

UNIVERSITY OF TARTU
Institute of Computer Science
Computer Science Curriculum

Oliver Viljamaa

Measuring the Effect of User-Perceived
Load Metrics on Conversion Rate in the
Context of TransferWise

Bachelor's Thesis (9 ECTS)

Supervisor: Prof. Marlon Dumas

Supervisor: Mihkel Nõges, MSc

Tartu 2018

Measuring the Effect of User-Perceived Load Metrics on Conversion Rate in the Context of TransferWise

Abstract:

The user experience of loading a web page can have an effect on business metrics. One of these metrics is conversion rate: the percentage of visitors performing a desired action when reaching a given state such as visiting a web page. The effect of the page load experience on conversion rate has been widely studied. However, the observed effects vary significantly across studies due to the use of different load metrics and due to differences in the populations. Therefore, to ensure reliable results, this effect should be measured for each site separately. The problem addressed by this work concerns TransferWise not having measured this effect and therefore lacking insights on whether or not to optimise page load experience, to what extent, and with respect to which metrics.

In this setting, the contributions of this work are threefold: (i) a precise definition of a set of user-perceived load metrics for TransferWise, together with a solution for collecting these metrics within TransferWise production environment; (ii) a dashboard application for measuring the effect of load metrics on conversion rate; and (iii) an initial analysis of the relations between load metrics and conversion rate for different acquisition channels.

The results of the analysis confirm the expected correlations between increase in load time and decrease in conversion rate, and show that the extent of this relation varies by load metric and by acquisition channel.

Keywords:

Page load time, user-perceived load metrics, web performance, conversion rate, TransferWise

CERCS: P170 Computer science, numerical analysis, systems, control

Kasutaja tajutava laadimise mõõdikute mõju konversioonimäärale mõõtmine ettevõtte TransferWise kontekstis

Lühikokkuvõte:

Kasutaja kogemus veebilehe laadimisel võib mõjutada ärimõõdikuid. Üks nendest mõõdikutest on konversioonimäär: protsent külastajatest, kes sooritavad antud olukorda jõudes, näiteks veebilehte külastades, soovitud toimingut. Veebilehe laadimiskogemuse mõju konversioonimäärale on laialdaselt uuritud. Täheldatud mõju erineb aga uurimuste vahel märkimisväärselt, tingituna erinevatest laadimismõõdikutest ja uuritud populatsioonide erinevustest. Seetõttu tuleks usaldusväärsete tulemuste saamiseks mõõta mõju iga lehekülje jaoks eraldi. Antud töös käsitletakse probleemi, et TransferWise ei ole seda mõju mõõtnud ning seetõttu puudub ülevaade, kas lehe laadimiskogemust optimeerida või mitte, missuguses ulatuses ja missuguste mõõdikute osas.

Selle taustal on käesoleval tööel kolm panust: (i) kasutaja tajutava laadimise mõõdikute defineerimine TransferWise'i kontekstis koos lahendusega nende mõõdikute kogumiseks TransferWise'i tootmiskeskkonnas; (ii) näidikulaua rakenduse implementatsioon laadimismõõdikute mõju konversioonimäärale mõõtmiseks; ja (iii) esialgne laadimismõõdikute ja konversioonimäära seose analüüs erinevate kanalite lõikes.

Analüüsi tulemused kinnitavad oodatud korrelatsioone laadimisaja suurenemise ja konversioonimäära vähenemise vahel ning näitavad, et selle seose ulatus erineb laadimismõõdikute ja kanalite lõikes.

