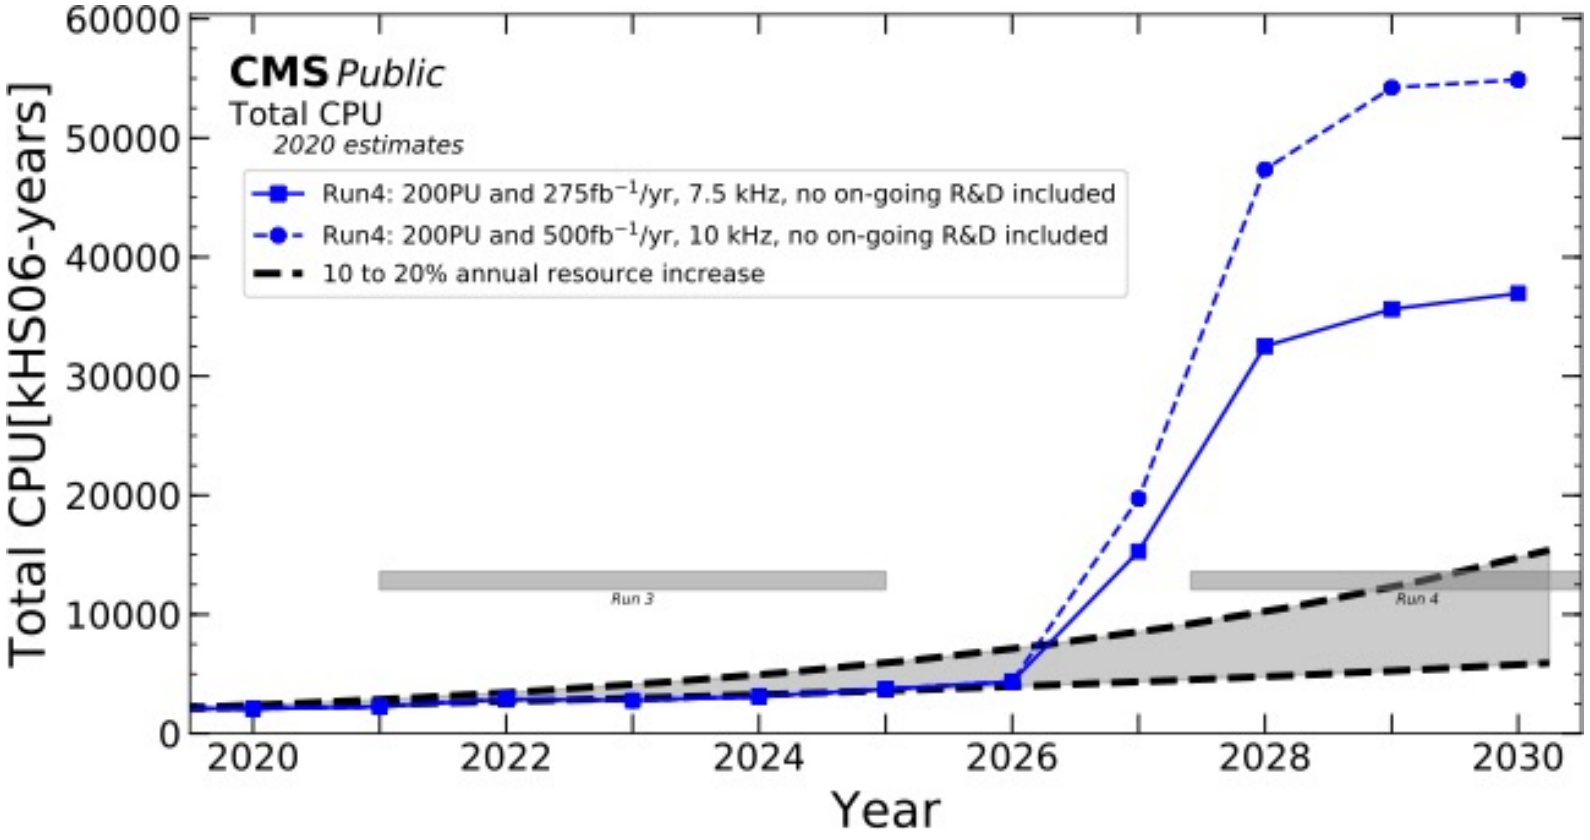

CHF/HS06 Future Projections

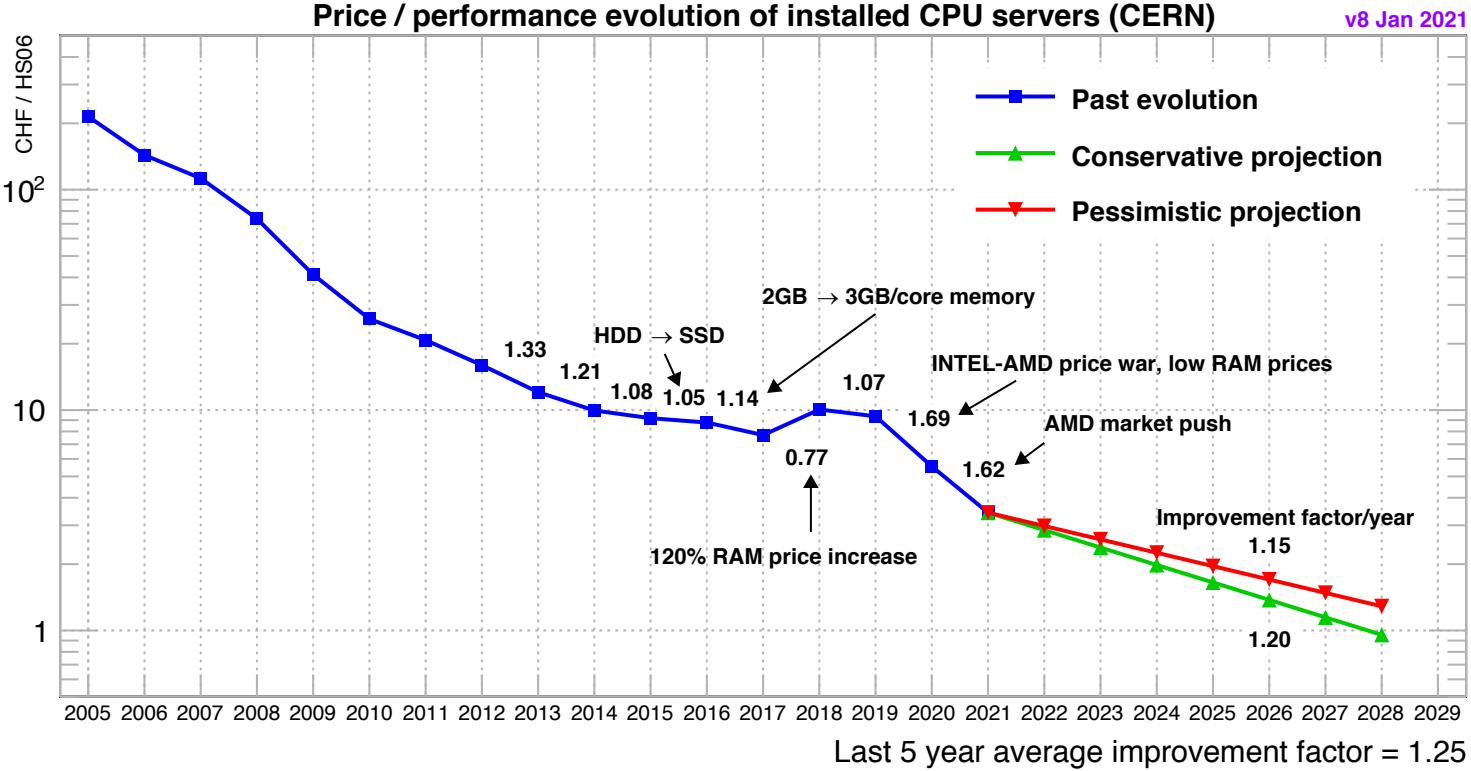

Last 5 year average improvement factor = 1.25

Machine Learning in HEP: Jet Generation

Generative Adversarial Networks (GANs)

- Generator + discriminator in tandem

Hadronic jets

- As calorimeter images
  - Convolutional NNs
  - But: sparse and unordered
- As point-cloud-style datasets
  - Graph NNs
  - But: physics features and variable particle number

