# Analyzing Users in Parallel Computing: A User-Oriented Study


Stephan Schlagkamp
Robotics Research Institute
TU Dortmund University
stephan.schlagkamp@udo.edu

Rafael Ferreira da Silva
Information Science Institute
University of Southern California
rafsilva@isi.edu

Johanna Renker, Gerhard Rinkenauer
Leibniz Research Centre for
Working Environment and Human Factors
{renker, rinkenauer}@ifado.de



***Abstract*—The performance evaluation of parallel computing environments is crucial for the design of parallel job schedulers, as well as policy definitions. The analysis of user behavior is fundamental to unveil individual behaviors and reactions to different system performances (e.g., scarce resources, low throughput, etc.). In this paper, we present an analysis of parallel computing users based on responses to the Questionnaire for User Habits of Computer Clusters (QUHCC). The survey is composed of 7 measures and 53 items, and was answered by 23 users of computer clusters at TU Dortmund University. We investigate several influences on working behavior, including the influence of slow responses on working times, strategies to cope with high contention and poor performance, user's experience, and user satisfaction. Analysis results reveal that user satisfaction is negatively correlated to the application slowdown; users tend to work after hours to improve their efficiency; informal agreements between users are established to coordinate executions and reduce the system load; and scientific experiments may include several clusters, thus the user submission behavior should be seen from a multi-dimensional perspective. We then compare and discuss the analysis results with conclusions obtained from statistical trace analysis to reveal unknown and hidden correlations and feedbacks between system characteristics and the subsequent job submissions. Our findings indicate that the user characteristics together with the historical information (traces) are crucial to build a concise understanding of feedback effects between the user satisfaction, their job submission behavior, and the system performance. Additionally, this paper also provides a first overview of which user reactions may be the most relevant for dynamic performance evaluation.**



***Keywords*-User behavior analysis, user satisfaction, parallel computing**


## I. Introduction

Modeling and simulation of high performance computing (HPC) and high throughput computing (HTC) systems have gained much attention over the past decades in terms of performance evaluation. Several studies have focused on the analysis and optimization of HPC and HTC systems by using workload traces collected at the infrastructure level [1]. These traces often include job characteristics and performance metrics (e.g., runtime, I/O, memory usage, etc.) of scientific applications that ran in such systems. However, none or few information about the underlying reasons of when and why jobs were submitted are provided. As a result, methods and techniques were developed to attempt to model the user behavior from these data [2]–[4]. Revealing user behavior in parallel computing is seen as an important aspect to rate the quality of computing infrastructures, model submission behavior, and therefore evaluate new scheduling techniques [1], [5], [6].

In this paper, we aim to enrich the knowledge about *human* user behavior in parallel computing by studying and characterizing user responses by means of a questionnaire. In particular, we apply the Questionnaire for User Habits of Computer Clusters (QUHCC) [7] to a group of 23 distinct users of two different computer clusters hosted at TU Dortmund University. We then identify, compare, and discuss aspects of the user submission behavior, revealed in this analysis, to conclusions obtained from statistical trace analyses in previous works [8], [9]. Our goal is to investigate the following questions: (1) which are the most common ways for users to adjust their working times and how does this influence satisfaction/waiting time satisfaction? (2) Does dissatisfaction lead to adjustments in user submission behavior? (3) Are experienced users using the systems more efficiently? (4) Is job cancellation an important aspect regarding subsequent user behavior and satisfaction? Additionally, we aim to unveil the most relevant aspects that should be focused on in future trace analysis, specially in studies targeting users' submission behavior modeling. The ultimate goal of this study is to provide insights to answer the following questions: (1) do the previous methods of modeling user submission behavior from traces suffice? (2) which aspects should be emphasized when analyzing workload traces?

Our findings indicate that (1) user satisfaction is negatively correlated to the application slowdown, however expert users are more likely to be satisfied; (2) users tend to work on weekends to cope with long completion times and to improve their efficiency, although they constantly apply strategies to exploit the given resources; (3) informal agreements between users are established to coordinate executions and reduce the system load, even though a usual consequence might hazard the system or other users, or violate policies; and (4) scientific experiments may run across several clusters optimized to different analyses (e.g., computation or visualization), thus the user analysis should identify and correlate the behaviors on each system.

Although we are analyzing local, smaller compute clusters (e.g., compared to NERSC[1]), our results are still of interest. The aspects of sharing resources among scientists and resource allocation are similar.

This paper is structured as follows. In Section II, we present an overview of the related work. Our methodology, scales of the QUHCC, and participants are described in Section III. Section IV presents the data analysis and discussion, where we first provide a general overview on the data basis of this study, and then a descriptive analysis of the answers provided by users, and finally a correlation analysis between different scales. Section V presents a discussion of how our findings imply changes in policies, increase user satisfaction, and increase focus in trace analyses. Section VI concludes this paper and presents future works.

