The Sky Is Not the Limit: Multitasking Across GitHub Projects

Bogdan Vasilescu[†], Kelly Blincoe[§], Qi Xuan[‡], Casey Casalnuovo[†], Daniela Damian[♯],
Premkumar Devanbu[†], Vladimir Filkov[†]

[†]Dept. Computer Science, University of California, Davis. Davis, CA 95616, USA
[§]Dept. Electrical & Computer Engineering, University of Auckland, New Zealand
[‡]Dept. Automation, Zhejiang University of Technology, Hangzhou 310023, China
[♯]Dept. Computer Science, University of Victoria, Victoria, BC, Canada
{vasilescu, ccasal, ptdevanbu, vfilkov}@ucdavis.edu
k.blincoe@auckland.ac.nz, xuanqi@zjut.edu.cn, danielad@uvic.ca

ABSTRACT

Software development has always inherently required multitasking: developers switch between coding, reviewing, testing, designing, and meeting with colleagues. The advent of software ecosystems like GITHUB has enabled something new: the ability to easily switch *between* projects. Developers also have social incentives to contribute to many projects; prolific contributors gain social recognition and (eventually) economic rewards. Multitasking, however, comes at a cognitive cost: frequent context-switches can lead to distraction, sub-standard work, and even greater stress. In this paper, we gather ecosystem-level data on a group of programmers working on a large collection of projects. We develop models and methods for measuring the rate and breadth of a developers' context-switching behavior, and we study how context-switching affects their productivity. We also survey developers to understand the reasons for and perceptions of multitasking. We find that the most common reason for multitasking is interrelationships and dependencies between projects. Notably, we find that the *rate of switching* and *breadth* (number of projects) of a developer's work matter. Developers who work on many projects have higher productivity if they focus on few projects per day. Developers that switch projects too much during the course of a day have *lower* productivity as they work on more projects overall. Despite these findings, developers perceptions of the benefits of multitasking are varied.

CCS Concepts

•**Information systems** → **Data analytics;** •**Human-centered computing** → **Empirical studies in collaborative and social computing;** •**Software and its engineering** → **Open source model;**

Keywords

Multitasking; GITHUB; productivity

1. INTRODUCTION

Multitasking is a staple of high performing professionals, programmers included. It is the ability to stop working on a task, switch to another, and return eventually to the first one, as needed or as scheduled. The goal is to optimize human resource allocation, while reprioritizing tasks dynamically [17]. When done well, or at least in a disciplined way, multitasking can yield dividends [1, 4]. Clearly, if a task in the queue has a higher priority than the current one, switching them can improve performance. For example, programmers encounter this when starting a new coding task, while their previous, urgently needed fix undergoes testing. If testing uncovers bugs, the urgent fix will require attention, and the new coding task will be put on the back-burner.

Multitasking comes at a cost though [9]. Humans, programmers included, have a certain, limited amount of *cognitive flexibility*, the mental ability to switch from thinking about one concept to thinking about another. Limitations apply to both the number of concepts we can juggle, as well as to the difficulty in switching between them. For example, coding simultaneously 7 simple functions is easier than coding 8 of them. Likewise, switching between two small matrix multiplication problems is more difficult than switching between two small integer additions. As we can imagine, reaching our innate limitations can result in decreased performance on all tasks and perhaps even diminished quality. It is unknown how far multitasking can be pushed safely, although some anecdotal evidence is available.

Software developers have long been pushing the limits on multitasking [27] because of the innate modularity of the development process and the independence of module processing (*e.g.*, one can code while tests are being executed). In open-source software, developers also commonly contribute to multiple projects in the same time period, bridging different communities [20, 30]. With the advent of social coding tools like GITHUB, this has intensified. It is not uncommon to find prolific developers contributing code to 5-10 GITHUB projects in the same week. In fact, contributing to as many GITHUB projects as possible is an accomplishment, valued by peers and employers alike [32].

There are various reasons why developers are more prolific on GITHUB compared to other platforms. The features and usability provided by GITHUB play a big role [34]. It is one example of how novel technology benefits programmers, and it empowers them; GITHUB's platform is responsible for the inception of many projects, that otherwise wouldn't have existed [34]. And the payoffs are also substantial [34]. With

so many drivers for multitasking, it is easy to see how one could cross the limits from safe, productive multitasking into an overloaded mode, where code output falls and bugs start to multiply. The question really becomes where are those limits, and what are their determinants?

Multitasking is a complex phenomenon, with costs along different dimensions. This paper reports on a mixed methods study of multitasking and focus switching. Through analysis of longitudinal data, we investigate how productivity (*i.e.*, outputs produced per unit time) of prolific GITHUB programmers is determined by the number of projects they work on, how much they focus on each (relative to the others), and how diverse the projects they contribute to are in terms of programming languages. Notably, the very platform (GITHUB) that has introduced this multi-project multitasking phenomenon, also gives us all the tools we need to understand when programmers are at risk to approach their limits because of it. Further, a survey of 128 of these prolific developers was used to better understand the reasons for, and perceptions of, multitasking. The highlights of our findings are that:

- 98% of prolific programmers have contributed to multiple projects per day at least once, but the patterns are nuanced, and the reasons vary.
- Within limits, multitasking is associated with higher productivity, but developers' perception is mixed.
- Productivity does not decrease so long as developers work on 4 or fewer projects per week. Yet, developers often want to contribute to more projects regardless of their current workload.

We introduce our research questions and related works in Sec. 2, followed by our research methods (Sec. 3) and results (Sec. 4). We discuss the significance of our contributions in Sec. 5 and offer some concluding remarks in Sec. 6.

2. BACKGROUND

Multitasking and Performance. Multitaskinghas become increasingly common among knowledge workers [17, 31, 37]. While there is some disagreement between studies, it is generally believed that multitasking has non-monotonic (concave) effects on individual productivity [1,4], with plenty of theoretical evidence from Psychology, Management, and Organizational Behavior (*e.g.*, [1,37]).

Positive effects are attributed to several factors. First, multitasking may increase productivity since inevitable lulls in one project (*e.g.*, waiting for information from customers and colleagues, waiting for the build process to finish) can be filled with tasks related to other projects [1, 4, 37]. By keeping multiple projects active at the same time, one can switch focus between those projects during periods of relative down-time, enabling them to utilize their time more efficiently, therefore increasing their productivity.

Second, multitasking may increase productivity through cross-fertilisation and learning. If knowledge and skills required to work on a project are transferrable to other projects as well, one may experience decreasing costs of contributing to multiple projects simultaneously, as they are able to realize such knowledge transfers [29]. Similarly, exposure to more projects brings about exposure to different environments, providing opportunities to develop transferrable skills. We have found evidence for such positive effects of knowledge transfers on productivity (past experience with a programming language) during prior work on GITHUB [10].

On the other hand, several theories explain the cognitive processes that account for decreased performance in multitasking conditions, *e.g.*, memory-for-goals [2]. Performing a task may require maintaining some information mentally. This information is stored in the central part of the working memory, an area of the brain with high access speed but single-task capacity [3]. Memory-for-goals explains how initiating a task requires strengthening its goal ("a mental representation of an intention to accomplish a task" [2]) in memory to the extent to which its activation rises above other competing goals [49]. Therefore, in a multitasking context, switching to another task involves first retrieving its goal from memory (especially if multiple tasks require storing state [9], which is not uncommon for programming), and this takes time. With more multitasking, one starts to incur cognitive switching costs for interrupting one task and resuming another [9]. Ultimately, too much multitasking leads to mental congestion, which negatively affects productivity. As a result, workers start to experience "project overload" [59], and become increasingly slower and more error prone [8, 16, 46, 48]. The duration of the interruption, its complexity, and the moment when it occurs, all play an important role in how disruptive task switching is [2, 18].

In addition to potential multitasking (*e.g.*, when the technological context changes), switching between projects on GITHUB also involves a social component: different projects may involve different social contexts and different teams, all of which require adjusting to. Research has found that the more diverse the "working spheres" associated with each team are, the more disruptive switching between those teams becomes, and the more it hinders productivity [31]. GITHUB is known to be a particularly diverse [53], social [14] ecosystem, wherein many social attributes become salient [52] and can impact impression formation and collaboration [33].