Message Passing GAN (MPGAN)

- Outperforms point cloud GANs
- https://arxiv.org/abs/2106.11535

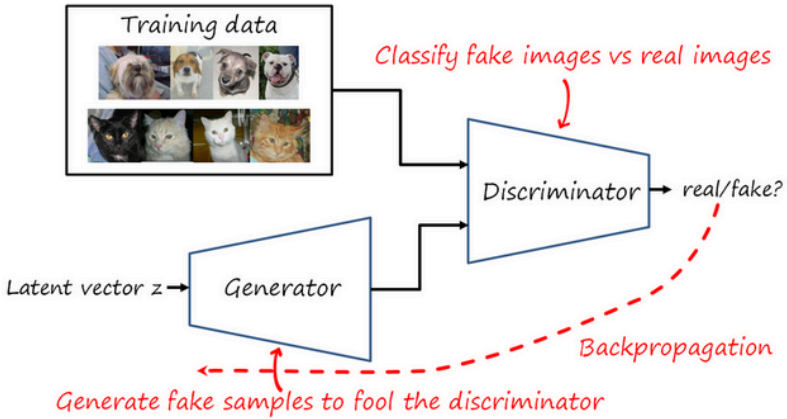

Machine Learning in HEP: Jet Generation

Generative Adversarial Networks (GANs)

- Generator + discriminator in tandem

Hadronic jets

- As calorimeter images
  - Convolutional NNs
  - But: sparse and unordered
- As point-cloud-style datasets
  - Graph NNs
  - But: physics features and variable particle number

Message Passing GAN (MPGAN)

