CSA 2012 Group Discussion Transcripts

May 31 and June 1, 2012
Westin Arlington Gateway; Arlington, Virginia
http://csa.cs.pitt.edu/

Compiled by Timothy J Parenti
7 July 2012

Discussion 1: A single view of today's infrastructure

[0:00] **JONES:** …individual groups and to try to come up with a unified vision that will put us ready to do our second breakout sessions, so I'll have each of the group leaders from the last sessions present their vision and their— we should all have this diagram.

[0:27] Some of them have been put into slides; it looks like some of them are still on the easel.

Group 1-A

[0:37] **JONES:** So, I was leading the discussion for Group A, and… unfortunately, the slides are advancing themselves.

[0:49] We started out by trying to take the existing tools that we're all familiar with and sort of the leading tools out there and categorize them in various ways and look at sort of where they sat along axes of various metrics.

[1:11] So this was our first attempt; we sort of categorized things, going from, for example data-centered, down to the board, down to the chip, inter-chip, even component-level… we sort of started with physical description, and eventually this became component-level.

[1:33] So this was our attempt to come up with those metrics, axes that we could use to create some kind of chart… and we decided that there were so many of these that we needed to really abandon that a little bit. And so what we ended up with is an x-axis that's sort of the level of integration across the whole system that we would be able to support in a tool, versus its scalability, and categorized where the tools would fit on this kind of chart.

[2:10] As a result, we had several different axes that we couldn't really include. The first one was a fidelity concept; flexibility; performance, which had some relationship to the y-axis; and interoperability, which had some relationship to the x-axis.

[2:27] So this is what we ended up with; we had the class of tools that were trying to accomplish a lot of the system altogether, although we were pushing and pulling a little bit to say, "Where do these fit?" and we had some more specific component solutions, and then we had individual component solutions, and various other locations.

[2:52] I won't list every single one, but we tried to place things according to where they would fit on this chart, and where we ended up with was we didn't really feel like this represented the gaps to these solutions. If you look at the chart, we sort of said, "Well, actually if our goal is to be in the top-right corner, from this chart it looks kind of like we're almost there," which is really not the case, and so they asked me to put "log scale" on the chart. [laughter]

[3:29] So here are the gaps that we identified: We talked about the level of interoperability, how many levels in the hierarchy that could be supported and what did this mean toward the integrability and the time-to-solution that something like this would provide. How large of a capacity of a system can we model in this metric? Were these validated tools? How could we test to make sure that these were giving us real, trustable results?

[3:59] What are the modularity elements of the tool? So this is more like thinking about "Are we coming from a standard or an API that everybody is comfortable with?" and so we can get more modular capability of the tools.

[4:16] Again, fidelity came back, and then there was a distinction between absolute fidelity — in other words, "Is a tool giving me exactly the same result as an implementation?" — versus relative fidelity, which is to say, given two different types of architectures, is the relationship between them reasonable, even though the true fidelity may be off by a constant or something of that nature.

[4:43] And then, the maintainability and, in particular, extension for new technologies as they develop. Then we kind of got onto a little tangent about benchmarks, but we decided that was its own breakout session, so we kind of left it at that.

[4:58] So that's the summary of what the A Group did. Let's see. Who was in charge of the B Group?

Group 1-B

[5:28] **LANGE:** Okay, so we have half our diagram; the other half got lost in the…

[5:34] But essentially, what we did was try to figure out sort of the taxonomy, diagram out sort of how things were designed and how they're coupled together with tools that exist now.

[5:43] What we ended up doing was sort of taking a tangent more of what we thought things would look like, but what seemed to be apparent to us is that we're sort of on the way there, so things are getting kind of close, we're just sort of slowly building out to get to where we think we want to be in the future.

[5:58] So the basic thing we did is … we came up with three sort of general components. You have, when you break up a simulation, you have sort of the target, the simulator itself, and the host that it is running on. The target here is essentially what you want to simulate, you know, a CPU or a system, a collection of systems… what it is you want to determine the behavior of.

[6:21] But then, inside the simulator itself, there are a bunch of different components that sort of act as different roles throughout the system, so the topic has sort of a meta-layer that sort of coordinates different simulation… I don't want to say "components," but different parts of the simulator together and glues them together and lets them communicate. An example would be SST.

[6:41] Below that, you have sort of more complex, system-level simulators such as gem5 or VMM, or SimNow, QEMU box, etc.; you have a huge list of these, which we don't have, unfortunately.

[6:53] But basically the idea here is that we couldn't really come up with a definition, but there was something about this where you had multiple components, you could actually run possibly an OS, something… what you would sort of determine as a system in real-time. It was basically determined as a full spectrum of solutions out there, and it's kind of hard to identify what they would be.

[7:13] The lowest systems here are different components such as DRAMSim, SimpleScalar, other examples that essentially looked at a single component and weren't really concerned with actually running full operating systems, full applications, or full workloads; they were just looking at a very small component for a given workload defined in some way.

[7:35] And then, sort of below this, we have this idea of metrics. So we have all of these simulation components and systems simulators and meta-simulators, all dump out some form of metrics. For right now, this is pretty much a complete mess. There's two approaches that seem to be predominant, that is, essentially aggregating everything into a single text file and parsing that when you have time, or if you're really lucky, you have sort of a statistical network that can look at the outputs as they're coming out and make determinations based on, essentially, analysis during run-time.

[8:10] Yeah, so I think that's pretty much it. So, Group C?

Group 1-C

[8:47] **CHO:** Yeah, just, I thought that the picture on the screen is too small; I decided to prepare some slides.