Võtmesõnad:

Veebilehe laadimisaeg, kasutaja tajutava laadimise mõõdikud, veebijõudlus, üleminekumäär, TransferWise

CERCS: P170 Arvutiteadus, arvutusmeetodid, süsteemid, juhtimine (automaatjuhtimisteooria)

Contents

1	Introduction	5
2	Background	6
2.1	User-perceived load metrics	6
2.2	Conversion rate	9
2.3	Effect of the load experience on conversion rate	10
2.4	Organisational context	12
3	Tracking load metrics	13
4	Measurement dashboard	14
5	Effect of load metrics on conversion rate	15
6	Conclusion	16
	References	17
	Licence	22

1 Introduction

An average person visits tens to hundreds of web pages a day, including news sites, work applications, search engines, and social networks. The interactions with these pages have created expectations for the users on how a page should look, feel, and act. These expectations also regard load experience. Furthermore, the expectations for a faster load experience are increasing year by year [Aka15; Eve13]. In addition to the benefit of fulfilling the user expectations, the load experience can affect business metrics such as conversion rate, the percentage of users who take the desired action. As conversion rate optimisation results in more desired actions for the same amount of traffic, companies such as Walmart [CKR12] and Mozilla [Cut10] have studied the effect of the load experience on conversion. However, the different metrics used for measuring the load experience, lack of user segmentation, and different populations do not allow the results to be generalised.

This thesis frames this problem in the context of TransferWise – a fast-growing financial technology company with a diverse customer base. As the effect of the load experience on conversion rate has not been studied in the organisation, there is a lack of knowledge on whether or not the teams should prioritise optimising the load experience compared to other feature work, which metrics to optimise for, and which user segments the optimisation would affect and by how much.

To address the problem, three contributions are made in this work:

1. a set of user-perceived load metrics are defined specifically for TransferWise and a solution for collecting these metrics within TransferWise production environment is implemented;
2. a dashboard application for measuring the effect of load metrics on conversion rate is implemented;
3. an initial analysis of the relations between different load metrics and conversion rate for different acquisition channels is conducted;

The remainder of the thesis is structured as follows. Chapter 2 explains the background of user-perceived load metrics, conversion rate, the effect of the load experience on conversion rate, and TransferWise as organisational context. In Chapter 3, the load metrics are defined for TransferWise pages and the implementation of the data collection solution for the production environment is explained. Chapter 4 lists the requirements and explains the implementation of the dashboard application created to measure the effect of load metrics on conversion rate. The results, implications, and limitations of the initial analysis are presented in Chapter 5. Finally, Chapter 6 is the conclusion of the thesis.

2 Background

This chapter explains the background of the concepts worked with in this thesis. The first section defines the load experience, compares metrics to represent it, and gives an overview of the methods and tools to measure it. The second section defines conversion rate, shows why it is an appropriate metric for measuring the success of web pages, and gives a comparison of the tools to analyse it. In the third section, the methods, results, and limitations of previous studies on the effect of the load experience on conversion rate are presented, along with an overview of tools with the capability to measure the effect. The final section describes TransferWise, its methods for measuring conversion rate, and how the effect of the load experience has been measured in the past.

2.1 User-perceived load metrics

Page load is the whole experience of waiting to load a web page. That includes the time spent in the server to send the requested HTML to the user and the time spent in the client to process the received HTML and the assets it requests [Newb]. According to the *performance golden rule* formulated based on data from the HTTP Archive, ‘80-90% of the end-user response time is spent on the front-end’ [Sou12]. Therefore, measuring response times on the server would not give a realistic view of the load experience, and it is mostly measured on the client-side. To make the experience quantifiable, it is often reduced to a single metric of page load time defined as the time it took to load the page. However, that time can have different meanings.

The browsers dispatch events to signal occurrences such as network activity or user interaction [Ms2+15]. These events can be observed using event listeners. In general, studies and tools rely on the values of either the `load` or the `DOMContentLoaded` event [Newb]. The `load` event marks the time at which all the required resources have been downloaded and processed [Gri]. `DOMContentLoaded` marks when the HTML document has been completely loaded and parsed [MDNa]. The reasons for the popularity of these events are historical and related to their ease of use. In the past, when web pages were mostly HTML and images, `window.onload` that triggered on `load` was appropriate for measuring the user-perceived load experience [Sou13]. The event has been standard across all browsers since early versions [MDNb], and `DOMContentLoaded` has good support as well [MDNa], making them easy and safe to use. However, for modern web pages making heavy use of JavaScript, AJAX, and progressive rendering, `DOMContentLoaded` fires too early and `load` happens too late [Den]. The fact that `load` does not correlate with what

the user perceives as loaded is demonstrated well by comparing the load experiences of Amazon and Gmail [Sou13]. Amazon is almost completely rendered at 2 seconds, but load does not happen till 5.2 seconds, as seen on Figure 1. On the other hand, Gmail load happens after 3.3 seconds, but the page is mostly rendered only at 4.8 seconds, as seen on Figure 2. Thus, these traditional metrics do not reflect the user-perceived load experience accurately.




