Experience and Learning from an NSM Nodal Center for Training in HPC and AI

Rupesh Nasre.
IIT Madras

***Abstract*—Training is crucial to develop the required human resource in HPC and AI. A well-trained community fuels the growth of the country, and of the scientific community in general. The government of India envisioned training of tens of thousand of learners to fulfill the future needs of the country in these domains. It was part of a nation-wide large project called the National Supercomputing Mission. Accordingly, four Nodal Centers were established in different parts of the country, whose mandate was to impart quality training in HPC and AI to the interested people. The nodal centres received funding and support from the ministry and the implementation agencies. This report narrates the experience by one of the nodal centers at IIT Madras, highlights lessons learned from this experience. We also suggest avenues for future training programmes. In particular, we discuss various certificate courses, themed workshops, industry workshops, internships, along with the various HPC resources created along the way. We also brief about the feedback we have received from selected participants. We hope that various efforts in these and similar directions would benefit from this learning.**

I. Introduction

In 2015, Indian government announced the National Supercomputing Mission (NSM) budgeted for Rs.4500 crore (over 500 million USD) over a period of seven years. The vision was to bring together the industry, academia, and government agencies working in this domain, to address nation's problems emanating from various scientific, strategic, and societal applications [1]. Two agencies, Indian Institute of Science (IISc), Bengaluru and Centre for Development of Advanced Computing (CDAC), under the government bodies, Department of Science and Technology (DST) and Department of Electronics and Information Technology (DeitY) were identified as the implementation agencies. A goal of this mission was to provide impetus to the *culture of supercomputing* in India. In particular, the objectives [2] were:

- to make India one of the world leaders in Supercomputing and to enhance India's capability in solving grand challenge problems of national and global relevance
- to empower Indian scientists and researchers with state-of-the-art supercomputing facilities and enable them to carry out cutting-edge research in their respective domains of expertise and work
- to minimize redundancies and duplication of efforts, and optimize investments in supercomputing
- to attain global competitiveness and ensure self-reliance in the strategic area of supercomputing technology

The goals were modest, but the effort required was huge, especially due to the fragmentation of efforts by individuals thus far. To achieve these goals, efforts were needed at multiple focal points, from the computing infrastructure all the way to the supercomputer users, covering hardware, compilers, middleware, operating systems, and applications. Thus, a full-stack wide-ranged holistic development was essential. Toward this multi-foci objective, NSM [2] envisaged several deliverables:

1) to install supercomputers of varying capacities (up to 3 PF) in different parts of the country
2) to connect these new and existing supercomputers by the National Knowledge Network (NKN) [3]. NKN offers a high-speed (1..10 Gbps) network connecting R&D labs and academic institutions in the country, and is provided by the central government.
3) to build supercomputers indigenously as a multi-phase project to push the country towards self-reliance.

To build this ecosystem, NSM was planned under four pillars: infrastructure, R&D, applications, and human resource development (HRD). Infrastructure pillar involved the development and installation of supercomputers. R&D provided financial support for projects toward exascale computing, such as architecture, compilers, and middleware. Applications pillar identified specific thrust areas of national importance, and supported applications in these areas such as medicine, weather, computational chemistry, geophysics, CFD, and computational physics. HRD pillar developed HPC-aware manpower. Together, the four pillars were supposed to help improve the supercomputing climate in the country.