[8:55] So, I realize that this is a really hard task, given that we have such a nice variety in the background of the participants here. You know, there are so many different needs and so many different problems that we try to solve we the simulation methodology, even if we have so many tools and we have some ideas, it's very hard to converge on the single thing that we can nicely present. I believe everybody else also feels the same.

[9:24] Anyways, we tried to be organized here and can we intelligently list all we've got? So we decided to come up with a model; we named it the "3M" model. We tried to list the things, based on three different aspects: Module —

that's covered, module here is a component in a system; and Methods; and Metrics.

[9:47] So what are they? Methods: We have different techniques here, analytical modeling, statistical simulation, discrete event-based simulation, or some kind of function emulation; these… we could list the things based on the methods.

[10:01] Metrics: Traditional instructions per cycle type of metrics, or more a little bit high-level throughput, or clock-cycle time as hardware design, goal in area or power or energy…

[10:17] And modules, as I said, is a CPU core pipeline, is a GPGPU and accelerators, DRAM, disk, network, the whole cluster…

[10:26] Based on this "3M" model, we tried to list the things that we know of, and we gave up soon because it made our picture so messy. So we decided to focus on the well-known tool, just one tool, or one framework, which is gem5.

[10:45] So, okay, where is gem5, according to them? This is our picture, actually, and this picture clearly captures how hard this task is. [laughter] We tried to list the modules, and then this is the metrics.

[11:09] Then we had all kinds of discussions on fidelity… You know, we need to spend some time to agree on the language first. What is fidelity, right? What kind of fidelity are you talking about? Michael from IBM was saying, "When we talk about floating-point unit at bit-level accuracy… bit-level cycle-level accuracy, that's fidelity in that project," and, you know, you may mean something very different with the same term "fidelity" here. So we had a big discussion on that one.

[11:42] Anyways, so okay, let's focus on gem5 here. So this is the same thing, okay? We translated this, real-time, to this table. We've got modules and metrics, and gem5 is right here; it's a discrete event-based simulator framework and it nicely covers a good chunk of the modules that we thought of… felt interesting and important, and it can be expanded further down to even model large-scale clusters as well.

[12:17] And then we found that it's a pretty nice framework where we could actually plug in the existing tools, at least in the mind of one person — you know, everybody knows who that is — these things can be plugged in nicely, possibly by tweaking the existing interface; maybe it just takes one day, right?, to plug in things like, given some accelerator models or GPGPU kind of simulators, possibly could be plugged in.

[12:47] But then after that, after we found that this is a pretty nice framework, that we have these kind of things in our community, and then we realized that, you know, there may be some important things that are still missing or not, you know, the modes to report for other problems that they want to solve.

[13:04] Then we had this discussion of coverage. So we classified different platforms into datacenter and cluster-scale kind of things, mobile — these are two extremes — datacenter/cluster, mobile, and then everything in between, right?, servers and PCs. And we got a bunch of things, especially in the middle.

[13:27] We started to see some nice tools and frameworks and attempts to cover the datacenter and cluster-scale — examples were BigHouse from Michigan and SST, Sandia — and the mobile: gem5 could be also, and is used in the mobile platforms, to model mobile platforms. And so there are other tools that are somewhat geared toward developing software systems, not really cycle kind of tools, but there are tools coming up.

[13:59] So, compared with the servers and PCs, kind of single-node or few-nodes kind of situations, we have many, many tools, but we saw that we may need more things in the datacenter or in the mobile levels.

[14:14] Timescale: We talk about "cycles," nanoseconds, possibly some milliseconds, but as we go along in minutes and days if we look at large-scale or in some systems issues, solid-state drives, storage systems, networking… you know, there are, as we move along that line, we see some limitations there.

[14:38] Deployment: Do we want to run our simulation on parallel platforms or on host platforms, sequential host, and FPGA?

[14:47] And we had a lot of different discussions… you know, we had endless questions. So I didn't know when we'd have to end, but anyways: Interface and modules, we had that discussion in the context of gem5; fidelity, I already mentioned that; and some interesting thing here, automatic correlation between different models, different simulations, and different layers — for example, can we actually "synthesize" analytical model from more low-level physical data and models?

[15:19] So one example that we've got is, okay, now the AMD produces a processor, and then they find some gap between their real physical implementation and the simulator and they try to fix the simulator based on the physical, you know, the processor they actually got. That takes a lot of time and effort, so can we actually translate whatever we've got upward and downward that will improve the productivity?

[15:45] And then there was a discussion of "functional" versus "real data-driven." Here the functional means, "Okay, here's a packet at cycle 100," and the real data-driven is "here's your packet at cycle 100 and here's the data as well," so that, you know, your item will actually work on the build and so on.

Session 1 Conclusion

[16:07] **JONES:** Okay, so our task is to take these three pictures and to integrate them. Can you put your picture back up on the slide?

[16:17] We've got the Group B picture and I've got… here's our… could you hold this? Can you put up your slide? Your slide was very attractive.

[16:45] So, we have… so here are our pictures. So now we need one picture.

[16:55] Actually, I think there's a lot of commonality between some of what was presented. In particular, I see a lot of commonality between this and this. At least one of the axes is very similar, this module axis is very similar. And some of the axes we couldn't include, they included, so what's clear is that we're talking about a multidimensional chart that includes a lot of different components together to sort of look at the design space for where are we in terms of simulators.

[17:31] Maybe the best questions to ask are, "What are the problems that our current simulation environments don't solve?" and "Why is it that they don't solve these problems successfully?" and "How can we design something that can target and solve some of those problems?" Bruce?

[17:55] **CHILDERS:** Could the third question be in the next breakout… for the next breakout session?

