







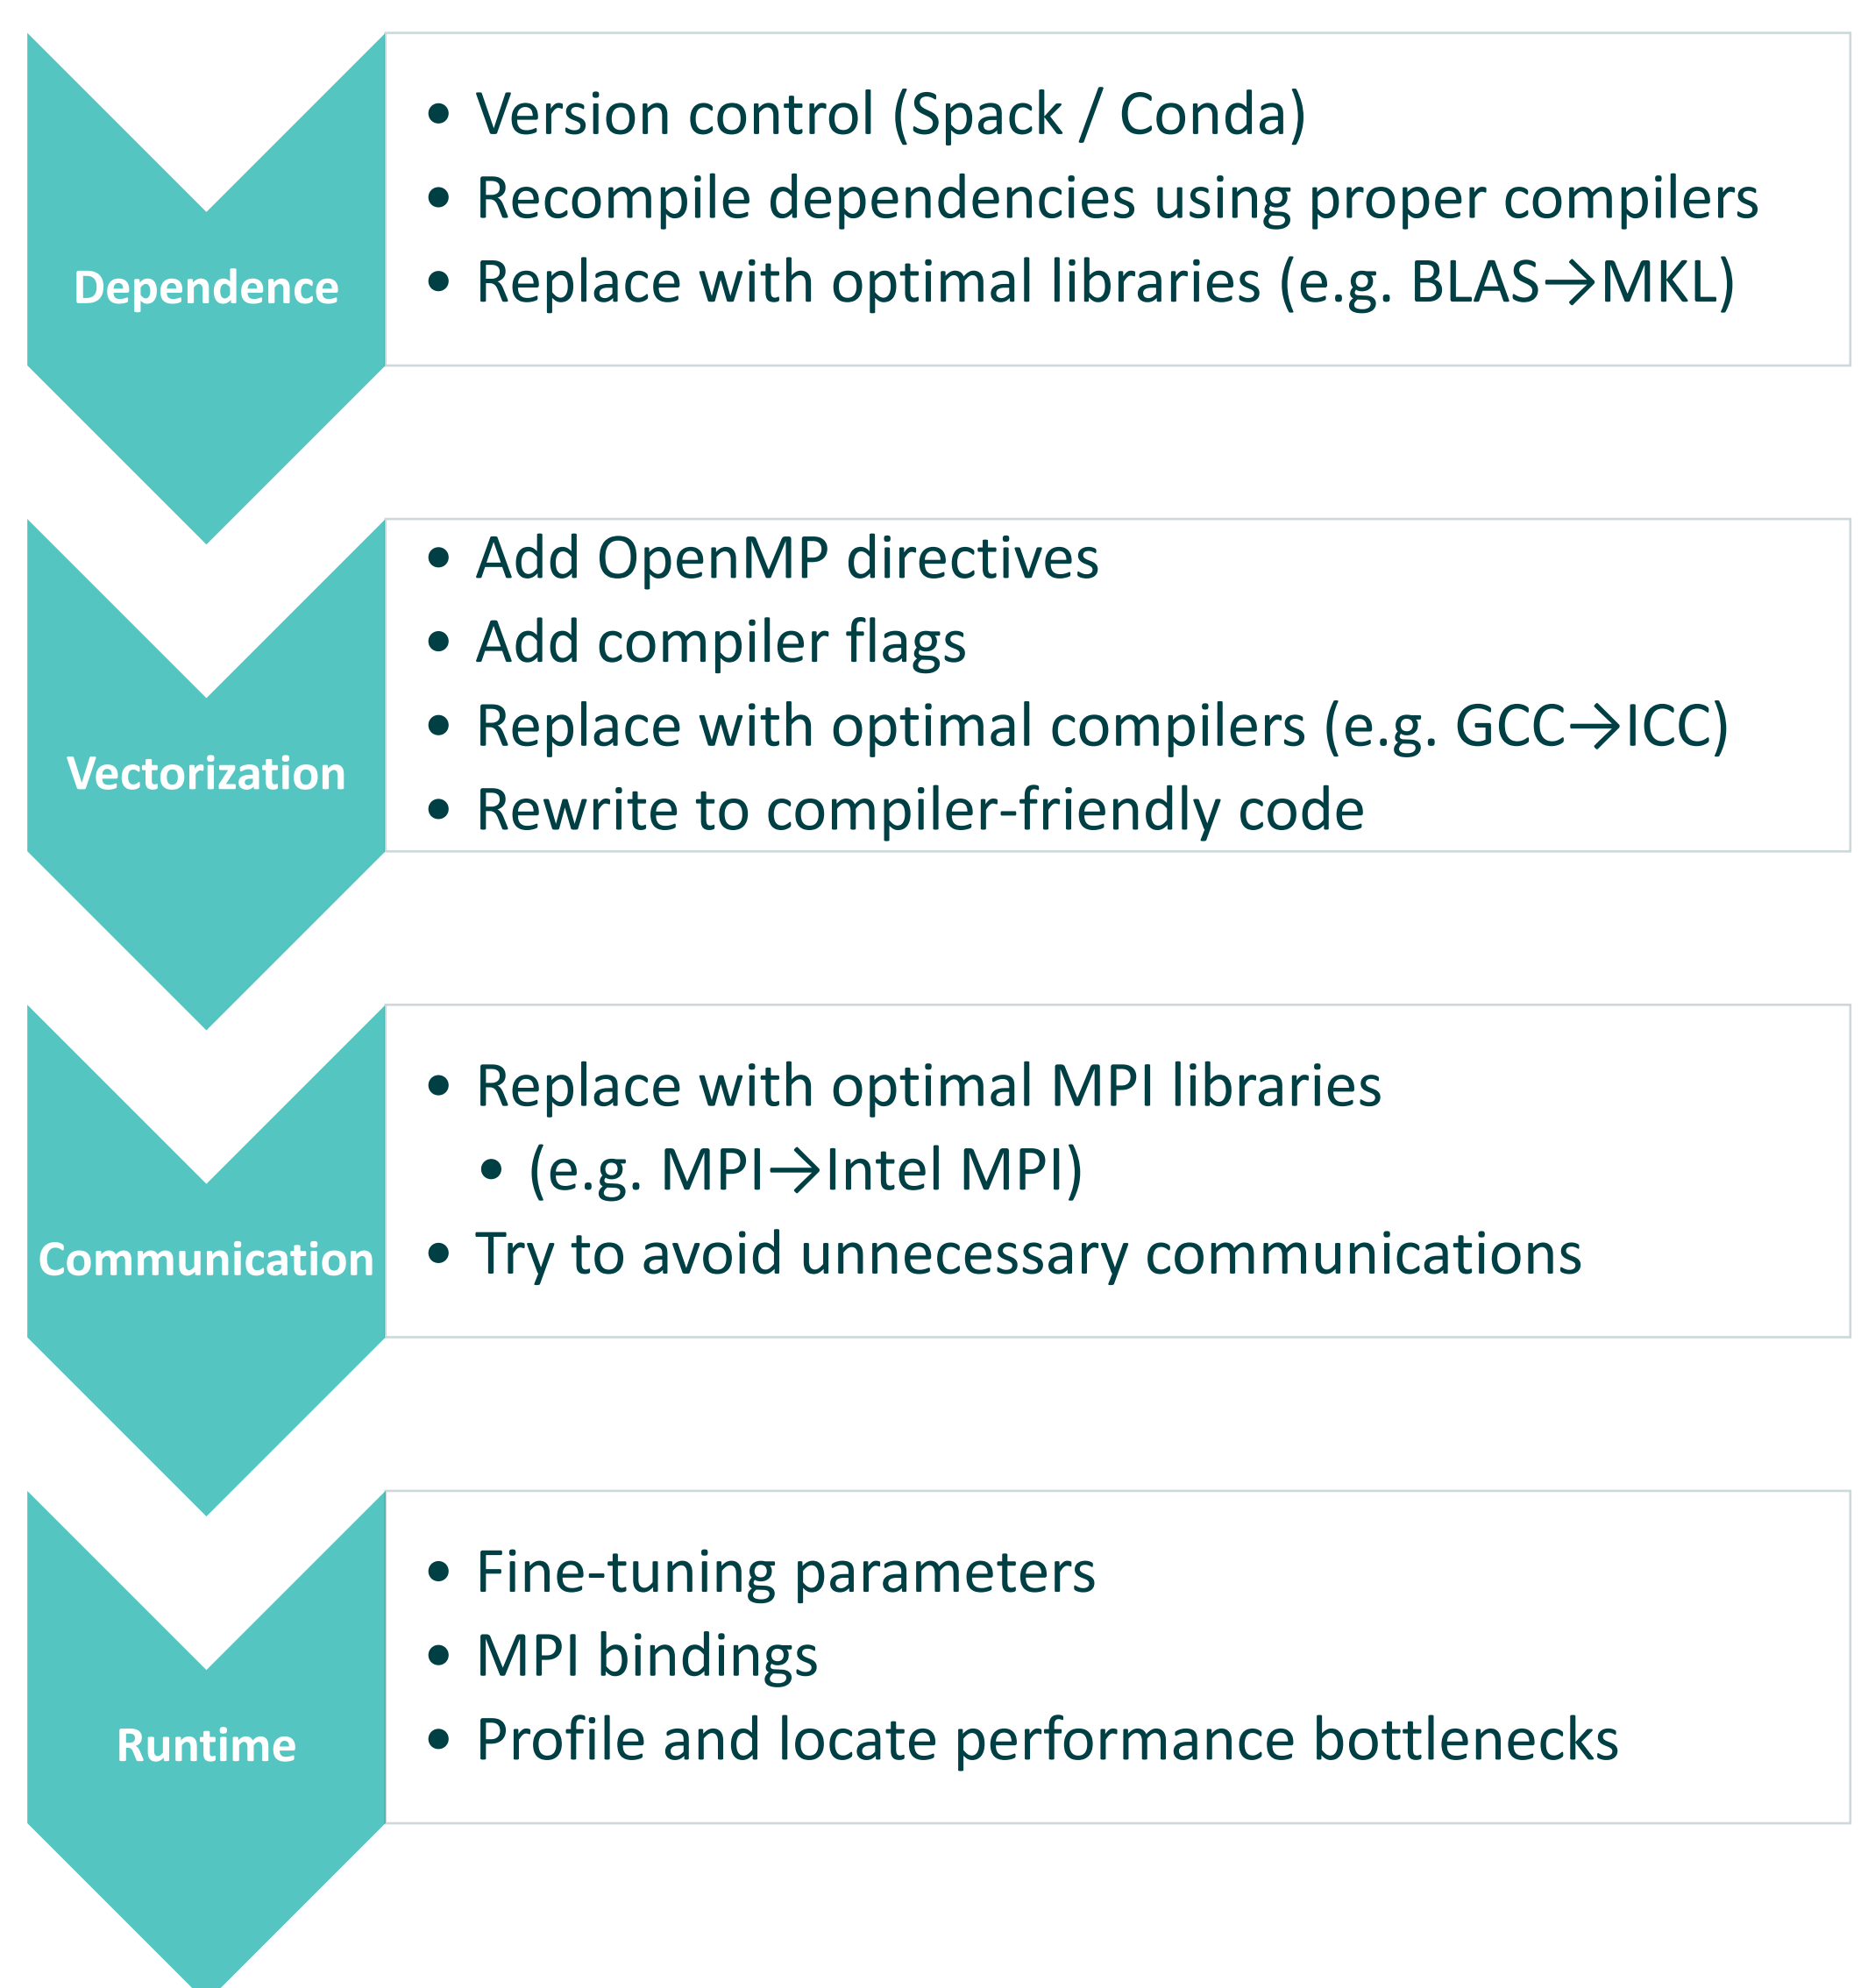
Team Members

 <p>Tongzhou Gu Year 4 @ Computer Science Maintenance & Captain HPC / MLSys / Embedded</p>	 <p>Yiqi Zhang Year 4 @ Biology Maintenance Neuroimaging</p>
 <p>Botian Xu Year 4 @ Computer Science Reproducibility Challenge ML / AI</p>	 <p>Tong Yuan Year 4 @ Microelectronics Mystery Application Tensor / Computational media</p>
 <p>Yingwei Zheng Year 2 @ Computer Science Cardioid CG / Compiler</p>	 <p>Bingzhen Wang Year 2 @ Computer Science Quantum ESPRESSO CG / Logic</p>

We, SUSTech Supercomputing Team (SST), are sponsored by the Center for Computational Science and Technology (CCSE) at SUSTech, advised by Dr. Fan Jing, the chief director of CCSE. The center is equipped with multipurpose supercomputing platforms, including NVIDIA Ampere, Intel Cascade Lake, Knights Mill, AMD EPYC Zen 2, and FPGAs. They are all open for our team to practice skills in real scenes virtually without any limitation.