- Outperforms point cloud GANs
- https://arxiv.org/abs/2106.11535

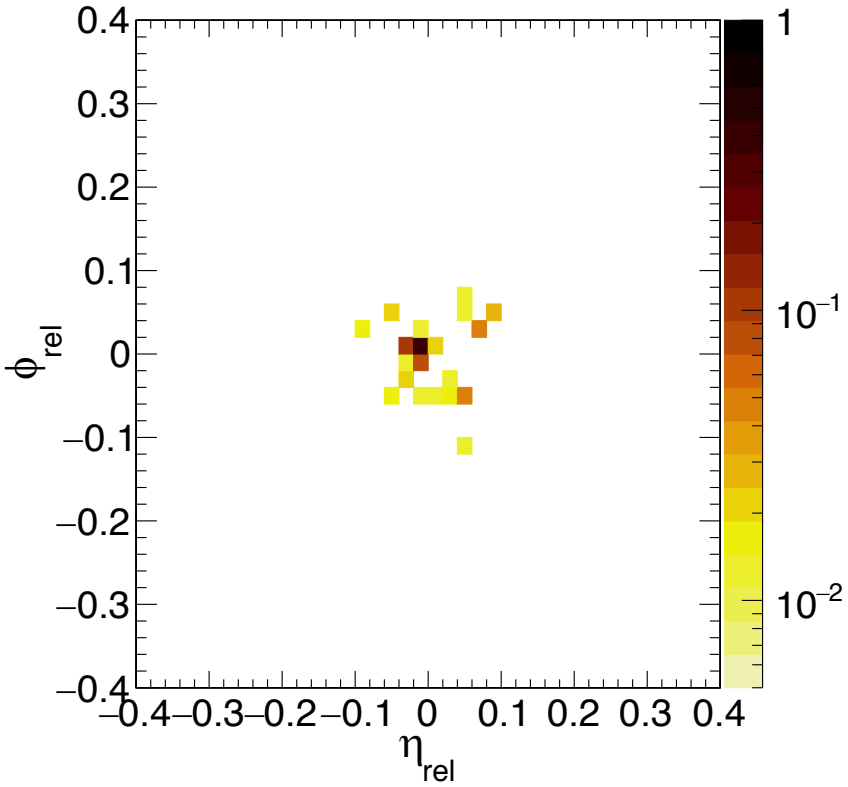

Machine Learning in HEP: Jet Generation

Generative Adversarial Networks (GANs)

- Generator + discriminator in tandem

Hadronic jets

- As calorimeter images
  - Convolutional NNs
  - But: sparse and unordered
- As point-cloud-style datasets
  - Graph NNs
  - But: physics features and variable particle number

Message Passing GAN (MPGAN)

- Outperforms point cloud GANs
- https://arxiv.org/abs/2106.11535

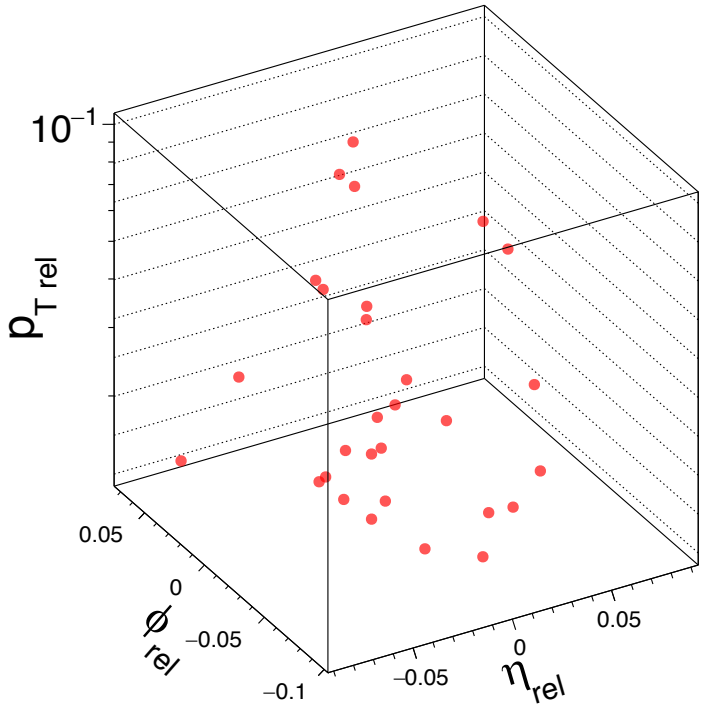

Machine Learning in HEP: Jet Generation

Generative Adversarial Networks (GANs)

- Generator + discriminator in tandem

Hadronic jets

- As calorimeter images
  - Convolutional NNs
  - But: sparse and unordered
- As point-cloud-style datasets
  - Graph NNs
  - But: physics features and variable particle number

Message Passing GAN (MPGAN)