[17:58] **JONES:** Right, so I'm trying to motivate what we're going to do for the next breakout. So, since Bruce basically tasked me with combining the pictures together, I'm declaring victory [laughter] that we have come up with several pictures that are somewhat compatible, and that we can now move forward with the second breakout session… unless we want to discuss this… yeah?

[18:26] **WOOD:** It strikes me that a lot of the discussion in the Group B was… So a lot of the discussion in Group B revolved around the interfaces between the different modules and different layers of meta-system, system, component, how the metrics were combined and such of that, and I think that focusing on just positioning where they fit in a hyperspace ignores that interoperability, which I think is really one of the things that the next breakout session really needs to draw upon.

[19:07] **JONES:** Yeah, so I totally agree. That was one of the things that we identified as big issues in our group as well. Other comments?

[19:23] **MOSSÉ:** So the other thing that we talked about also in Group B is…

[19:26] **MAN 1:** Where's Derek?

[19:27] **MOSSÉ:** Huh?

[19:28] **WOOD:** Derek's the best person to comment on this.

[19:29] **MOSSÉ:** Where's Derek?

[19:30] **WOOD:** Right there.

[19:31] **MOSSÉ:** Maybe Derek can talk about it.

[19:34] **CHIOU:** Okay, so one of… we separated their entire simulator into three layers, which is the target layer, which is what you're trying to simulate; the simulator layer, which is the simulator that is used to simulate that; and the host layer, which is whatever the simulator runs on. And so this is a…

[19:55] **MOSSÉ:** And I think that the other… what I was trying to say is the other two seem to… it seems that the other two are focusing on the middle of the page of that picture, and so there are commonalities, direct commonalities between those, and those two are, that look like there's commonalities with the center of the third picture over there.

[20:28] **MAN 2:** One other question I'm getting, particularly from the last diagram, Group C, is that, in trying to go for one picture, we may be forcing something that isn't necessarily the outcome that is wanted. The diagram which represented all the different features represents multiple axes, and with gem5 on it, it was nice and neat, but would be a complete mess if we included multiple other tools.

[21:00] But the interoperability point comes up if you take a diagram like that which represents the axes and features you want, and then ask yourself, "What does this particular one… what is it instantiated with?" If it's one tool with all components it has that are integrated, or is it the, let's say, all tools that are interoperable.

[21:22] The point being that the goal you'd strive for is a diagram that covers all the goals that you want out of a simulator or all the target systems you want out of a simulator, and it gives you something to shoot for with regards to what's been said.

[21:41] **JONES:** Okay. So we… other comments before we…?

[21:50] So we feel comfortable now moving to the next phase, where we're going to discuss, "How do we move forward, given what we have, in various different forms, and what we need?" which we know… we wouldn't all be here if we thought we had what we need, so we need to… let's go ahead and move forward into the second breakout session.

[22:18] Same groups, same locations.

Discussion 2: A single vision of the roadmap forward

[0:25] **LANGE:** Okay. So for this session we worked on especially developing roadmaps to try to see what all the groups came up with, so… okay.

Group 2-A

[0:38] **SAIDI:** Alright, so I'm going to talk about what we worked on. We didn't definitely get through all of our kind of models… so we didn't get through all of our models, but we did get through at least some of them.

[0:51] I think… we decided that the first thing we really need to do is come up with what's our needs going forward, some kind of, "What do we want to be able to do with these tools?" Do we want to simulate one core or a thousand cores? And that has some bearing on what we should be trying to… how we should be mutating them.

[1:16] Close on the list is some model of governance for how to make decisions with this body. While a huge group is great, at some point you have to come up with a consensus and how do you do that? And then to go forward, our kind of idea was, well, we should create a set of APIs that describe how various components talk to each other. This could be events, time, how communication is done in the system, how statistics are presented, how output is generated, how to migrate things like architectural state from one thing to another, and kind of one bigger question that's not as defined is the physics models, power and area and how those would be integrated and if there needs to be feedback loops there. In case you're doing power stuff, maybe you want to feedback your current thermal temperature into the system which makes you make decisions about what frequency you run at.

[2:25] And so, with step 3, creating these different APIs and the documentation that goes with them, the fourth step then would be to try these out on some limited scale to see if do we actually manage to solve the problems that we're aiming for.

[2:42] And if we do, then wrap existing tools in some API, and hopefully all the new additions and components will be in the new API and then we'd have interoperability between tools. Hopefully the APIs will be written in such a way that parallelization/VM execution/execution on FPGAs will be possible to abstract at the different API levels.

[3:12] There's a whole bit that we can't… kind of controlling some central repository where all the stuff is put and how to deal with that, how to validate it, how to do the testing. Ultimately, the only thing we have to say there is it really needs fulltime support staff to do, and don't expect grad students or other people to do it. [laughter]

[3:35] I think that's kind of it for Group A.

[3:41] **LANGE:** So, Group B?

Group 2-B

[3:47] **MOSSÉ:** This is group B. We're talking about the roadmap for the future.

[3:55] So this is the roadmap, not in chronological order, or at least not in strict chronological order. Define layers, it seems similar to the other group; define interfaces, seems similar to the other group; implement some basic components and runtimes and then look at what the inputs and configurations and metrics would be; do some functional validation of this sort of infrastructure; and then validation of different metrics. So the metrics are yet to be defined. So this is the big picture of…

[4:32] During the discussion… We came up with this big picture after talking a lot about the different issues. So we started with the 3M model from the last session which is similar to… is 3M that, uh…?

[4:49] **CHILDERS:** It's a trademark.

[4:50] **MOSSÉ:** Uh-oh. Shhhhh! [laughter] We started with some model, the "thrM" model, which is similar to the Target-Simulation-Host model that we had, and started talking about interfaces between the component modules

and the interface between the meta-modules and the components.