The HRD pillar was headed by CDAC, responsible for generating enough human resources that can spearhead supercomputing activities in the country. This was planned via nodal centres who could conduct training activities to achieve the tentative goal of 20,000 HPC-aware experts and professionals. There were four nodal centers at four Indian Institutes of Technology (IIT): Goa, Kharagpur, Madras, and Palakkad. IITs are centrally-funded autonomous institutions in the country for engineering and technology education and research, and are legally termed as Institutes of National Importance [4]. From the coordinating agency – CDAC – A Senior Director handled this effort and a faculty member from each IIT was the coordinator of the nodal center. CDAC, which received funds from the government, provided funds, guidance, and support to each center to conduct the training activities. In particular, each nodal center received the following:

- A compute cluster, named *Param Vidya*, having a CPU and two A100 GPUs, along with the required hardware and software

- Funds to conduct the training programs
- A licensed version of an online meeting platform to conduct online events (*gotowebinar*)
- Guidance and monitoring from CDAC

The people expected to be involved under the NSM activities were everyone interested in HPC and AI. On the organization side, the original teams spanned primarily academicians and personnel from the government agency CDAC. Industry personnel were not part of these teams. This did not forbid industry from getting involved in various junctures, as the academicians are closely connected with their industry counterparts (see also Section IV). On the participation side, on the other hand, the trainees were envisaged to be from any category: students, faculty members, industry professionals, freelancers, entrepreneurs, etc. The knowledge was expected to be available for any learner. We also found regular participation from various industrial agencies in our training programs for both the established and cutting-edge areas.

By the time the training sessions began (July 2020) we were into the pandemic. Our earlier plans of hosting participants locally had to be abandoned. While this limitation put brakes to our efforts in one direction, it also worked as a blessing in disguise and allowed us to improve the reach of our sessions. It was also easy for us to utilize various online tools to record the sessions and disseminate the relevant resources. While personal touch of the offline sessions was missing, the logistics were getting simplified: we did not need to worry about participants' accommodation, food, and safety. Further, since the number of *seats* was not fixed, we did not require implementing participant categories and seat reservation, needed when the funding is from the government.

Overall, in the last four years, the nodal centers have conducted a range of training programs, from one day sessions to 3-month-long courses, internships, as well as themed workshops. Some of these training programs were conducted together, while others were organized individually by a nodal centre [5]. Along the way, we learned several aspects of HPC expertise in India, lessons in organizing, and challenges in effective conduct of the training sessions. This report details our experience and learning from organizing and conducting these training programs, which can help in planning future training activities under NSM or otherwise.

A. Intended Audience

This report should potentially be helpful for:

- Large-scale training organizers: NSM Nodal Centre worked tirelessly for over four years, conducting numerous training programs, with participation from experts from within the country and abroad. The lessons learned by us would be directly relevant when one plans to conduct such training sessions at a large scale for a variety of audience.
- Students: While a student may not directly observe the back-stage preparation, this writeup would provide glimpses into the planning that goes into arranging such sessions, and would provide avenues toward a more involved participation. It can not only help improve the learning, it can also help better the overall training program.
- Industry: Industry professionals and technologists can directly benefit with participation in such training programs. This can be two-fold: one by asking their team members attend specific sessions of interest, and two by being a trainer imparting practical education augmenting the theoretical studies from academia.

B. Outline

The sequel is organized as follows. Section II describes about the short-term and the long-term courses conducted. Section III describes about the various themed workshops (e.g., on CFD or Mechanics). Section IV describes about various workshops conducted in collaboration with the industry. Section V describes our model about NSM internships. Section VI lists various resources we created as part of this effort, which are likely to have lasting impressions. Section VII summarizes the experience of the convener from CDAC, Ashish Kuvelkar, in coordinating the activities across nodal centres. Section VIII provides a glimpse into the feedback provided by participants across the country listing how it helped them in their professional life. We conclude in Section IX and summarize our learning for the future training sessions.

II. Certificate Courses and Workshops

One of the most relevant events we conducted toward satisfying the mandate given to the nodal centres was online certificate courses and workshops. These were planned very similar to online courses conducted worldwide – with a registration phase, followed by scheduled online live lectures, and concluded with an online exam (conducted with the help of Mettl Secure Browser). Almost all our events provided e-certificates to the participants, based on varied criteria (attendance and exam score as decided by the instructors).