- Outperforms point cloud GANs
- https://arxiv.org/abs/2106.11535

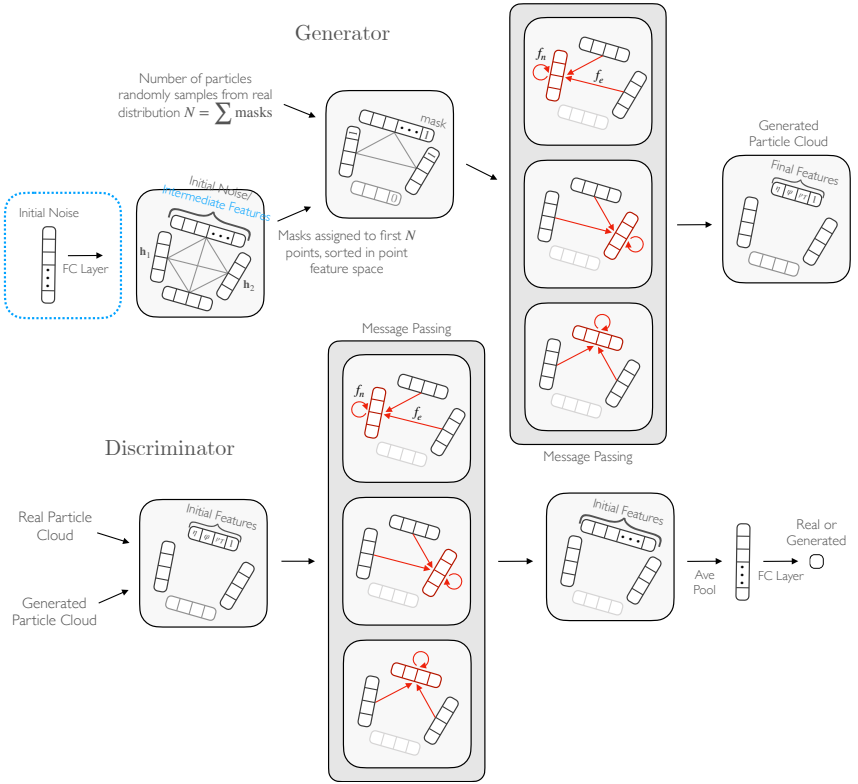

Machine Learning in HEP: Tracking

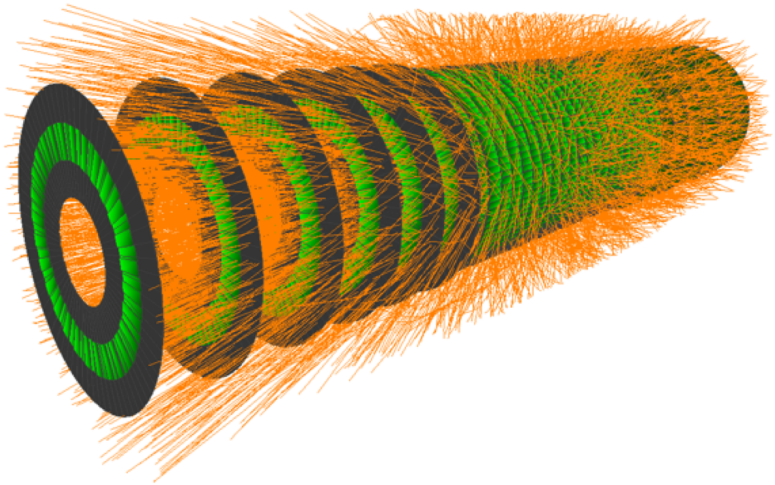

One of the most challenging HL-LHC computing tasks

- Exponential complexity with n_{hits}
- Track-ML challenge in 2018

Joining efforts with Exa.TrkX

- Geometric Deep Learning
  - Metric learning
  - (Attention) Graph NNs
- https://arxiv.org/abs/2103.06995

Machine Learning in HEP: Tracking

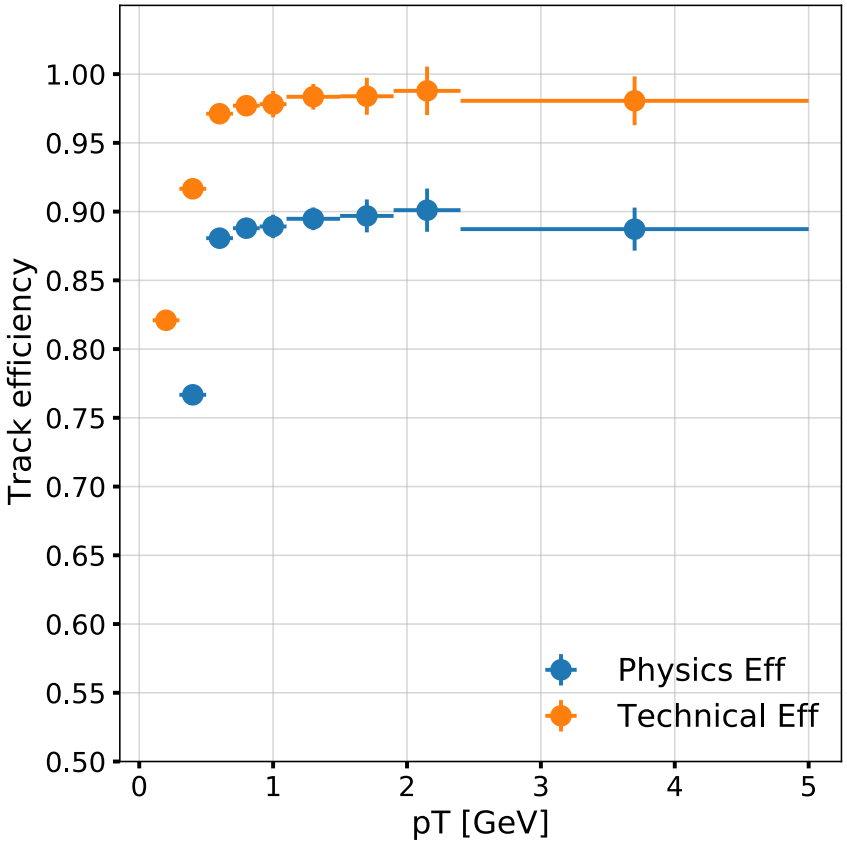

One of the most challenging HL-LHC computing tasks

- Exponential complexity with n_{hits}
- Track-ML challenge in 2018

Joining efforts with Exa.TrkX

- Geometric Deep Learning
  - Metric learning
  - (Attention) Graph NNs
- https://arxiv.org/abs/2103.06995

