

MIC Tuning WG - 2015-2016

This file contains notes on tuning working group meetings during the fall of 2015 and all of 2016.

See the MIC Tuning WG file leading up to ISC15 [here](#).

See materials on data preconditioning [here](#).

All in the MIC community who wish to contribute to this effort are welcome. Contact [CJ Newburn](#) to be added to invites and communications.

Conventions within this file

- Most recent meetings first
- Participants in italics not confirmed, not in italics are confirmed

May 12 MIC Tuning WG mtg

Alessandro Rigazzi (Cray), Ariel, Cameron (RPI), JohnM, Fabio Baruffa (LRZ), JamesT, Alex Breuer, Kent, Srinath, Subhash Saini, CJ

Brief review of QMSprof, a tool for triaging workloads' suitability for manycore based on SDE and Sniper

- Ariel: Often overwhelmed with data that he got from standard tools. Found ITAC to be useful. Likes this overview, to understand different kinds of workloads and their characteristics. Has contact with some folks with RSNDA, through NESAP (JackD), and could do a public vs. internal comparison a couple months from now.
- John: Does this compute reuse distances? Can you zero in on particular data structures? No, not in what we offer here. It could be integrated.
- Public vs. Intel-internal version of Sniper
 - Public version doesn't support AVX512, no KNx models
 - Intel-internal version is available under RS-NDA
- Intel is looking to try this out with app engineers and technical computing engineers and in prep for customer dungeons
- Purpose is NOT perf projection, not useful for absolute numbers, just the shape of the curve and some indication of potential bottlenecks

[IXPUG @ISC 2016](#)

Target for BoF submissions

Monday the 16th. Selection about June 3rd.

Fabio and Luigi considering making a booth talk. Got early access to KNL. Expecting to have a draft in about a week. Encouraged to submit to BoF as well. Not much update since Ostrava.

Ariel not attending, so not available to submit what he's been working on.

AlexB not attending, but Alex Heinecke is presenting a paper on their Seisol/KNL work. They were encouraged to submit to the BoF.

Alessandro won't be going. Asked to reach out to folks at Cray to encourage a submission.

Apr 14 MIC Tuning WG mtg

Michael, AlexB, Momme, Andy, Antonio, Ariel, Leo, Damian, David, Emily, Fiona, Georg, Gilles, Guillaume, Heinrich, John, Kent, Kirti, Klaus-Dieter, Luigi, Dima, Ruchira, Srinath, Estella, Suryanarayanan Natarajan

- Upcoming event: IXPUG/ISC BoF and workshop
 - see ixpug.org for details
 - For workshop (Thu, times TBD): Extended abstract Apr 15, full paper Apr 29, can continue to revise until meeting date
 - For BoF (Wed June 22, 8:30-9:30): Use IXPUG ppt template for lightning talks
 - Special MIC Tuning sessions over the next two weeks
- News: KNL SDP and remote access
 - KNL SDP
 - Can pre-order SW development platform that includes a KNL. Basic system is <\$5K. Expected delivery time frame in the next month.
 - Some signed up for these already, or planning to, e.g. ICHEC, RTWH Aachen
 - Was announced in HPC Wire yesterday. This will be posted to ixpug.org shortly.
 - See the announcement on the ixpug.org [resources page](#) and an [announcement](#).
 - KNL remote access
 - **Submit NOW for Knights Landing (KNL) Access:** please complete the form by clicking [HERE](#). Compute time of 4 hours or less will be prioritized over longer time-period requests. Once access is approved, you will be notified in a separate email. Please be aware that KNL results need to be reviewed by Intel (James Reinders). That can happen as part of the submission to EasyChair for the ISC workshop and/or BoF.
- Process, timing: Sharing KNL results
 - Can be included in submission to EasyChair site
 - ICHEC, TACC, other IPCCs like Aachen, ZIH, Moscow State, ZIB
- Update: progress in various working groups
 - See links on IXPUG working groups page
 - MPI WG - Michael, Heinrich
 - Good kickoff, broad involvement, initial overview from Heinrich
 - Focused on advanced MPI
 - Building a pipeline of topics, including Jeff Hammond, persistence, taskification, KNL, non-blocking collectives
 - First Tuesday of every month