A. Execution and Observations

Table I lists various certificate courses we conducted. The course topics varied from introductory HPC course to GPU programming and Deep Learning. The duration of the courses also ranged from 6 days all the way to 4 months. Most of the courses were live, except for the GPU Programming course from February 2021 which was run with recorded videos. Lectures were delivered by experts from various academic institutions from India (such as IITs), industry (Intel, KLA, NVIDIA, Qualcomm), and a government agency (CDAC).

The course fee for Introduction to Deep Learning (marked as Paid*) was waived for nominated international delegates. While free courses attracted more registrations, the topic of the course mattered considerably, followed by the timing. For instance, a 4-month course in GPU Programming witnessed 1500+ registrations in February 2021. A similar course on CUDA Programming run over 40 days received less than one-third (472) registrations. Both the courses were online and free, and were offered by IIT Madras. We attribute this reduction

Event	Start Month	Duration	Instructors	# Registrants	Type
Scientific Computing on GPUs with OpenACC	November 2023	6 days	IIT Madras	200	Free
Mini-course on Concurrent Programming	July 2022	6 days	Mahindra Ecole Centrale, Intel Labs, IITs Ropar, Roorkee	1577	Free
CUDA Programming	May 2022	20 lectures over 40 days	IIT Madras	472	Free
GPU Programming with OpenACC	February 2022	10 days	IIT Madras	55	Paid
GPU Programming with CUDA	February 2022	50 lectures over 14 weeks	IIT Madras	472	Free
Introduction to Deep Learning	June 2021	7 weeks	IITs Goa, Madras, Delhi and NVIDIA	55	Paid*
Introduction to GPU Programming	June 2022	6 lectures over 6 weeks	IIT Madras, KLA	50	Paid
Introduction to Machine Learning	March 2021	2 months	IITs Goa, Madras	598	Paid
GPU Programming	February 2021	4 months	IIT Madras	1500+	Free
Introductory HPC Course	November 2020	3 months	IITs Goa, Kharagpur, Palakkad, Kanpur, IIT Tirupati, CDAC, NVIDIA	878	Paid

TABLE I
CERTIFICATE COURSES

to two reasons: (i) The first course started in February when students are undergoing their studies, while the second one started in May when many students have exams in India. (ii) Several interested learners (students, faculty, industry personnel) got their learning from the first instance of the course, and therefore, the second course witnessed relatively smaller footfall. We observed that the learners have a much larger demand for CUDA than other programming languages – even if the other programming language works with NVIDIA GPUs. This is also affected by the visibility and popularity of the language. For instance, OpenACC is relatively simpler to code, and works with NVIDIA GPUs as well as CPUs. But for a GPU Programming with CUDA course, the registrations were 472 in number, while for OpenACC, the number was 200 (both the courses were free). This was more stark considering that the CUDA course was run for a semester, while the OpenACC course was for only six days (and in our experience, typical learners prefer shorter courses).

Another observation, which now seems obvious, is that the more generic courses received more registrations compared to the specialized ones. Thus, an introductory HPC course, although a paid one and running for 3 months, had 878 registrations. On the other hand, paid GPU Programming courses hovered around 50–55 registrations. Similarly, introductory ML course had 598 paid registrants, while for DL, the number reduced 10-fold. Between free and paid versions, clearly the free versions were crowd-pullers. But we found that the difference can be considerably larger (4–10×). For instance, for OpenACC, the difference was 200 vs. 55, while for CUDA, it was 50 (paid) vs. 472 (free) vs. 1500+ (free) depending upon the duration and the time of the year.