## II. Related Work

Workload archives are widely used for research in distributed systems to validate assumptions, to model computational activity, and to evaluate methods in simulation or in experimental conditions. An overview of workload modeling and their usage in performance evaluation is provided by Feitelson [1]. Available workload archives [10]–[15] mainly capture information about job executions, but lack information on, e.g., job dependencies or reasons related to job submission decisions, which would lead to better understanding of the user behavior. Therefore, some efforts have been made to identify correlations in user submission behavior from traces. Researchers have tried to find interpretations of data recorded in workload traces and model human submission behavior [2], [5], [8], [9]. These analyses have found application in job submission behavior models to evaluate parallel job schedulers [4], and to improve user satisfaction through user-aware job scheduling [3]. Schwiegelshohn claims that understanding user behavior will support more convincing evaluations of parallel job schedulers and therefore increase the potential of practical usability [6]. Although some works have focused on characterizing user behavior from workload traces, they lack knowledge about user interactions external to the system (which are only indirectly present in the recorded data) that could significantly impact the user behavior.

## III. Methodology

Analyses presented in this work are based on answers to an online questionnaire among users from computer clusters at TU Dortmund University. The questionnaire was developed with the help of a focus group, i.e., the system administrators of the computer clusters at the Physics Department of TU Dortmund University.[2] They provided expertise knowledge on user habits, and user satisfaction and requirements, from which questions were derived for the questionnaire. The Questionnaire on User Habits in Computer Clusters (QUHCC) is composed of 53 questions divided into seven measures, of which some are divided into sub-measures. We focus our analysis on the measures of: (1) level of expertise, (2) waiting for jobs, (3) influence on working times, (4) usage of strategies, (5) general job adjustment, (6) user-centered job adjustment, (7) job cancellation, and (8) user satisfaction. In the following, we refer to these measures as *scales*[3]. In the following subsections, we briefly describe each of these scales, the participants, and the computational resources.

### A. Scales overview

The questions developed for the questionnaire target different HPC and HTC topics, which aim to provide broadly understanding of user behavior in parallel computing infrastructures. A detailed description on the process and interactions to develop QUHCC can be found in [7]. In questionnaire design, several *items*, i.e., user answers provided to statements or questions on the same topic are combined into a scale value [16]. Therefore, a scale aims to measure an averaged value for each participant on a certain topic. Questions in QUHCC have different answer categories, e.g., binary (yes/no), or multiple answers values to weight the answer (e.g., strongly disagree, disagree, somewhat disagree, somewhat agree, agree, or strongly agree). For the sake of clarity and because the possible answers are symmetrically, we map answers to values in $[0, 1]$ or $[1, 6]$, respectively. The answers to each item $\mathrm{s}_i$ of a scale $s(u)$ for a participant $u$ consisting of $n$ items are then normalized to a scale value in $[0, 1]$:

$$\mathrm{s}(u) = \frac{\sum_{i=1}^{n} s_i(u)}{n \cdot \max\{A\}} \in [0, 1], \tag{1}$$

where $A$ is the answer possible values for any item $s_i$. In QUHCC, $A := \{0, 1\}$ or $A := \{1, \ldots, 6\}$ holds for binary decisions and multiple decision, respectively. The scales of QUHCC considered in this study target (1) a descriptive analysis of each item to capture the specific user behavior, and (2) a map to values between zero and one to assess how strongly a user agrees to the considered scale. Below, we provide an overview of each of the scales:

***Level of Experience*** (LE): rates the user's experience with parallel computing. This is a generic scale, i.e., it is not bound to any computing paradigm (e.g., HPC, HTC, etc.). It consists of four items, and evaluates the user's confidence when using computational resources.

***Waiting for Jobs*** (WJ): identifies dependencies between two consecutive experiments. Job dependency is detected if the user requires the completion of a set of jobs in order to trigger the following analysis.

***Influence on Working Times*** (IWT): focuses on user reactions or strategies to seek for the system lower usage in terms of adjusting daily working time patterns. It consists of five items, and aims to identify alternative working times users typically adopt to circumvent poor system performance.

[1] http://www.nersc.gov

[2] http://www.phido.physik.uni-dortmund.de

[3] Since this is the first time users are analyzed with QUHCC, scales are not yet validated.

***Usage of Strategies*** (US): identifies strategies users commonly use to obtain faster results and better resource allocation. Such strategies include submission of bags of tasks, job prioritization, moving computation to another resource, or informal agreements with other users.

***General Job Adjustments*** (GJA): focuses on adjustments of job requirements to use the parallel computing infrastructure more efficiently. This scale consists of six items.

***User-centered Job Adjustments*** (UJA): focuses on *self-centered* adjustments of job requirements. Typically, these job adjustments are not required from a global point of view, they target specific jobs tuning so that resources may be allocated/consumed in the user's benefit. This scale is composed of five items, and aims to identify the most common adapting strategies users perform to burst their executions.

***Job Cancellation*** (JC): measures the percentage of submitted jobs that are voluntarily cancelled by the user. Since this scale is defined by a number, we have also asked for the main reasons leading to this action. We classified them as *Useless Result*, *Configuration Error*, *Programming Error*, and *Using other Resource*.

***User Satisfaction*** (USF): measures how participants perceive system performance in terms of waiting times. This scale is composed of four items, and evaluates the impact of job response times in the user's expectance of the system performance.

