

Supplemental Material for “Visualizing a Moving Target: A Design Study on Task Parallel Programs in the Presence of Evolving Data and Concerns”

Katy Williams, Alex Bigelow, and Kate Isaacs

1 INTRODUCTION

We provide further detail on our collaborators involved with Atria’s design and the participants in our evaluation. We then provide more figures showing how Atria was used in practice and our evaluation sessions.

2 COLLABORATORS AND PARTICIPANTS

	Runtime Team	Performance Analysis Team	Gatekeeper	Front-line Analyst	Fellow Tool Builder	Evaluation participant	Prior familiarity	Influenced design	Deployment use
PM			•	•	•				D1, D7
P1		•	•	•	•		•	•	D2
P2		•		•	•		•	•	D1
R1	•		•	•	•		•	•	D1
R2	•		•	•	•		•	•	D5, D6
R3	•			•	•	•	•	•	D1
R4	•			•	•	•	•	•	D3
R5	•			•	•	•	•	•	
R6	•			•	•	•	•	•	
R7	•			•	•	•	•	•	
R8	•			•	•	•	•	•	
R9	•			•	•	•	•	•	
R10	•			•	•	•			
R11	•			•	•		•	•	D4, D6

Fig. 1. Team members involved in design or evaluation and their roles. Everyone had multiple roles, acted as a front-line analyst, and was a fellow tool builder.

The Phylanx project is divided into three teams: the Runtime Team, the Performance Analysis Team, and the Visualization Team. The Runtime Team is the largest with a principal investigator (PI), local project manager, several scientists, and student researchers. The Performance Analysis and Visualization teams each had a PI, student, and a scientist.

Each team resides at a different institution and time zone. In-person meetings occurred a few times a year, usually at the Runtime Team’s institution or the Supercomputing conference. There were weekly teleconferences with the team regarding all project goals (not only

- Katy Williams, Alex Bigelow, and Kate Isaacs are with the University of Arizona. E-mails: kawilliams,alexrbigelow,kisaacs@email.arizona.edu.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

visualization). Visualization-focused video conference meetings could be scheduled on request, usually when the Visualization Team would demonstrate new features (new design iterations).

Figure 1 lists all Runtime and Performance Analysis team members who were involved in Atria’s design process, used Atria deployments, or who were involved in evaluation sessions. The Phylanx program manager is also listed.

Most non-student team members have formal training in computer science or significant experience in scientific computing. Student team members were enrolled in either engineering or computer science graduate programs. Most team members, including the program and project managers, contributed code to the project, with the exception of newer students, such as R10, who were ramping up and new to the specific technologies of the project.

3 DEPLOYMENTS

Atria was deployed in several different ways, with design variations adapted to diverse, evolving needs.

Initially our data was output to the application’s standard out along with other debug output. Users had to extract Atria-specific data from the text dump manually. The first deployment (D1) was designed to serve data piped from the command line for generic *ad-hoc*, as-needed visualization. This deployment required users to clone and keep updated their own version of Atria. This version has been used extensively by R3, but also on occasion by P2 and R1. We also hosted our own version with fixed data as a demonstration site to communicate with the rest of the team. Additionally, PM originally used this version sporadically as a communication tool to describe how Phylanx worked to others outside the project.

A variant with a date selector was designed and deployed (D2) to augment a series of nightly regression tests that P1 had generated, to enable retroactive inspection, shown in Figure 2. Atria was then integrated into the nightly regression reports, allowing users to examine the previous night, any already collected regression data, or a comparison between two dates.

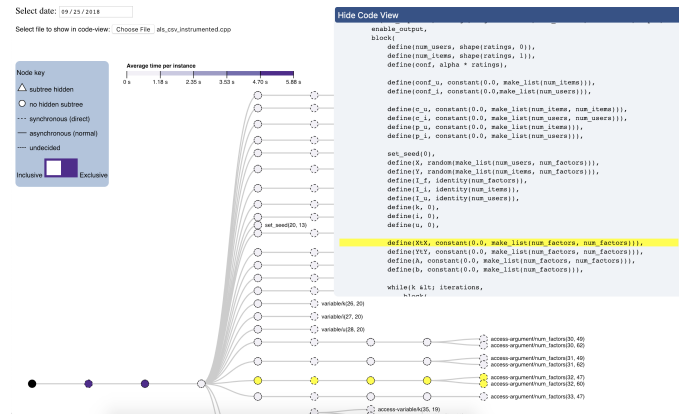


Fig. 2. A view of the deployment created for retroactive inspection of nightly regression tests (D2), including a calendar selector at the top-left.

