

ORBIT: Oak Ridge Base Foundation Model for Earth System Predictability

Xiao Wang, Aristeidis Tsaris, Siyan Liu, Jong-Youl Choi, Ming Fan,
Wei Zhang, Junqi Yin, Moetasim Ashfaq, Dan Lu, Prasanna Balaprakash
Oak Ridge National Laboratory, Oak Ridge, United States
{wangx2, tsarisa, lius1, choij, fanm, zhangw3, yinj, mashfaq, lud1, pbalapra}@ornl.gov

Abstract—Earth system predictability is challenged by the complexity of environmental dynamics and the multitude of variables involved. Current AI foundation models, although advanced by leveraging large and heterogeneous data, are often constrained by their size and data integration, limiting their effectiveness in addressing the full range of Earth system prediction challenges. To overcome these limitations, we introduce the Oak Ridge Base Foundation Model for Earth System Predictability (ORBIT), an advanced vision-transformer model that scales up to 113 billion parameters using a novel hybrid tensor-data orthogonal parallelism technique. As the largest model of its kind, ORBIT surpasses the current climate AI foundation model size by a thousandfold. Performance scaling tests conducted on the Frontier supercomputer have demonstrated that ORBIT achieves 230 to 707 PFLOPS, with scaling efficiency maintained at 78% to 96% across 24,576 AMD GPUs. These breakthroughs establish new advances in AI-driven climate modeling and demonstrate promise to significantly improve the Earth system predictability.

I. PROBLEM OVERVIEW

Earth system predictability, including tasks such as accurate prediction of extreme weather and climate events, is crucial for safeguarding society and protecting economic stability, as it enables timely preparedness and mitigation strategies. The simulation of such events, however, is challenging as it demands a large ensemble size to accurately represent the diversity of possible scenarios, especially those with a low probability of occurrence. Consequently, there is a pressing need for rapid and accurate Earth system models to ensure that forecasts are timely without compromising prediction accuracy for effective decision-making.

In the Earth system, what unfolds at sub-daily to daily weather scale eventually averages out and can be characterized as sub-seasonal to seasonal climate variations. Gradual shifts in weather and seasonal patterns determine the longer-term climate changes in response to internal and external forcings. Although the Earth system operates seamlessly across different scales, the

current physics modeling approach involves using separate prediction systems for each scale, with one system used for weather forecasting, another for sub-seasonal to seasonal prediction, and a third for long-term decadal to centennial climate projections.

In response to the above challenges, AI-based models are emerging as a promising approach and aim to integrate prediction at different scales into a single system. These models leverage advanced machine learning techniques to handle the intricacies of the Earth system data, potentially revolutionizing our ability to forecast and respond to climate change. Notable models like FourCastNet [1], GraphCast [2], and PanguWeather [3] primarily use consistent datasets such as ERA5 [4] for weather forecasting, which limits their flexibility and applicability across varied tasks. Their reliance on homogeneous training and testing data is a significant challenge, restricting their utility for diverse Earth system predictability applications. Moreover, Earth system modeling benefits from abundant simulation data from multi-model resemble projects such as CMIP6 [5], which provide deep insights into the system’s structure, function, and dynamics. However, these rich datasets are underutilized in advancing Earth system modeling capabilities. Addressing this gap by integrating such extensive data can markedly improve AI model robustness and applicability, pushing the boundaries of what these tools can achieve in Earth sciences.

Recently, an AI foundation model based on the vision transformer (ViT) architecture, called ClimaX [6], has been proposed for both weather and climate modeling. It enables pre-training on CMIP6 multi-model simulation data and can be fine-tuned on limited labeled data for a variety of prediction tasks, including weather forecasting, sub-seasonal to seasonal prediction, and multi-decadal climate projection. Despite that ClimaX shows promising prediction performance, it has a relatively small model size (115 million parameters) and considers a small number of variables from limited atmospheric layers, which limits its capacity for accurate prediction and scalability

potential on large leadership-class HPC systems.

Scaling ViTs at the pre-training phase is more challenging than natural language processing (NLP) transformer models due to several key factors: 1) ViTs process high-dimensional image data that requires significantly more computational power and memory compared to the one-dimensional text data handled by NLP; 2) ViTs need to capture complex spatial dependencies within images, which is more computationally intensive than managing sequential dependencies in text; 3) The higher resolution and complexity of image data necessitate greater computational resources and sophisticated processing techniques to scale effectively; 4) Efficiently leveraging hardware for ViTs involves more complex parallelization and optimization strategies due to the intricate nature of image processing; and 5) Existing scaling strategies are mainly optimized for NVIDIA GPUs with fast interconnects, with limited support for AMD or other non-NVIDIA hardware. This results in suboptimal performance when using alternative platforms. Furthermore, optimizations are heavily biased towards NLP transformers, neglecting the specific needs of ViTs on these platforms.

We develop ORBIT, Oak Ridge Base Foundation Model for Earth System Predictability. We designed ORBIT to (1) enable scaling the model size up to 113 billion parameters, (2) incorporate 91 channels of climate variables, and (3) pre-train on 10 different CMIP6 datasets [5] with 1.2 million observation data points. Our reported 113-billion model represents the largest dense ViTs to date, surpassing the existing largest dense vision model [7] by five-fold, and a thousand-fold larger than the largest AI foundation model for climate—ClimaX [6]. ORBIT achieves 707/230 PFLOPS for the 10/113 billion parameter models on 24,576 Frontier AMD GPUs using mixed-single BFLOAT16 precision.