### B. Participants and computational resources

The questionnaire was applied to users from both HPC and HTC environments at TU Dortmund University. In total, 23 users from different science domains (including mathematics, statistics, chemistry, physics, and computer science) took part in this study. Users mainly use two computer clusters from the the TU Dortmund University: LiDO[4], a cluster composed of 432 nodes and 3,584 CPU cores; and a cluster at the Physics Department composed of 114 nodes and 912 CPU cores. While LiDO is mostly used to compute tightly-coupled jobs, jobs from the Physics' cluster are mostly embarrassingly parallel jobs submitted as bags of tasks. Additionally, two participants work at the Statistics Department using a cluster composed of 85 CPU cores, which are unequally distributed among 11 nodes[5]. Since all surveyed users belong to different science domains and use both computation paradigms, we argue that our findings are independent of a particular scientific domain or parallel computing paradigms, but represent a general picture of working with distributed computing resources.

## IV. Data Analysis and Discussion

In this section, we analyze and discuss the data obtained from the application of QUHCC. First, we present an overall analysis of the answers provided per scale. Then, we use this aggregated data to conduct a descriptive statistics analysis of the items within scales, and to seek for correlations among different scales. Finally, we discuss the importance of our findings, and limitations and challenges of conducting workload trace analysis.

### A. Overview of the collected data

Fig. 1 shows the distribution of scale values for the user's answers provided to QUHCC. Each scale is computed according to Equation 1, and represents the weighted average of the items within a scale. Below, we present the general interpretations of the analysis of the distributions. In the following subsections, we will expand and discuss each of these findings. We present the seven scales according to the order from the scales overview in Section III-A.

1) Most of the users have significant experience with parallel computing environments with an average scale value of $\mu = 0.82$, and standard deviation $\sigma = 0.15$ (Fig. 1a);
2) Users tend to seek for alternative ways to circumvent poor system performance. There is no major action that prevails the behavior of a single user ($\mu = 0.50$ and $\sigma = 0.26$, Fig. 1b);
3) Most of the users ($\mu = 0.46$, $\sigma = 0.25$) use at least one strategy to improve jobs execution (Fig. 1c). Note that these strategies include both computational methods and human decisions/interactions;
4) Users adjust the job requirements according to the available capacities, and are aware of the efficient use of the computing infrastructure ($\mu = 0.67$ and $\sigma = 0.09$, Fig. 1d);
5) Users frequently use techniques to improve their executions ($\mu = 0.60$, $\sigma = 0.14$). Nevertheless, we experience outliers of highest (0.8–1.0) or lowest (0.0–0.2) values (Fig. 1e);
6) Most of the users cancel less than 20% of their jobs ($\mu = 0.12$, $\sigma = 0.15$). We experience only one outlier, a single user mentioned that about 75% of the jobs are canceled;
7) Users are often satisfied with the overall turnaround time of an experiment ($\mu = 0.64$), however the standard deviation of $\sigma = 0.28$ implies that some users also experience longer waiting times than expected (Fig. 1g).

To determine whether users often wait for job completion in order to proceed their analyses, we defined jobs length into three categories: (1) *small*–job runtime up to four hours; (2) *medium*–job runtime up to three days; and (3) *large*–longer than three days (we also asked the participants for how long they are willing to wait for results). Fig. 2 shows the distribution of answers ranging from strongly disagree to strongly agree. For job categories small and medium, the median answer is *somewhat agree*, while for large jobs is *disagree*. This result suggests that job lengths ranging from several hours up to three days have more influence on consecutive working. According to the provided answers, results longer jobs are not as crucial for consecutive working. In a general setup, schedulers and allocation strategies handle users and jobs equally, i.e., scheduling properties are limited

[4] http://lidong.itmc.tu-dortmund.de/ldw/status.html
[5] https://www.statistik.tu-dortmund.de/rechnerdoku/tutorial/Computecluster.html

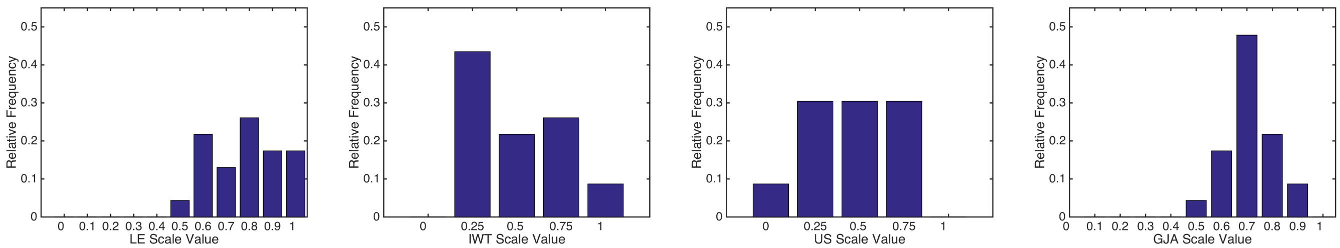


(a) Level of experience (LE) (b) Influence on working times (IWT) (c) Usage of strategies (US) (d) General job adjustments (GJA)

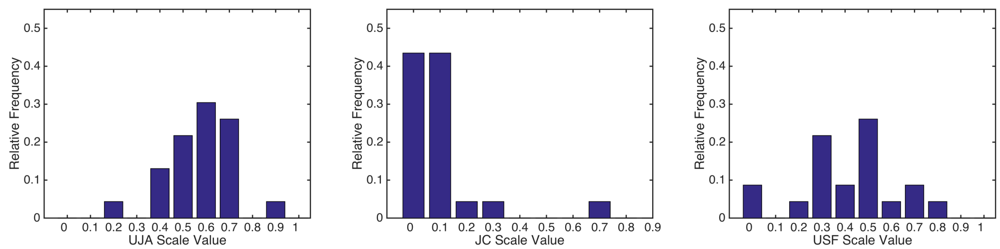


(e) User-centered adjustments (UJA) (f) Job cancelation (JC) (g) User satisfaction (USF)

Fig. 1: Distribution of scale values according to answers provided by 23 users of computer clusters at TU Dortmund University. Answers from *strongly disagree* to *strongly agree* are represented between $[0, 1]$.