To streamline the setup burden of D1, we created another deployment (D3) which we hosted at visualization team’s institution. This version allowed users to paste data from the standard out data dump into a text area, shown in Figure 3. R3 used this version when heavily examining one of the Phylanx applications, but as their work shifted from that application, is not using it now. Around the same time, R11 requested a self-contained, serverless HTML version of D1 with embedded data (D4).

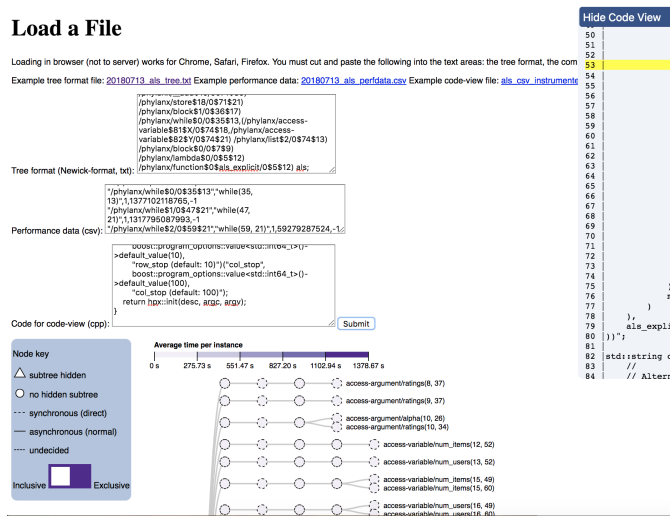


Fig. 3. A view of the deployment hosted at the authors’ institution, for cases where ports can not be opened publicly (D3), including text areas for pasting data.

At the request of the project manager, R2, we created a version with a specific dataset to work with Jupyter Notebook (D5) for demonstrations at the 2018 Supercomputing conference (SC18). This version contained only the main tree view, no code view or tooltip. In addition to R2, we also used this version, as well as D1, for demonstrations at SC18.

A persistent goal had been a more general Jupyter Notebook deployment. Having overcome technological considerations to create the SC18 demo, we iterated on the design, moving the code view and tool tips, resulting in (D6), shown in Figure 5 in the paper. We incorporated this version into a new demonstration notebook, both for the team members who give general demonstrations (R2, PM) and as an example for potential Phylanx users.

Finally, PM created a variant for use and demonstration at a secure facility (D7). This version is undergoing review and could not be shared with the authors.

4 ATRIA FIGURES FROM EVALUATIONS

At the start of each evaluation session, we introduced the participant to Atria by walking them through a tree of a factorial code snippet, shown in Figure 4. The familiarity of the code and the small size of the tree made it a tractable example of Atria’s features.

After the factorial demonstration, we asked the user to perform a series of tasks, as listed in Section 7.2.1. The single tree view is shown in Figure 5.

For task L1 (“Find a primitive that takes a lot of time”) the expected workflow was to visually scan for the nodes with the most saturated color. To answer the follow-up questions (“How long does it take without its children? With?”) the user needed to hover the mouse over the node to read the tooltip, which displayed the exclusive time (“without its children”) and the inclusive time (“with”).

In task L2 (“Find a primitive that is executed synchronously”) the user visually scanned the tree for a node with a dashed border. There are many instances of nodes with dashed borders in the tree so participants quickly found a correct node. Some participants had difficulty distinguishing between the close-dashed and spaced-dashed lines.

Task L3 (“Find a primitive that is executed asynchronously”) was similar to task L2. However, there was only one such node in this example, so some participants had difficulty finding this single node with a solid border.