To scale ORBIT, we propose a novel Hybrid Sharded Tensor-Data Orthogonal Parallelism (Hybrid-STOP) by taking advantage of unique mathematical property for matrix chain multiplication and distributing model parameters among GPUs in alternating row and column shards. Through this innovation, the Hybrid-STOP algorithm combines tensor parallelism and fully sharded data parallelism together to achieve better scalability that a single parallelism cannot achieve. Importantly, this approach is designed to be architecture-agnostic, not requiring specialized hardware, which broadens its applicability across various computing platforms.

Our work represents the first instance of pre-training a large ViT model on a non-NVIDIA platform, achieving impressive scaling efficiency. This accomplishment is particularly significant given the limited interconnect




Fig. 1: Architecture of ClimaX foundational model.

capabilities of the Frontier platform, especially when compared to the faster interconnects of more recent NVIDIA platforms, and the less advanced software stack typical of non-NVIDIA environments. The successful scaling of the ORBIT model under these conditions not only highlights its robust design but also underscores the potential of alternative hardware platforms for training large AI models, thanks to the innovative Hybrid-STOP scaling approach that we developed.

II. STATE OF THE ART

Fig. 1 shows the architecture for the ClimaX foundational model with multiple input channels [6]. Each input channel is a 2D image corresponding to different climate variables, such as humidity, wind speed, and temperature. After independent patch tokenization for each channel input, the embedding for the tokenized image patches is aggregated together across different channels using cross-attention. The aggregated embeddings are then




Fig. 2: Fully Sharded Data Parallelism (FSDP) forward and backward pass.

used in the transformer training block, with each layer of the training block consisting of a self-attention and a feed-forward sub-layers. Then, predictions are computed by projecting the embedding output from the feature space to the image space.

Unfortunately, ClimaX only has 115 million model parameters. The small model size, in turn, limits ViT’s capability to learn complex data patterns from diverse and chaotic earth system information. In fact, the limitation of small vision models is not unique to climate modeling. The largest foundational dense ViT model to date has 22 billion parameters [7], and the size trails far behind those of the NLP models, which are now demonstrating capabilities in the trillion parameters scale [8], [9], [10].

Scaling large ViT has unique challenges, as mentioned before in Sec. I. The memory requirement for training a ViT model surpasses that for an NLP model at the same model scale due to the high dimensionality nature of image data. For climate modeling, ViT can have large channel dimensionality (up to 91 channels in our experiments) and each channel represents a unique state of different climate variables. Large input channels are advantageous, leading to more climate variables to be used for prediction. Unfortunately, the memory requirements for training ViT models also escalate rapidly with both the image resolution and the number of input channel variables, resulting in significantly more computations and memory use compared to NLP models.

To reduce computations and memory requirements, an effective parallelism technique is the Fully Sharded Data Parallelism (FSDP) [11], [12]. It enables distributed computations and memory by sharding both data batches and model parameters. Fig. 2 explains its operations using a 2-GPU example. In Fig. 2(a) forward pass, each GPU receives a different data batch and a different model shard. Then, the two GPUs perform an all-gather operation on their model shards to collect the full model in order to compute the full model activations. Since

each GPU is only interested in the activations relevant to its assigned model shard, each GPU in this example discards half of the activations, with GPU 1 keeping the left half and GPU 2 keeping the right half. In Fig. 2(b) backward pass, the model shards are all-gathered again to compute the gradients for the full model. Note that since each GPU is assigned a different data batch, the magnitude of the full model gradients is also different across GPUs. To synchronize the gradients, the model gradients are averaged together and then scattered to each GPU through a reduce-scatter operation. Finally, half of the model shard gradients are discarded, in the same way as how activation shards are discarded in the forward pass. Note that in Fig. 2, a temporary copy of the full model has to be collected during both forward and backward passes. Therefore, the performance of FSDP is limited by its peak memory use when gathering the full model parameters.

Alternative to FSDP, tensor parallelism [13], [14] does not gather model parameters during training. Instead, each GPU keeps model parameters sharded throughout training, but the GPUs are required to communicate frequently on the activations. In contrast, pipeline parallelism [15], [16], [8], [17] does not shard model parameters. It partitions a model instance into stages and distributes stages across GPUs, where both activations and gradients are communicated across stage boundaries. Nevertheless, there are two challenges for tensor and pipeline parallelisms. First, the scalability of both tensor and pipeline parallelisms is limited by the specific model architectures. The scalability for tensor parallelism is limited by the number of multi-attention heads [13], while the scalability for pipeline parallelism is limited by the number of model layers [18]. Second, vanilla FSDP does not support tensor and pipeline parallelisms to achieve further speedup and memory reduction. Therefore, constructing a solution that can integrate FSDP with these alternative parallelisms can significantly improve the scalability for ViT models.

III. INNOVATION REALIZED

A. Hybrid Sharded Tensor-Data Parallelism

To address the challenges discussed in Sec. II, we developed a novel Hybrid Sharded Tensor-Data Orthogonal Parallelism (Hybrid-STOP). It takes advantage of unique mathematical property for matrix chain multiplication and divides model parameters in alternate row and column shards. Through this innovation, the Hybrid-STOP algorithm combines the tensor parallelism and FSDP together to achieve better scalability than any single parallelism. In addition, our proposed parallelism technique is a general approach and is not limited by specific model architectures. Furthermore, the Hybrid-STOP algorithm avoids the peak memory use problem as in FSDP and leads to better memory reduction capability, by keeping the parameters sharded throughout training.