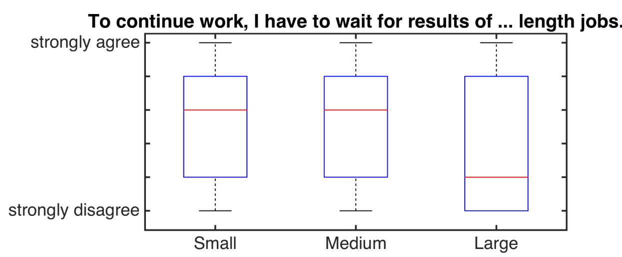


Fig. 2: Distribution of answers provided for the waiting for jobs (WJ) scale.

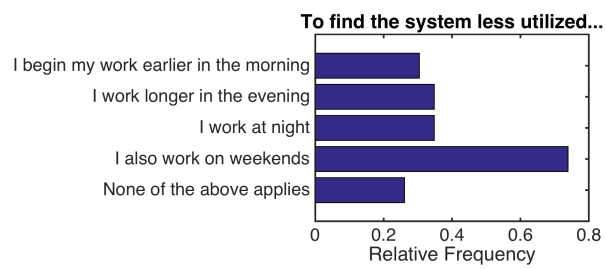


Fig. 3: Relative frequency of answers in the influence on working times (IWT) scale.

to basic job characteristics. Analysis results suggest that there is a need to prioritize jobs by the importance of continuing work. In Section IV-C, we show how these jobs may impact the user behavior.

Note that no user reported severe dissatisfaction. A possible explanation is that the questionnaire was conducted with users who are currently researching in a scientific environment, i.e., users that already have experience with the system (see also the descriptive analysis of the USF scale below). The analysis of users who left the system is out of the scope of this paper, since we target the understanding and modeling of continuous users and their behavior.

### *B. Descriptive analysis of scales*

The aggregated data shown in the previous subsection allows the inference of the general user behavior according to a scale. Nevertheless, it does not unveil knowledge related to, for example, specific user decisions or patterns. Therefore, in this subsection we present a descriptive analysis of the scale items observing the relative frequency of answers.

Fig. 3 shows the relative frequency of different participants' reactions to poor system performance. About 75% of the users tend to work on weekends, followed by working longer in the evening, and working at night ($\sim$35%). About 30% of the users work earlier in the morning, and only about 25% do not change their habits when facing poor system performance. This result indicates that users need to employ alternative strategies (besides working hours) to detour the low performance of the system and optimize their executions.

Fig. 4 shows the relative frequency of participants using a certain strategy to improve their executions. Besides working after hours, users also seek for improvements at the infrastructure and social (human interaction) level. Most of the users submit bags of tasks (above 60%), even though only a few (7 out of 23) mainly work with the cluster of the Physics Department (which instantiates an HTC environment). This result demonstrates that the capabilities of the HPC environment are underutilized, which may be due to users that do not have proper knowledge of the usage guidelines.

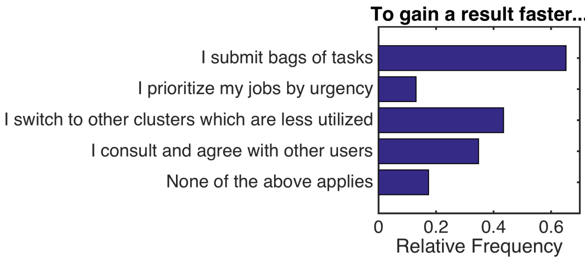


Fig. 4: Relative frequency of answers in the usage of strategies (US) scale.

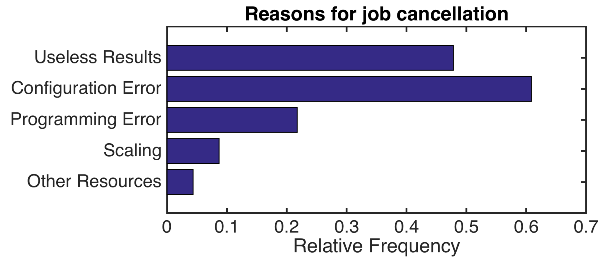


Fig. 5: Relative frequency of answer categories in the job cancellation (JC) scale.

In a previous work [8], we observed that this is not common in supercomputers, where policies are stricter, and allocations and usage are actively monitored. Our data suggests, that job prioritization is not common (less than 20% of the users). Since users have significant experience with the system and most of the executions are embarrassingly parallel jobs, jobs within an experiment have the same importance. About 40% of the users move their computations to other resources when facing high workload contention. This behavior clearly indicates a metric of success (in the form of feedback), and should also be considered in trace analyses. Surprisingly, several users ($\sim$35%) establish informal agreements to coordinate their executions. This fact arouses the interest in investigating inter-user co-operation that could unveil *on-the-fly* deviations to the user behavior, or violations to the service level agreements. This result also indicates that the current system policies do not satisfy the users' requirements. Last, less than 20% of the users do not use any of the listed strategies. This analysis demonstrates the need to further investigate user-strategies to improve parallel computing environments, either by conducting trace analysis, workload modeling, or scheduler design. Each of these research interests must be aware of these fluctuations and behavior traits. Trace analysis must be aware of this hidden knowledge, which are only indirectly represented in recorded data. Therefore, a better workload modeling, and results on importance of jobs can lead to better satisfaction.