Figure 1. Amazon 90% rendered (2 seconds) and load (5.2 seconds) [Sou13].

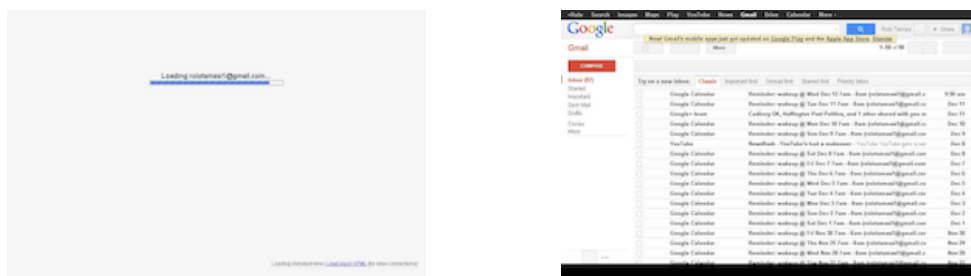


Figure 2. Gmail load (3.3 seconds) and 90% rendered (4.8 seconds) [Sou13].

There have been attempts to define a metric that effectively represents the user-perceived load experience. However, creating one such metric that can be standardised across browsers has been deemed difficult, and a set of metrics has been recommended instead [Pan16]. The furthest advancements on this topic have been made by engineers at Google, who have created a guide to what they call *user-centric performance metrics* [Wal]. The guide formulates questions for different stages in the load experience and brings out what metrics answer them, shown in Table 1. As Google has participated in creating the W3C Paint Timing working draft [Pan17], the definitions for the metrics in the guide and draft align. *First Paint* (FP) is defined as the time when the browser renders anything other than the default background paint. *First Contentful Paint* (FCP) marks the time when the browser renders any text, image, non-white canvas, or SVG. Since *First Meaningful Paint* (FMP) and *Time to Interactive* (TTI) metrics depend on what is considered important on a page, defining the rules for the browsers for these metrics is difficult

and they have not been added to the Paint Timing draft yet. Accordingly, the guide leaves defining First Meaningful Paint to web developers knowing what is meaningful in the context of the page. It loosely defines Time to Interactive as the point at which the application is visually ready and capable of responding to user input. The metrics are shown as frames of a load experience for an example page on Figure 3.

Table 1. Questions and corresponding metrics [Wal].

Question	Metric(s)
Is it happening?	First Paint, First Contentful Paint
Is it useful?	First Meaningful Paint
Is it usable?	Time to Interactive

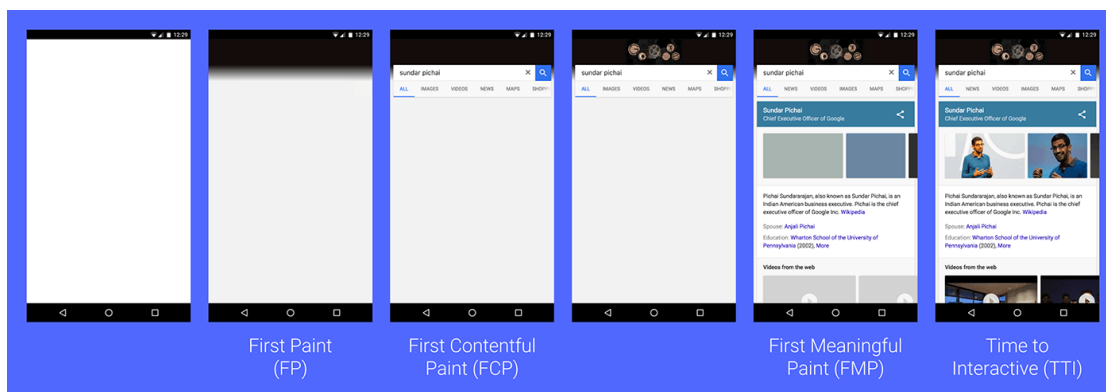


Figure 3. Load metrics as frames of a load experience [Wal].