To understand the Hybrid-STOP algorithm, let us first explain the key mathematical property that the algorithm utilizes. If we performs two matrix multiplications in order, namely $y \leftarrow xAB$, where x, A, B are all matrices and y is the output, then the above computations are mathematically equivalent to:

$$y \leftarrow xAB = x[A_{*,1}, A_{*,2}] \begin{bmatrix} B_{1,*} \\ B_{2,*} \end{bmatrix}, \quad (1)$$

where $A_{*,1}$ and $A_{*,2}$ are two sub-matrices of A , split along the column direction. $B_{1,*}$ and $B_{2,*}$ are two sub-matrices of B , split along the row direction. In a more general term:

$$y \leftarrow xAB = \sum_{k=1}^K xA_{*,k}B_{k,*}, \quad (2)$$

where K is the number of sub-matrix shards. $A_{*,k}$ is a column shard for matrix A and $B_{k,*}$ is a row shard for matrix B . The same property can also be applied to gradient computations. If we denote the gradient of y with respect to x as $\frac{\partial y}{\partial x}$, then:

$$\frac{\partial y}{\partial x} = \sum_{k=1}^K B_{k,*}^T A_{*,k}^T. \quad (3)$$

The above equations are directly relevant to the core operations of the ViT training block. Each layer of the transformer training block, as shown in Fig. 1, consists of a self-attention and a feed-forward sub-layer. For both sub-layers, the primary computations can be summarized as performing two matrix multiplications in sequence. For the self-attention sub-layer, the primary computations are $\text{softmax}(QK^T)V$, where Q is a query matrix, K is a key matrix and V is a value matrix [19]. For the feed-forward sub-layer, the primary

computations are $\text{GeLU}(xA)B$, where x is input to the feed-forward sub-layer. A and B are model parameters for linear transformation [19], [13]. Notice that even if we include the additional computations for softmax and GeLU, the key operations for both the self-attention and the feed-forward sub-layers are still conformed to the general form: $y \leftarrow xAB$. The Hybrid-STOP algorithm recognizes this pattern between the feed-forward and self-attention sub-layers. It uses Eqns. (2) and (3) to distribute model parameters and computations for both self-attention and feed-forward, while also integrating tensor and FSDP parallelisms.

Fig. 3 explains the Hybrid-STOP operations using an example with 4 GPUs. Fig. 3(a) shows the forward pass. GPUs 1 and 2 are an FSDP group. GPUs 3 and 4 are another FSDP group. Both FSDP groups are outlined by red rectangular boxes and each FSDP group shards both data batches and model parameters. Meanwhile, GPUs 1 and 3 are a tensor-parallel group, highlighted by purple dash lines, and GPUs 2 and 4 are another tensor-parallel group. Each tensor-parallel group shares the same data batches, but shards different model parameters.

At time steps T_2 and T_3 in Fig. 3(a), the first FSDP group, consisting of GPUs 1 and 2, gathers the left half columns of the model parameters A from Eqns. (2) and 3). Then the activations corresponding to the left half columns are computed at time step T_4 . Similarly, the second FSDP group, consisting of GPUs 3 and 4, gathers the right half column shard of the model parameters A and computes half activations. Then at time step T_6 , the first FSDP group gathers the model shards for the top half row shard for the model parameters B from Eqns. (2) and 3), and the second FSDP group gathers the bottom half row shard for B . By multiplying the half columns activation from time step T_4 with the half row shard for B from time step T_6 , each GPU computes an individual activation output, $xA_{*,k}B_{k,*}$ at time step T_7 . Then GPUs in each tensor-parallel group add their individual activations together, based on the mathematical principle defined in Eq. (2), and compute the final activation and training loss.

In the backward pass in Fig. 3(b), each FSDP group gathers their row shards for model parameters B at time step T_1 . Then each FSDP group performs a gather and a reduce-scatter operations to compute the gradients for the parameter row shards at time step T_2 , in the same manner as the FSDP algorithm operates in Fig. 2(b). The only difference is that hybrid-STOP does not gather the full model. Instead, GPUs 1 and 2 gathers the first half row shard, and GPUs 3 and 4 gathers the second half row shard. Then, each FSDP group gathers column shards of model parameters A at step T_3 and compute the




Fig. 3: Hybrid-STOP forward and backward pass. GPUs 1 and 2 are an example FSDP group, highlighted by a red rectangular box. GPUs 1 and 3 are an example tensor-parallel group, highlighted by purple dash lines.

gradient for its column shard at step T_4 . At time step T_5 , the gradients for each GPU’s activation are computed. To synchronize gradients, GPUs in the same tensor-parallel group add their activation gradients together through all-reduce to compute the final gradients, following the mathematical principle in Eq. (3).

Note that in Fig. 3, the Hybrid-STOP does not gather a temporary copy of the full model like the FSDP algorithm does in Fig. 2. Instead, the Hybrid-STOP keeps the model sharded, thereby leading to lower peak memory footprint and scaling to larger model size. In addition, since the proposed Hybrid-STOP algorithm is based on the general form of matrix chain multiplications in Eqns. (2) and (3), Hybrid-STOP can be applied to for both feed-forward and self-attention, irrespective to transformer architecture. This is in contrast to tensor parallelism, whose scalability is limited by the number of self-attention heads, or pipeline parallelism, whose scalability is limited by the number of layers.

B. Optimization

This section discusses various optimization techniques that were used to aid training convergence, scalability,

computing efficiency and memory use.

Architecture Optimization. Prior work on training 22-billion ViT model reports divergent training loss due to extremely large attention logits with near-zero entropy [7]. To solve this issue, we adopted the same approach from [7], by applying additional layer normalization to the queries, Q and keys, V , before computing the self-attention scaled dot product. Through these additional layer normalization, the attention logit growth is contained and the training loss divergence is prevented.