[5:13] We talked about the interface of the target, the simulation, how to define each one of these interfaces, but one of the consensus was that it needs to be simple and needs to be hidden at a certain level of abstraction. So it's similar to… somebody alluded to compilers. You don't know what's going on but you can still do a bunch of stuff. You don't know what register allocation is, you don't know what one component may be doing, but you can still do whole system simulation.

[5:45] So what are these interfaces and needs? Do we need to define if it is a discrete event or, if we are going to put visualization, at what level you're going to define these interfaces and what you start with. Defining the data format is something that is needed — the data format for the interface — but also the semantic of this data format and also the adaptability; you have to allow for future changes because nothing we do is static.

[6:16] What's the interface description language we are going to use? Is it a single language? Is it multiple languages? In our eventual elimination which languages we talked about. Is a port-style interface sufficient? So there are some issues with this that some people are working on, but this may be sufficient.

[6:36] Looking at what levels, layers of stacks, so if you look at network protocol stack, you have to talk about different layers — the hardware layer, the port layer, the protocol layer, the metric report layer, or the statistics gathering layer — all of this. In order to define all this, then we need different levels.

[6:59] **SAIDI:** So do we need to define something like the OSI seven-layer network stack for hardware simulation?

[7:06] **MAN 3:** Doesn't SystemC do that?

[7:09] **SAIDI:** I will readily confess, I don't know enough about…

[7:11] **MAN 4:** …well you don't need all seven layers…

[7:16] **MOSSÉ:** When we were talking about it, we said "What is it?" So are we just making some kind of parallel whether doing this layering would be a good thing or would not be a good thing.

[7:28] I think we came up with these in like maybe four minutes or three minutes, so we really don't know what the answer is, but it seemed a good idea at the time that if we define this at this layer and we can define in a uniform manner that everybody that likes it, or everybody abides by it, that it would make the creation of the simulation much easier, right? Now is this the right level or the right interfaces? We don't know. But something that would…

[8:03] We also talked about parallelization to address scalability. So do you parallelize, not only at the input level and the workload level, but also the simulation level? Do we need to define some rules of rollbacks and the modules of different fidelities and different accuracies so that you can abstract instead of using cycle-accurate simulators for one part, you just want functional simulator or something like that?

[8:31] We talked extensively for the time frame about repositories; how to deal with this unified infrastructure. Should we create a repository that you give your inputs and these inputs for this particular problem are stored in a database and then, when you run your algorithm, you get to run everybody else's algorithm on every input that is stored?

[9:02] So you can offer workloads and algorithms, or you can offer an algorithm and end up running against everybody else's workloads, and then you get all the outputs and outcomes and so you don't have to implement somebody else's algorithm, but you can validate against their algorithm. You get to see their algorithm because, if this is an open source project, then you get to see and you get to validate somebody else's algorithm if you want.

[9:30] One of the things that was brought up is this similar to scientific workloads, and also that we need some kind of taxonomy or ontology, but similar to the duodecimal system, that classifies the inputs and the algorithm so that you

can hone in and you can zone in to the particular section that you want to deal with.

[9:53] Also, it's not here, but we talked about this being also a good thing as a teaching tool or as a modification… so if I looked at an entire simulation infrastructure and I know that I want to do page-replacement algorithms in the cache I can just look at this and go into that particular piece of code instead of having to guess or to read the entire simulation infrastructure.

[10:18] Then we talked about validation and Group A also talked about some kind of review board that would be able to validate. So how do you do this validation? Regression tests, unit tests, valid and correct interfaces, or formal verification and model checking? We didn't really discuss this too much; we just said "This is needed", "What are we going to test against?", "How are we going to do this?", but this is something that is needed, as we said in the roadmap; validation is something that is in there.

[10:55] So that's it; what we talked about.

[10:59] **LANGE:** So, so far it seems like there's a large amount of agreement between groups, so… I see Group C is getting set up.

Group 2-C

[11:18] **CHO:** Okay, Group C. This is a partial roadmap for the infrastructure.

[11:27] The key word that we came up with is "interoperable simulator." That's the thing that we want to pursue here, so naturally we spent 99.9 percent of the time on the issues of interface. We want an "awesome" interface that captures some of the good characteristics I'm going to talk about a little bit in the next slide.

[11:55] We thought about performance, accuracy, traditional metrics that we'd like to somehow optimize for simulators, and support for fast, correct interpretation of results. As I said, 99.9 percent of the time was spent on "awesome interface."

[12:13] So what is "awesome interface"? We want a common-component interface that has the following characteristics. For example, intuitive, simple interface, because there will be a lot of people working on different components, and we want to have interface that works with — nicely and simply… intuitively works with — those people and the existing infrastructure we've already developed.

[12:45] Because our systems are evolving, we want our interface to be extensible. For example, the interface is simple, but sometimes people have crazy ideas — nice ideas — about how to get some more information from different parts of the chip — for example, CPU chip — to design a new thing. Does the caching controller have to look into the register file? Not likely, but we may have a crazy idea — great idea — about why that is a good idea, then how can we extend our interface to allow that?

[13:20] We also thought that this interface should enforce a structural decomposition of a system. And this interface should also include some of the features that's needed for global management of simulation tasks themselves. And also, this interface… to have this interface, we need to agree on the notion of "timing" inside of a simulation, and we kind of like the discrete event-based simulation framework that could promote this kind of interface.

[13:59] I also heard that the common infrastructure is going to be something important to get together these different pieces of infrastructures — statistics, for example, report generation scripts, visualization support — these things must be somehow defined so that we can effectively integrate different components in the simulation infrastructure.