Although most of the users have a low rate of job cancellations (Fig. 1f), we have classified the main causes why jobs are deleted (Fig. 5). Most of the jobs are often removed due to configuration errors or useless results. Note that configuration and programming errors imply a useless result. In case an answer could not be explicitly classified into one of the categories, we accounted it for all possible interpretations (e.g., if an answer states that mistakes were done during job preparation, it would be accounted as configuration and programming error). The high percentage for useless results is mainly due to the non-convergence of iterative methods (see Section IV-A).

Fig. 6 shows the distribution of answers provided to the six items in the general job adjustment scale (GJA). Whiskers are defined as 1.5 IQR—interquartile range, i.e., the distance between the upper and lower quartile. Users mostly adjust their jobs according to the available capacities of the system (median: *somewhat agree*, upper quartile: *agree*). Therefore, when analyzing and comparing workload traces, one needs to be aware that different system capacities may lead to distinct shaped jobs. On the other hand, most of the users do not tend to scale their jobs (e.g., increase the number of cores, or the number of jobs within a bag of tasks) if resources are idle—lower quartile: *strongly disagree*. In most parallel computing systems where queuing systems do not hold jobs longer than necessary, timing the job submission may significantly improve the system performance (in terms of waiting times), and as a result user satisfaction. However, optimized job submission is not frequently used. Answers range from *strongly disagree* (lower quartile) to *somewhat agree* (upper quartile), with median *somewhat disagree*. This result indicates that users are not aware of (or are not willing to explore) other system capabilities regarding job scheduling (e.g., advance reservation, queue priorities, and among others), since they often work after hours and/or establish informal agreements among them (Fig. 4). Conversely, users do care for not disturbing the system performance, since they do interrupt jobs that would not produce useful results. However, it is not clear whether they constantly (or have the capability to) monitor jobs outcomes. Users *somewhat agree* (with whiskers ranging across the whole spectrum of answers) that having system load information is helpful to determine job submission and their expectations, which could also influence user behavior. Choosing systems suiting the job requirements is also common among participants (the median answer is *agree*). This result shows that users are aware of their jobs, systems, and care for the specifications.

Fig. 7 shows the distribution of answers provided to the items of the user-centered adjustment scale (UJA). Some of the questions target similar aspects as the GJA scale, but they are user-centered. Although users tend to monitor the status of the system (Fig. 6), high contention do not prevent users to submit their jobs. In Section IV-A, we suggested that users may plan in advance their executions, or there is no need for instant results. However, the analyses performed in this subsection revealed that users seek for alternatives to avoid contention.

### C. Correlation analysis between scales

The analysis conducted in the previous subsections showed that the data in all scales are widely spread. Although corre-

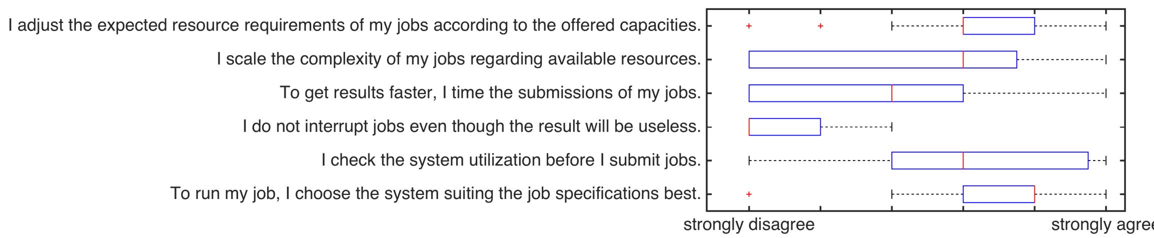


Fig. 6: Boxplots of users' answers to the general job adjustment (GJA) scale.

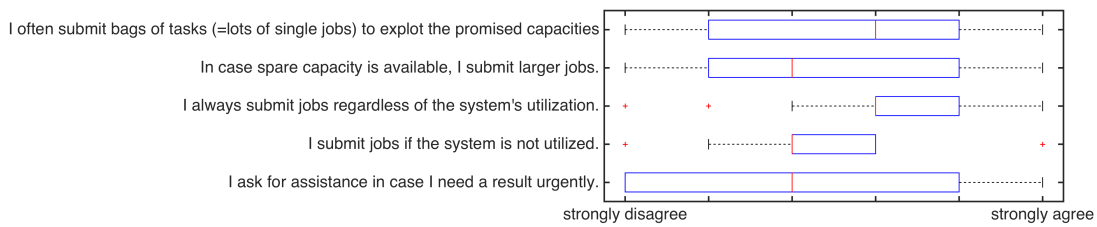


Fig. 7: Boxplots of users' answers to the user-centered job adjustment (UJA) scale.

lations do not necessarily describe causal dependencies, they allow us to infer dependencies (e.g., there is a relationship between user satisfaction and the adjustments of working times). Therefore, we use the Spearman's correlation coefficient to identify statistical relationships between the scales. We assume a $p$-value of less than $0.05$ (in some cases we also consider values below $0.1$) to rate correlations as of statistical significance. An advantage of Spearman's correlation coefficient is that it is independent of the ordinal scale values.