Registration, attendance, and completion have witnessed a large variation – in a way similar to MOOCs. We have seen that the number of attendees is often 20–40% of the registration, and the number of participants completing the course (based on attendance and/or exam) is often less than 10% of the registration. That way, most of the online lectures have student participation similar to a university class. This is true for most of our courses – whether free or paid – where the registration is open to everyone. Both students and industry professionals start enthusiastically but often get busy with their academic assignments or company work. From this perspective, physical presence in offline training sessions may be useful, but the registration numbers may dwindle for offline training. However, interestingly, the attendance is far better (above 75%) and the completion rate is high (above 70%) when the participants are nominated by their guides. We had the opportunity to conduct such an online course on GPU Programming twice, with invitations sent to the guides from IITs and they nominating one or two students from their groups. We must note that there was no fee for these courses (participants are more serious if they have paid for a course). The attendance was very good and a large majority of them received e-certificates (based on attendance). We believe the same inference can be extended for nomination of industry professionals by their managers. On the other hand, restricting all the training programs to an invited audience may not be as per the government funding agency's expectations.

Table II lists various HPC workshops conducted. While these also had a training element similar to the certification courses, these were based on different topics chosen by the speakers rather than following a coherent build-up as needed in a course. Two of these were for 2 and 3 days, while the winter school was for a longer duration, permitting more involved and gradual treatment. We also observe that a 2-day workshop in March received 2× number of registrations than a 3-day workshop in July.

B. Lessons Learned

- Duration of a course does not predictably impact the number of learners interested. Even for longer-term courses, we have found a large number of registrations.
- Paid courses can have an order of magnitude fewer registrations than the free ones. However, the seriousness and focus of the participants is much higher with the paid courses – even with a nominal fee (our student fees were typically 500–1000 Indian rupees which is USD 6–12), almost resembling an elective at a university.

Event	Start Month	Duration	Instructors	# Registrants	Type
HPC Workshop	March 2021	2 days	from IIT Dharwad and IIT Palakkad	781	Free
Computer Architecture Winter School	December 2021	2 weeks	18 speakers from IITs, IISc and Industry	111	Free
HPC Workshop	July 2020	3 days	10 from IIT Madras	382	Free

TABLE II
HPC WORKSHOPS

- Making a certain amount of attendance mandatory towards certification can result in a win-win situation. The downside of this is that an admin needs to followup with queries on not receiving certificate, logging in via a different account, logging in for a few minutes to qualify for attendance, etc. We had to answer these queries for several months after the events. Therefore, it is helpful to have a staff maintaining these logs and answering the relevant queries.
- The month of the course, and its timing matter quite a lot for the participants. There are professionals who prefer to attend sessions in the evenings and weekends, and then there are students with exactly opposite preferences. We have found a majority in the first lot.
- Having a quiz at the end of the course is very helpful to provide feedback to the learners. It also helps in filtering the top learners. Giving different certificates to these top learners can help boost their motivation. If the quiz marks affect certification, a large amount of energy and preparation needs to be spent on the mock test and the final exam – even if the grading is automated.
- Providing a teaching or mentoring opportunity to the interested course toppers can create a positive feedback loop. Two of our one-week courses were conducted by toppers from the previous training sessions. Providing access to editable slides and codes was crucial for them toward a smooth delivery of the material.

III. THEMED WORKSHOPS

To cater to the interests of participants from specific domains, we organized themed workshops. Since these were based on a certain focus of interest to the corresponding organizers, these were more research-oriented, in contrast to the other training sessions. The organization flow was similar to the certificate courses with a registration phase, followed by the live sessions with attendance mandate etc. However, these themed workshops did not conclude with an exam. Another difference was that the themed workshops were organized by a group of colleagues who drafted the program by involving their colleagues and research scholars, whereas the certificate courses were planned by us (nodal centre coordinators) involving a few teachers.

A. Execution and Observations

Table III lists various themed workshops we conducted. The topics had a varied range: Computational Fluid Dynamics, Civil Engineering, Material and mechanics, AI and Biology. The speakers spanned Indian institutions, various foreign universities, as well as industry. The HPC CFD Workshop also had paper submissions (via easychair.org) – the only one among all our conducted events. The workshop on ML for Construction Automation experienced a large number of registrations – primarily due to its publicity; the free event was co-organized by IIT Madras and International Association of Automation and Robotics in Construction. The event was also conducted in the summer vacation.