[14:29] The awesome interface shouldn't hurt simulation performance. For example, when we define the interface, we need to consider whether such an interface will not hurt the possibility of parallelization and so on. And also we thought that this interface should allow different modeling strategies using different languages, for example.

[14:56] So, roadmap. Short-term, next one, two years: Define the glue logic, basically define the interface, such that, you know, we can develop glue logic and wrappers for components and so

on. Implement a logical notion of time. For the short-term, we also want a set of common libraries to define and develop for visualization, statistics, and so on. And for the medium-term we thought that practical parallel simulation strategies and techniques will be needed to be developed. And that's it.

Session 2 Conclusion

[15:27] **LANGE:** Okay, so it appears that pretty much everyone agrees that there's probably two main issues that need to be addressed here. So one is this idea of how do we have interfaces between different simulation environments, different simulation tools, and have them talk to each other? Then how do we actually make sure that they're talking to each other in the right way and what the result is coming out is accurate.

[15:52] I guess what struck me as either the good news or the bad news is that the issues don't seem to be technical. They seem to be more organizational; getting people to agree on how things are going to work how they're going to get together and have some sort of, just, committee or some group that gets together and decides these things. So I really don't think that's probably the purview of today, but something to think about… Go ahead…

[16:19] **CHTCHELKANOVA:** I would like to give a comment. So it seems like this group, your group, is really something similar to, let's say, nanoHUB. You guys know what nanoHUB is? It's a repository for simulation of raw materials.

[16:37] Why am I mentioning is they created a structure for content management called HUBzero. So they basically can take it and create whatever content you want. And they also, if you take a look at nanoHUB — and NSF spent a lot of money to help them create it — they have actually interfaces to help people to do simulations. They have access to exceed now, before it was terabytes, so you can actually run your job through this portal. Let's put it this way: It's a web-based portal. So something like that is coming to good benefit. That's number one observation.

[17:25] Number two observation is you may consider having something like a social organization. I would suggest that you take a look… Right now I am on detail in Geo Directorate, working with EarthCube. So if you Google "EarthCube" and take a look what we helped the community to create in terms of bits, there is a blog, there is a membership, there is all kinds of discussion groups, there is a place to keep materials.

[18:00] So we have really active working groups, and I think that if you find it interesting, there are some people who would like to take leadership roles. It's a very good example, I think, of how a community can self-organize. And if, as a community, people go forward form a software institute, which is the point of all these discussions, I think it would be really cool if, in your proposal, you go handle maybe all these components and you will show that you are very serious.

[18:37] **LOURI:** Form a CRI.

[18:39] **CHTCHELKANOVA:** Even CRI, yeah.

[18:40] **LOURI:** A community-oriented CRI.

[18:41] **CHTCHELKANOVA:** Yeah we should also come with CRI and… because a software institute is basically a CRI for NSF. So there are definitely some opportunities, but it's beside the point.

[18:59] So basically you need probably a social site… in something like content management, which would be, an example, HUBzero. And as we're promoting this, it's just an example, you know; decide on what you want.

[19:18] **LANGE:** Are there any other comments? Okay, so I think it's… time for dinner?

[19:26] **CHILDERS:** Yeah so dinner will be at 6:00 in F. Scott Fitzgerald Salon D. I think it's out this door and around the corner, on the other side of right here, but you can't go…

Discussion 3: Achieving the vision and next steps

[0:00] **CHILDERS:** …time to get back together.

[0:13] **MOSSÉ:** [singing] Hello! Is it me you're looking for? [laughter]

[0:19] **CHILDERS:** The reason he was doing that is because I don't do singing.

[0:25] So, this is the final group discussion session, and so we had three breakouts; the first one was trying to identify what's this role of central coordination, assuming this is kind of what we want. The second thing is incentivizing and rewarding participation and then trying to get people to buy in and take action and do something. And the third one, then, is to identify sort of the tangible next things to do.

[0:53] We'll go… just in order, since that probably makes sense, so… Group A, Alex?

Group 3-A: Role of centralized coordination

[1:02] **JONES:** So, yeah, once again, we're working in the Dark Ages with flip charts.

[1:07] But what we were able to do was sort of identify steps that the central coordination would need to take and some of the pieces of what those steps would look like moving forward.

[1:21] So Step 0 is sort of what we're doing now and what we're going to be doing moving forward, which is to establish the working group of people that are interested parties in digging into some of these problems and creating the infrastructure that we need in order to move forward with this goal.

[1:41] Step 1, this working group would then have a process for taking proposals for a couple of things concurrently. So, needs and requirements from people, what they're simulating… what they would like their simulation infrastructure to be, what to do; as well as proposals for what the standard layers and then, eventually, the standard interfaces that would come from these standard layers would look like.

[2:11] And this is imagined to be a very iterative process. We're thinking multiple workshops and discussion around initial proposals in order to make something the community is comfortable with.

[2:25] And our goals are to come up with an implementation of the interface; to create a centralized repository of information, potentially modeled on some of the existing repositories that some of the more popular tools work from; and to create an example workflow of how one can get an actual simulation accomplished within that… within their target, working on their host, and that sort of thing.

[2:59] We also need to determine a mandate, and create a reasonable mandate that is limited to something that can be accomplished in an effective manner, as well as get everybody speaking the same language and maybe create a glossary of terms of what we mean. We actually had a little bit of this going on in the meeting that… you know, we would say one thing, and multiple people would think different things related to that word. So these are important components of what we need to do.

[3:33] In terms of the needs documents, the number one item to take away from this is the workflows. So this, in some sense, encapsulates everything on this list, but determine what the workflows are that are of interest and that need to be handled, and potentially what would eventually need to be handled in the future, so we can think ahead in terms of making the interfaces and workflows extensible and in keeping up with current technologies.