Task L4 (“Find a primitive that is repeated in the code”) required more interaction with Atria besides simple visual scanning. Since the yellow lines that connect repeated primitives are only visible when the user hovers over the primitive, the users used their mouse to manually scan over nodes in the tree and to discover if the primitive was reused.

Once the user completed tasks on a single tree, we presented them with the comparison tree view, as show in Figure 6. We asked the user to perform a series of tasks that focused on comparing two datasets, as listed in Section 7.2.1.

For task C1 (“Which run was slower?”) the user visually scanned the entire tree to collect information about the hue and saturation of the nodes as a whole. A tree with more saturated purple nodes indicates that the net change in time increased from the first run to the second run, therefore the differences in the second code led to a faster run. Likewise, a tree with more orange nodes indicates that the first run of the code was faster than the second run.

We only presented task C1* to users with a performance analysis background. This question (“Why might it have been slower?”) required the user to investigate highly saturated nodes of the slower run’s color and reason about the primitive represented by the node.

Users accomplished task C2 (“Find a primitive that changed execution mode”) by visually scanning the tree for nodes with a pink border. Most users confirmed their answer by reading the tooltip.

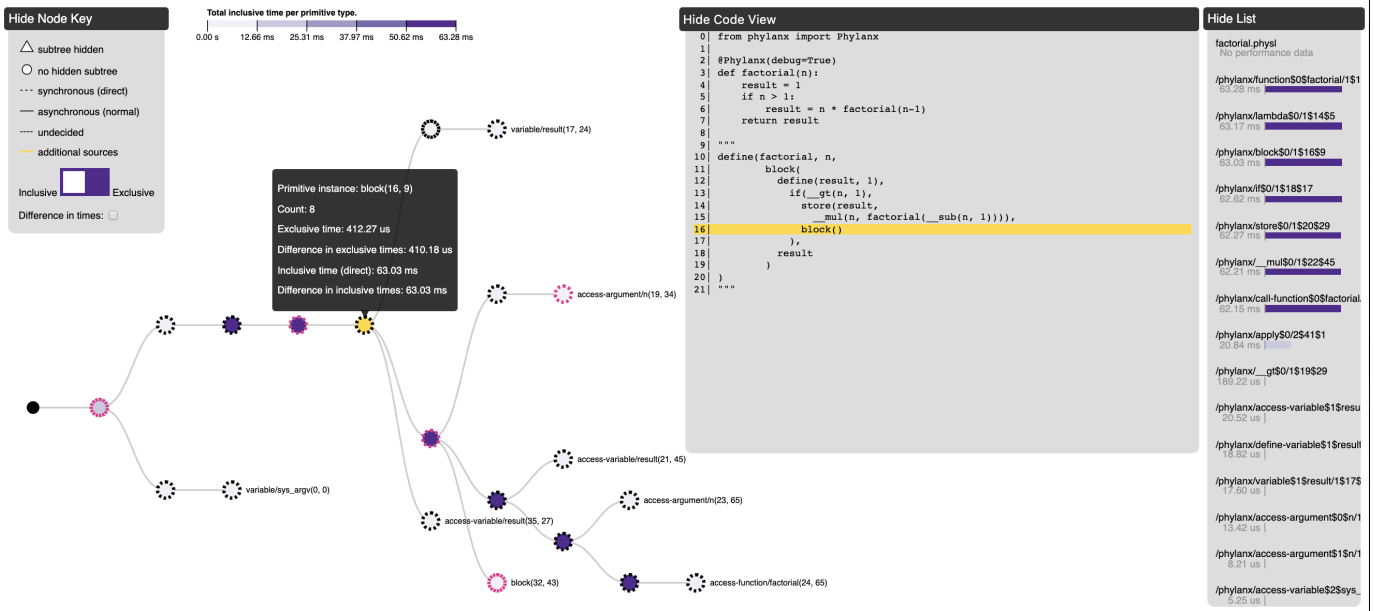


Fig. 4. A view of the demo tree we showed participant at the start of each Evaluation Session.