- Vectorization - Georg
 - Detailed walk throughs of ideas
 - Filling SIMD lanes
 - Contributions of data, implementations posted to github repo
 - Upcoming: SIMD template data layout library
 - Mondays twice a month
 - Good involvement of compiler engineers
- New memory types - Ruchira
 - Some presentations on MCDRAM
 - Upcoming: sub-NUMA clustering
 - Once a month

IXPUG EMEA Conference, Czech Republic: March 14th-18th, 2016

Video Links

- <http://livestream.com/hpvideostudio/ixpug/videos/115792658>
- <http://livestream.com/hpvideostudio/ixpug/videos/115887555>

Details

This is a reminder that we are looking forward for submissions for the IXPUG Workshop held in conjunction with the Intel Parallel Computing Center EMEA Meeting and Tutorials.

- Event information on IXPUG.org: <https://www.ixpug.org/events/ixpug-ostrava> (don't forget to sign-up your IXPUG membership)
- Abstract Submission (via EasyChair): <https://easychair.org/conferences/?conf=ipccixpug2016>
- Conference
Registration: <https://events.it4i.cz/indico/event/0/registration/register#/register>

Mar 10 MIC Tuning WG mtg

Andy Mallinson, Antonio Gomez, Ariel Biller, Ben Han, Benoit, David Mackay, Gilles, Jerome, Marcel Erhardt, Mukundhan Selvam, Kevin O'Leary, Nikola Tchipev, Hideki, Ruchira, Sergi, Thomas, Josh Tobin

Action items: Hideki to point to a presentation on compiler support for stride 2 that can be relevant for complex.

ICHEC submission review - Michael

- Complex split into 2 int arrays
- Increase in cache misses for compiler case

- Please bring the 6x compiler vs. intrinsic on MIC into the Vect WG
- Please double check that data is aligned

Hartree - Sergi

- Set perf goals, analyze progress toward them
- Analyze vectorization effectiveness with tools, look at remaining bottlenecks

Technical Univ Munchen - Nikola

Feb 11 MIC Tuning WG mtg

Attendees: ThomasS, RuchiraS, ZakharM, Kirill Rogozhin, JamesR, MikeLee, Damian, Luigil, FabioB, SrinathV, KevinO, HidekiS, Shailen Sobhee, Timothy Stitt, Sergi, Luke Mason, GeorgZ, Karthik, MeenaA, Subhash Saini, Ariel Biller, KentM, LeoB, EmilyM, Guilherme, CJ

- ISC and SC planning
 - BoF submissions
 - ISC, Feb 15
 - IXPUG on MIC Tuning
 - NERSC and new memory candidates were folded into this
 - Sandia/KNL on cluster experiences
 - SC, July 31
 - IXPUG on MIC Tuning
 - Possible: new memory types:
 - <https://drive.google.com/folderview?id=0B5oMsTn7u4anUnVUVWJmRmRwVWM&usp=sharing>
 - Tutorial submission candidates for ISC - due Feb 15
 - Cluster Checker - application, API that comes with Cluster Edition
 - VTune Memory Access Analysis - new features, e.g. memory structures
 - Python & DAAL tutorial - preview available now
 - Snapshot Tools (MPS, APS, SPS) - key metrics
 - SIMD Data Layout Template library - C++ AoS to SoA vectorization help
 - Vectorization / OpenMP 4.X / v17 Compiler features - simd and tasks
 - Vectorization Advisor (incl. AVX512) - new version
 - *Please offer additional candidates and feedback on desired content*
 - Workshops
 - SC - Feb 14
 - No planned submissions from IXPUG
 - ISC - submission accepted
 - Application Performance on Intel Xeon Phi – Being Prepared for KNL & Beyond (see [proposal text](#))
- [IPCC/IXPUG](#) Ostrava (Mar 14-18) prep report