Meanwhile, we are always seeking opportunities to improve our diversity. Team members major in various disciplines, such as Computer Science, Microelectronics, as well as Biology. Some of them come from minority Autonomous Regions, such as Guangxi Zhuang Autonomous Region and Inner Mongolia Autonomous Region. Moreover, we encourage females and non-binary to join us and strictly prohibit any form of discrimination. We are privileged to have alumni, Shuxin Zhen and Zhao Li, who participated in SC20 VSCC, as our consultants. They are pursuing greater successes in academia and industry now, as an outstanding female and transgender.

Optimization Approaches



The bottlenecks of applications and benchmarks vary from case to case.
HPL: Hunger for **computation FLOPS**. It is one of the ideal workloads for GPUs.
HPCG: Hunger for **memory and network bandwidth**. It is ideal to run on machines with GPUs and InfiniBand.
IO500: Consume **disk and network bandwidth**. It is a suitable workload for machines with high-speed SSD and InfiniBand.
Cardioid: According to webinars, its performance is limited by both **memory bandwidth and computation FLOPS**. Thus, It will work with GPUs well.
Quantum ESPRESSO: According to slides on the user guide, it creates **peer-to-peer communications** a lot, while it doesn't benefit from GPUs greatly. Thus, it should run on CPU servers with InfiniBand.

Cloud Configuration

 <p>AMD</p> <p>HB120rs v3</p> <ul style="list-style-type: none"> • 120x Zen3 Core • 0.4TB/s RAM BW • 7T+ FP64 Flops • 200 IB HDR • 7GB/s SSD • Good at <ul style="list-style-type: none"> • IO500 • CPU-intensive workloads 	 <p>NVIDIA</p> <p>ND96asr v4</p> <ul style="list-style-type: none"> • 8x A100 GPUs • 12.4TB/s VRAM BW • 70T+ FP64 Flops • 200G IB HDR • NVLink • Good at <ul style="list-style-type: none"> • HPL / HPCG • GPU accelerated workloads
--	--

Conda	Spack	HPCG	HPL
Intel oneAPI HPC	NVIDIA Docker	Slurm	IO500
NVIDIA HPC SDK	Singularity		
BeeGFS / NFS (over RDMA / IBofIP)			DAOS
Ubuntu 20.04 HWE			

Workloads	SKU	# of Nodes	Est Time (hrs)	Cost
HPCG, HPL	ND96asr v4	64	1	\$1741
IO500	HB120rs v3	48	1	\$173
Cardioid	ND96asr v4	4	3.5	\$381
Quantum ESPRESSO Reproducibility	HB120rs v3	32	7	\$806
Mystery	ND96asr v4 HB120rs v3	4 32	3.5	\$381 \$403

OS: **Easy to use**. There are rich resources including tutorials and troubleshooting. Also, it supports relatively new kernel (HWE) **with the latest hardware support**.
Job Scheduler: Capable of **dynamically scaling Azure Cluster**. Besides, it helps to bind CPU cores.
Compiler: Capable of generating **hardware-efficient assembly** and performing hardware-based **profiling**. Meanwhile, it bundles **optimized MPI libraries**.
Container: NVIDIA optimized HPCG and HPL and distributed their binaries as Singularity Images. Moreover, keeping the **environment isolated** could avoid programs from crashing due to contaminated environment.
Package Manager: Good at managing **common HPC / Python applications and libraries** and **isolating Python environments**, while they don't require excessive permissions so that users are less likely to damage the system.
Distributed File System: NFS is **simple, stable and easy to configure**, while DAOS is High-Performance Distributed File System developed by Intel. It is The **2nd best file system** in IO500 ranking. Thus, it is capable of improving IO500's score.