Hierarchical Parallelism. Integrating the FSDP and tensor parallelisms does not provide a sufficient source of parallelism to scale to a large supercomputer. To address that, additional parallelism is needed. In addition, how to map the Hybrid-STOP algorithm to the architecture of a large supercomputer also requires special consideration.

To scale to large supercomputer, we add Distributed Data Parallelism (DDP) [20] to Hybrid-STOP as an orthogonal level of parallelism. Fig. 4 explains the interactions among different parallel groups and how they are mapped to the hierarchical structure of a supercomputer. The tensor parallelism requires frequent fine-grain communications to perform activation reduction




Fig. 4: Hierarchical parallelism of the Hybrid-STOP. Each horizontal purple rectangle represents a tensor-parallel group. Vertical red rectangles represent FSDP groups. Green rectangles represent DDP groups.

in each training layer. Therefore, GPUs in the same tensor-parallel group are mapped to GPUs in the same node of a supercomputer to take advantage of low latency peer-to-peer network in a node. For illustration, the tensor-parallel groups are highlighted by horizontal purple rectangles in Fig. 4. Meanwhile, each GPU is connected with multiple CPUs to load input data across different climate variable channels.

In contrast, FSDP parallel groups are mapped to GPUs from different nodes, which has slower node-to-node network than the tensor-parallel group’s peer-to-peer network within the same node. We design it this way because the FSDP parallel group requires coarser communication by performing gather and reduce-scatter communications on a model shard. In the example of Fig. 4, GPUs 1 and 5 is an FSDP group and is highlighted by a vertical red rectangle. GPUs 2 and 6, 3 and 7, 4 and 8 are other FSDP groups.

The DDP groups require least communication and they require only one gradient reduction for each global data batch. Therefore, each DDP group is mapped to a sub-cluster of the supercomputer and each sub-cluster consists of multiple nodes. In addition, different DDP groups are assigned with different data subsets to train. In the example of Fig. 4, nodes 1 and 2 form a DDP group, and nodes 3 and 4 is another DDP group. Each DDP group consists of two tensor-parallel groups and four orthogonal FSDP groups.

Activation Checkpointing. We used PyTorch activation checkpointing technique [21], [22] to trade compute for memory saving on large models. Instead of keeping activation outputs from the forward pass in memory for the expected use in backward pass, the activation checkpointing technique omits saving them in memory. Then during the backward pass, the needed activation outputs are recomputed by rerunning the forward pass segment for the activation outputs.

Mixed-Precision. We used BFLOAT16 mixed-precision for faster computation. One issue with using BFLOAT16 mixed-precision is that gradients with small or large magnitudes are not always representable in BFLOAT16, and are either flushed down to zero or exploded to infinity. To address the above mentioned issues, we applied the PyTorch dynamic gradient scaling mechanism [23] during our training. It automatically detects the value range of the computed gradient magnitude. If the values are beyond the value range of BFLOAT16, the gradient scaling will scale the gradient magnitudes to the proper value range representable in BFLOAT16, and then reverse the scaling during the parameters update.

Layer Wrapping and Prefetching. To further reduce communication overhead, the Hybrid-STOP does not shard model parameters from all training block layers at once. Instead, it shards model parameters layer by layer to reduce the size of communication messages. At each layer, the FSDP groups communicate on shards of the model parameters from the same single layer each time before moving on to the next layer.

In addition, the Hybrid-STOP algorithm prefetches the gathering for the model shards before the model shards are needed, so that cross-GPU communication can occur asynchronously and the communication overhead can be hidden in the computations. For example in Fig. 3(a) forward pass, the gathering for the row shards of model parameter B is needed at time step T_6 . However, the Hybrid-STOP algorithm can prefetch the gathering when computing the activations at time step T_4 , thereby hiding the gathering communication overhead in the computations. Similarly in the backward pass in Fig. 3(b), the gathering for the model column shards at time step T_4 can be prefetched at time step T_2 .

IV. HOW PERFORMANCE WAS MEASURED

Model Configuration. All scalability performance numbers presented in the Result Sec. V are based on five different ViT model sizes: 115 million (1024 embedding, 8 layers, 8 attention heads), 1 billion (3072 embedding, 8 layers, 8 attention heads), 10 billion (8192 embedding size, 11 layers, 32 attention heads) and 113 billion parameters (12288 embedding size, 56 layers, 32 attention




Fig. 5: The maximal model size that each parallelism can scale to at different numbers of GPUs.

TABLE I: 113 billion model walltime per observation data point using 512 GPUs, with and without optimizations from Sec. III-B.

Layer Wrapping	×	✓	✓	✓	✓
Mixed Precision	×	×	✓	✓	✓
Prefetching	×	×	×	✓	✓
Activation Checkpoint	×	×	×	×	✓
Walltime (secs)	OOM	1.83	0.49	0.47	0.33

heads). The architecture was the same as the ClimaX AI foundational model [6], with a figure illustration in Fig. 1. The only exception is that we used additional layer normalization to self-attention queries and key vectors, as described before in Sec. III-B, to prevent training loss divergence.

Pre-training Dataset. We used pre-training data from CMIP6 [5], which is an international effort across different climate modeling groups to compare and evaluate climate models. The data can be accessed from the CMIP6 archive [24]. We used ten different pre-training data sources from CMIP6 (MPI-ESM, AWI-ESM, HAMMOZ, CMCC, TAI-ESM, NOR, EC, MIRO, MRI, and NESM). These sources provide a dataset spanning 65 to 100 years of simulation, yielding more than 1.2 million observation data points with a time difference of 6 hours between two consecutive observation data points.