In summary, we expect productivity benefits for those who engage in multitasking (we focus on multitasking across GITHUB projects in this paper), through load balancing, more efficient work practices, learning, and cross-fertilisation. However, there will be a point of diminishing returns, after which these benefits will be offset by cognitive overload.

Focus Switching in Software Engineering. Multitasking in software engineering is prevalent, to the extent that rules of thumb have been proposed to discourage contributing to multiple projects at a time [55]. Perhaps more so than in other pursuits in life, multitasking for programmers amounts to switching away from one well defined technical task, *e.g.*, coding, before it has been finished, and onto another. Coming back to such unfinished tasks necessitates some retention of precise technical detail, then refocusing on a new task, and later a recall of the previous one.

Developers switch focus frequently due to interruptions [22,27,60], either external or self-inflicted [13]. For example, technical dependencies may exist between tasks, and this can cause developers to interrupt one another for coordination purposes [5, 11, 21]. Developers' focus switching patterns within a project are complex [57]. Focus switches occur when changing from one task to another, or when stopping work to coordinate or talk to someone else on the team [56].

When developers resume a task after an interruption, they must reconstruct the task's context [38, 40–42]. Cues and tools can help developers pick up where they left off, making it easier to remember task details [26, 39], but even so there is considerable overhead involved when switching. For

example, editing many files concurrently is found to negatively impact software quality [58]. Similarly, fragmented within-project work drives down developer productivity [50]. In this paper, we look beyond individual projects and investigate focus switching patterns and their impact on productivity when developers contribute *across* multiple projects.

Furthermore, a recent study found that developers associate minimal interruptions and context switches with higher productivity; however, they also reportedly feel productive when significant task switching occurs [35]. A possible explanation for this disparity is that software developers are not always able to provide an accurate, retrospective report on the amount of interruptions they face in a day [44]. Our study continues this investigation of *perceived* effects of multitasking on productivity.

Research Questions. We frame our study in the context of multitasking *across* GITHUB projects, *i.e.*, switching back and forth between multiple projects that one contributes to, within a short period of time (our operationalization below is one week). Our study was guided by two research questions. First, we seek a deeper understanding of multitasking and its effects for software developers active on GITHUB.

RQ1. What are the trends of, reasons for, and effects of multitasking and focus switching on developer productivity in GitHub?

Second, having identified the important multitasking and focus switching effectors on programmer productivity, we proceed to investigate how these dimensions interact, and which tradeoffs between them exist.

RQ2. Are there limits on multitasking (and what are they) before productivity is impacted?

3. METHODS

To answer our research questions, we followed a mixed-methods approach characterized by a sequential explanatory strategy [15]. We analyzed development activity and perceptions of prolific GITHUB developers. We combined: (1) an analysis and regression modeling of repository data, to quantitatively examine the effects of multitasking and focus switching on productivity; and (2) a user survey, to garner additional qualitative insight into the developers' perceptions of multitasking, focus switching, and their effects.

3.1 Data Set

Prolific Developers. We utilized a GITHUB dataset collected during prior work that contains information on prolific developers with a long and active contribution history [10]. Such developers were identified using GHTorrent [19], as those with: at least 5 years between their earliest and latest recorded commits; and at least 500 commits in total, made to at least 10 different non-fork repositories. The dataset contains details about the date, size (LOC added/removed), and contents of all commits authored by 1,255 developers across 58,092 public repositories accessible on GITHUB at the time of mining. Commit data was obtained by cloning each repository locally and parsing its git logs, and was used for our repository analysis. Multiple aliases used by a single developer were resolved using a series of heuristics [54] (we found 395 developers, or 31.5%, with two or more aliases; the maximum number of aliases per developer was eight). The languages of source code files were determined using filename extensions and some contextual information [10].
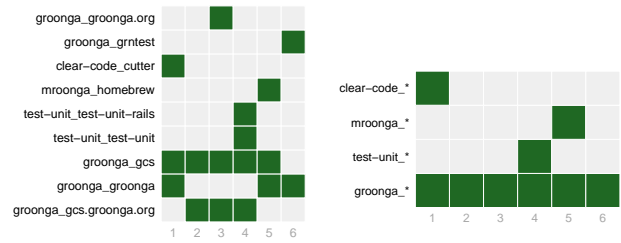

Figure 1: *Left*: Sample from one developer's daily contributions during a week in 2012. *Right*: Aggregation of the data into projects. The *x*-axis is the day index.

In addition, we invited developers in the most active three-quartiles (by total count of days active) for the user survey.

Conceptualization of a Project. Since our analysis focuses on multitasking across projects, we need to conceptualize how we define the boundaries of a software project. The naive approach would be to consider each GITHUB repository as its own separate project. However, we observed that in some cases, software projects are organized into multiple separate repositories on GITHUB.[1] Such repositories are conceptually and technically related, or even interdependent [7]. It is arguably less costly to switch between such related repositories (in terms of context switches) than other, less related repositories on GITHUB.

To account for this technicality, we conservatively group all repositories owned by the same GITHUB user or having exactly the same name[2] into sets of repositories, which we hereafter refer to as *projects*. All data about individual contributions to these repositories is then correspondingly aggregated at project level. This transformation ensures that we do not overestimate the number of focus switches during short time intervals. For example, the left part of Figure 1 depicts a sample from a developer's daily activity across different repositories in a week in 2012 (a day [column] – repository [row] cell is highlighted if the developer has contributed commits to that repository that day; days without any activity have been omitted). In total, the developer has contributed commits to 9 different repositories over the course of that week. On the busiest day, she contributed to 4 of these. The resulting aggregation of this data, using our conceptualization of a project, depicted in the right part of Figure 1, shows contributions to only 4 projects.

Note that our definition of *project* as a collection of repositories owned by the same GITHUB user is broader than the definition of project as a main repository together with all its forks, proposed in the literature [25].

3.2 Multi-project Multitasking Productivity

For the quantitative analysis, we developed a model to investigate the relationship between outputs produced per unit time (as a proxy for productivity) and a multitude of factors relating to multitasking and focus switching. Here we describe the factors and the model itself.

3.2.1 Temporal Resolution

An important question when studying multitasking and focus switching is what time interval to consider. During a longer period of time, be it objective (*e.g.*, week, month, year) or context-dependent (*e.g.*, project duration, length

---

[1] *e.g.*,, Ruby Standard Library `https://github.com/rubysl`
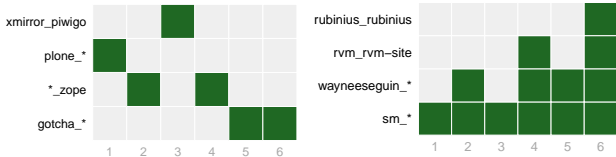
[2] Repos are identified by the `user_login/repository_name` stub.

Figure 2: Two developers contributed to 4 projects in a given week. *Left*: Sequentially; *Right*: Multitasking.
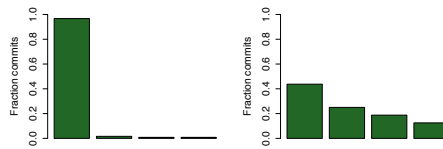

Figure 3: Distributions of commits over projects for two developers who contributed to 4 projects in a given week. Both worked sequentially, on not more than one project each day. *Left*: $S_{\text{Focus}} = 0.25$. *Right*: $S_{\text{Focus}} = 1.85$

of employment, major release), a developer may experience many focus switches without her necessarily also "juggling" many tasks over a shorter interval (*e.g.*, hour, day, working session). For example, suppose Alice and Bob both work on four different projects in a given week, and are free to schedule their tasks on each of the projects. They both interleave their tasks but in different ways. Each day, Alice only interleaves tasks on a single project, so she never works on more than one project in a day. Bob interleaves his tasks such that he contributes to multiple projects each day. When viewed at this resolution, both switch contexts (focus) between their four projects over the course of a week.

We modeled the interplay between focus switching at two temporal resolutions—daily and weekly,[3] and using different measures, as described next.