Fig. 8 presents an overview of Spearman's correlation values $\rho \in [-1, 1]$ between all scales. The correlation coefficient is expressed as an ellipse shape. The more elliptical it is, the higher the correlation coefficient. Otherwise, the more circular it is, the lower the correlation. Furthermore, the shape is tilted according to whether the correlation is negative ($\rho \to -1$, red) or positive ($\rho \to 1$, blue). The colors also indicate the strength of the correlation coefficient: pale colors indicate weak correlation values and therefore uncorrelated, while darker colors indicate increasing correlations. `WaitSmall`, `WaitMedium`, and `WaitLarge` represent the three jobs lengths from the waiting for jobs (WJ) scale. Although variables are not strictly correlated (which is expected), we notice several narrow elliptical shapes that may represent interesting correlations. Thus, we further analyze these correlations below.

We extract the most significant values according to the $p$-value. We assume that a $p$-value of 0.05 means significant correlation, while a $p$-value inferior to 0.1 will only serve the purpose of understanding and discussing a few parameters, but further study must be performed to underline or decline these correlations. Table I shows the Spearman's correlation values and $p$-values for the most significant correlations shown in Fig. 8. Correlations are ordered by significance, i.e., $p$-values < 0.05 in the top, and $p$-values $\in$ $[0.05, 0.1[$ at the bottom.

By means of the correlation table, we investigate the following claims, which were raised as part of the motivation of

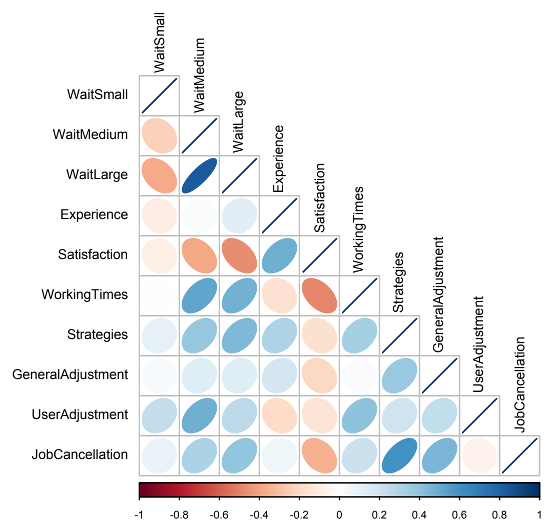


Fig. 8: Spearman's correlation map between scales. Two variables are full correlated if the ellipse is a line.

this work:

1) *Dissatisfaction with system response times is correlated with changes in working time behavior.* The correlation coefficient between the scales for USF and IWT indicates a negative correlation ($\rho < -0.48$, $p < 0.02$). This result confirms previous findings, which also suggest that satisfaction and experience are correlated [7];
2) *Completion of small- and medium-length jobs have more impact on the consecutive working.* The answers provided to scale WJ (Fig. 2) emphasize the importance of short- and long-running jobs on the user's consecutive work. However, in a correlation analysis we can only report significant correlations for medium and large jobs. Waiting for results of medium-sized jobs ($\rho > 0.52$,

TABLE I: Statistically significant correlations between scales with $p$-values $< 0.05$ (upper), and $p$-value $< 0.1$ (bottom).

| Correlated Scales | | | Spearman corr. coeff. $\rho$ | $p$-value |
|---|---|---|---|---|
| WaitMedium | - | WaitLarge | $\rho > 0.83$ | $p < 0.001$ |
| Strategies | - | JobCancellation | $\rho > 0.60$ | $p < 0.002$ |
| WaitMedium | - | WorkingTimes | $\rho > 0.52$ | $p < 0.02$ |
| Experience | - | Satisfaction | $\rho > 0.48$ | $p < 0.02$ |
| Satisfaction | - | WorkingTimes | $\rho < -0.48$ | $p < 0.02$ |
| WaitMedium | - | UserAdjustment | $\rho > 0.48$ | $p < 0.03$ |
| WaitLarge | - | Satisfaction | $\rho < -0.46$ | $p < 0.03$ |
| WaitLarge | - | WorkingTimes | $\rho > 0.47$ | $p < 0.03$ |
| GeneralAdj. | - | JobCancellation | $\rho > 0.45$ | $p < 0.03$ |
| WaitLarge | - | Strategies | $\rho > 0.44$ | $p < 0.04$ |
| WorkingTimes | - | UserAdjustment | $\rho > 0.40$ | $p < 0.06$ |
| WaitLarge | - | JobCancellation | $\rho > 0.39$ | $p < 0.07$ |
| WaitMedium | - | Satisfaction | $\rho < -0.38$ | $p < 0.07$ |
| WaitSmall | - | WaitLarge | $\rho < -0.37$ | $p < 0.08$ |
| WaitMedium | - | Strategies | $\rho > 0.38$ | $p < 0.08$ |
| Strategies | - | GeneralAdj. | $\rho > 0.37$ | $p < 0.09$ |
| Satisfaction | - | JobCancellation | $\rho < -0.35$ | $p < 0.10$ |