Each pre-training observation data point has size $128 \times 256 \times C_d$, where 128×256 is an image for a single climate variable with a spatial resolution of 1.40625 degree. C_d is the configurable input channel dimension and represents the number of climate variables to be used for training. At each model size (115 million, 1 billion, 10 billion, and 113 billion), we pre-trained two different ViT models, one with an input channel of 48 climate variables ($C_d = 48$) and one with an input channel of 91 variables ($C_d = 91$). The 48 climate variables are

chosen based on the variable settings in ClimaX [6]. The 91 variables are carefully chosen by climate scientists to represent an even broader range of physical variables, including 3 static variables, 3 surface variables, and 85 atmospheric variables spanning 17 pressure levels in the atmosphere. The full list of the 91 chosen variables will be provided in the final paper’s appendix.

Fine-Tuning Dataset. We utilized the ERA5 dataset [4], spanning from 1979 to 2020, for model fine-tuning. Specifically, data from 1979 to 2018 was allocated for training the fine-tuning data. Data from 2019 and 2020 was used for validation and evaluation, respectively. This partition of the fine-tuning dataset aligns with the current standard used by the Weatherbench2 [25]. The fine-tuning dataset has the same spatial resolution as pre-training and uses 91 variables, with an input observation data point size of $128 \times 256 \times 91$. The output variables selected were geopotential at 500 hPa (z500), temperature at 850 hPa (t850), 2-meter temperature (t2m), and zonal wind at 10 meters (u10).

System Details. Experiments were performed using the Frontier Supercomputer at the Oak Ridge Leadership Computing Facility. Each Frontier node has a single 64-core AMD EPYC CPU and 8 GPUs, with 64 GB of memory for each GPU. Among the 8 GPUs, every two of them share an MI250X graphics card and are connected with Infinity Fabric CPU-GPU. The four MI250X graphics card in the same node are connected with Infinity Fabric GPU-GPU of 50GB/s. The nodes are connected via a Slingshot-11 interconnect with 100GB/s. For the software, we used PyTorch v2.2, ROCm v5.7.0, MIOpen v2.19.0, RCCL v2.17 with libfabric v1.15.2 plugin.

Performance Metrics. The total number of floating point operations (FLOPs) of the systems was collected via the Microsoft DeepSpeed Profiler [26] and we only gathered the FLOPs on GPUs. Only the mixed-precision BFLOAT16 results were reported. The scalability of the Hybrid-STOP algorithm was measured on the pre-training datasets using the following three metrics:

- *Time-to-solutions.* We reported two time-to-solutions for pre-training, one for 48 channel variables and one for 91 channel variables. The size of each 48-channel observation data point is $128 \times 256 \times 48$. The size of each 91-channel data point is $128 \times 256 \times 91$. The time-to-solution is defined as the average wall clock time to process each observation data point. The time-to-solution also equals to the runtime per epoch, divided by the number of data points in an epoch.
- *Strong scaling efficiency.* Defined as the speedup for training an epoch at different numbers of GPUs, divided by the number of GPUs. We use the runtime at 512 GPUs (64 nodes) as the 100% efficiency baseline.




Fig. 6: Time-to-solution in (a) and memory footprint in (b) at 512 GPUs with different parallelism configurations.

- *Training loss.* We provided the pre-training loss for all 4 model sizes, with the training loss computed as the latitude-weighted Mean Squared Error (wMSE).

For the final model prediction performances on fine-tuning evaluation data, predictions were evaluated over the entire grid points for the year 2020. The area differences in the grid cells from the poles to the equator yield an inordinate bias to the polar regions, if all cells are weighed equally. Thus, we evaluated the prediction performance based on two latitude-weighted metrics. One metric is wRMSE, and the other metric is latitude-weighted Anomaly Correlation Coefficient (wACC). wACC is calculated as the Pearson correlation coefficient of the anomalies with respect to the climatology for each variable, with the values ranging from -1 (perfect anti-correlation) to 1 (perfect correlation). 0 represents predictions that are identical to the climatology. A higher wACC suggests that the model predictions closely follow the actual spatial variations observed in climate data, which is desirable. On the other hand, a lower wRMSE is indicative of better model accuracy.

V. PERFORMANCE RESULTS

A. Scaling Maximal Model Size

Fig. 5 shows the maximal model sizes that tensor parallelism, FSDP, and the proposed Hybrid-STOP can scale to from 1 to 512 GPUs (64 nodes). All runs in the figure used batch size of 2 and 48 channels of climate variables. Note that the model size for FSDP can only be scaled to 20 billion parameters. The model size for tensor parallelism is scaled to 73 billion parameters. The Hybrid-STOP algorithm, however, scales the model size to 143 billion parameters at 512 GPUs. Both FSDP and tensor parallelism have their limitations. The model size scaling performance for FSDP is limited by its peak memory use when collecting temporary copies of model parameters. Tensor parallelism, however, is limited by


the ViT architecture as the amount of tensor parallelism cannot be more than the number of attention heads. In comparison, the hybrid-STOP does not have either limitation and enable better model scaling capability.

B. Optimization Techniques


Table I shows the walltime comparison, with and without the optimizations techniques from Sec. III-B. All reported numbers in the table were for the 113 billion parameters model on 512 GPUs (64 nodes), using 48 channels of variables. The program was out of memory (OOM) when none of the optimization techniques was used. With layer wrapping alone, walltime was 1.83 seconds per observation data point. With both layer wrapping and BFLOAT16 mixed-precision, walltime was reduced to 0.49 seconds. With all four optimizations, the walltime was further reduced to 0.33 seconds.