3.2.2 Multitasking Dimensions

Projects per day. At the finest temporal resolution (day), we measured multitasking activity using the number of different projects contributed to that day. Contributions are measured by commits. Clearly, this represents a lower bound on the number of switches per day: a developer contributing to k projects in one day has to switch focus at least k times that day.[4] Similar count-based metrics are used in other studies to represent task switching [1].

At the coarsest temporal resolution (week), we used the average number of projects per day (AvgProjectsPerDay) as an aggregate measure of multitasking, as per [4] (inactive days excluded). For example, if Alice was active for 6 days during a particular week, with the distribution of projects day_1: 2 different projects; day_2: 3; day_3: 1; day_4: 1; day_5: 3; and day_6: 2, then we say that Alice contributed to $(2 + 3 + 1 + 1 + 3 + 2)/6 = 2$ projects per day on average that week.

AvgProjectsPerDay captures the distinction between developers who, over the course of a week, tend to work on projects sequentially day-to-day (AvgProjectsPerDay=1 in Figure 2 left), and those who interleave contributions to multiple projects each day, *i.e.*, multitask (AvgProjectsPerDay=2.2 in Figure 2 right).

Weekly Focus. While effective at identifying periods of sequential and interleaved contributions, the AvgProjectsPerDay measure is not useful to distinguish how evenly developers divide their attention among their projects. This differentiation is important because it has been found that more narrowly focused developers have less cognitive burden, resulting in higher productivity and quality [45, 57]. For example, Alice and Bob both contributed to 4 projects one

week, and both worked sequentially (*i.e.*, AvgProjectsPerDay=1); however, over the course of that week, Alice focused her efforts on only one of the 4 projects (Figure 3 left), *i.e.*, most of her commits went to one project; in contrast, Bob contributed more evenly to his 4 projects (Figure 3 right).

From an information theoretic point of view, Alice's weekly focus switching behavior is predictable: knowing the historical distribution of her commits to different projects, one can accurately predict the object of Alice's next focus; her next contribution is more likely to be in a project to which she contributes frequently. Stated differently, Alice is less likely to switch focus between projects, since she spends most of her time contributing to a single project. In contrast, there is more uncertainty in Bob's behavior: since he contributes more uniformly to his projects, he is approximately equally likely to contribute to any one of them in the next time period. Bob is more likely to switch focus than Alice.

We measured the uncertainty in a developer's focus switching behavior (or the diversity of focus switches) in a given week using the Teachman/Shannon entropy index, a commonly used diversity measure in many scientific disciplines [4, 6, 57]. We denote this measure S_{Focus}, defined as:

$$S_{\text{Focus}} = -\sum_{i=1}^{N} p_i \log_2 p_i, \qquad (1)$$

where p_i is the fraction of the developer's commits this week[5] in project i, and N is the total number of projects this week. S_{Focus} ranges between 0, when a developer contributes to a single project that week, and $\log_2 N$, when a developer contributes equally (*i.e.*, $p_i = 1/N$) to all N projects.

Similarly, we measured a developer's language entropy S_{Language}, defined analogously over the L different programming languages of the files touched that week:

$$S_{\text{Language}} = -\sum_{i=1}^{L} p_i \log_2 p_i \qquad (2)$$

S_{Language} is a proxy for the overall complexity of one's contributions to different projects: when writing code in multiple programming languages, in addition to switching focus between different projects (which involves restoring project-specific contexts), one also must switch focus between different languages (which may involve different skills).

Day-to-Day Focus. To capture the diversity of focus switches also at a finer temporal resolution (day), we adapted the focus shifting networks (FSNs) developed by Xuan *et al.* [57] for file-level focus. In our case, for a developer contributing to N projects in a given week, her FSN is a weighted directed graph over all projects. Two projects i and j are connected by an edge if they have been committed to by the developer on successive, but not necessarily con-

---

[3] One week is considered a relevant period in software development [37]. For convenience, our two temporal resolutions both use objective time intervals; subjective time intervals cannot be directly inferred from repository data.

[4] We assume that going from a neutral state, before contributing to any projects that day, to a state of focus imposed by the day's first project also counts as one switch.

[5] p_i can also be defined in terms of files touched, or LOC added/removed; all tend to be highly correlated.
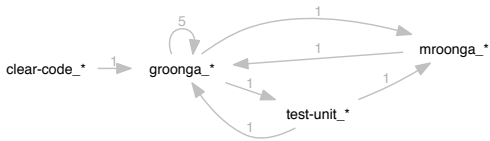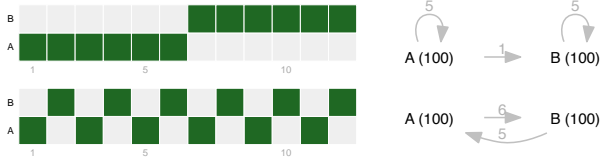
Figure 4: Corresponding FSN for Figure 1–right.

Figure 5: FSNs capture different aspects of one's focus switching behavior than AvgProjectsPerDay and S_{Focus}.

secutive, days. Edges point from the earlier-commit project to the later. The weight of the edge is the number of times the switch from i to j occurred.[6]

For example, recall the activity of the developer in Figure 1–right. Her corresponding FSN is depicted in Figure 4. There are 4 nodes corresponding to the 4 different projects she contributed to, over the 6 active days that week: `clear-code_*`, `groonga_*`, `test-unit_*`, `mroonga_*`. The day-to-day switch from `clear-code_*` to `groonga_*` occurred once (day$_1$–day$_2$); the switch from `groonga_*` to `groonga_*` is also recorded, since it occurred 5 times (all subsequent days except the last); other edges are constructed similarly.

FSNs capture additional information about a developer's day-to-day multitasking and focus switching behavior (relative to AvgProjectsPerDay and S_{Focus}), since they are based on time-series of commits rather than just aggregate counts. To illustrate this, consider the two scenarios in Figure 5. Alice and Bob both contribute equally to two projects A and B (100 commits to each; $S_{\text{Focus}} = 1$), and they never multitask (AvgProjectsPerDay = 1). However, in the first scenario Alice finishes all her tasks on project A before starting work on project B, while in the second scenario Bob *alternates* between days focused on A, and days focused on B. While neither AvgProjectsPerDay nor S_{Focus} distinguish between these scenarios, the corresponding FSNs (shown on the right in Figure 5) capture the two behaviors.

To leverage this, we considered focus switching as a Markov process—the next state is completely determined by the current one, and we measured the diversity of a developer's *day-to-day* focus switches (as opposed to aggregated at week-level with S_{Focus}) using Markov entropy [57]. We refer to the measure as S_{Switch}, defined as:

$$S_{\text{Switch}} = - \sum_{i=1}^{N} \left[p_i \sum_{j \in \pi_i} p(j|i) \log_2 p(j|i) \right], \qquad (3)$$

where π_i is the set of outgoing neighbors of node i (*e.g.*, $\pi_A = \{A, B\}$; $\pi_B = \{B\}$ for the top FSN in Figure 5); $p(j|i)$ is the conditional probability that the developer switches focus from i to j, defined as $p(j|i) = \frac{w_{ij}}{\sum_{k \in \pi_i} w_{ik}}$; w_{ij} is the weight of the edge $i \rightarrow j$; and p_i is defined as above.

S_{Switch} can be seen as a measure of the repetitiveness of one's *focus switches* from one day to the next: the lower

---

[6]Xuan *et al.* [57] used slightly different weights, including the number of files changed in a commit in the computation. Our definition reduces the collinearity with commit size.
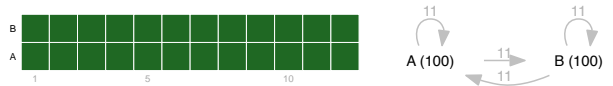

Figure 6: Example of a repetitive day-to-day behavior.

the value, the less repetitive one's day-to-day behavior is. Bob (Figure 5–bottom; $S_{\text{Switch}} = 0$) has the least repetitive day-to-day switches: he never contributes to the same project on two consecutive days. To illustrate the extreme opposite, consider the example in Figure 6, where Charlie contributes to exactly the same projects A and B on consecutive days ($S_{\text{Switch}} = 2$). Alice's behavior ($S_{\text{Switch}} = 0.325$) is in between the two extremes.