Heterogeneous Computing

Prototype: Phase-2 CMS HLT

- https://cds.cern.ch/record/2766559
- 37.3 MHS06 farm (PU = 200)
- Process events at 750 kHz
- High efficiency, low output rate

General-purpose GPUs

- Cheaper CHF/HS06
- Structure of Arrays
- Vectorized loops

CMSSW with GPUs

- Run-3: pixel tracks, calo local reco
- Phase-2: HGCAL reconstruction

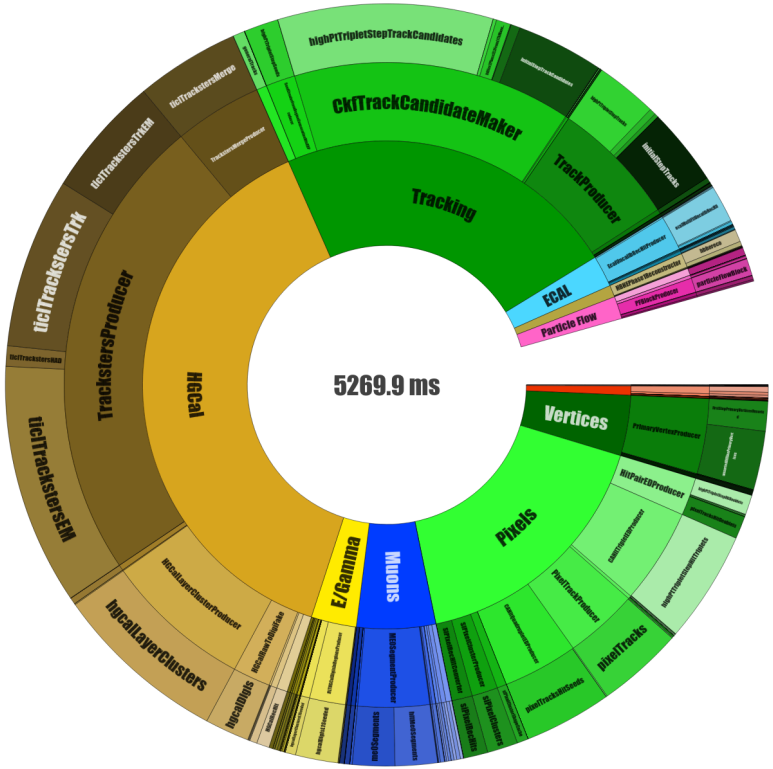

Heterogeneous Computing

Prototype: Phase-2 CMS HLT

- https://cds.cern.ch/record/2766559
- 37.3 MHS06 farm (PU = 200)
- Process events at 750 kHz
- High efficiency, low output rate

General-purpose GPUs

- Cheaper CHF/HS06
- Structure of Arrays
- Vectorized loops

CMSSW with GPUs

- Run-3: pixel tracks, calo local reco
- Phase-2: HGCAL reconstruction

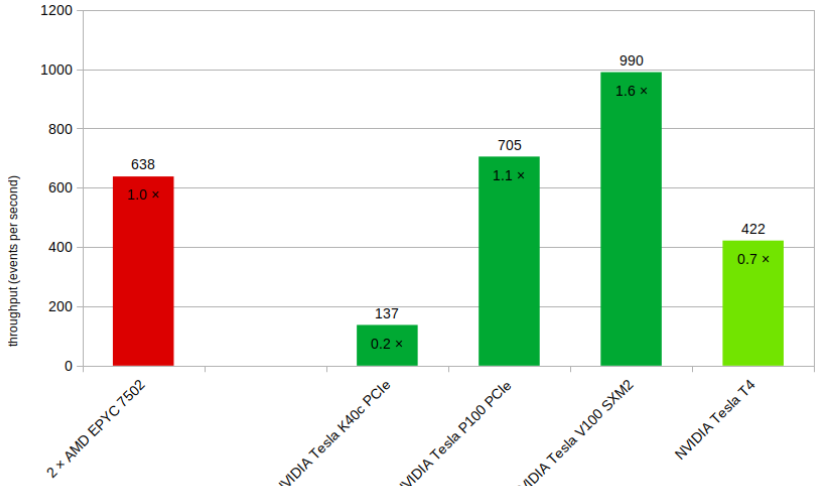

Heterogeneous Computing

Prototype: Phase-2 CMS HLT

- https://cds.cern.ch/record/2766559
- 37.3 MHS06 farm (PU = 200)
- Process events at 750 kHz
- High efficiency, low output rate