[4:06] This would talk about things like, "What are the components we need to incorporate?", "What kind of experiments?", "What types of hosts would we be running on?", "What type of interfaces?", "What are the needs of those interfaces, dependent on the different workflows that we'll need to have?"

[4:23] So let me skip ahead to the next… before I talk about the bottom half of that page. This is our vision for what this repository would look like. And this is not going to be all created in version 1, but in essence, these are the things

that we think need to be included in order to make this successful.

[4:44] Obviously, we need to be able to retain all the different models that are of interest for the system. We also need to include tools, workloads, example workflows — potentially these are releases of simulators that do a potential widely-used kind of system that people are looking at a lot.

[5:11] There needs to be a skeleton or skeletons of infrastructures that can be built upon, that have sort of the same fundamental components, or the same model for parallelization, the same model for events, and these sorts of things, potentially that would work on the same kinds of hosts, et cetera.

[5:37] So there needs to be something here that can be built upon; we want to store the target configurations of interest. We want to have some way of doing regression tests for things that are checked into the repository to make sure that they meet a standard and are not extremely buggy code that doesn't run and it's just without documentation.

[6:00] And… let's see what I meant by this… Oh! This is a reproducibility element. We want to make sure that there is a way to keep either a data vault or some kind of signature of an experiment so that one, when reading a paper, can find the signature and reproduce those results, either by going to the repository and getting those results out, or by checking out a workflow that would actually create the results locally on your system.

[6:31] This makes it possible to do real, apples-to-apples comparisons in papers to verify results, to make sure that the community is really moving forward, as opposed to sort of not having a good way to do real comparisons.

[6:50] Finally, the plan in terms of stages, not necessarily in terms of a timeline, would be we would like to look at revision 0 of this repository as soon as possible. This might be collecting the most widely-used tools — the suggestions in the group were maybe some combination of gem5 and SST — to build from, even before we have all the interfaces defined, just so we can start to collect a centralized sort of location for a lot of the infrastructure that is out there now, and we can have a… get people thinking and looking towards this, before we get to what we'd call revision 1, which is an initial release of the core interface standard.

[7:41] This would be, again, maybe looking at gem5, SST, with using these APIs that we've created and potentially some other basic component tools that might be standard and widely-used at this point, and then revision 2-plus is to continue to add the features and capabilities of the repository that I mentioned a few moments ago.

[8:06] So, that's a summary. With that, I would… Anybody from Working Group A want to mention anything I missed? We got pretty much everything.

[8:18] Do we want to talk about this, or just go to B? Okay, so we'll go on to B.

Group 3-B: Incentivizing and rewarding participation

[8:27] **MOSSÉ:** So, we're from B. We were talking about incentivizing a single infrastructure, and this is the promise.

[8:37] We already talked about the pros *ad nauseam*: Intellectual growth, changes in architecture possible, components are reusable, augmentation of infrastructure. And the same way, we talked about having an initial infrastructure, so one of the things that the discussion revolved around was "What does it look like?"

[8:57] So this initial infrastructure would have to be funded, we thought about two or three years to develop the initial infrastructure, and then three to five years' maintenance, and if we could get… even though Bruce said, "Don't talk about funding," we assume funding is going to be forthcoming for this.

[9:15] And so the question is how to do it. One way of doing it is having different levels. Even though here's bottom-up, I'm going to talk about it top-down, because this is what I wanted to put in the slide to start with and since I

have the microphone, I'll do it anyway. [laughter]

[9:33] One… So… Okay, so I'll talk about it this way. [laughter] So you start from the bottom, and you build the little parts and then you integrate that into components, and then you go for the distribution. So the model here is this tells you… a group or an entity or somebody can say, "Here is my distribution. This contains that part, that part, that part, that part, and that part, and I'm going to aggregate that and I'm going to distribute this."

[10:04] And there could be several different distributions of this. So we can think about one single infrastructure, one distribution, or several distributions. So think about the Linux model.

[10:14] And these components would be you build a component and you put it in the repository; and you build a component and you don't want to put it in the repository, but you keep it on your website; and you build a component and you don't want to distribute it, and so it doesn't go into the distribution, but there are these components.

[10:29] So the first level is what you guys were talking about, something very simple that just shows what the layout of the land is, and then it can start being populated more and more by having that.

[10:46] Do you guys want to add something?

[10:48] **MAN 5:** I think the key point with the kernel would be to separate the simulation framework from the simulator. So the kernel is the thing that's got the components, the interfaces, and discrete-event model and such, and then the components are the things that actually do stuff. The kernel also has all the tracing debug, and I think that distinction might be a useful line to draw. Maybe not.

[11:11] **WOOD:** I like that. It's very important.

[11:13] **MOSSÉ:** So we talked a lot about incentivizing people, and we separated it into two parts: One is "incentivize use of other tools… other people's tools and code," and the other is incentivizing people to contribute and releasing their own.

[11:30] To start with that, the key issue is to provide value, right? So it has some… I'm only going to use somebody else's code if it is well-documented, if I have a good validation of it, if I can trust these results, if I have some kind of infrastructure that gives me benchmarks or inputs or types of applications or traces, gives me visualization tools, gives me statistics and debugging infrastructure.

[11:59] We could create "committees" that create these distributions. These words are loose, right? So we talked about these distributions committees, groups, companies, whatever you want. And individuals, or — "individuals" is also a loose term — groups, smaller groups create these components.

[12:16] So with that, if this means that some people contribute, some people maintain, you have a good infrastructure, so you're providing more value, and therefore people will use this infrastructure.