3.2.3 Regression Analysis

We modeled the variability in outputs produced (*i.e.*, LOC added) per unit time—our productivity proxy, as dependent on control measures and the three dimensions of multitasking: projects per day, weekly focus, and day-to-day focus. For each developer, our data consists of measurements of the different variables across multiple multitasking weeks (only weeks with Projects > 1 were modeled).

As customary in regression modeling with skewed predictors, as some of ours are, we first removed outliers. Whenever one of our variables x was well fitted by an exponential distribution, we rejected as outliers values that exceeded $k(1 + 2/n)\text{median}(x) + \theta$ [43], where θ is the exponential parameter [47], and k is computed such that not more than 0.5% of values are labeled as outliers, for each variable. Table 1 presents summary statistics for our filtered data set.

Then we fit a linear mixed-effects model with a random-effects term for developer. This allows us to capture developer-to-developer variability in the response (LOC added), (*e.g.*, some developers being naturally more productive than others), rather than assessing the contributions of specific developers, which we are less interested in. Additionally, we allow for deviations in slope of a developer's time trend from the population values (*i.e.*, we allow for the possibility that, for example, developers with higher initial productivity may, on average, be less strongly affected by time passing). We also tested, but found insignificant, the inclusion of a random-effect term for the time window in which the measurement was taken, to capture longitudinal, week-to-week variability (*e.g.*, some weeks developers may be more productive due to other, unobserved variables). All other variables were modeled as fixed effects. We used multiple linear mixed-effects models, as implemented in the functions `lmer` and `lmer.test` in R. Coefficients are considered important if they were statistically significant ($p < 0.05$). Their effect sizes are obtained from ANOVA analyses. We evaluate our model's fit using a marginal (R_m^2) and a conditional (R_c^2) coefficient of determination for generalized mixed-effects models [23, 36], as implemented in the `MuMIn` package in R: R_m^2 describes the proportion of variance explained by the fixed effects alone; R_c^2 describes the proportion of variance explained by the fixed and random effects together.

To check collinearity among the predictors we use the VIF, or variance inflation factor; all were below 4, except for those between the interaction terms and their comprising factors, which is expected. We conclude collinearity between our variables is not an issue. To ensure compliance with `lmer`'s modeling assumptions, we also checked the QQ-plot for our model, which showed good match with a normal distribu-

Table 1: Summary statistics for week-level data (78,552 rows; 1,193 developers; outliers removed).

Statistic	Mean	St. Dev.	Min	Median	Max
GlobalTime	1,142.76	74.24	991	1,150	1,263
UserTime	308.05	194.12	1	271	2,307
Repositories	3.70	2.22	2	3	18
Projects	2.57	0.98	2	2	9
DaysActive	3.98	1.63	1	4	7
Languages	3.19	1.43	1	3	14
Commits	23.89	30.06	2	15	943
FileTouches	88.56	174.31	2	39	2,737
LOCAdded	4,151.36	15,691.15	1	635	265,702
LOCDeleted	2,282.19	9,217.39	0	247	148,977
S_{Language}	0.82	0.52	0.00	0.84	2.92
AvgProjectsPerDay	1.44	0.48	1.00	1.33	6.00
S_{Focus}	0.92	0.44	0.02	0.92	3.01
S_{Switch}	0.64	0.53	0.00	0.69	2.83

tion. The residuals between the observed and model fitted values for LOC added showed no difference in variance variability across the range.

The following describes our regression variables:

Response. Our response is **LOCAdded**, as a measure of outputs produced per unit time (proxy for productivity). It sums the count of lines of code added per commit, over all commits by a developer in a given week. We also experimented with Commits and FileTouches. All are highly correlated (rho $\simeq$ 0.82 in both cases).

Main predictors. We include **AvgProjectsPerDay** (Multitasking), S_{Focus} (Weekly Focus), and S_{Switch} (Day-to-Day Focus), discussed above, and their interactions.

Controls. Our controls are:
- **GlobalTime**: Week index of the current week, with respect to 1990-01-01 (chosen arbitrarily for simplicity). Controls for potential generic environment changes.
- **Projects**: Total number of projects contributed to in a given week.
- **Languages**: Number of different programming languages across all files touched that week.
- S_{Language}: Controls for the overall "linguistic" complexity of one's contributions.
- **UserTime**: Per developer index of the current week, relative to one's first ever recorded GitHub contribution. Controls for changes in developer experience with time, that may have affected individual productivity.

3.2.4 Hypothesis Testing

To test for a difference in the medians between two populations we use the non parametric Wilcoxon-Mann-Whitney test, for unpaired samples, and the Mann-Whitney paired test for paired samples. We use the p-values to determine statistical significance, and supplement those with the Hodges-Lehmann point estimate for effect sizes, $\hat{\Delta}$. We also use the multiple contrast test procedure $\widetilde{\mathbf{T}}$ [28] for Tukey-type contrasts at 95% confidence level, to distinguish significant from non-significant differences between all pairs of distributions.

3.3 User Survey

We sent an online survey to 851 GitHub users selected from the set of prolific developers described earlier. The survey included multiple choice, Likert scale, and open-ended questions. We asked users about their software development experience in general, and with GitHub; which factors influence their contributing to new repositories; what makes
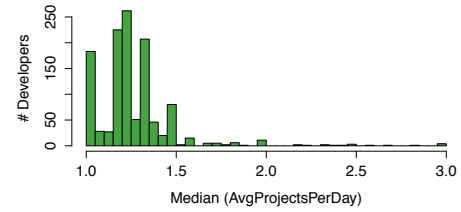

Figure 7: Distribution of median(AvgProjectsPerDay).

them switch between projects; and their perceptions of the impacts of contributing to multiple projects.[7] We received 128 responses (15% response rate). We iterated through the open-ended responses using grounded theory methods [12], to categorize them and identify themes. The survey participants reported development experience was 17.2 years on average (median 15; range 7 to 40), while their GitHub experience was 5.9 years on average (median 6; range less than 1 to since GitHub was founded).

4. RESULTS

4.1 RQ1: Multitasking, summarized.

4.1.1 Amount of Multitasking

Do developers multitask? We examined developers' commit activity and also asked survey participants to report the number of projects they contribute to in an average day and week. From the repository analysis, we found that multitasking across projects over the course of a week is not uncommon: developers contributed to multiple projects in 37% of the developer–weeks in our dataset (or 78,562 out of 132,277). The answer is not quite as clear for daily multitasking. Consider the distribution of developers' median AvgProjectsPerDay (Figure 7): for each developer, the median is computed across all weeks when they contributed to multiple projects. For example, a developer with median = 1 contributes to a single project per day on average, at least half of the weeks (since AvgProjectsPerDay cannot be smaller than 1 by construction), *i.e.*, she has a tendency to work sequentially. The distribution is right-skewed ($\gamma_1 = 3.08$). Even when contributing to multiple projects during a week, some developers still frequently focus on only few projects each day: 15% (183 out of 1,193) had median = 1; 95% (1,129 out of 1,193) did not contribute (on average) to more than 1.5 projects each day. Still, almost all developers (98%; 1,165 out of 1,193) contributed to multiple projects per day at least once.

Considering all projects, not just those hosted on GitHub, our survey participants reported contributing to an average of 2.7 projects per day (median 2; range 0–10). Over the course of a week, respondents stated that they contribute to an average of 6 projects (median 5; range 0–30).