One of the flagship events of the NSM Nodal Centre was HPC Research Week. The plan was to have one day of talks on one theme by various researchers in the country. The organizers chalked out six themes: Aerospace, Biology, Chemical, Computers, Mathematical, and Mechanics. Each day had 10-15 talks from Indian experts working in that theme. This was envisaged as a research showcase of various researchers working in these areas within the broad framework of HPC. We also thought that prospective Masters and PhD scholars would find this research glimpse useful to choose a particular group / university to apply to. There was a decent number of registrations (511), with participants interested in multiple themes: Aerospace (242), Biology (171), Chemical (157), Computers (380), Mathematical (262), Mechanics (196).

B. Lessons Learned

- Willingness and enthusiasm of the instructors is better for a themed event, since they are working in the relevant field and the event provides an opportunity to collaborate.
- If feasible, themed workshops should be conducted offline, where the relevant focused group comes together, interacts, and learns together.
- When having peer-reviewed workshops, the quality of submissions may not match the organizers' expectation. Hence, especially in the initial years, it helps to have a mix of accepted workshop papers and invited talks.
- The number of registrations does not reflect the amount of participation. The organizers' expectation of importance to a theme may mismatch with that of the participants.

IV. INDUSTRY WORKSHOPS

While industry was not part of the original team of organizers, it is an integral part of the overall HPC and AI ecosystem. Hence, we naturally collaborated with industry partners in conducting training sessions. Since IIT faculty members have industry connect, it was an easily-formed collaboration.

A. Execution and Observations

Table IV lists various industry workshops held over a period of four years. The rationale behind these workshops was twofold: (i) to provide industry-centric view of HPC and AI to the participants, and (ii) to provide the know-how of the

Event	Start Month	Duration	Instructors	# Regs	Type
HPC Research Week	November 2023	6 days	71 across Indian institutions	511	Free
HPC Symposium on AI and Biology	January 2022	4 days	Indian Institutes, NCSU, Google, Intel, UBrasilia	95	Paid
HPC Workshop on Materials and Mechanics	July 2021	3 days	IITs, IISc, UMich, TU Dresden, NCSU, NITT	649	Free
ML for Construction Automation	June 2021	5 days	IIT Madras, USharjah, UCambridge	1241	Free
HPC CFD Workshop	December 2020	3 days	6 from IITs	132	Free

TABLE III
THEMED WORKSHOPS

leading companies. The companies were Qualcomm, ARM, KLA, and Intel. All the industry sessions were free and ranged from 1 day to 6 spaced-out lectures. The Intel Workshop was conducted offline at IIT Madras and contained students from the institute. Two of the sessions stood out in the number of registrations. ARM Workshop received over 1000 registrations and KLA Workshop received over 28000 registrations. The ARM Workshop was conducted by a German engineer from ARM, while the KLA Workshop had five speakers from US, Europe, and India. The extremely large number of registrations are credited to the wide publicity the organizers gave to the event – using LinkedIn as well as via IIT Madras's Communications Office. The number of participants who simultaneously attended the KLA Workshop hit a maximum of 4500+, which was supported by the KLA infrastructure with Zoom and YouTube's live-streaming.

B. Lessons Learned

- A good publicity (institute LinkedIn, alumni forums, NPTEL mailing lists) can significantly change the number of registrations.
- Having foreign speakers seems to play a positive role in participants' interest.
- Associating an event with the students' curriculum at a university can fetch more participation. This does not mean the student interest is more.
- An industry workshop is a win-win situation for the company and the participants to know of the philosophy, product line, research focus of the company. Academia is poised well to coordinate neutrally between the two parties.
- NSM pillars (infrastructure, R&D, applications, and HRD) involved primarily academicians, apart from the government agency CDAC. It would be helpful if industry representatives are involved from the beginning in such activities.