Metrics representing the user-perceived load experience have been labeled *user-centric performance metrics* [Wal], *progressive web metrics* [Iri], *milestone timings* [Web], and *user-perceived wait time* [MAC17]. To put the emphasis on user perception and load, this work uses the term *user-perceived load metrics*. As these metrics reflect the user-perceived load experience more accurately, they are preferred to the traditional `load` and `DOMContentLoaded` events.

In order to gain actionable insights from load metrics on a web page, the metrics need to be monitored [Fri18]. Monitoring methods and tools generally fall into one of two categories: *synthetic monitoring* and *real user monitoring* (RUM) [Sta17]. Synthetic monitoring, also known as *active monitoring*, consists of running *synthetic tests* simulating the path an end user would take on the site at set intervals and monitoring the results [Eve15b]. Synthetic testing allows stress testing the application before launch [Sta17], competitive benchmarking [Eve15b], and digging deep into the load experience details [Zem]. Yet, due to a limited combination of

configurations, synthetic monitoring can not reflect the whole spectrum of user experiences [Eve15b]. Alternatively, real user monitoring, also known as *passive monitoring*, gathers data from every user across every network and geographic location as they move through the site [Eve15b]. Thus, it fully represents the spectrum of real user experiences. Real user data can give insights into user properties such as acquisition channel [Eve15b] and how users contribute to business metrics [Zem]. Active monitoring cannot give similar information due to its synthetic nature. Real user experiences are ultimately what matters, but synthetic monitoring can keep the bad experiences from happening in the first place [Wal17]. Therefore, the benefits of synthetic monitoring and RUM are complementary [Aka12], but only real user monitoring can be used when measuring the effect on business metrics.

Various tools for real user monitoring exist, with some of them supporting synthetic testing in addition to RUM. The tools include SOASTA mPulse [SOAb], SpeedCurve LUX [Spe], AppDynamics [App], DynaTrace [Dyn], NewRelic [Newa], and Pingdom [Pin]. In addition, the metrics can be measured for real users using open-source tools like Boomerang [SOAa] or manually. The metrics can then be sent as events to analytics tools such as Google Analytics [Goo] or Mixpanel [Mixa]. A number of the RUM tools also have built-in support for custom metrics using the *User Timing* API, a performance measurement interface. Alternatively, the metrics can be marked using the client-side snippet of the tool. Moreover, some of the tools allow adding custom front-end business metrics, such as conversion rate.

2.2 Conversion rate

There is a large number of business metrics to monitor and optimise for business success, for example registrations, active users, purchases, revenue, and conversion rate [Dat16]. The last, conversion rate, is a popular indicator for the success of web pages. It has been popularised by being the main metric measured when *A/B testing*: showing different variants of a page to different visitors to find out which variant is the most effective [SK13].

A *conversion* happens when a visitor takes the desired action that has a measurable value to the business. For example, this can mean a clickthrough to another page, a registration, a download, or a purchase [AGP12]. The definition of *conversion rate* is the percentage of visitors that take the desired action. It is calculated by dividing the number of converted visitors by the number of visitors in the same timeframe. Therefore, the formula for calculating conversion rate is the following:

$$\text{conversion rate} = \frac{\text{number of visitors that take the desired action}}{\text{number of visitors}} \quad (1)$$