- IPCC - Mon pm, Tue am
 - NDA sessions, invite only
 - Presenting latest results
 - Feb 22 deadline: initial draft of program created after that
- IXPUG - Tue pm - Fri
 - Submissions open for talks, tutorials; submit at ixpug sites
 - Feb 22 deadline: initial draft of program created after that
 - Submission at ixpug.org
- Working group report outs
 - Sign up at <https://www.ixpug.org/working-groups>
 - New memory types - Ruchira
 - High bandwidth memory (HBM, MCDRAM) and non-volatile memory (3D Cross Point)
 - Reviewed SC15 Tutorial content on how to enable MCDRAM (Ruchira to add link)
 - White paper in the works, will be posted to IDZ
 - Looking to gather more usage models
 - Vectorization - Georg
 - Next mtg Feb 15
 - Worked through a problematic case from IPCCs
 - Case study getting posted to github.com/ixpug, under vectorization WG
 - Life Sciences
 - Had a kickoff, gathering steam for broader inclusion

Discussions

- KNL access and results reporting
 - **KNL book** forthcoming, targeted for ISC
 - Varied access to HW; agreements for working with pre-production machines require vetting of results reporting, but approval is fairly likely
- Calendaring
 - James: Recommended method is to send both Outlook and Google

ISC Abstracts

Proposal for ISC'16 BOF Session

Gearing Up for Intel Xeon Phi Based Supercomputers

- Dr. Richard A. Gerber, National Energy Research Scientific Computing Center, Lawrence Berkeley National Lab. (NERSC)
- Dr. Chris J. Newburn, Intel Corp.
- Dr. Thomas Steinke, Zuse Institute Berlin (ZIB)

The BOF seeks to build community among those developing HPC applications for systems incorporating the Intel Xeon Phi many-core processor and builds on the successful BOFs at SC15, ISC2015, and SC14, which drew up to 150 participants each. Within months of ISC2016, large Xeon Phi Knight's Landing processors will make their appearance in large HPC systems. Taking advantage of the processor's full capabilities requires tuning and optimizing using programming techniques and tools targeted at a combination of CPUs and the Xeon Phi. Threading, vectorization, memory contiguity and alignment, and data locality are important for performance. The BOF is a place to share experiences and gain insights. The BOF will start with Lightning Talks that share key insights and best practices. It is followed by a moderated discussion among all those in attendance, which will include representatives from Intel and other software and tool providers. The BOF will close with an invitation to an ongoing discussion through the Intel Xeon Phi Users Group (IXPUG).

Topics include Many-core systems, Programming models, Performance evaluation and tuning, Application scalability, Status of performance analysis and optimization toolchain

Targeted audience: application developers on many-core platforms, experts in code optimization for Xeon Phi platform, **Estimated attendance:** 100

Proposal for ISC'16 BOF Session

The First Intel® Xeon Phi™ Processor is Landing: Early Experiences with an Integrated HPC Stack

Sandia - Si Hammond, Jim Ang/Intel - Rajesh Agny

In 2015, pre-production Intel® Xeon Phi™ processors (codenamed Knights Landing) and Intel® Omni-Path fabric products were integrated by Penguin Computing to create the Bowman advanced architecture testbed cluster at Sandia National Laboratories. In this BoF, we will hold an informal community discussion on the early experiences of porting applications to utilize with utilizing the key elements of the Bowman cluster such as the first generation of the Intel® Omni-Path system interconnect, updated "HPC software components stack" from Intel and the new architectural features of the Intel Xeon Phi™ processor including up to 72-cores, high bandwidth on-package memory, dual 512-bit vector units per core, four-way multi-threading and binary compatibility with Intel® Xeon® processors. This session will provide an overview of the Bowman design and discuss initial code porting and modernization efforts that are underway at Sandia including the introduction of scalable threaded and vectorized algorithms, the benefits of the "Intel® Scalable System Framework" as well as the injection of new methods to utilize multiple memory spaces. We show that these efforts enable performance improvements on both Intel® Xeon® and Xeon Phi™ processors. Sandia is building on their algorithms using the production Intel® Xeon Phi™ coprocessor (codenamed Knights Corner) and will show compelling performance and energy efficiency improvements over their existing deployments utilizing Knights Landing for range of applications in finite-element analysis, explicit solvers and hydrodynamics. Finally, Sandia will highlight their future plans for utilizing Intel® Xeon® and Xeon Phi™ processors and Omni-Path fabric in the forthcoming supercomputing deployments and how these learnings/efforts can contribute to wider Exascale code development readiness for Exascale.