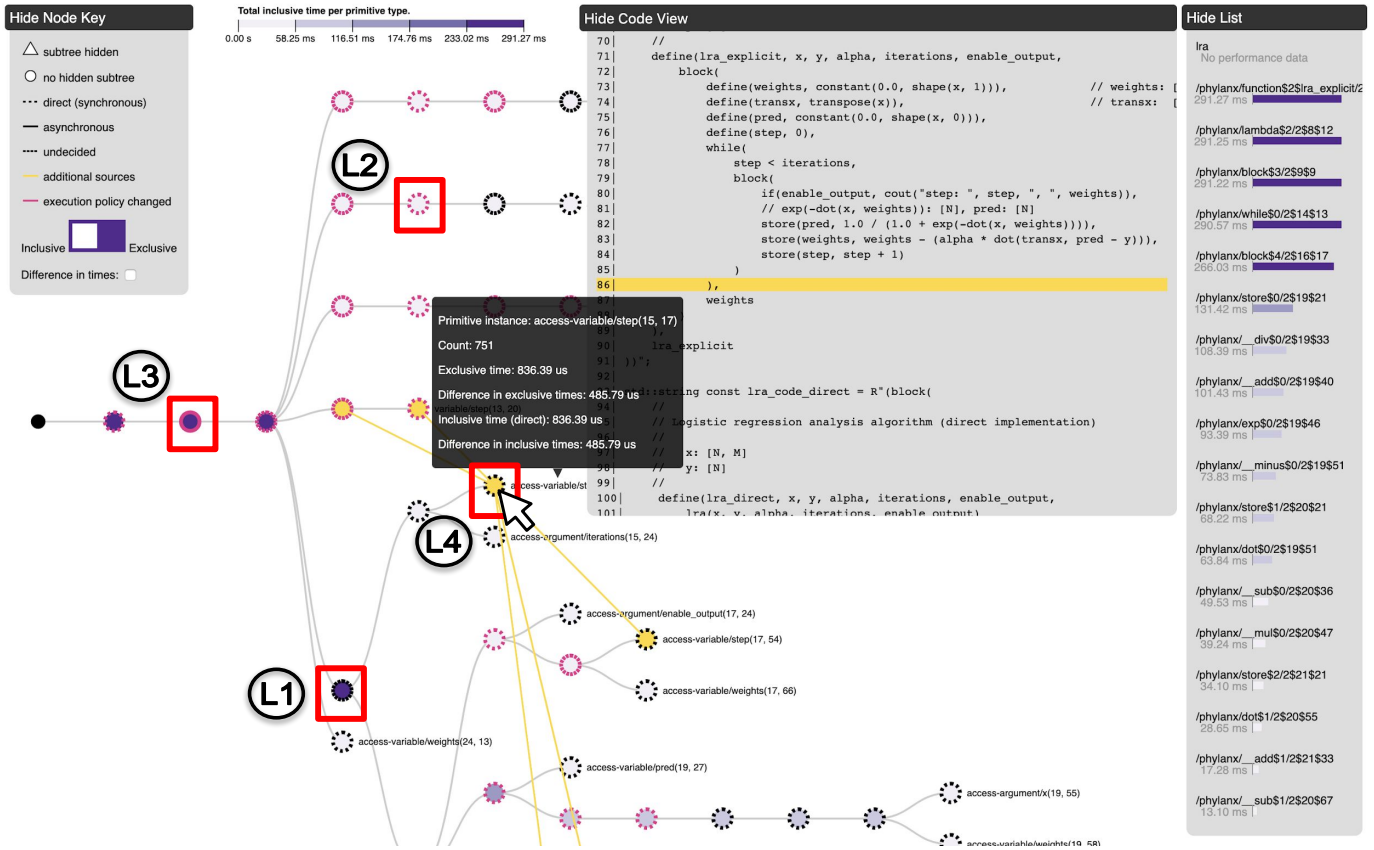


Fig. 5. A view of the single tree that participants used to perform our tasks during the evaluation sessions. The red boxes show a possible answer to each question presented in the task. Each box is labeled with its corresponding task, as they are enumerated in Section 7.2.1 (Tasks L1 - L4).