As maximising the conversion rate results in a larger number of desired actions without having to attract more visitors, it is fitting metric to optimise for business success [Qua]. Compared to registrations, active users, purchases, and revenue, it has the benefit of taking the baseline traffic into account. For this reason, it can be used to measure the performance of the page in isolation. In addition, its flexibility allows finding out the effect for any subset of visitors, any desired action, and any conversion funnel leading to that action. The conversion funnel is the sequence of steps a visitor needs to go through before they can reach the conversion [Bul]. When investigating conversion with an end goal of a click on the landing page, a funnel may consist of two steps, the visit and the click. For an e-commerce site investigating conversion to a purchase by a new customer, the funnel could consist of page visit, registration, product page view, and purchase steps. In the last example, the registration step ensures that only new customers are reflected in the conversion rate to the last step. Based on the above, conversion rate is an appropriate business metric to use for measuring the success of web pages.

As conversion rate is an easily measured and effective metric, it is well-supported by analytics tools. For example, Google Analytics [Goo], Mixpanel [Mixb], and Heap Analytics [Hea] all support conversion rates and funnel analysis. The complications with analysing conversion often lie in gathering the data and analysing the results for different user segments. This may be a reason why companies such as Airbnb opt for building an internal analytics and A/B testing platform [Par17]. As the custom frameworks are built specifically for the needs of the businesses, they can be used for investigating the interplay of any metrics by any variable. This is needed because metrics such as conversion rate are determined by many factors.

2.3 Effect of the load experience on conversion rate

Load is one of the numerous factors contributing to the experience users have with a page. Hence, the effect of the load experience on business metrics such as revenue has been studied in the industry. As conversion rate optimisation ultimately leads to increasing revenue and reaching other business goals, the effect of the load experience on conversion rate has been studied. There are not many recent academic studies on this topic. Instead, the research has mostly been conducted by online businesses aiming to understand the trade-off between optimising the load experience and implementing other features [Bru09]. In addition, studies have

been carried out by companies selling performance tools and services with the goal of proving the importance of the load experience [Eve15a].

The research methods and the published results have evolved noticeably over time. One of the earliest reports on the effect of load originates from Amazon in 2006 [Lin06]. The company reported that every 0.1 second delay costs 1% in revenue. The figure is often quoted and has therefore become the standard expectation in the industry, albeit the method and result not being explained in detail. Conversely, Akamai published their research on the state of online retail performance in 2017 [Aka17] comprised of roughly 10 billion user visits from leading retail sites and presented as a comprehensible study.

Two approaches prevail in terms of study methods. Firstly, the studies bring out the difference between the peak conversion rate and the conversion rate for an experience that took x seconds longer. Whether the difference is relative to the peak conversion rate of the higher load time conversion rate differs by study and is not always clarified. Using that method, Firefox found a 1 second increase in the load time to decrease conversions by 2.7% [Cut10] in 2010, and for Walmart a 1 second improvement increased conversions up to 2% in 2012 [CKR12]. In a case study by SOASTA in 2015, pages 1 second faster experienced a 27% conversion rate increase [Eve15a], and findings by Akamai in 2017 showed a 1 second delay resulting in conversion rates up to 21.8% lower [Aka17]. The other method is comparing the average or median load time for conversions and non-conversions. Firefox found that the median for non-conversions was 75% slower than for conversions [Cut10] (2.1s vs 3.67s), Walmart saw a difference of 87% between averages (3.22s vs 6.03s) [CKR12], and the most recent study by Akamai presented converting pages being 26% faster than non-converting [Aka17].

These studies have mostly been carried out using traditional load metrics or simplistic heuristics, not a set of user-perceived load metrics. The fact that the same traditional metric may mark a different user-perceived point in a load experience for different pages makes the studies are hard to compare. As the studies use one metric instead of multiple, they also do not show which metric affects conversion rate most and what part of the experience should therefore be optimised. In addition, users are either not segmented or segmented only by device. Hence, the effect on conversion for segments defined by other parameters such as acquisition channel remains unstudied. Furthermore, the differences in the results prove the effect of the load experience can vary significantly between populations. This suggests that the most reliable way of finding out the effect of the load experience on conversion rate for a specific page is performing measuring for that page.