C. Configurations for Hierarchical Parallelism

The Hybrid-STOP algorithm uses three orthogonal parallelism groups, which are FSDP, tensor, and DDP. The configuration of the parallel group size for each level of parallelism has a noticeable impact on both runtimes and memory use. Fig. 6 shows the walltime and memory use for the 113 billion model to process each observation data point. All numbers in the figure were run on 512 GPUs and used a DDP group size of 1. Performance numbers were presented with different combinations of FSDP and tensor parallel group sizes. The program ran out of memory when using either FSDP or tensor parallelism alone. Using an FSDP group size of 64 and a tensor parallel group size of 8, the program has the fastest computation speed with 0.33 walltime second per data point when using a batch size of 3. This runtime is 25 times faster than the performance when using a FSDP group size of 2 and a tensor parallel group size of 256. The hierarchical parallelism configurations, however, have less impact on memory



(a) Strong Scaling for 48 Climate Variables



(b) Strong Scaling for 91 Climate Variables

Fig. 7: The strong scaling efficiencies and time-to-solutions for different model sizes, with 48 channels of climate variables in (a), and with 91 channels of variables in (b). T represents the walltime to process each observation data point. E represents the strong scaling efficiencies, compared to runtimes at 512 GPUs.

use. A mild memory footprint increase can be observed when increasing the FSDP and decreasing the tensor parallel group size.

D. Strong Scaling and Time-to-Solution

Fig. 7 shows the strong scaling efficiencies, denoted as “E” in the figure, and walltime per observation data points, denoted as “T” in the figure, from 512 GPUs (64 nodes) to 24,576 GPUs (3,072 nodes). Experiments in Fig. 7(a) used 48 channels of climate variables and experiments in Fig. 7(b) used 91 channels of climate variables. We observed that the Hybrid-STOP algorithm retained excellent strong scaling efficiency when scaling up the model size and the input channel size. With 48 input channels, all four model sizes (113 billion, 10 billion, 1 billion, and 115 million parameters) maintained strong scaling efficiencies in the range of 78% to 83% at 24,576 GPUs, compared to the baseline performance at 512 GPUs. For the performance with 91 input channels, all model sizes retained strong scaling efficiencies in the range of 81% to 96% at 24,576 GPUs.

For time-to-solution, the 113 billion parameter model takes 8E-03 second for each 48-variable observation data point at 24,576 GPUs with a sustained computing throughput at 230 PFLOPS. One pre-training epoch takes 2.7 wall-clock hours on 24,576 GPUs for the 113 billion

parameter model to see all 1.2 million observation data points. The 10 billion parameter model takes 3E-04 second for each 48-channel observation data point at 24,576 GPUs, with a sustained computing throughput at 707 PFLOPS. In contrast, due to increased memory use for more input channels, processing each 91-channel observation data point takes more walltime, requiring 1E-02 second for the 113 billion parameter model, and 4E-04 second for the 10 billion parameter model at 24,576 GPUs.

E. Pre-Training Loss

Pre-training loss for all four model sizes is presented in Fig. 8. All four models used a fixed global data batch size of 2,880 and trained for 2.5 epochs. Despite of high initial loss, the 10 billion and 113 billion parameter models converged rapidly and outperformed the smaller 100 million and 1 billion models after seeing 2 million observation data points.

F. Model Prediction

Fig. 9 presents a comparative analysis of prediction performance, measured by wACC and wRMSE, between fine-tuned models with different parameter scales - 100 million and 1 billion. The upper panels show the model performance for z500 and t850, while the lower panels




Fig. 8: The pre-training loss comparison among different model sizes with 48 input channels.

illustrate outcomes for t2m and u10. The results reveal that the 1 billion parameters models demonstrate superior prediction skills by consistently producing higher wACC and lower wRMSE across all evaluated lead times and meteorological variables. The bar graphs, depicting the wRMSE show a reduction in prediction errors for the 1 billion parameters models as opposed to the 100 million parameters models, demonstrating the advantage of employing models with increased capacity. Moreover, the line plots for wACC highlight that despite a general decline in wACC over extended forecast periods, particularly for measures such as z500 and u10, the 1 billion parameters model exhibits great stability in predictive performance, even for longer lead times up to 10 days. The sustained level of performance of the 1 billion parameters model over 100 million parameters model suggests that larger parameter scale models possess more robust and reliable forecasting capability for medium to long-range predictions and display greater potential in advancing Earth system predictability.

In this study, our main objective was to address the scaling challenges inherent in the pre-training ViT foundation model for Earth system predictability, a critical bottleneck that has not been sufficiently tackled in previous research. Our ORBIT model’s design specifically targets these scaling issues; preliminary prediction accuracy results after a few epochs already demonstrate promising results with large model sizes. This underscores the potential of our approach in effectively managing large-scale model pre-training. In the current phase of our research, we have concentrated on evaluating models with up to 1 billion parameters to establish a robust baseline for prediction accuracy across various scales and meteorological conditions. The decision to not yet compare the 10 billion and 113 billion parameter models is primarily

due to time constraints and the progressive nature of our scaling experiments. This stepwise approach ensures that our findings are both scalable and reproducible within the limits of our allocation, setting a solid foundation for future expansions in model complexity. It’s important to note that models such as ClimaX, Pangu-Weather, and FourCastNet leverage high-resolution data and longer sequence processing, presenting additional complexities in direct comparison due to the variations in data scale and intricacy. In our future work, we will homogenize the data and experimental settings to compare ORBIT with the smaller parameter scale weather forecasting models in a rigorous evaluation, ensuring a fair assessment of its performance relative to existing solutions. This strategy allows us to focus initially on overcoming scalability challenges and the scaling advancements necessary for pre-training large ViTs.