ISC BoF Title: IXPUG Working Group: General Vectorization
Organizer: Georg Zitzlsberger; Speaker: TBD

The Intel Xeon Phi User's Group (IXPUG) is an independent users group that provides a forum for the free exchange of information and ideas that enhance the usability and efficiency of scientific applications running on large Xeon Phi-based High Performance Computing (HPC) systems. Last year (2015), different working groups (WG) have been created within this international forum [1]. One of those is the "General Vectorization" WG. Its goal is to address common problems with SIMD vectorization, providing solutions, and also discussing gaps in current technologies and standards. It is concentrated at a single location and also offers the chance to meet compiler engineers and key contacts to standards like OpenMP, C++, ..., around SIMD vectorization. It is open for everyone interested and the discussions are public [2]. For this BoF session we'll provide a report out of recent key findings in real world applications, such as patterns that are hard to vectorize with current compilers, solutions or workarounds, and the current status in discussing those. We highly encourage everyone in the HPC community to contribute, exchange knowledge and help improving the ecosystem via this "General Vectorization" WG.

Jan 14 MIC Tuning WG mtg

Attendees: O'Leary Kevin, Han Benedict, Ariel Biller(WIS/NESAP);, Bockhorst Heinrich, Cantalupo Christopher, Damian Alvarez, John Linford, Kent Mifeld, Kogut Jaroslaw, Kulakowski Krzysztof, Mallinson Andrew, Oertel Klaus-Dieter, Srinath V, Antonio Gomez, Stephan Ethier (Princeton U), Michael Lysaght (ICHEC), Gilles Civario (ICHEC), Yount Chuck, Judy Qiu, Chuck Yount, Georg Zitzlsberger, Ruchira Sasanka, Emily McCallum, Gerard Gorman, Benoit Scherrer, Ignacio Hernandez, Fabio Baruffa (LRZ), Luigi Iapichino (LRZ)

Logistics

- Times

Charter

- Overall coordination
- Report out on progress in WG, also in tools
- Prep for events
- Influencing tools features and bugs

Upcoming events

- IXPUG/IPCC event in Ostrava, Mar 14-18
 - Mon-Tue for IPCCs under NDA
 - Tue-Wed for IXPUG user forum, open to public
 - Thu-Fri for tutorials
- IXPUG workshop and probably BoF at ISC, June 19-23, Frankfurt

- IXPUG event at Argonne, Chicago, Sep
- IXPUG BoF at SC in SLC
- Possible Colfax event; they are exploring IXPUG co-sponsorship

Working groups

- Long running; ready to pause?
 - Data Preconditioning for Locality - taxonomy decorated
- New
 - Life Sci - by science and by optimization technique; opportunity for joint funding
 - New memory types - high-bandwidth memory, non-volatile (3D Cross Point); focused on analyses and tools for MCDRAM; usage model; prioritizing features; VTune white paper in progress; participation by architects (Ruchira)
 - General Vectorization - tough cases; systemic and strategic issues; language interfaces; participation by compiler architects; all relevant tools including Vector Advisor (Georg)
- Proposed
 - Floating point precision (Jim Demmel)
 - Nested parallelism (MichaelL)
 - Performance portability - Open standard alternatives (Gilles)
 - **New MPI features** - MPI 3, non-blocking collectives; experiences, feature requests (MichaelL); StephanE; JohnL; Srinath; Jaroslaw
 - Code modernization (John Linford) - addressed by other WG? SW carpentry of how to approach it
 - Surviving legacy code (Ariel Biller) - unit testing, Fortran
 - DevOps (Guilherme Amadio)
- Deferred
 - Vector packing and dynamic scheduling - may be spun off from gen vectorization

github.com/ixpug

- SrinathV
- Collaborators can share codes that are useful
- Benchmarks, stress tests, reproducers
 - vectorization stress tests from Argonne
- Community codes - HARMONIE (ICHEC)? mini-PARSEC (Ariel)?