[12:29] On the other side, why would you release that? So we had to do a lot of education, right?, that people notice that if you do a little part, you get a whole, and if you don't do the part, you still get the whole, so what's the incentive? So there is the incentive, the altruistic incentive, or the community incentive, right?

[12:51] So you get also visibility. Yesterday or the day before, I used the word "notoriety" and people said, "No, notoriety is a bad thing; visibility is a good thing." Here, the good thing is visibility. You could also get notoriety by releasing your code, right? Oooh… [laughter]

[13:12] One of the things is if you have this infrastructure and you release your single part to the code, and if somebody complains about another part, you can blame somebody else. You only have to maintain your own code; you don't have to do maintenance of the other parts. But if you build the entire infrastructure, then you have to maintain the entire infrastructure.

[13:32] We could have the seal of approval, or what Bruce calls "Good Housekeeping… Stamp…"

[13:38] **CHILDERS:** Seal of Approval, Stamp of Approval…

[13:40] **MOSSÉ:** "…Stamp of Approval," of the certain distribution, or a certain module, or a certain component. You could claim that you're using this infrastructure, which is validated and is common, so you could have a better "broader impact" if you're going to contribute to this infrastructure.

[14:02] We talked about also contributions, and if you contribute your module immediately, then you don't get leverage of your own research, but we talked about some models where you can publish a paper or you can do something, and then six months later, you release this, or three months later or twelve months later. So then, there is incentive for you to keep it private for a little bit, but also incentive from the community for you to release it a little later on.

[14:30] And we talked about having a survey — and probably this group is a good group to start with — that says, "Which tools do you use? Why do you use these tools?" so that we can get a little sense of what would be a good incentive for people to use, and "What would this CSA need for you to be able to, or to want to use it?"

[14:55] And so, that's it. Group C?

Group 3-C: Next actions and steps to take as a community

[15:42] **CHIOU:** That wasn't supposed to be the first slide, but… [laughter]

[15:47] Alright, so I think we all are in agreement, basically. I saw components of what we discussed in all of the first two presentations, so… This is "Next actions and steps."

[15:57] So clearly the first action item's going to be a working group, simulator writers from academia and industry, people who have stakes in all of this. The first thing is the working group would define the scope — you know, network layers is one of them — and then examine what's out there. I think that there is a lot of simulator activity that's out there; a lot of it is represented in this room.

[16:23] And so once we examine what's out there, decide on some common, extensible interface and a common infrastructure that would be necessary for us to be able to interact with each other.

[16:36] Education and training, so a portal, I think both other groups talked about that as well, we can start this right now. We could have online tutorials, video lectures, slides, webinars, these kinds of things, and the repository. Coding standards is going to be very important. A lot of this stuff is going to be as we progress further on in, but the intention is that we would be able to send a new grad student to go and read the portal, watch the videos, and be able to get started in this important work.

[17:09] One thing that was pointed out is that a lot of the stuff that's out there is actually pretty reasonable, it's just not used the way it should be used, which is we should configure and appropriately set the right parameters, this kind of thing. And so perhaps even training for a scientific method, best practices for computer architects, and how they should use simulators is something we should consider.

[17:29] We also, at some point, are going to need implementers. One possibility is to write a CRI proposal and try to get funding for some sort of center, which would then essentially have engineers whose job is specifically to maintain this interface, enforce coding standards, this kind of thing.

[17:49] Or we could potentially hopefully… our hope is that we could convince industry that having these appropriate interfaces is worthwhile enough for them that maybe they might want to dedicate some engineering resources to this. And of course we have the always-present grad students who will be included in doing some of this, as well.

[18:08] **WOOD:** So you're excluding yourself from actually contributing? [laughter]

[18:14] **CHIOU:** Uh, no. I have not done any implementation work since I've become a faculty

member. So yeah, I'm sure anyone who wants to do it, absolutely, is going to be welcome, but these are the people we thought initially, and so, of course, the first action item is working group volunteers, and clearly people in this room are people who would like to be part of the working group. So any and all who are interested in volunteering: This is your opportunity.

[18:43] I don't know how we want to do this. Bruce, do you want to take emails for people who want to be in the working group, or…?

[18:48] **CHILDERS:** So, we have a mailing list. I can send out an email today or tomorrow to encourage people to start communicating and maybe solicit membership from that.

[19:02] **CHIOU:** Okay.

[19:04] **MAN 6:** So, to your question, it seems like there probably… or there might be more than one working group. So… some way of sort of organizing a way of volunteering for a particular type of activity?

[19:19] **WOOD:** So that's the job of the first working group is to figure out what other working groups… [laughter]

[19:25] **CHILDERS:** So, we'll do this, some of this is part of pulling all these materials together into a final report. Maybe we can identify some of the common themes that we've talked about. As we've been talking the whole time, I've been taking notes, and there are certain things that keep coming up over and over again, and those may be natural things to form two or three — three is probably too many — two working groups around and try to get people involved.

[19:51] I'll take the responsibility for trying to do some of that and then send out an email. That email won't be tomorrow, but at some point in the not-too-distant future, I will solicit this "Here's some common things we saw; are these good working groups?" and then we'll see if we can go from there.

Session 3 Conclusion

[20:12] **CHILDERS:** So, Ahmed, did you want to say a few things about NSF programs?

[20:18] **LOURI:** Oh! It's a good thing that Tracy is here. [laughter]

[20:25] **CHILDERS:** So the reason that I pointed to is, well, we're talking about what are the next steps, so maybe NSF can inform us about what are some of the opportunities where this would fit to help move some of this along.