General-purpose GPUs

- Cheaper CHF/HS06
- Structure of Arrays
- Vectorized loops

CMSSW with GPUs

- Run-3: pixel tracks, calo local reco
- Phase-2: HGCAL reconstruction

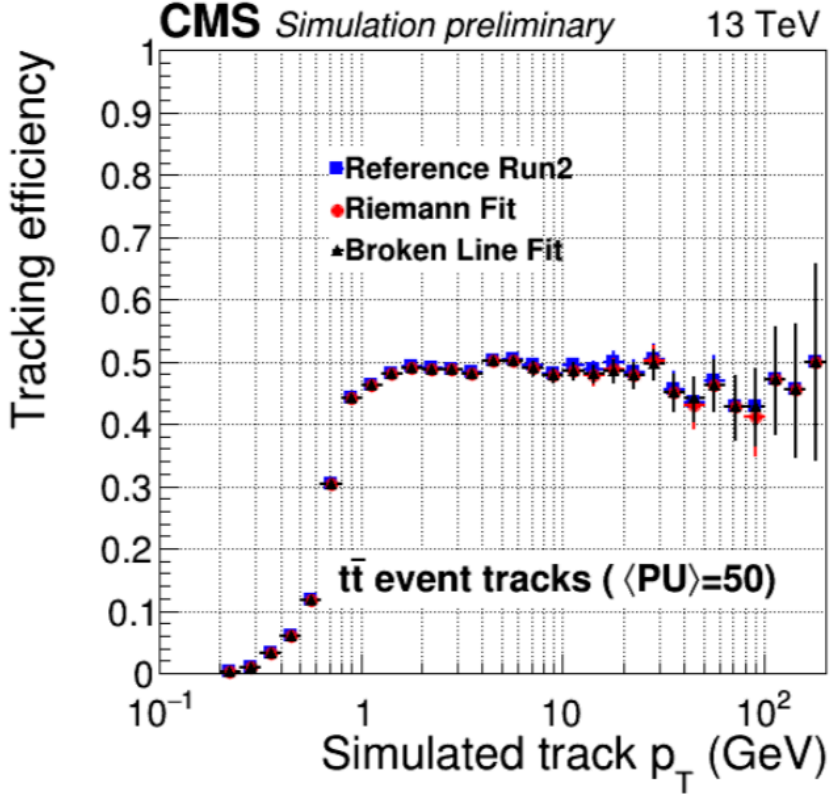

Networking – LHCONE

Private network connecting

Tier1s and Tier2s:

- Serving any LHC sites according to their needs and allowing them to grow
- Model: sharing the cost and use of expensive resources
- A collaborative effort among Research & Education Network Providers

LHCONE services

- L3VPN (VRF): routed Virtual Private Network – operational
- P2P: dedicated, bandwidth guaranteed, point-to-point links – in development
- Monitoring: monitoring infrastructure – operational

LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, NOvA, XENON, JUNO)