Measuring the effect of the load experience on conversion rate is mainly attempted

by performance tools. For example, conversion metrics are natively supported in SOASTA mPulse [SOAb], SpeedCurve LUX [Spe], AppDynamics [App], and DynaTrace [Dyn]. These metrics are defined on the front-end, for example a purchase button click or a HTTP request to the transaction URL. However, these metrics are not as reliable as back-end events that can be specified to trigger only on real validated occurrences. Although measuring user-perceived load metrics on universal analytics tools such as Google Analytics or Mixpanel requires more setup, it ensures the quality and alignment of tracked business metrics. Hence, analytics tools should be considered for measuring the effect of user-perceived load metrics on conversion rate.

2.4 Organisational context

TransferWise is a UK-based financial technology company founded in 2010 by Estonian expatriates Taavet Hinrikus and Kristo Käärman [Tra]. The mission of the company is ‘money without borders: instant, convenient, and eventually free’ [Tra17]. To achieve the mission, TransferWise offers a series of products, such as international money transfer, borderless accounts, and business payouts. The service has three million customers, with £2 billion transferred every month and £2 million in bank fees saved each day [Kää18]. TransferWise powers more than 750 currency routes [Tra18] and is run by over 1,000 employees in nine offices across the world [Tra18]. The organisational structure is flat, with most of the company organised as autonomous independent teams focused around specific customer problems [Pei15].

[content omitted due to containing confidential information]

3 Tracking load metrics

[content omitted due to containing confidential information]

4 Measurement dashboard

[content omitted due to containing confidential information]

5 Effect of load metrics on conversion rate

[content omitted due to containing confidential information]

6 Conclusion

The aim of this work was to solve the problem of teams at TransferWise not knowing the effect of user-perceived load metrics on conversion rate. The problem was solved by the three contributions made by the author. Firstly, user-perceived load metrics were defined specifically for TransferWise pages and collected within TransferWise production environment. Secondly, a dashboard application to measure the effect of the metrics on conversion rate was implemented. Thirdly, an initial analysis was conducted focused on the differences between load metrics and acquisition channels.

As a result of the metric collection and dashboard creation, teams at TransferWise can measure the effect of load metrics on conversion rate. Thus, they can learn whether to prioritise improving the load experience, which metrics to optimise for, and which user segments optimisations would affect.

The findings of the initial analysis of 5.6 million visitors supported the existing research by showing a connection between the load time increase and conversion rate decrease. The median load time for converted users was up to 23% lower and a load time increase of 1 second was found to increase conversion rate by up to 40% for certain metrics and channels. In terms of acquisition channels, social channel was shown to have the biggest effect of load time on conversion rate and the smallest effect was observed for direct channel. Between user-perceived load metrics, First Paint had the biggest effect on conversion rate.

However, the results across studies and different pages at TransferWise vary and the metrics used in the analysis were simplistic. Hence, further research is required for transferable and generalisable results. In addition, opportunities for future work include A/B testing load improvements and using logistic regression to quantify the relations in a more fine-grained manner.