[20:40] **LOURI:** Well, first of all, I'd like to thank you. Thank you very much for coming and doing this; this has been incredibly rewarding to me personally, but also I'm sure it is going to be very valuable to the community at large.

[20:54] As far as funding, I think, obviously, we talked about CRI Atlanta, and CRI has mechanisms for community infrastructure that can be sustained and that can be actually renewed, and we can actually try that, and there are mechanisms where we can have grad students plus professional programmers, if you will, as post-docs will have you to do that, so this is certainly one venue.

[21:23] The other one I'd like to you to focus on is perhaps if a set of you get together and maybe think about an expedition in computing. I can personally see that could be as an expedition because of this hope of the work and some of you said maybe it's not on the table, but that's what expedition is good for.

[21:46] The third one, I think, is maybe coming up down the pipe is the ACI. I brought a brochure here, just in case, so I'll give you the brochure number and maybe… Tracy is here, which is the Deputy Division Director of CCF, can talk a little bit more… but it is "NSF: 12-051".

[22:08] David has been involved in some whitepaper writing for this purpose. There's a select set of people, volunteers, that they volunteer to write whitepapers; it's actually ready, it's in CI's hands. I don't know… Did you release it to the public?

[22:23] **WOOD:** Uh, yeah. There's… I can send out the URL. There is a URL for it.

[22:29] **LOURI:** Right. I think CCC is also…

[22:30] **WOOD:** Yeah, it's on the CCC website.

[22:33] **LOURI:** It's on CCC website as well; you can read that. That is basically a set of scientists just like you that get together and write, I think it's called "Computer Architecture for the 21st Century," if I'm not mistaken.

[22:45] **WOOD:** Yeah, I think if you Google "21st Century Computer Architecture," it's the first hit.

[22:49] **LOURI:** Oh, yeah, yeah.

[22:50] **WOOD:** Let me just double-check that… Yeah. So if you Google "21st Century Computer Architecture," first hit's the CCC blog and the second hit is the PDF.

[23:01] **LOURI:** Okay. So that's the whitepaper. At NSF, we're taking that seriously, and people are working on that, and my understanding is for this ACI Advanced Computing Infrastructure that I would encourage you to read, it's on CI's website, if you want to look it up quickly, it's as I said, "NSF: 12-051." Yes.

[23:27] The document, it calls for a lot of things about… basically, it's HPC, harnessing parallelism, but also this type of modeling simulation tools, it calls for tools, and calls for all kinds of… This fits within that, so actually, I think there may be some solicitations coming. This is just a document, but there could be some solicitation coming up as soon as this summer, but watch out for summer plus fall for that, and hopefully you will respond to that, and take advantage of that as far as to solve immediate source of funding to continue this.

[24:06] I really like the idea of community building again, and continuing, as Alex says, maybe another set of workshop or workshops… Jean-Luc is here, so we are also thinking about closing the loop with a design implementation type of community that fits easy into this thing, so let us continue talking on that, and if you have any questions… Chris?

[24:35] **MINEO:** I would just add that the SI^2 program is a…

[24:39] **LOURI:** SI^2, yeah, the software institute yesterday, that we talked about, that's also a very good source of funding for that; thank you for reminding me of that. So that's another good source for this.

[24:53] If you have any… once you get going and you can get some volunteers for the working group and you want to get in touch with me or Almadena or Chris or whatever, so… you are invited to do so.

[25:08] Thank you very much. I'll give it back to you, Bruce.

[25:13] **CHILDERS:** So I want to open it up, if there's anything else that anybody wants to contribute to the workshop before we close out, any final parting shots, final parting thoughts on things that we should be doing, things we missed in the discussion, anything of this sort.

[25:33] Wow. That's hard to believe that there's nothing more to be said about simulation…

[25:40] **RODRIGUES:** So, I just want to say, from Sandia and the DoE perspective, this is of course a… simulation is something we consider very important and which we've invested a lot in, and hopefully we'll be continuing to do that.

[25:52] And so we have infrastructure at Sandia already sort of dedicated to things like regression testing, threshold software support, as well as actually working on simulators, and so we're very interested in partnering with others, supporting proposals going forward, and really trying to be a service for all of you… and kind of helping industry, because of course, we want them to succeed, helping academia, and sort of being a neutral party amongst all these groups to try to bring this forward, so we've very interested in working with you.

[26:23] DoE's actually holding a workshop kind of similar to this, more internally, sometime in August, so hopefully there will be some connections and some follow-ups from that as well.

[26:36] **CHILDERS:** Okay, as a final sort of administrative thing, if you have been making slides, taking notes, during the… I've also been taking notes and things. I don't think I have everybody's slides, particularly from the most recent breakout session, so if you made slides and things, please send that material to me.

[26:54] We're going to pull it together in some form or another and write a final report. We'll try to distribute a draft of the final report prior to releasing it publicly through the web site so that people can comment and add stuff to it.

[27:08] I know everybody also sent responses to the questionnaires. We would like to include those in the final report as well, so if you want the opportunity to go back and do an editing before we do this, feel free to do that and send me an updated file. Of course, if you're happy with the questionnaire as originally sent, no problem; we'll get it included.

[27:33] I think that's all I want to say. There's boxed lunches out there. There's a lot of boxed lunches, so if you want to take two boxed lunches if you're on the road today, feel free. Unfortunately, we were unable to change the ordering with the hotel. The hotel is relatively inflexible about things. But feel free.

[27:55] Thank you so much for coming. I actually think this was an extremely productive meeting and it's a great seed activity, but as a community, if we want to make this happen, we do have to take the next steps, and I'll try to facilitate some of that, but ultimately, it's up to the community to actually participate and do something.

[28:11] So thank you so much, and have safe travels home. [applause]