V. INTERNSHIPS

A key aspect of the nodal centre was to be a facilitator. NSM Internships was a concept based on this aspect.

A. Execution and Observations

In the year 2022, the nodal center coordinated the HPC/AI summer internships between the interested faculty colleagues and the student applicants. The nodal center took care of the internship stipends and permitted the faculty-student pair to focus on the technical aspects of the work. This was viewed well by the participating faculty colleagues, who did not have any travel plans. For the 7 projects floated by colleagues across various IITs, we received 573 applications. The numbers suggest that the demand for internships is quite high. The applications were collected via a google form (including each applicant's CV). The most popular project was titled "Identifying biological circuits capable of specific functionalities using ML-based ODE solutions" with 185 applications, while the least popular was "Parallel implementation of IIF solver for stiff Advection-Reaction-Diffusion equations" with 24 applications. We made the project-wise applicant details available to the mentors, who selected their mentees based on the project-specific background required.

B. Lessons Learned

- Nodal centers can effectively facilitate the academic internships. It would be worthwhile to invest into industry internships in future.
- Since interns are not marked as *staff* or *speakers*, the nodal centres should receive the flexibility towards such payments.
- When done at the India level, the number of applications can be huge. One needs an easy filtering and ranking mechanism to deal with the scale. This is especially relevant when the ranking mechanism can be queried under RTI (right to information).

While the internship program was conducted successfully, we could not resume it in the following years as we were organizing other events. However, we note that there is a large potential in the nodal centres in facilitating such collaborations. The idea can also be extended to industry internships, as well as winternships and in-semester internships.

VI. RESOURCES

In the process of conducting this variety of events, the nodal centers also developed various resources. These resources range from a mailing list to a video repository. These resources were not originally planned, but were needed as we needed to communicate with the relevant participants.

A. Execution and Observations

In the early days, we emailed colleagues and collaborators to announce various events. As each event happened, we included the registration email ids in the future event publicity. At one stage, we had over 30,000 email ids to be emailed to. While these mass emails were scripted, the email publicity

Event	Start Month	Duration	Instructors	# Regs	Type
Intel Workshop on oneAPI	April 2024	1 day	Intel Leadership India	100+	Free
Intel Workshop on oneAPI	February 2023	1 day	Intel Leadership India	100+	Free
AI and HPC in Semiconductor Manufacturing	September 2021	5 days	KLA Leadership India, US	28000+	Free
ARM Workshop	March 2021	2 days	ARM, Germany	1042	Free
Qualcomm Lecture Series	December 2020	6 lectures over 6 weeks	Qualcomm Researchers India	365	Free

TABLE IV
INDUSTRY WORKSHOPS

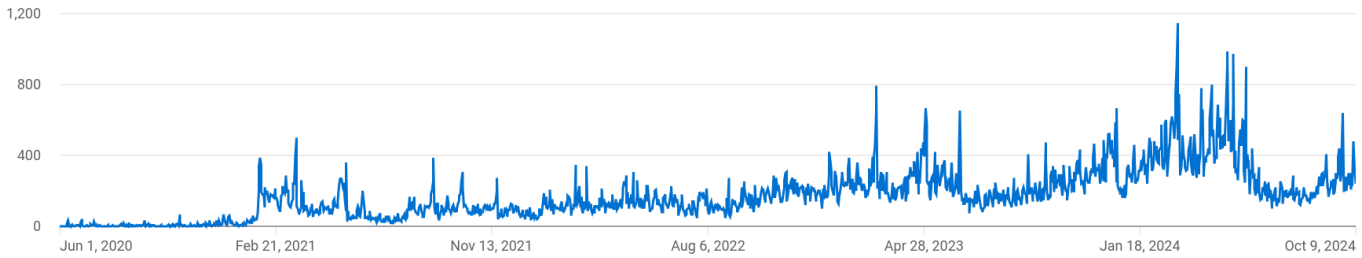

Fig. 1. Views of HPC Education YouTube Channel over four years

was time-consuming and we knew that we were spamming several of them. We did not have an unsubscribe mechanism.