Does within-day multitasking scale with the number of projects per week? We investigated whether contributing to more projects each week is also associated with contributing to more projects daily (therefore with more within-day focus switches), as one would expect. Figure 8 summarizes the distribution of AvgProjectsPerDay for different values of Projects (each beanplot shows the distribu-

---

[7] A complete list of questions is available at `http://kblincoe.github.io/survey/Focus_Shifting_Survey.pdf`.
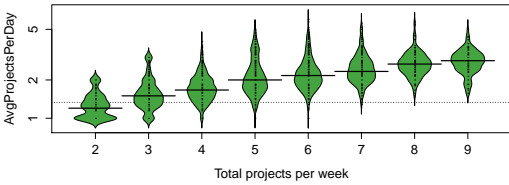
Figure 8: Beanplots showing the distributions of AvgProjectsPerDay, for different values of Projects. The solid horizontal lines represent the medians in each group. The dotted horizontal line is the overall median. Log y-axis.
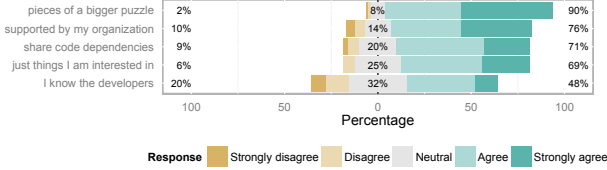

Figure 9: Reasons for contributing to multiple projects. Survey responses on how projects switched between are related.

tion of AvgProjectsPerDay for all users over all weeks). Our finding is that *developers do significantly more daily context switches as they participate in more projects per week.* All groups are significantly different pairwise using the $\widetilde{\mathbf{T}}$ [28] procedure at 95% confidence level, and higher Projects corresponds to higher AvgProjectsPerDay. For example, developers contributing to a median of 5 projects per week switch between a median of 2 each day, although in fact they could have focused on only one each day.

4.1.2 Reasons for Multitasking

Our survey investigated reasons for working on multiple projects. We asked the respondents to indicate, using a five-point Likert-type scale, their agreement to a number of statements about how the projects they contribute to are related. As shown in Figure 9, interdependencies, personal interest, and social relationships were all stated as strong reasons for contributing to multiple projects. Participants were also able to provide other ways the projects they contribute to are related: common responses were that they are an end user of the software tool and want to fix bugs impacting it, and that they contribute to a mix of projects due to their job as well as their personal interests.

We also asked about reasons for switching between projects when working on multiple projects. The strongest were all related to a need being identified in another project, either because of dependencies, newly created issues, or a request from another developer. The summarized responses are shown in Figure 10. Again, there was space for participants to describe other ways they schedule their work on multiple projects. The most common write-in answer, mentioned by five respondents, was that they simply liked to change focus for the sake of working on something different. For example, one participant said "I find it easier to produce quality code if I break up long stretches working on a single codebase with time spent on a completely different problem. It freshens up the mind." [P57] This response indicates that increased motivation could be a positive reason to multitask.

4.1.3 Productivity Effects

Is multitasking associated with more outputs produced per unit time? As a preliminary quantitative
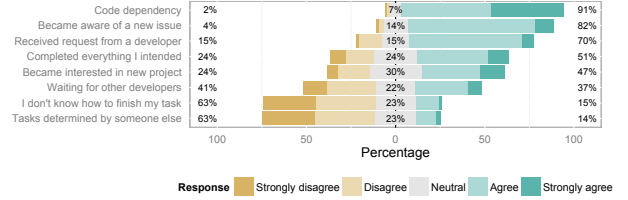

Figure 10: Reasons for switching between projects. Survey responses on why developers switch from one project to another when working on multiple projects.
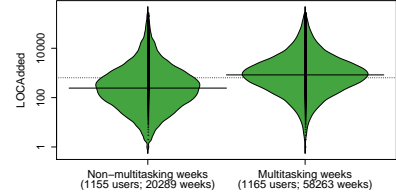

Figure 11: Beanplots comparing the distributions of LOCAdded during periods of sequential working (left) and multitasking (right). Solid horizontal lines: medians in each group. Dotted horizontal line: overall median. Log y-axis.

analysis, we found that more focus switches are associated with higher LOCAdded. Splitting the data into two groups, weeks of sequential work (AvgProjectsPerDay=1) and weeks with multitasking (Figure 11), the WMW rank-sum test reveals that periods of multitasking are more productive in terms of LOCAdded per week ($p < 2.2$e-16; $\hat{\Delta} = 379$).

Still, this analysis does not consider many confounds. As a refinement, we turn to the multiple regression analysis. Table 2 shows the effects of our independent and control variables for the multitasking productivity model; the response is log(LOCAdded). In addition to the model coefficients and corresponding standard errors, the table shows the sum of squares, a measure of variance explained, for each variable. All predictors have been z-transformed to reduce potential collinearity in the interaction terms. The statistical significance is indicated by stars. The model fits the data well; it explains 35% of the variability in the data using only the fixed effects ($R_m^2 = 0.35$), and 55% when considering both the fixed and random effects ($R_c^2 = 0.55$).

We start by discussing the effects associated with our controls. As expected, Projects and Languages have significant, positive effects. *Developers who contribute to more projects and who use more programming languages are associated with more lines of code added per week.* Together, these two effects explain about 60% of the variance explained.

More interestingly, S_{Focus} has a strong, negative effect, accounting for 12% of the variance explained. Recall that S_{Focus} captures a developer's overall project focus at week level (the lower the value, the more of one's time goes into a single project; the higher the value, the more evenly one divides their attention across projects). *Developers keeping fewer projects as their focus, rather than spreading themselves thin, are associated with more LOC added per week.* S_{Language} paints a similar picture: focusing on fewer languages requires less context switching between them, and is associated with more LOC added per week.

AvgProjectsPerDay has a significant positive effect on LOCAdded (1% of the variance explained). *More within-day multitasking is associated with more LOC added per week.*

	Coeffs (Errors)	Sum Sq.
(Intercept)	$0.069\ (0.012)^{***}$	
GlobalTime	$-0.037\ (0.005)^{***}$	15.71^{***}
Projects	$0.263\ (0.006)^{***}$	2566.87^{***}
Languages	$0.549\ (0.004)^{***}$	11505.20^{***}
S_{Focus}	$-0.300\ (0.004)^{***}$	2757.75^{***}
S_{Language}	$-0.231\ (0.003)^{***}$	2354.39^{***}
AvgProjectsPerDay	$0.046\ (0.004)^{***}$	215.29^{***}
S_{Switch}	$0.225\ (0.004)^{***}$	2255.09^{***}
Projects:S_{Focus}	$0.032\ (0.003)^{***}$	0.16^{***}
Languages:S_{Language}	$-0.045\ (0.002)^{***}$	345.96^{***}
Projects:S_{Language}	$-0.021\ (0.003)^{***}$	68.66^{***}
Projects:AvgProjectsPerDay	$-0.024\ (0.003)^{***}$	128.60^{***}
S_{Language}:AvgProjectsPerDay	$0.026\ (0.003)^{***}$	59.86^{***}
Projects:S_{Switch}	$-0.104\ (0.003)^{***}$	404.80^{***}
AvgProjectsPerDay:S_{Switch}	$0.058\ (0.003)^{***}$	203.12^{***}
S_{Language}:S_{Switch}	$-0.029\ (0.003)^{***}$	35.99^{***}
AIC = 171919; BIC = 172105; LogLik = -85939		
Num. obs. = 78552; Num. groups: fUserID = 1193		

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 2: Multitasking productivity model. The response is log(LOCAdded) per week. $R_m^2 = 0.35$. $R_c^2 = 0.55$.

Day-to-day focus switching (S_{Switch}) also has a strong, significant effect (10% of the variance explained). Recall that S_{Switch} captures the diversity of day-to-day focus switches (higher values indicate more repetitive behavior). The positive coefficient suggests that *repetitive, predictable day-to-day switches are associated with more LOC added per week.* Psychology research suggests that repetitive work is associated with lower levels of arousal [51] and decreased productivity, with switching tasks being one source of stimulation [24] to counteract this effect [51]. Our result suggests that switching between projects on consecutive days is more "expensive" than it is stimulating: a more repetitive day-to-day work pattern is associated with more outputs produced.

4.1.4 Perceived Impacts of Focus Switching

While our repository analysis is limited to LOCAdded per week, our survey responses touch also on other impacts.