VI. IMPLICATIONS

Earth system predictability confronts challenges due to the complex nature of environmental systems and the myriad of influencing variables. This complexity demands robust, adaptable, and highly scalable computational models to enhance our predictive capabilities and deepen our understanding of the Earth system processes. To that end, AI foundation models for climate are emerging as a promising approach in recent years. The contributions of the ORBIT model significantly impact both the field of AI for climate science and the future of HPC systems and applications.

With the ORBIT model’s capacity to use large ViT architecture, incorporate 91 variables/channels, and pre-train on 10 model datasets from the CMIP6 project, a substantial advancement in AI-based Earth system modeling has been achieved. The model’s ability to process large ViT model and a wide variety of input channels allows it to capture more complex interactions and feedback mechanisms within the Earth’s climate system. This enhanced model complexity is crucial for improving the accuracy of predictions related to weather forecasting, sub-seasonal to seasonal predictions, and long-term climate projections. As climate change continues to affect global ecosystems and human societies, precise predictions become increasingly vital for planning and adaptation strategies.

The ORBIT model’s effective scaling up to 113 billion parameters on a non-NVIDIA platform, such as the Frontier supercomputer, also marks a considerable advance in the use of HPC resources for complex AI tasks. With up to 113 billion parameters, ORBIT represents the largest dense vision transformer to date, surpassing the existing largest dense vision transformer by five-




Fig. 9: Latitude-Weighted Anomaly Correlation Coefficient (wACC) and Latitude-Weighted Root Mean Squared Error (wRMSE) scores for prediction results in the test year 2020, comparing the ORBIT models with 100 million and 1 billion parameters across four output variables.

fold [7], and a thousandfold larger than the largest foundational climate model [6]. This success demonstrates that the innovation in scaling algorithms can overcome challenges such as less advanced software stacks and limited interconnect capabilities. It also promotes a more inclusive approach to HPC system development that does not depend exclusively on one type of hardware/software stack, thereby enabling diversity in hardware utilization and encouraging the development of more adaptable and robust HPC systems.

Furthermore, the introduction of Hybrid Sharded Tensor-Data Orthogonal Parallelism in ORBIT showcases the deeper integration of AI and HPC. This approach optimizes the distribution of model parameters across multiple GPUs, boosting the efficiency and scalability of AI models. Such progress is essential for the future of AI applications in science and industry, where the demand for processing large-scale, complex datasets continues to grow. Efficiently scaling AI models on diverse HPC platforms will likely drive further innovations and a broader adoption of AI technologies across various domains. Importantly, this approach is not reliant on specialized HPC architectures, making it applicable on a broad range of HPC systems.

Besides climate modeling, ORBIT provides a template for other scientific fields that process large image datasets. Fields such as astrophysics, material science, bi-

ology, and complex computational fluid dynamics simulation could benefit from applying the scaling techniques developed for ORBIT. This methodological exchange enhances the overall capability of HPC systems to address a broader array of scientific questions, extending the frontiers of AI and HPC research.

In a nutshell, the achievements of the ORBIT model has the potential to advance the field of Earth system predictability with improved methodological AI and HPC innovations and broaden the capabilities of AI and HPC for climate modeling. These advances promote a more cohesive approach to addressing some of the most urgent computational and environmental challenges of our era.

ACKNOWLEDGMENTS

The authors thank Verónica G. Melesse Vergara and Mallikarjun (Arjun) Shankar for their support of high performance computing resources. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory (ORNL), which is supported by the Office of Science of the U.S. Department of Energy (DOE) under Contract No. DE-AC05-00OR22725. This research was supported by the AI Initiative Project sponsored by the Director’s Research and Development Program in ORNL and by the BER-ASCR SciDAC Program in the DOE. This manuscript has been co-authored by ORNL, operated