Dec 7 MIC Tuning WG mtg

Chuck, MikeB, Andy, Klaus-Dieter, John, Georg, David, CJ

- Publicity
 - Opt in - send mail, give them an easy way to accept
 - Try to get all past attendees of IXPUG events and IPCCs signed up
- SC15 BoF review
 - ~180 attendees near the beginning
 - Lots of lightning talks: tales from the trenches and tools highlights
 - 6 formal talks out of 20 submissions - 30% acceptance rate

- Looking ahead
 - Univ of Ostrava/IT4I, Czech Republic: March/April, training with hands-on sessions, IXPUG user forum, day for NDA sessions for IPCCs
 - Colfax training event in the works
 - ISC16 - submissions for BoF and full-day workshop
 - Argonne National Lab - Sep 26-30'16 in Chicago; annual US mtg; use Theta
 - Topics of interest
 - Tools for debugging and tuning, automation; more on perf but also functionality (David)
 - Getting ready for KNL, MCDRAM/HBM, booting configs, SDE, memkind - try to identify experiences that we can get approved to share (Klaus-Dieter, Chuck, Georg)
 - OpenMP perf portability wrt GPUs and trade-offs nesting levels (John)
 - Georg/Mike - LAMMPS intrinsics vs. additional loop nests
 - John - Pearls 1 book: Jason Sewall, Phalanx/AWE
 - Klaus-Dieter - MD codes
 - ICHEC?
 - MPI vs. OpenMP trade-offs; NUMA effects
- Related efforts
 - HPC Developers Conference
 - ISTEP - Intel training effort
- Working groups
 - Starting up: LifeSci, New memory, general vectorization
 - We suggest to defer the others for now
 - Help to hear directly from customers, establish priorities, share insights

Nov 12 SC15 BoF PC mtg

Rakesh, Lisa, Georg, David, Kent, Lisa, Jack, Karthik, CJ

- Zhengji
- Overall agenda
 - Welcome (5) Richard
 - Going parallel with efficiency on the future Knights family (10) CJ
 - Tales from the Trenches - lightning talks (35) (CJ)
 - Open discussion - tuning focus (10) (Thomas)
 - Tools (20) KevinO, MarkO, MikeL, Antonio? (TBD)
 - Compare/contrast Allinea vs. Intel tools - Kevin can help (3)
 - Richard soliciting TotalView?, Tau, Vampir? **(3 ea, depending on novelty)**
 - Antonio will also try to do a compare/contrast (3)
 - Zhengji (4-5)
 - Intel tools - Kevin O'Leary/Mike Lee (3)
 - Open Discussion (5-10) (Thomas)
 - Wrap up (1) (David)

- Mixed
- First cut

Nov 11 WG mtg

Alvaro, Andy, Jason, Chuck, Roland, Luigi, Martin, Alexander, Georg, JimB, Klaus-Dieter, Rakesh, Kevin, Karthik, JohnE, CJ

Please post your foil updates (pdf) [here](#). Since I asked and none of you had objections to posting your work, I'll take that as an approval for sharing. For the sake of reviews, it's also best to re-upload to EasyChair.

Alvaro, RTM

- Make offload between Xeon and MIC clearer
- Could show impact of memory reduction on subdomain size and/or data xfer amount with OFFLOAD_REPORT
- Add stepwise perf improvements and correlate with PerfMon data from VTune

Alexander Moskovsky

- Perhaps pick an area to focus on
- HDF5?
- Will try to update tomorrow before flying to US, or perhaps later

Luigi

- Eff of 92% and 3.7x was an Advisor estimate; eff was actually 55-60% and speedup was 2.5x. So it turns out that there is potential for improvement, whereas that wasn't clear before.
- Georg: We have a reproducer which Luigi already shared, let's try to work that with Advisor engineers to root-cause the issues. Luigi to re-share the code.
- Updated for HSW, previously was IVB. Advisor estimates were more accurate for HSW than IVB. That was a surprise to Kevin.
- Kevin: could you add a brief slide about how Vector Advisor helped?
- Luigi: can offer feature requests; Intel guys can disposition those
- We got insights from an email thread about inlining

GROMACS

- Missing: improvement on MIC vs. unoptimized MIC, can't easily get that
- Karthik: Can check the FLOP counts with SDE:
<https://software.intel.com/en-us/articles/calculating-flop-using-intel-software-development-emulator-intel-sde>
- Can make some updates by

QE and PFARM

- Great insights in mapping work to a hetero platform, and using tools and characterization data to direct how to optimize and to correlate perf improvements with supporting data
- Streamlining this so that it can stand alone is a challenge, but a good investment.