Positive Impacts. We asked developers about the impacts of contributing to multiple projects in parallel (responses summarized in Figure 12). The impact with the highest level of agreement (47%) was that contributing to multiple projects increases the chances of each project succeeding. There was a weak correlation between the number of projects a developer reported contributing to each week and their perceived impact of this behavior on project success (rho = 0.25, $p = 0.005$). This shows that the more cross-project focus switching a developer does, the more they believe this will increase the project's success. This could indicate that some developers, especially those who do multitask frequently, are aware of the knowledge transfer benefits [29]. In fact, many respondents described this benefit, *e.g.*, "Working on different projects, especially when the ecosystem (programming language, runtime environment) varies greatly, increases the chance to take advantage of things learned from one project in another." [P7]

40% of respondents also agreed that when working on multiple projects they are able to resolve more issues. Respondents noted several reasons for this, such as the reduced wait time when working on multiple projects. For example, "I can switch to another project when I get stuck and need some time to work out a good solution" [P96], and "I don't get blocked [waiting] on a dependency (by contributing to the
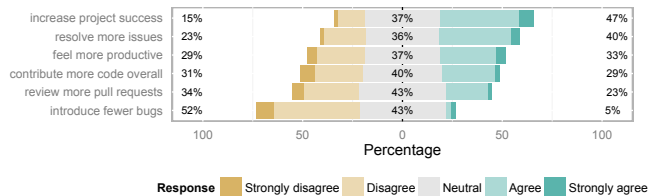

Figure 12: Perceived impacts of multitasking. Survey responses on the impacts of contributing to multiple projects in parallel (compared to focusing on fewer projects).

	Day			Week		
Projects	> 1	1	U	> 2	1-2	U
n	102	19		104	17	
Feel more productive	0.11	-0.37	1242^{*}	0.13	-0.59	1269.5^{**}
Contribute more code	0.02	-0.63	1377^{**}	0	-0.59	1219^{**}
Review more pull reqs	-0.04	-0.68	1348.5^{**}	-0.05	-0.71	1236^{**}
Resolve more issues	0.23	0.05	1091.5	0.27	-0.24	1177^{*}
Introduce fewer bugs	-0.49	-0.74	1144	-0.49	-0.76	1084
Increase project success	0.43	-0.05	1259^{*}	0.46	-0.29	1309.5^{***}

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 3: Impacts of multitasking segregated based on respondents who reported multitasking (>1 project per day or >2 projects per week) compared to all other responses. Table shows the mean response (and Mann-Whitney U) where responses were coded as follows: strongly disagree (-2), disagree (-1), neutral (0), agree (1), strongly agree (2).

dependency [myself])." [P6] Other positive impacts noted in the open-ended responses include an increased social network and reduced overall effort, for example, due to fixing a bug in the correct location rather than implementing multiple workarounds in many dependent projects.

Negative Impacts. However, participants were not very positive about the effects of multitasking on the quality of their code: only 5% agreed that they introduce *fewer* bugs when switching focus. In the open-ended responses, the most commonly cited negative impact was related to the cost of context switches; *e.g.*, one participant stated "Working on multiple projects also requires more 'context switching' mentally, which can be a drag on productivity." [P107]

Confusion. For three of the statements relating to the impact on productivity ("contribute more code"—the response variable we studied also quantitatively; "review more pull requests"; and "feel more productive"), there was a high amount of variance in the responses: some believe they are more productive while working on multiple projects, while nearly as many believe that they are not. We did observe a difference in responses based on how many projects developers reportedly contribute to. Comparing the responses of those who switch focus across many projects, with those that do not (1 project per day, or 1-2 projects per week), we find that the more frequent focus switchers are more likely to report positive benefits from this multitasking (Table 3).

4.2 RQ2: Limits on multitasking.

4.2.1 Quantitative Analysis

While the discussion above paints a general picture of the directionality of the various multitasking effects on outputs produced per unit time, the model in Table 2 also yields interesting two-way interactions. We illustrate the interpretation of such interaction terms with the example Projects—
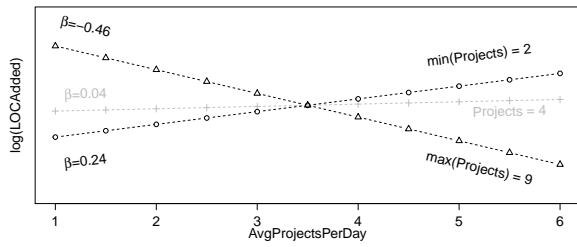
Figure 13: Interaction Projects (week) : AvgProjectsPerDay

AvgProjectsPerDay. Keeping all other variables constant, the interaction can be described by:

$$\text{LOCAdded} = \alpha_1 \cdot \text{Projects} + \alpha_2 \cdot \text{AvgProjectsPerDay}$$
$$+ \alpha_3 \cdot \text{Projects} \cdot \text{AvgProjectsPerDay}, \quad (4)$$

where α_1, α_2, and α_3 are the coefficients in the model.

To understand this interaction, we follow the standard practice of varying one term while holding the other constant, and vice versa. For example, fixing Projects allows us to examine the effects of increasing AvgProjectsPerDay on the response (*what happens to LOC added per week as there is more daily multitasking?*), for different levels of Projects. We illustrate this process in Figure 13, where the two regression lines shown in black correspond to the extremes of Projects (2 and 9 in our data set). When Projects is at its lowest (2), increasing AvgProjectsPerDay is always associated with an increase in the response, since the slope is positive ($\beta = 0.33$). In other words, more daily multitasking is associated with more LOC added per week, for developers that only take on 2 projects per week. At the other end of the scale, when Projects = 9, AvgProjectsPerDay is always associated with a decrease in the response, since the slope is negative. In other words, the more daily multitasking, the fewer LOC added per week, if developers take on too many different projects that week. Resolving $\alpha_2 + \alpha_3 \text{Projects} = 0$ allows us to estimate the point of diminishing return, *i.e.*, Projects = 4 (horizontal regression line shown in gray). Combining the two findings, we conclude: *higher average daily multitasking is associated with more outputs produced per week (albeit with diminishing returns), as long as developers don't take on more than 4 projects per week.*[8]

Similarly, we express tradeoffs between the other dimensions of multitasking and focus switching by investigating the interaction between two other pairs: (1) AvgProjectsPerDay and S_{Switch} and (2) Projects and S_{Switch}.

The interaction between AvgProjectsPerDay and S_{Switch} reveals that *higher average daily multitasking is associated with more LOC added per week when one's day-to-day focus switching is more repetitive.* This suggests that one can handle more complexity each day, as long as one doesn't have to also handle more complexity day-to-day.

The interaction between Projects and S_{Switch} shows that *increasing the weekly total number of projects is associated with increases in outputs produced, when developers don't follow too repetitive day-to-day patterns.* This suggests that taking on many projects can prevent boredom and be beneficial; however, if switching between them becomes too repetitive, productivity decreases.

---

[8]We find a similar result at month level: increases in AvgProjectsPerDay are associated with increases in LOCAdded per week, below 5 projects per month in total.
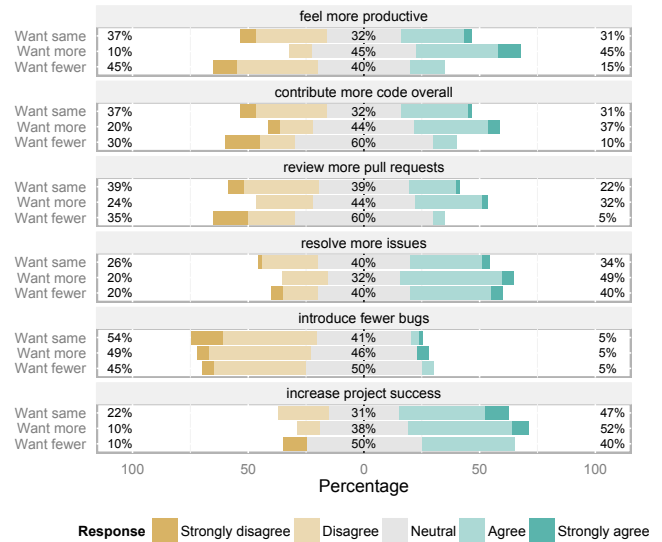

Figure 14: Survey responses on the impacts of multitasking, grouped based on the respondents' desire to increase/decrease multitasking. Those who wish to increase multitasking are more positive about the benefits.