by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

REFERENCES

- [1] J. Pathak, S. Subramanian, P. Harrington, S. Raja, A. Chattopadhyay, M. Mardani, T. Kurth, D. Hall, Z. Li, K. Azizzadenesheli, P. Hassanzadeh, K. Kashinath, and A. Anandkumar, “Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators,” 2022.
- [2] R. Lam, A. Sanchez-Gonzalez, M. Willson, P. Wirmsberger, M. Fortunato, F. Alet, S. Ravuri, T. Ewalds, Z. Eaton-Rosen, W. Hu, A. Merose, S. Hoyer, G. Holland, O. Vinyals, J. Stott, A. Pritzel, S. Mohamed, and P. Battaglia, “Graphcast: Learning skillful medium-range global weather forecasting,” 2023.
- [3] K. Bi, L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, “Pangu-weather: A 3d high-resolution model for fast and accurate global weather forecast,” 2022.
- [4] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, A. Simmons, C. Soci, S. Abdalla, X. Abellan, G. Balsamo, P. Bechtold, G. Biavati, J. Bidlot, M. Bonavita, G. De Chiara, P. Dahlgren, D. Dee, M. Diamantakis, R. Dragani, J. Flemming, R. Forbes, M. Fuentes, A. Geer, L. Haimberger, S. Healy, R. J. Hogan, E. Hólm, M. Janisková, S. Keeley, P. Laloyaux, P. Lopez, C. Lupu, G. Radnoti, P. de Rosnay, I. Rozum, F. Vamborg, S. Villaume, and J.-N. Thépaut, “The era5 global reanalysis,” *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020. [Online]. Available: <https://rmetsonline.wiley.com/doi/abs/10.1002/qj.3803>
- [5] V. Eyring, S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor, “Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization,” *Geoscientific Model Development*, vol. 9, no. 5, pp. 1937–1958, 2016. [Online]. Available: <https://gmd.copernicus.org/articles/9/1937/2016/>
- [6] T. Nguyen, J. Brandstetter, A. Kapoor, J. K. Gupta, and A. Grover, “Climax: A foundation model for weather and climate,” 2023.
- [7] M. Dehghani, J. Djolonga, B. Mustafa, P. Padlewski, J. Heek, J. Gilmer, A. Steiner, M. Caron, R. Geirhos, I. Alabdulmohsin, R. Jenatton, L. Beyer, M. Tschannen, A. Arnab, X. Wang, C. Riquelme, M. Minderer, J. Puigcerver, U. Evci, M. Kumar, S. van Steenkiste, G. F. Elsayed, A. Mahendran, F. Yu, A. Oliver, F. Huot, J. Bastings, M. P. Collier, A. Gritsenko, V. Birodkar, C. Vasconcelos, Y. Tay, T. Mensink, A. Kolesnikov, F. Pavetić, D. Tran, T. Kipf, M. Lučić, X. Zhai, D. Keysers, J. Harmsen, and N. Houlsby, “Scaling vision transformers to 22 billion parameters,” 2023.
- [8] S. Dash, I. Lyngaas, J. Yin, X. Wang, R. Egele, G. Cong, F. Wang, and P. Balaprakash, “Optimizing distributed training on frontier for large language models,” 2023.
- [9] W. Fedus, B. Zoph, and N. Shazeer, “Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity,” 2022.
- [10] Z. Ma, J. He, J. Qiu, H. Cao, Y. Wang, Z. Sun, L. Zheng, H. Wang, S. Tang, T. Zheng, J. Lin, G. Feng, Z. Huang, J. Gao, A. Zeng, J. Zhang, R. Zhong, T. Shi, S. Liu, W. Zheng, J. Tang, H. Yang, X. Liu, J. Zhai, and W. Chen, “Bagualu: targeting brain scale pretrained models with over 37 million cores,” in *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 192–204. [Online]. Available: <https://doi.org/10.1145/3503221.3508417>
- [11] Pytorch, “Introducing pytorch fully sharded data parallel (fsdp) api,” <https://pytorch.org/blog/introducing-pytorch-fully-sharded-data-parallel-api/>, 2022, accessed: 2023-07-31.
- [12] Y. Zhao, A. Gu, R. Varma, L. Luo, C.-C. Huang, M. Xu, L. Wright, H. Shojanazeri, M. Ott, S. Shleifer, A. Desmaison, C. Balioglu, P. Damania, B. Nguyen, G. Chauhan, Y. Hao, A. Mathews, and S. Li, “Pytorch fsdp: Experiences on scaling fully sharded data parallel,” *Proceeding of VLDB Endowment*, vol. 16, no. 12, p. 3848–3860, August 2023. [Online]. Available: <https://doi.org/10.14778/3611540.3611569>
- [13] M. Shoenybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, “Megatron-lm: Training multi-billion parameter language models using model parallelism,” 2020.
- [14] Y. Xu, H. Lee, D. Chen, B. Hechtman, Y. Huang, R. Joshi, M. Krikun, D. Lepikhin, A. Ly, M. Maggioni, R. Pang, N. Shazeer, S. Wang, T. Wang, Y. Wu, and Z. Chen, “Gspmd: General and scalable parallelization for ml computation graphs,” 2021.
- [15] C. He, S. Li, M. Soltanolkotabi, and S. Avestimehr, “Pipetransformer: Automated elastic pipelining for distributed training of transformers,” 2021.
- [16] Y. Huang, Y. Cheng, A. Bapna, O. Firat, M. X. Chen, D. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, and Z. Chen, “Gpipe: Efficient training of giant neural networks using pipeline parallelism,” 2019.
- [17] C. Kim, H. Lee, M. Jeong, W. Baek, B. Yoon, I. Kim, S. Lim, and S. Kim, “torchpipe: On-the-fly pipeline parallelism for training giant models,” 2020.
- [18] D. Narayanan, M. Shoenybi, J. Casper, P. LeGresley, M. Patwary, V. A. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro, A. Phanishayee, and M. Zaharia, “Efficient large-scale language model training on gpu clusters using megatron-lm,” 2021.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS’17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 6000–6010.
- [20] S. Li, Y. Zhao, R. Varma, O. Salpekar, P. Noordhuis, T. Li, A. Paszke, J. Smith, B. Vaughan, P. Damania, and S. Chintala, “Pytorch distributed: Experiences on accelerating data parallel training,” *Proceedings of the VLDB Endowment*, vol. 13, 2020. [Online]. Available: <https://doi.org/10.14778/3415478.3415530>
- [21] A. Gruslys, R. Munos, I. Danihelka, M. Lanctot, and A. Graves, “Memory-efficient backpropagation through time,” 2016.
- [22] T. Chen, B. Xu, C. Zhang, and C. Guestrin, “Training deep nets with sublinear memory cost,” 2016.
- [23] Pytorch, “Automatic mixed precision package - torch.amp,” <https://pytorch.org/docs/stable/amp.html#torch.cuda.amp>. GradScaler, 2024, accessed: 2024-04-05.
- [24] L. L. N. Laboratory, “Cmip6 - coupled model intercomparison project phase 6,” <https://pcmdi.llnl.gov/CMIP6/>, 2024, accessed: 2024-03-31.
- [25] S. Rasp, S. Hoyer, A. Merose, I. Langmore, P. Battaglia, T. Russel, A. Sanchez-Gonzalez, V. Yang, R. Carver, S. Agrawal *et al.*, “Weatherbench 2: A benchmark for the next generation of data-driven global weather models,” *arXiv preprint arXiv:2308.15560*, 2023.

[26] Microsoft, "Microsoft deepspeed," <https://github.com/microsoft/DeepSpeed>, 2024, accessed: 2024-04-05.