References

- [AGP12] Tim Ash, Maura Ginty and Rich Page. *Landing page optimization: The definitive guide to testing and tuning for conversions*. John Wiley & Sons, 2012.
- [Aka12] Akamai. *Situational Performance Measurement with a Splash of RUM*. Dec. 2012. URL: <https://blogs.akamai.com/2012/12/situational-performance-measurement-with-a-splash-of-rum.html> (visited on 23/04/2018).
- [Aka15] Akamai. *Akamai Technologies - 2014 Consumer Web Performance Expectations Survey*. May 2015. URL: <https://www.akamai.com/us/en/multimedia/documents/content/akamai-performance-matters-key-consumer-insights-ebook.pdf> (visited on 23/04/2018).
- [Aka17] Akamai. *Akamai Technologies – The State of Online Retail Performance – Spring 2017*. Apr. 2017. URL: <https://www.soasta.com/wp-content/uploads/2017/04/State-of-Online-Retail-Performance-Spring-2017.pdf> (visited on 23/04/2018).
- [App] AppDynamics. *Browser RUM Metrics*. URL: <https://docs.appdynamics.com/display/PRO44/Browser+RUM+Metrics> (visited on 08/05/2018).
- [Bru09] Jake Brutlag. *Speed Matters*. June 2009. URL: <https://research.googleblog.com/2009/06/speed-matters.html> (visited on 20/04/2018).
- [Bul] Zach Bulygo. *How Marketing Funnels Work*. URL: <https://blog.kissmetrics.com/how-marketing-funnels-work/> (visited on 08/05/2018).
- [CKR12] Cliff Crocker, Aaron Kulick and Balaji Ram. *Walmart - Real User Monitoring*. Feb. 2012. URL: <https://www.slideshare.net/devonauerswald/walmart-pagespeedslide> (visited on 20/04/2018).
- [Cut10] Blake Cutler. *Firefox & Page Load Speed – Part I*. Mar. 2010. URL: <https://blog.mozilla.org/metrics/2010/03/31/firefox-page-load-speed-part-i/> (visited on 21/04/2018).
- [Dat16] Databox. *18 Metrics Every SaaS Company Should Track*. Nov. 2016. URL: <https://medium.com/databox-insights/18-metrics-every-saas-company-should-track-8a5045a09190> (visited on 08/05/2018).

- [Den] Artem Denysov. *Improve your web application using Progressive Web Metrics*. URL: https://docs.google.com/presentation/d/1D4foHkE0VQdhcA5_hiesl8JhEGeTDRrQR4gipfJ8z7Y/ (visited on 23/04/2018).
- [Dyn] Dynatrace. *Real user monitoring (RUM)*. URL: <https://www.dynatrace.com/capabilities/real-user-monitoring/> (visited on 08/05/2018).
- [Eve13] Tammy Everts. *How to Find Your Website's Performance Poverty Line*. June 2013. URL: <https://blog.radware.com/applicationdelivery/applicationaccelerationoptimization/2013/06/web-performance-poverty-line/> (visited on 21/04/2018).
- [Eve15a] Tammy Everts. *Case study: Mobile pages that are 1 second faster experience up to 27% increase in conversion rate*. Sept. 2015. URL: <https://www.soasta.com/blog/mobile-web-performance-monitoring-conversion-rate/> (visited on 22/04/2018).
- [Eve15b] Tammy Everts. *Performance Monitoring 101: A beginner's guide to understanding synthetic and real user monitoring*. Apr. 2015. URL: <https://www.soasta.com/blog/synthetic-real-user-measurement-monitoring-rum/> (visited on 23/04/2018).
- [Fri18] Vitaly Friedman. *Front-End Performance Checklist 2018*. Jan. 2018. URL: <https://www.smashingmagazine.com/2018/01/front-end-performance-checklist-2018-pdf-pages/> (visited on 23/04/2018).
- [Goo] GoogleAnalytics. *Adding analytics.js to Your Site*. URL: <https://developers.google.com/analytics/devguides/collection/analyticsjs/> (visited on 08/05/2018).
- [Gri] Ilya Grigorik. *Analyzing Critical Rendering Path Performance*. URL: <https://developers.google.com/web/fundamentals/performance/critical-rendering-path/analyzing-crp> (visited on 23/04/2018).
- [Hea] Heap. *eCommerce*. URL: <https://heapanalytics.com/solutions/ecommerce> (visited on 08/05/2018).
- [Iri] Paul Irish. *paulirish/pwmetrics: Progressive web metrics at your fingertipz*. URL: <https://github.com/paulirish/pwmetrics> (visited on 23/04/2018).
- [Kää18] Kristo Käärmann. *TransferWise Mission Report Q1 2018*. Apr. 2018. URL: <https://transferwise.com/gb/blog/transferwise-mission-report-q1-2018> (visited on 09/05/2018).