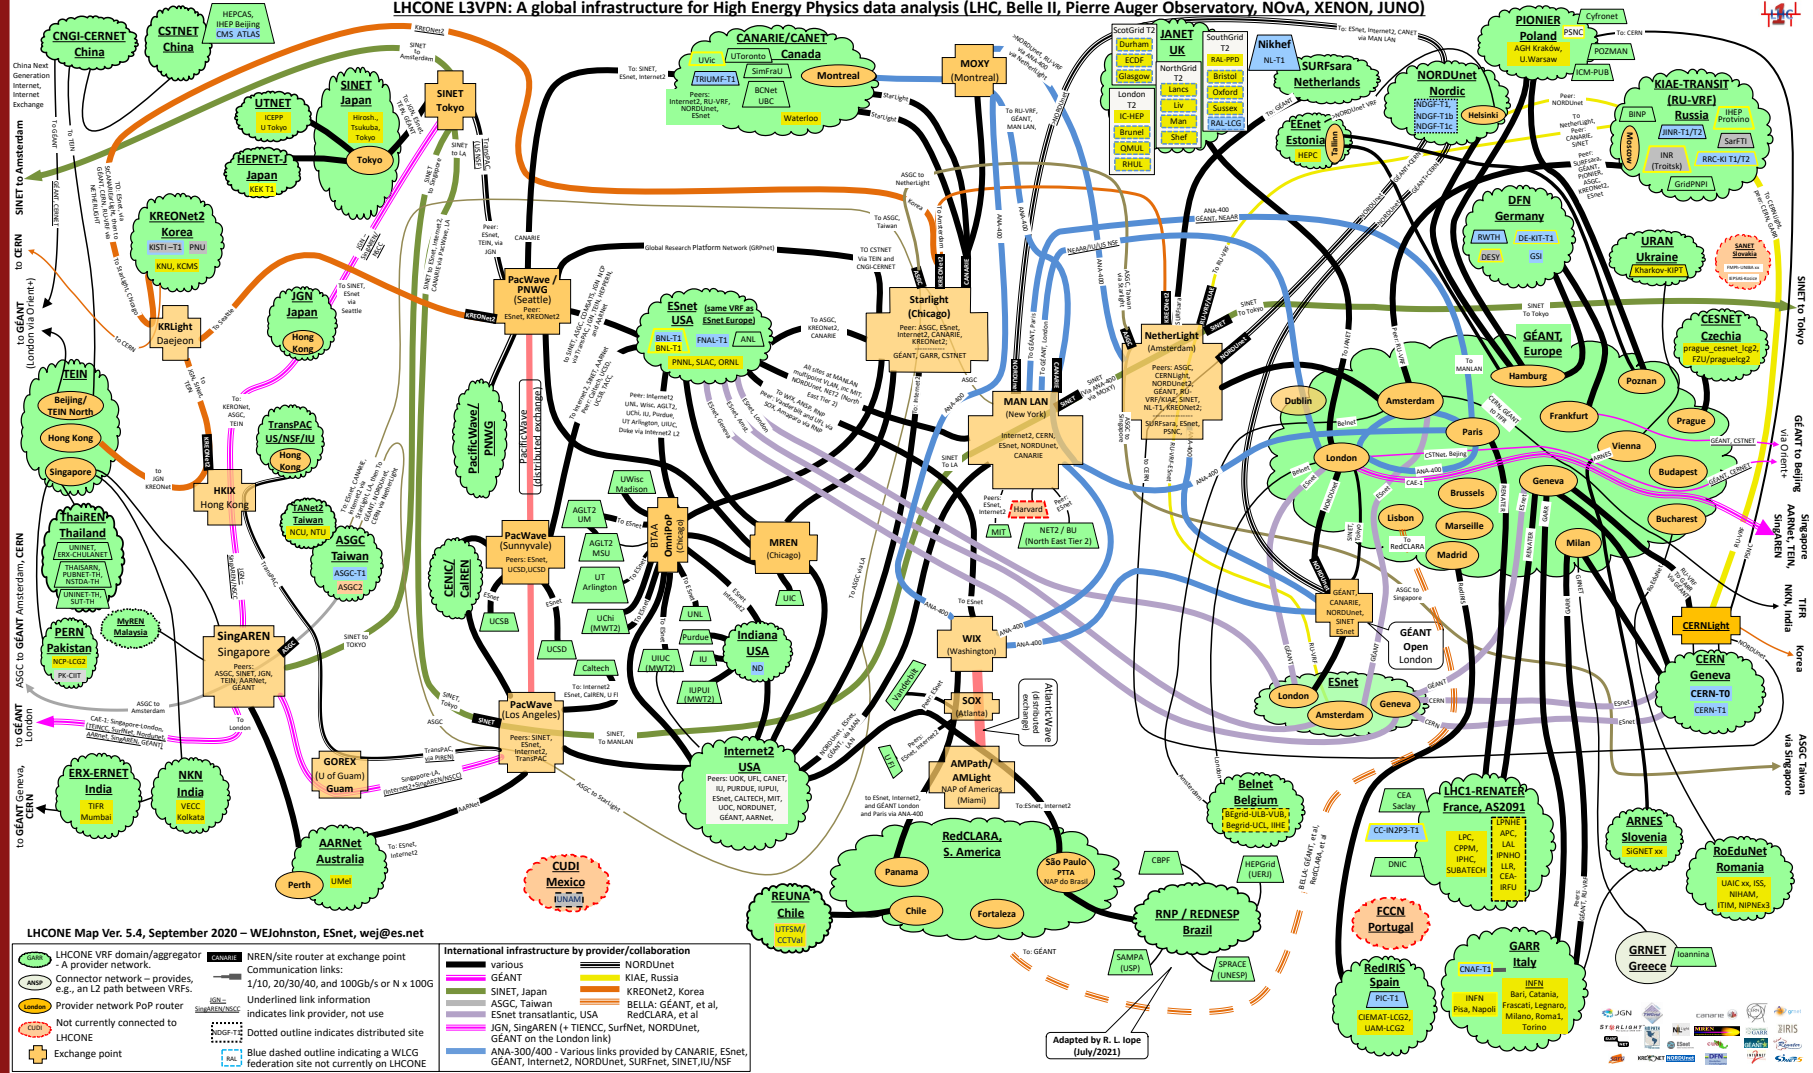

LHCONE Map Ver. 5.4, September 2020 – WEJohnston, ESnet, wej@es.net

- LHCONE VRF domain/aggregator - A provider network.
- Connector network – provides, e.g., an L2 path between VRFs.
- Provider network PoP router
- Not currently connected to LHCONE
- Exchange point
- NREN/site router at exchange point
- Communication links: 1/10, 20/30/40, and 100Gb/s or N x 100G
- Underlined link information indicates link provider, not use
- Dotted outline indicates distributed site
- Blue dashed outline indicating a WLCG federation site not currently on LHCONE

International infrastructure by provider/collaboration

- various
- GÉANT
- SINET, Japan
- ASGC, Taiwan
- ESnet transatlantic, USA
- JGN, SingAREN (+ TIENCC, SurfNet, NORDUnet, GÉANT on the London link)
- ANA-300/400 - Various links provided by CANARIE, ESnet, GÉANT, Internet2, NORDUnet, SURFnet, SINET,IU/NSF
- NORDUnet
- KIAE, Russia
- KREONet2, Korea
- BELLA: GÉANT, et al, RedCLARA, et al