4.2.2 Developer Perceptions on Limits

Complementary to this repository analysis, we asked the survey respondents if they would prefer to contribute to more projects, fewer projects, or their current number of projects. Participants are most likely to want to increase the number of projects they contribute to. Only 14.6% of respondents indicated that they would prefer to contribute to fewer projects (46.9% to more; 38.5% current amount).

However, a Spearman's rank correlation test shows the number of projects a developer already contributes to in an average day and week do not correlate with their desire to contribute to more, fewer, or the same number of projects (day: rho = 0.10, p = 0.29; week: rho = 0.01, p = 0.86). Developers who contribute to the most projects are no more likely to want to contribute to fewer and vice versa. This could imply that developers do not have a good sense of when to stop multitasking, despite the negative effects to productivity outlined by theory and confirmed by our model.

However, we find that *participants who want to contribute to more projects are more optimistic about the impacts of multitasking.* The exception is the impact to code quality (bugs introduced), to which all participants responded equally negatively. Figure 14 shows the survey responses relating to the impacts of multitasking, grouped by the participants' reported desire to contribute to more, fewer, or their current number of projects. For each statement, three bars show the aggregated responses for each of these three groups of participants (want to contribute to more/fewer/same number of projects). The responses for the "want more" group are skewed to the right compared to the other two groups (for all statements other than "introduce fewer bugs"), showing the responses from this group are more positive.

5. DISCUSSION
5.1 Implications for Software Practice

We found that developers do a significant amount of multitasking across projects over the course of a week. Devel-

opers indicated that the most common reason for this multitasking is the interrelationships between projects. Focus switches often occur because of a need identified on another project, indicating that such focus switching is not planned, but occurs because of necessity.

Multitasking is beneficial. Yet, we found that despite these unplanned interruptions, developers who switch focus across projects generate more outputs than those who do not. Switching projects allows developers to make more efficient use of their time when one project reaches a lull, and provides them with opportunities for learning and knowledge transfer, as discussed in the Background section. However, perceptions towards multitasking and its effect on productivity are varied. Those who do multitask are more likely to believe multitasking can be beneficial. Software teams could benefit from practices and tools that support team/work assignment that account multitasking.

There are limits on multitasking. However, we found that limits exist on the amount of focus switching one can participate in before their productivity begins to decline. Yet, developers lack a good sense of when to stop multitasking. Nearly half of all survey participants would prefer to contribute to more than their current number of projects. We saw no differences in response based on the number of projects a developer already contributes to, indicating they do not seem to be aware of the limits of multitasking. Hence, developers could benefit from tools which predict and notify of approaching such limits based on their work patterns.

5.2 Research Agenda

While our investigation resulted in significant new insights on multitasking and focus switching across projects, there are still many unanswered research questions, opening some interesting avenues for future research.

Predicting productivity impacts. Our model can identify the effects of various features of multitasking on developer productivity. A logical next step would be to use these models to predict the productivity of individual developers based on their past contribution patterns. A predictive model could be the foundation of a tool that can monitor developer focus switching behavior, and provide warnings to developers (or their managers) when risky patterns emerge that could indicate productivity will be negatively impacted.

Further investigation of non-multitaskers. Our survey participants who reported no or very little multitasking, were more pessimistic about the impacts of multitasking on their productivity and project success. Future work can investigate why developers choose not to multitask. Do non-multitaskers contribute to projects with fewer dependencies to other projects, resulting in a reduced need for multitasking? If the projects do not have fewer dependencies, how do non-multitaskers avoid switching focus when these dependencies necessitate changes in another project? Do they simply create an issue on that project and wait for someone else to make the change, resulting in reduced productivity? The answers to these questions can result in the conceptualization of methods and tools designed to change the perceptions and habits of these developers, to promote an increase in multitasking to reap the associated productivity benefits.

Further investigation of reasons for focus switches. We identified limits to multitasking. Contributing to more projects per week is associated with more outputs produced, as long as developers do not work on too many projects in a single day. Future work can investigate reasons for this finding. Perhaps the reasons for focus switching is different for those who contribute to many projects per day, compared to those who contribute to only a small number of projects per day. The differences between these two groups may provide additional insight, revealing which focus switching reasons are more likely to result in increased productivity. With a greater understanding of the beneficial focus switches, tools can be devised to help developers distinguish between the various interruptions they face.

5.3 Threats to Validity

All developers in our study are highly prolific contributors, and our survey respondents were self-selected. Thus, our results may not generalize beyond this research setting. The thematic analysis of the survey was performed by only one person, introducing a possible risk of unreliable results.

We define projects by aggregating repositories related by name. Given the ease with which different projects can depend on each other on GITHUB, determining the exact boundaries of a project is difficult. Our definition may overly combine repositories, but we believe it makes context switches more likely to be between distinct entities. The results may be different if a coarser or finer grained definition of project were to be used.

The aliased user identities refer to the original author of the commit, while the dates chosen are from the committer fields, which mark when the change was finally added to the main repository. These fields can be different if the commit was originally created in a fork. We found the author and committer to be different in only 1.78% of commits, so we do not believe this will significantly affect the results.

We considered contributions as commits only. Other types of contributions could also be considered in future work.

The projects we extracted were based on data from GITHUB and GHTorrent, which limits the generalizability.

Finally, measuring productivity as LOC added per week, or any other count of user interaction with the artifact, is certainly somewhat naive and incomplete. However, it is neither unprecedented nor unjustified by prior research [57].

6. CONCLUSIONS

Programming is multi-faceted, inherently involving context-switching. With the advent of ecosystems like GITHUB, another tier of context-switching becomes possible: switching between projects. Using a mixed-methods approach (survey+quantitative analysis of mined data) we study this phenomenon. We find that up to a certain point, switching between projects is associated with productivity increases; but beyond that, it appears that overall output of developers declines. Yet, the survey responses indicate that developers seem to believe it's always better to work on ever-more projects, regardless of how many they already work on! Our work is *actionable*: managers and developers can take advantage of our findings to manage both the number of projects they work on and the intensity of their context-switching.

Acknowledgements

7. REFERENCES

[1] R. F. Adler and R. Benbunan-Fich. Juggling on a high wire: Multitasking effects on performance. *International Journal of Human-Computer Studies*, 70(2):156–168, 2012.

[2] E. M. Altmann and J. G. Trafton. Memory for goals: An activation-based model. *Cognitive Science*, 26(1):39–83, 2002.

[3] J. R. Anderson. Human symbol manipulation within an integrated cognitive architecture. *Cognitive science*, 29(3):313–341, 2005.

[4] S. Aral, E. Brynjolfsson, and M. V. Alstyne. Information, technology, and information worker productivity. *Information Systems Research*, 23(3-part-2):849–867, 2012.

[5] E. T. Barr, C. Bird, P. C. Rigby, A. Hindle, D. M. German, and P. Devanbu. Cohesive and isolated development with branches. In *FASE*, pages 316–331. Springer, 2012.

[6] R. Benbunan-Fich. An entropy index for multitasking behavior. In *ICIS*. AIS, 2011.

[7] K. Blincoe, F. Harrison, and D. Damian. Ecosystems in GitHub and a method for ecosystem identification using reference coupling. In *MSR*, pages 202–211. IEEE, 2015.

[8] J. P. Borst, N. A. Taatgen, and H. van Rijn. The problem state: a cognitive bottleneck in multitasking. *Journal of Experimental Psychology: Learning, memory, and cognition*, 36(2):363, 2010.

[9] J. P. Borst, N. A. Taatgen, and H. van Rijn. What makes interruptions disruptive? a process-model account of the effects of the problem state bottleneck on task interruption and resumption. In *CHI*, pages 2971–2980. ACM, 2015.

[10] C. Casalnuovo, B. Vasilescu, P. Devanbu, and V. Filkov. Developer onboarding in GitHub: The role of prior social links and language experience. In *FSE*, pages 817–828. IEEE, 2015.

[11] M. Cataldo and J. D. Herbsleb. Coordination breakdowns and their impact on development productivity and software failures. *IEEE TSE*, 39(3):343–360, 2013.