$p < 0.02$) is slightly stronger correlated to the influence on working times than large jobs ($\rho > 0.47$, $p < 0.03$). Regarding satisfaction, waiting for large ($\rho < -0.46$, $p < 0.03$) and medium length ($\rho < -0.38$, $p < 0.07$) jobs are nearly equal negatively correlated. Similar results are observed for the usage of strategies, where positive correlations are observed: $\rho > 0.44$, $p < 0.04$ for large, and $\rho > 0.38$, $p < 0.08$ for medium. Nevertheless, waiting for medium-sized jobs are most relevant for adjustments ($\rho > 0.48$, $p < 0.03$). These results indicate that short and medium jobs should be prioritized when optimizing consecutive work, however large jobs may also negatively impact the satisfaction, and therefore lead to the use of strategies;

3) *User experience increases satisfaction.* Although one could argue that experienced users have a deeper understanding of how a system could perform better, and they are therefore dissatisfied with the system, the Spearman correlation coefficient reveals a positive relationship between experience and satisfaction ($\rho > 0.48$, $p < 0.02$). The analysis however does not reveal whether this is due to the experience and coping with waiting times, or that experienced users actually use the system better towards their own needs. We then argue that teaching researches about the underlying mechanisms of the resources, as well as best practices, is beneficial for user satisfaction;
4) *Job cancellation has significant influence in the user behavior and satisfaction.* Our analysis unveiled strong correlations that support this claim: (1) the user behavior shows positive correlations to the usage of strategies ($\rho > 0.60$, $p < 0.002$), and to general job adjustments ($\rho > 0.45$, $p < 0.03$). Furthermore, the waiting for results from large jobs is also strongly correlated to the job cancellation scale ($\rho > 0.39$, $p < 0.07$); (2) the user satisfaction presents negative correlation between the JC scale and the USF scale ($\rho < -0.35$, $p < 0.10$). This result highlights a need for autonomic tools that seamlessly enable the execution of parallel computing applications on high performance systems to rule out this aspect, in case there is a causal relationship between the results from these two scales. Examples of such tools include MapReduce [17] and scientific workflows [18], [19], however limitations at the infrastructure level may narrow their efficiency (e.g., two-factor authentication, or limited network connectivity).

## V. Discussion on Opportunities and Challenges on Understanding User Submission Behavior

The study of *human* user behavior conducted in this paper have extended the understanding of the users' working behavior in parallel computing. Although several aspects investigated in this paper may be common knowledge[6], our findings unveil factors that are often hidden in workload traces, and human elements that cannot be modeled from traces due to the lack of knowledge on users and their submission behavior. Below, we discuss the challenges and limitations of trace-based analysis, and underline current open questions:

1) *Users' expertise.* Participants of the study consider themselves as experienced users. However, subsequent analyses revealed that users are not aware of all capabilities of the system (e.g., advance reservations). Thus, there is a need to define methods and techniques to identify users' expertise level in parallel computing systems. In order to build this knowledge, direct interview or questionnaires would be required along with active monitoring of users reactions to application and system issues. Even though this knowledge cannot be extracted from workload traces in a first instance, results from the user-centered study would aid the identification of such patterns in available traces. Additionally, a study on how less experienced users would impact the work of expert users may lead to the development of experience-aware scheduling algorithms;
2) *Modeling interactions with other systems.* This study revealed that users tend to use different strategies to cope with issues and low system performance. Although bags of tasks can be estimated [5], [20], and job priorities can be obtained [14], [15] from workload traces, user strategies (e.g., use another cluster) cannot be modeled from them. This result indicates that user behavior analysis should include workloads from different systems that could be potentially used by the user. A typical example is the use of a cluster optimized for computations, and another optimized for visualization [8];
3) *Unveiling user agreements (and possibly violations).* Users from the same institution or research group tend to establish informal agreements to improve their performance. Although these agreements might not be

[6]To the best of our knowledge, this is the first work that provides scientific data to support this common knowledge in the field of distributed computing.

hazardous to the system, they may violate policies (e.g., a user provides spare resource allocation to another user to run an experiment that is not listed in the allocation request). Questionnaires may unveil this practice, however it would not prevent the misuse of the system. Therefore, there is a need of methods to automatically detect such practices (e.g., through job modeling based on past executions) and prevent or alert users of possible misconducts.

## VI. Conclusion

Understanding user behavior is crucial to the improvement of user satisfaction, while promoting optimal resource utilization in parallel computing systems. This paper presented the analysis of answers provided to the Questionnaire on User Habits in Computer Clusters (QUHCC), to unveil human user behavior in parallel computing. The study was conducted with 23 users of different science domains from three computer clusters at the TU Dortmund University. Analysis results shows that although most of the users have substantial knowledge of the system, they still need to work after hours to obtain better response from the systems. Additionally, users tend to establish informal agreements between them to improve their efficiency. These human interactions, for example, cannot be captured in workload traces and therefore limit the analysis. On the other hand, workload traces can provide information on previous users that abandoned the system. Although this information may not explicitly highlight the causes related to leaving the system (e.g., user dissatisfaction, or the user had no more experiments to run), relations could be inferred from, for instance, the number of successful executions and experienced waiting times, which may lead to dissatisfaction. Therefore, we argue that the deeper understanding of users together with the historical information (traces) are crucial to build a concise understanding of feedback effects between the user satisfaction and the system performance.