NCC – NAP Network Connection

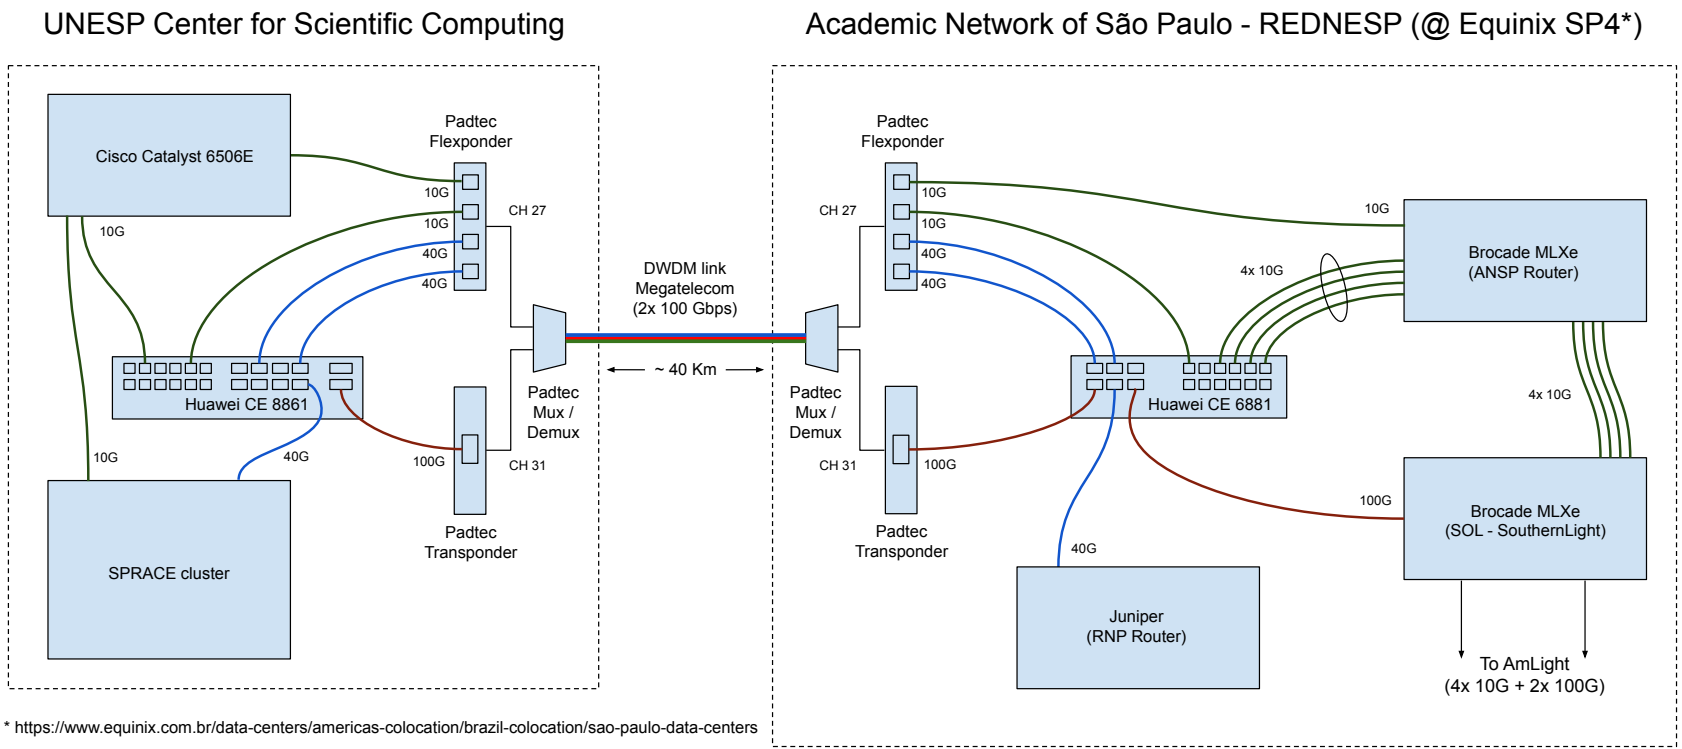

* https://www.equinix.com.br/data-centers/americas-colocation/brazil-colocation/sao-paulo-data-centers

Conclusions and Outlook

The HL-LHC represents a monumental step in computing

- Similar to the Tevatron → LHC jump
- Computing power evolution alone will not save us

Two approaches to Research and Development

- Machine Learning
- Heterogeneous computing

SPRACE will be part of the HL-LHC computing era

- Upgrades foreseen for 2021, 2023 – and beyond.
- Involvement in R&D activities

Computação no SPRACE na era do HL-LHC

Infrastrutura de Computação

- Upgrade (Storage + WNs) 2021: (92k + 150k) = 242k USD
- Upgrade (Storage + WNs) 2023: (92k + 150k) = 242k USD
- Instalação de GPUs na Tier-2: 108k USD
- Upgrade para HL-LHC (2026): 242k USD
- Upgrade para HL-LHC (2031): 242k USD
- Infrastrutura do datacenter: 40k USD

P&D para Computação

- ML TT-V: 2020-2025: 178k BRL
- HetComp TT-V: 2020-2025: 178k BRL
- ML TT-V: 2026-2027 178k BRL
- HetComp TT-V: 2026-2027 178k BRL

Cluster de Desenvolvimento

- 2 servidores, 8 GPUs: 59k USD

Azul: previsto no Projeto Temático FAPESP 2018/25225-9

Vermelho: ainda não custeado