To account for it, we built a googlegroup called nsmevents to announce various events and disseminate information to only the interested participants. It has 1278 members currently. This resource has been useful to disseminate HPC related events within the country – even if the event is not related to NSM. We would be able to utilize this resource very effectively for future purposes too.

Due to the online mode of events, we also had an easy opportunity to record most of the sessions. We made the free session videos (and on occasions, paid ones too) publicly available at a new HPCEducation YouTube channel [6]. This has become a popular resource to learn various HPC and AI related material. It has 27 playlists housing 3-to-55 videos each in them spanning topics from basic parallelism concepts to HIP Programming for AMD GPUs. The channel has over 4500 subscribers, averaging 3 subscribers per day. It has been viewed 0.3 million times. Figure 1 shows the views by visitors over the years. It clearly indicates that the channel has been increasing in viewership till April 2024 – time until when various events were conducted. The average view duration is 4:32 minutes, which indicates that viewers spend time to learn concepts. In fact, the top viewed video is 7-minute explanation of OpenMP `parallel for` with about 23,000 views. Top five visitor countries are India, US, Pakistan, Germany, and Egypt in that order.

Apart from the longer session videos, we have also curated short 2-to-5 minute videos explaining various concepts. Three such playlists are OpenMP, HPC, and CUDA. Each video explains one concept with animation and voice-over, taking a simple example, mentioning corner-cases, and explaining the variations. For instance, topics in the OpenMP playlist include Amdahl's Law, Parallel For, Shared and Private Variables, Critical Section, Atomics, Barrier, Single, Reduction, Environment Variables, etc. These playlists have become quite popular and permit learners to refer to them for quick reference. We have also built an Android app for viewing these concept videos.

B. Lessons Learned

- Planning the resources for a longer-term can outlive the funding, and can continue to benefit the community.
- A balance is needed between online and offline sessions. General, high-level, breadth-oriented, introductory, non-hands-on sessions can be conducted online. Focused, advanced, hands-on sessions are more effective in an offline setting.
- Policy decisions are needed to make the paid events' recordings available to public.
- While a large number of viewers are interested in short conceptual videos, large videos based on quality content continue to be useful for the community.

VII. EXPERIENCE OF THE CONVENER

The four nodal centers were guided, supported, and monitored by Ashish Kuvelkar, Senior Director, HPC Technologies, CDAC, India. We summarize below his experience as the convener of the NSM HR vertical.

The idea behind setting up of NSM Nodal Centres for training in HPC and AI at four IITs was to leverage upon the expert faculties available at these institutes not only in the areas of HPC but also in the various other disciplines of engineering and science. The challenge posed by the Covid pandemic was effectively negated by the coordinators at these nodal centers, who worked together to run courses in online mode. This helped in reaching out to participants on pan-India basis. It was heartening to note the feedback given by some of the participants of such courses, who mentioned that, they could get a chance to be tutored by faculties from IITs, which otherwise they could have only dreamt of. The recorded lecture sessions were converted by one of the nodal centers (IIT Palakkad) into an NPTEL course. This course was well

received by the student community and based on the positive feedback given by the participants, NPTEL has repeated the course multiple times.

The connects that the Nodal Centre coordinators had with their fellow faculty members within and outside their institutes across departments, as well as in the industry (some of whom were the alumni of these institutes), helped them engage in conducting the courses to suit needs of a variety of audience. The courses conducted not only witnessed captive participants from the respective institutes, they were also able to reach out to the participants in the region. The Technical Assistants associated with the faculty members were very helpful in providing support to run the online as well as in-person courses and workshops.