Nov 4 WG mtg

Kengo Nakajima, MichaelK, Luigi, Nikolay, Andy, Antonio, Kevin O'Leary, JohnP, Chuck, Rakesh, JasonS, Martin, Antonio, Jason, Damian, CJ

Tools subsession

- Kevin can talk about VTune, Vector Advisor - many changes, things to try on AVX512 even without a KNL
- Include Allinea

Agenda

- Welcome (5)
- Going parallel with efficiency on the future Knights family (15) CJ
- Tales from the Trenches - lightning talks (50) CJ and customers
- Tools (10-15) KevinO, MarkO, MikeL, Antonio?
 - Compare/contrast Allinea vs. Intel tools - Kevin can help
 - Richard soliciting TotalView, **Tau**, **Vampir**
 - Antonio will also try to do a compare/contrast
- Open Discussion (35-40)

Nov 2 WG mtg

Florian Wende, Michael Lysaght, John Eblen, Antonio Gomez, Jim Browne, Estella, Damian, Heinrich, Andy, David, Chuck Yount, Jan, DimaP, Kent, Georg, JohnM, MartinP, MichaelB, Roland, ScottF, Florian, Karthik, JohnP, Hideki, Klaus-Dieter, CJ

BoF SC15 plan

- IXPUG highlights
 - IXPUG and HPC Dev Con talks on data preconditioning
 - Working group signups
 - Tools page
 - github - Srinath Vadlamani
 - Discussion forum
 - HPC Dev Con HiPC in Bangalore in December
- Lightning talks
- Tools highlights
- KNL HW updates and SW enabling for KNL and beyond
- Q&A on tuning and enabling issues

BoF lightning talk reviews

- Florian, Dynamic SIMD Scheduling
- MichaelL, Enabling the Quantum Collisions 'PFARM' code on Xeon Phi - Refactoring MPI communications for Symmetric Mode
- John/Roland, Vectorization of pair interactions in GROMACS
- John/Roland, Strategies to Optimize Offloading to the Xeon Phi
- Jim/Antonio, PerfExpert.- Workflow for Vectorization and Parallel Scalability
- Luigi, Vectorisation efficiency in a Gadget kernel: dealing with conditionals and data access
- Others present
- CJ
- Kent: Interested in perf, offload

Oct 5 WG mtg

Jason Sewall, Jan Zielinski, Stephane Ethier, Srinath Vadlamani, Glenn Brook, Tom Henderson, CJ

- Review of IXPUG 15 workshop
 - What went well, what could be improved
 - Scheduling conflict was problematic
 - Lightning talks were good, even more variety would be good. Specific techniques, learning what worked and what didn't.
 - Monday hands on - a couple more sessions would help. Work through work flow on a simple app or two. It's also helpful to consider the implications of extrapolating from a mini-app to the full app, particular wrt cache and MCDRAM capacity pressure.
 - Really useful, info packed sessions
 - New [tools page](#), off of ixpug.org/documents
 - Consider more 1:1, detailed technical interaction via a poster session (Peter Boyle)
 - New github site, <https://github.com/IXPUG> - Srinath Vadlamani
 - Purpose: collaborative space for community to iterate on creating examples, test codes
 - Current contents: autovec stress test
 - Response: compiler team interest in this, particular where cases include metadata on business impact. Can observe institutional affiliation of those who check in codes, who forks. Can add comments on codes. Could consider having a side doc that offers a table of pointers from biz issues to codes.
 - Possible additions: streams, for memory
 - Can add discussion pages
 - Increased visibility from posting codes that are submitted to Intel or other tools vendors as reproducers. Also include the work arounds and status of fixes.
- Working groups and proposals (see [list](#))

- Data preconditioning - presentation at IXPUG 15, also HPC Dev Con < SC15
- Life sciences
- New types of memory
- Any volunteers to create Google forms for signup?
- Planning for SC15 BoF
 - Wed Nov. 18, 5:30-7pm, [Paving the way for Performance on Intel® Knights Landing Processors and Beyond: Unleashing the Power of Next-Generation Many-Core Processors](#)
 - Suggested focus topics - MCDRAM methods
 - Format - include HW highlights
 - A qualitative notion of perf impact for KNL is interesting, even if not absolute