Our findings lead future research into directions of different granularity. Future research should focus to investigate workload traces according to findings of this work. For example, participants of our study answered, that slow responses of the system correlate with adjustments of working times. Since the distribution of job submissions is an important factor in workload modeling, it is important to reveal these feedbacks between performance and submissions in recorded traces. Furthermore, additional research is necessary on the correlations reported including user satisfaction. Causal dependencies are of interest and can be accessed by path analysis and structural equation modeling. Afterwards, they can be exploited towards user-oriented scheduling strategies and policies. For example, minimizing the necessity for users to adjusting working times, either on weekdays to the night, or early morning or to the weekend then is analyzed as an important step to increase user satisfaction. Simulations of increased system sizes and improving performance metrics are possible. In the broad picture, future work can exploit results from this work and work on the question in how far algorithms (in this case scheduling and allocation decisions) have influence on our daily lives and their social impact.

**Acknowledgments.** This work was partly funded by DOE under the contract number ER26110, "dV/dt – Accelerating the Rate of Progress Towards Extreme Scale Collaborative Science".

## References

[1] D. G. Feitelson, *Workload Modeling for Computer System Performance Evaluation*. Cambridge University Press, in press.

[2] ——, "Looking at data," in *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*. IEEE, 2008, pp. 1–9.

[3] E. Shmueli and D. G. Feitelson, "On simulation and design of parallel-systems schedulers: are we doing the right thing?" *Parallel and Distributed Systems, IEEE Transactions on*, vol. 20, no. 7, pp. 983–996, 2009.

[4] S. Schlagkamp, "Influence of dynamic think times on parallel job scheduler performances in generative simulations," in *JSSPP 2015 - 19th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP 2015)*, Hyderabad, India, May 2015.

[5] N. Zakay and D. G. Feitelson, "On identifying user session boundaries in parallel workload logs," in *Job Scheduling Strategies for Parallel Processing*. Springer, 2012, pp. 216–234.

[6] U. Schwiegelshohn, "How to design a job scheduling algorithm," in *Job Scheduling Strategies for Parallel Processing*. Springer, 2014, pp. 147–167.

[7] J. Renker, S. Schlagkamp, and G. Rinkenauer, "Questionnaire for user habits of compute clusters (QUHCC)," in *HCI International 2015-Posters' Extended Abstracts*. Springer, 2015, pp. 697–702.

[8] S. Schlagkamp, R. Ferreira da Silva, W. Allcock, E. Deelman, and U. Schwiegelshohn, "Consecutive job submission behavior at Mira supercomputer," in *ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2016.

[9] S. Schlagkamp, R. Ferreira da Silva, E. Deelman, and U. Schwiegelshohn, "Understanding user behavior: from HPC to HTC," in *International Conference on Computational Science (ICCS)*, 2016.

[10] "Parallel workloads archive," http://www.cs.huji.ac.il/labs/parallel/workload.

[11] A. Iosup, H. Li, M. Jan, S. Anoep, C. Dumitrescu, L. Wolters, and D. H. Epema, "The grid workloads archive," *Future Generation Computer Systems*, vol. 24, no. 7, pp. 672–686, 2008.

[12] C. Germain-Renaud, A. Cady, P. Gauron, M. Jouvin, C. Loomis, J. Martyniak, J. Nauroy, G. Philippon, and M. Sebag, "The grid observatory," in *Cluster, Cloud and Grid Computing (CCGrid), 2011 11th IEEE/ACM International Symposium on*. IEEE, 2011, pp. 114–123.

[13] D. Kondo, B. Javadi, A. Iosup, and D. Epema, "The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems," in *Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on*. IEEE, 2010, pp. 398–407.

[14] R. Ferreira da Silva and T. Glatard, "A science-gateway workload archive to study pilot jobs, user activity, bag of tasks, task sub-steps, and workflow executions," in *Euro-Par 2012: Parallel Processing Workshops*, ser. Lecture Notes in Computer Science, 2013, vol. 7640, pp. 79–88.

[15] R. Ferreira da Silva, W. Chen, G. Juve, K. Vahi, and E. Deelman, "Community resources for enabling and evaluating research on scientific workflows," in *10th IEEE International Conference on e-Science*, ser. eScience'14, 2014, pp. 177–184.

[16] P. Kline, *A Handbook of Test Construction (Psychology Revivals): Introduction to Psychometric Design*. Routledge, 2015.

[17] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.

[18] T. Fahringer, R. Prodan, R. Duan, J. Hofer, F. Nadeem, F. Nerieri, S. Podlipnig, J. Qin, M. Siddiqui, H.-L. Truong *et al.*, "Askalon: A development and grid computing environment for scientific workflows," in *Workflows for e-Science*. Springer, 2007, pp. 450–471.

[19] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger, "Pegasus, a workflow management system for science automation," *Future Generation Computer Systems*, vol. 46, pp. 17–35, 2015.

[20] A. Iosup, M. Jan, O. Sonmez, and D. Epema, "The characteristics and performance of groups of jobs in grids," in *Euro-Par 2007 Parallel Processing*. Springer, 2007, pp. 382–393.