[12] J. Corbin and A. Strauss. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Sage Publications, 2014.

[13] L. Dabbish, G. Mark, and V. M. González. Why do i keep interrupting myself? environment, habit and self-interruption. In *CHI*, pages 3127–3130. ACM, 2011.

[14] L. Dabbish, C. Stuart, J. Tsay, and J. Herbsleb. Social coding in GitHub: transparency and collaboration in an open software repository. In *CSCW*, pages 1277–1286. ACM, 2012.

[15] S. Easterbrook, J. Singer, M.-A. Storey, and D. Damian. Selecting empirical methods for software engineering research. In *Guide to Advanced Empirical Software Engineering*, pages 285–311. Springer, 2008.

[16] S. J. Gilbert and T. Shallice. Task switching: A PDP model. *Cognitive Psychology*, 44(3):297–337, 2002.

[17] V. M. González and G. Mark. Managing currents of work: Multi-tasking among multiple collaborations. In *ECSCW*, pages 143–162. Springer, 2005.

[18] A. J. Gould. *What makes an interruption disruptive? Understanding the effects of interruption relevance and timing on performance*. PhD thesis, UCL, 2014.

[19] G. Gousios and D. Spinellis. GHTorrent: Github's data from a firehose. In *MSR*, pages 12–21. IEEE, 2012.

[20] A. Hars and S. Ou. Working for free? Motivations of participating in open source projects. In *HICSS*, pages 9–pp. IEEE, 2001.

[21] J. D. Herbsleb and A. Mockus. An empirical study of speed and communication in globally distributed software development. *IEEE TSE*, 29(6):481–494, 2003.

[22] S. Jenkins. Concerning interruptions. *Computer*, (11):116–114, 2006.

[23] P. C. Johnson. Extension of nakagawa & schielzeth's r^2_{GLMM} to random slopes models. *Methods in Ecology and Evolution*, 5(9):944–946, 2014.

[24] D. Kahneman. *Attention and effort*. Prentice-Hall, 1973.

[25] E. Kalliamvakou, G. Gousios, K. Blincoe, L. Singer, D. M. German, and D. Damian. The promises and perils of mining GitHub. In *MSR*, pages 92–101. ACM, 2014.

[26] M. Kersten and G. C. Murphy. Using task context to improve programmer productivity. In *FSE*, pages 1–11. ACM, 2006.

[27] A. J. Ko, R. DeLine, and G. Venolia. Information needs in collocated software development teams. In *ICSE*, pages 344–353. IEEE, 2007.

[28] F. Konietschke, L. A. Hothorn, E. Brunner, et al. Rank-based multiple test procedures and simultaneous confidence intervals. *Electronic Journal of Statistics*, 6:738–759, 2012.

[29] A. Lindbeck and D. J. Snower. Multitask learning and the reorganization of work: from Tayloristic to holistic organization. *Journal of Labor Economics*, 18(3):353–376, 2000.

[30] G. Madey, V. Freeh, and R. Tynan. The open source software development phenomenon: An analysis based on social network theory. In *AMCIS*, page 247, 2002.

[31] G. Mark, V. M. Gonzalez, and J. Harris. No task left behind? examining the nature of fragmented work. In *CHI*, pages 321–330. ACM, 2005.

[32] J. Marlow and L. Dabbish. Activity traces and signals in software developer recruitment and hiring. In *CSCW*, pages 145–156. ACM, 2013.

[33] J. Marlow, L. Dabbish, and J. Herbsleb. Impression formation in online peer production: Activity traces and personal profiles in GitHub. In *CSCW*, pages 117–128. ACM, 2013.

[34] N. McDonald and S. Goggins. Performance and participation in open source software on GitHub. In *CHI*, pages 139–144. ACM, 2013.

[35] A. N. Meyer, T. Fritz, G. C. Murphy, and T. Zimmermann. Software developers' perceptions of productivity. In *FSE*, pages 19–29. ACM, 2014.

[36] S. Nakagawa and H. Schielzeth. A general and simple method for obtaining r^2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2):133–142, 2013.

[37] M. B. O'Leary, M. Mortensen, and A. W. Woolley. Multiple team membership: A theoretical model of its

effects on productivity and learning for individuals and teams. *Academy of Management Review*, 36(3):461–478, 2011.

[38] C. Parnin. A cognitive neuroscience perspective on memory for programming tasks. *Programming Interest Group*, page 27, 2010.

[39] C. Parnin and R. DeLine. Evaluating cues for resuming interrupted programming tasks. In *CHI*, pages 93–102. ACM, 2010.

[40] C. Parnin and S. Rugaber. Resumption strategies for interrupted programming tasks. *Software Quality Journal*, 19(1):5–34, 2011.

[41] C. Parnin and S. Rugaber. Programmer information needs after memory failure. In *ICPC*, pages 123–132. IEEE, 2012.

[42] C. J. Parnin. *Supporting Interrupted Programming Tasks with Memory-Based Aids*. PhD thesis, Georgia Institute of Technology, 2014.

[43] J. K. Patel, C. Kapadia, and D. B. Owen. *Handbook of statistical distributions*. M. Dekker, 1976.

[44] D. E. Perry, N. A. Staudenmayer, and L. G. Votta. Understanding and improving time usage in software development. *Software Process*, 5:111–135, 1995.

[45] D. Posnett, R. D'Souza, P. Devanbu, and V. Filkov. Dual ecological measures of focus in software development. In *ICSE*, pages 452–461. IEEE, 2013.

[46] C. Rosen. The myth of multitasking. *The New Atlantis*, 20(Spring):105–110, 2008.

[47] P. J. Rousseeuw and C. Croux. Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424):1273–1283, 1993.

[48] J. S. Rubinstein, D. E. Meyer, and J. E. Evans. Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4):763, 2001.

[49] D. D. Salvucci, N. A. Taatgen, and J. P. Borst. Toward a unified theory of the multitasking continuum: From concurrent performance to task switching, interruption, and resumption. In *CHI*, pages 1819–1828. ACM, 2009.

[50] H. Sanchez, R. Robbes, and V. M. Gonzalez. An empirical study of work fragmentation in software evolution tasks. In *SANER*, pages 251–260. IEEE, 2015.

[51] K. H. Teigen. Yerkes-Dodson: A law for all seasons. *Theory & Psychology*, 4(4):525–547, 1994.

[52] B. Vasilescu, V. Filkov, and A. Serebrenik. Perceptions of diversity on GitHub: A user survey. In *CHASE*, pages 50–56. IEEE, 2015.

[53] B. Vasilescu, D. Posnett, B. Ray, M. G. J. van den Brand, A. Serebrenik, P. Devanbu, and V. Filkov. Gender and tenure diversity in GitHub teams. In *CHI*, pages 3789–3798. ACM, 2015.

[54] B. Vasilescu, A. Serebrenik, and V. Filkov. A data set for social diversity studies of GitHub teams. In *MSR*, pages 514–517. IEEE, 2015.

[55] G. M. Weinberg. *Quality Software Management, 1: Systems Thinking*. Dorset House Publishing, 1992.

[56] Q. Xuan, P. T. Devanbu, and V. Filkov. Converging work-talk patterns in online task-oriented communities. *arXiv preprint arXiv:1404.5708*, 2014.

[57] Q. Xuan, A. Okano, P. Devanbu, and V. Filkov. Focus-shifting patterns of OSS developers and their congruence with call graphs. In *FSE*, pages 401–412. ACM, 2014.

[58] F. Zhang, F. Khomh, Y. Zou, and A. E. Hassan. An empirical study of the effect of file editing patterns on software quality. *Journal of Software: Evolution and Process*, 26(11):996–1029, 2014.

[59] A. Zika-Viktorsson, P. Sundström, and M. Engwall. Project overload: An exploratory study of work and management in multi-project settings. *International Journal of Project Management*, 24(5):385–394, 2006.

[60] L. Zou and M. W. Godfrey. An industrial case study of program artifacts viewed during maintenance tasks. In *WCRE*, pages 71–82. IEEE, 2006.