Sep 17 WG mtg

Attendees:

AndyM, Chuck, Guilherme, Sergi, Zhengji Zhao, Rakesh, Damian, Heinrich, Hideki, MikeB, Jack, Estella, SamW, CJ

Agenda

- Upcoming events
 - IXPUG in Berkeley
 - SC15 BOF?
- Review materials submitted to IXPUG'15

Aug 17'15 WG mtg

Attendees:

AndyM, CJ, Oliver Perks; [Wayne Gaudin](#); Luigil, LisaS, MichaelB, GeorgZ, John Michalakes, Zakhar, Hideki, ScottF, DimitryP, Rich, Heinrich Bockhorst, Klaus-Deiter, RichardG

Agenda

- Welcome, logistics
 - Monday 8am PDT seemed ok for those present; can't speak for those not here
- Report out on ISC15
 - Participation in BoF (60-65), workshop (65, turned 20 people away)
 - Materials posted at ixpug.org - including materials not presented
 - Great collaboration
 - joint reviews of submissions, communal feedback, connections with tools architects and experts
 - joint data collection and analysis (vectorization effectiveness, DRAM latency and bandwidth sensitivity), including by people not even there
 - collaboration in highlighting work to date and setting a research agenda - data preconditioning

- Great prework, very dynamic
- Upcoming events:
 - IXPUG, Sep 28-Oct 1 in Berkeley
 - SC15 BoF submitted
 - HPC Dev Con @ SC15 - Influence from IXPUG: OpenMP 4.x, data preconditioning sessions
 - HiPC15 industrial BoF submission in preparation
- Goals, relationship of topical working groups
 - Key areas of common interest to focus on
 - Share results and techniques
 - Building a community from diff institutions and experts and architects
 - Broad overview - this forum, and narrower focus - topical WG like on data preconditioning, DFT (JackD)
 - Prep for upcoming events: planning, discuss calls for submissions, review submissions, work logistics
- Relationship with IPCCs (Lisa Smith)
 - They are focused on modernizing codes that impact the community
 - That effort is a facet of the overall IXPUG effort
 - IPCCs are obviously directly sponsored by Intel
 - Monthly newsletter sent to ipcc@lotsofcores.com mailing list, which includes tips, tricks, case studies
- IXPUG-Sep15 prep
 - Format
 - Monday: Tutorials, e.g. MPI, Hybrid, OpenMP, Optimization case study (Richard Gerber)
 - Tuesday: Keynotes, general presentations, KNL update
 - Wednesday: Deeper dive on topics of greatest interest, experiences from the community (contact CJ Newburn)
 - Rolling opportunity to get feedback and present ongoing work
 - Thursday: library and tools (contact Richard Gerber/delegate); programming models, languages, runtimes (contact Scott French)
 - Separate tracks for part of this?
 - **Data reuse tool** - maybe frame the problem, gather requirements, present plans, gauge interest; reference RogueWave tool, ThreadSpotter, whose development has been frozen (follow up with Nicolai Pleszkun)
 - John: If we had a working and efficiency version of **Coarray Fortran**, we'd definitely be using it. Not performant yet. Georg: Compiler engineers are already working on this for the upcoming release.
 - No interest from those present in OCR or Legion
 - Thu/Fri: possible DFT workshop
 - Thu afternoon side sessions - maybe 3-5

- IPCC proposed
- Key focus areas
 - Vectorization effectiveness and memory tuning again?
 - Metrics: how you know if you have a problem and why it got better
 - Key challenges and their solutions
 - Submitted optimization case studies
 - Those remain the key 2 areas, no other suggestions for now
 - Special sessions
- Submissions
 - Putting instructions on ixpug web site for submission; easychair.org is being set up
 - Mechanism for indicating to broader community that you're interested and working on it, even if it's not done
 - Mix of talk lengths
 - lightning talks, 5 minutes lengthened to 7
 - longer talks, where we wish to share details completely enough to foster discussion
 - Luigi: short is ok, as long as there's enough time to discuss afterward; maybe topical sessions with discussion at the end across lightning talks in that session
 - Richard: long breaks, long lunches
 - Planning to solicit SC14 and ISC15 submitters
- Interest and registration
 - Register at ixpug.org