Thus, from C-DAC's and NSM's perspective, these Nodal Centers were very effective in meeting the goal of capacity building in HPC and AI and they will continue to do so, in the future.

VIII. Feedback

We also collected feedback from participants for knowing how various events have helped them.

A. Anecdotes

We list a few anecdotes verbatim below.

- PhD student from a private university: It is very helpful to my research and upcoming projects.
- PhD student from a public university: Being a computational chemist I didn't have great knowledge of computer science openmpi course help me to improve my parallel programming and also gave me insight of meaning of the programming that we use.
- Project associate from a public university who has attended five courses: Helped me to get a job a Project Scientist (HPC) in IITM Pune
- Employee from a private company: I have taken GPU based courses because of my work and interest more inclined towards that. Courses are really awesome
- Employee in a public sector company who has undergone four courses: NSM trainings have enabled me to switch jobs from software engineering to data science domain and have also empowered me to take the best architectural decesions while engaging with business stakeholders in the domains of Finance and Healthcare where i have worked with. All courtesy of the great programs that NSM team organizes time to time.
- Employee in a government organization: Opened new reserch avenues
- Faculty member in a private university: It helps me in train my students in HPC.
- Faculty member in a public university: The training has given insight about CUDA programming using C,C++. we have asked the students to go through the online lectures provided in the NSM Website

The above feedback clearly indicates that for a variety of learners, the NSM sessions have been useful for learning, advancing their careers, and sometimes, even in better teaching.

B. Lessons Learned

- Different participants attend such training sessions for various purposes – learning, certification, mandatory requirement, hands-on experience, etc. If the organizers know of these requirements, they may be able to conduct appropriate sessions customized for a group or tune the session towards such groups.
- The training sessions, originally envisaged primarily for Indian students, expanded the scope to cover a variety of learners (as indicated by the feedback givers). The organizers need to be watchful of such users and permit such a change of scope over time, in consultation with the funding and guiding agencies.
- A few participants do exceedingly well in multiple courses / workshops. There needs to be a mechanism to identify, appreciate, and utilize them effectively, as a positive feedback mechanism.

IX. Concluding Remarks and Future Plans

We narrated our experience with running the NSM Nodal Centre at IIT Madras. We conducted a variety of sessions on various topics related to HPC and AI, in collaboration with sister nodal centres and industry. Several colleagues helped us in this endeavour and we made many new connects as well. In the process, we developed multiple resources, which continue to help the learners even after the NSM activities are over. We also had our learning as narrated in this report, which can help the community when NSM 2.0 is implemented. We strongly believe that if the organizers make sincere attempt towards providing quality training to the community, it would help us grow together.

Acknowledgments

We thank Ashish Kuvelkar, C-DAC for his mentorship in running the nodal centres, and for providing his thoughts to be added to this writeup. We also thank coordinators of the sister nodal centres: Sandeep Chandran (IIT Palakkad), Sharad Sinha (IIT Goa), and Somnath Roy (IIT Kharagpur) for their help, inputs, and camaraderie. We acknowledge the help from our colleagues, collaborators, and students in conducting these training sessions. We are indebted to the reviewers for reviewing this work and help us improve the paper.

References

[1] M. of Science and Technology, "Announcement of National Supercomputing Mission," 2015.

[2] M. of Electronics and I. Technology, "National Supercomputing Mission," 2024. [Online]. Available: https://nsmindia.in/

[3] N. I. Centre, "National Knowledge Network," 2024. [Online]. Available: https://nkn.gov.in/

[4] Wikipedia, "Indian Institutes of Technology," 2024. [Online]. Available: https://en.wikipedia.org/wiki/Indian_Institutes_of_Technology

[5] R. Nasre, "NSM Nodal Centre at IIT Madras," 2024. [Online]. Available: http://nsm.iitm.ac.in

[6] N. N. Centre, "HPC Education YouTube Channel," 2024. [Online]. Available: https://www.youtube.com/c/HPCEducation