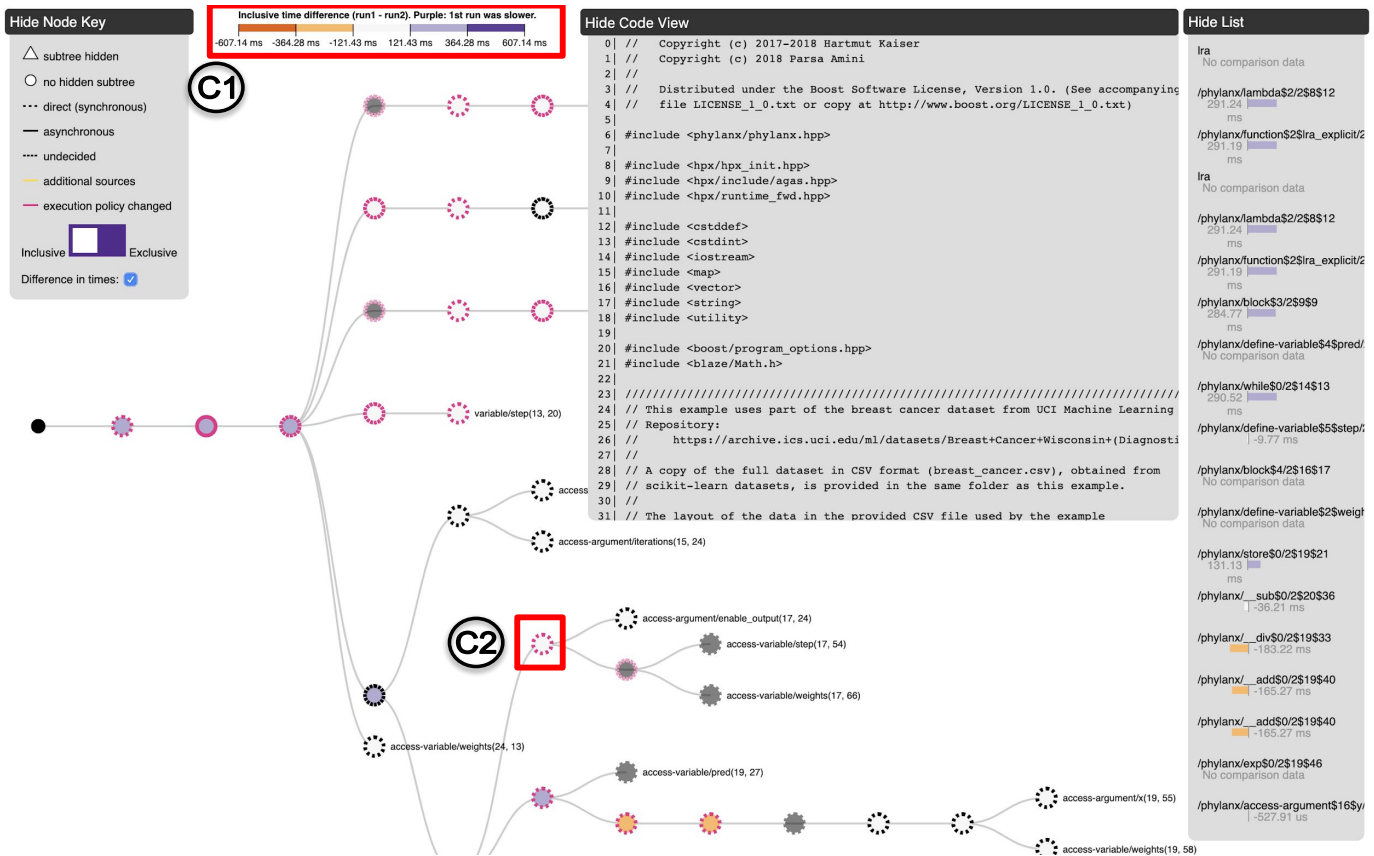


Fig. 6. A view of the comparison tree that participants used to perform our tasks during the evaluation sessions. The red boxes show a possible answer to each question presented in the task. Each box is labeled with its corresponding task, as they are enumerated in Section 7.2.1 (Tasks C1 and C2).