- [Lin06] Greg Linden. *Make Data Useful – Amazon, Findory*. Nov. 2006. URL: <http://www.gduchamp.com/media/StanfordDataMining.2006-11-28.pdf> (visited on 22/04/2018).
- [MAC17] Sam Meder, Vadim Antonov and Jeff Chang. *Driving user growth with performance improvements*. Mar. 2017. URL: https://medium.com/@Pinterest_Engineering/driving-user-growth-with-performance-improvements-cfc50dafadd7 (visited on 22/04/2018).
- [MDNa] MDN. *DOMContentLoaded - Event reference*. URL: <https://developer.mozilla.org/en-US/docs/Web/Events/DOMContentLoaded> (visited on 07/05/2018).
- [MDNb] MDN. *GlobalEventHandlers.onload - Web APIs*. URL: <https://developer.mozilla.org/en-US/docs/Web/API/GlobalEventHandlers/onload> (visited on 07/05/2018).
- [Mixa] Mixpanel. *Mixpanel JavaScript API Overview*. URL: <https://mixpanel.com/help/reference/javascript> (visited on 08/05/2018).
- [Mixb] Mixpanel. *Understand Every User's Journey*. URL: <https://mixpanel.com/> (visited on 08/05/2018).
- [Ms2+15] Ms2ger et al. *W3C DOM4*. W3C Recommendation. <http://www.w3.org/TR/2015/REC-dom-20151119/>. W3C, Nov. 2015. URL: <http://www.w3.org/TR/2015/REC-dom-20151119/>.
- [Newa] NewRelic. *Instrumentation for browser monitoring*. URL: <https://docs.newrelic.com/docs/browser/new-relic-browser/page-load-timing-resources/instrumentation-browser-monitoring> (visited on 08/05/2018).
- [Newb] NewRelic. *Page load timing process*. URL: <https://docs.newrelic.com/docs/browser/new-relic-browser/page-load-timing-resources/page-load-timing-process> (visited on 07/05/2018).
- [Pan16] Shubhie Panicker. *Why First Paint as Web Perf API?* Aug. 2016. URL: https://docs.google.com/document/d/1wdxSXo_jctZjdPaJeTtYYFF-rLtUFxrU72_7h9qbQaM/edit (visited on 23/04/2018).
- [Pan17] Shubhie Panicker. *Paint Timing 1*. W3C Working Draft. <https://www.w3.org/TR/2017/WD-paint-timing-20170907/>. W3C, Sept. 2017. URL: <https://www.w3.org/TR/2017/WD-paint-timing-20170907/>.

- [Par17] Jonathan Parks. *Scaling Airbnb's Experimentation Platform*. May 2017. URL: <https://medium.com/airbnb-engineering/https-medium-com-jonathan-parks-scaling-erf-23fd17c91166> (visited on 08/05/2018).
- [Pei15] Nilan Peiris. *We inspire smart people and we trust them*. Mar. 2015. URL: <http://tech.transferwise.com/we-inspire-smart-people-and-we-trust-them/> (visited on 09/05/2018).
- [Pin] Pingdom. *How to set up Real User Monitoring (RUM)*. URL: <https://help.pingdom.com/hc/en-us/articles/203611002-How-to-set-up-Real-User-Monitoring-RUM> (visited on 08/05/2018).
- [Qua] Qualaroo. *The Beginner's Guide to Conversion Rate Optimization*. URL: <https://qualaroo.com/beginners-guide-to-cro> (visited on 08/05/2018).
- [SK13] Dan Siroker and Pete Koomen. *A/B testing: The most powerful way to turn clicks into customers*. John Wiley & Sons, 2013.
- [SOAa] SOASTA. *Boomerang*. URL: <https://soasta.github.io/boomerang/> (visited on 08/05/2018).
- [SOAb] SOASTA. *mPulse | Performance Monitoring Tool for Websites & Apps*. URL: <https://www.soasta.com/performance-monitoring/> (visited on 08/05/2018).
- [Sou12] Steve Souders. *the Performance Golden Rule*. Feb. 2012. URL: <http://www.stevesouders.com/blog/2012/02/10/the-performance-golden-rule/> (visited on 22/04/2018).
- [Sou13] Steve Souders. *Moving beyond window.onload()*. May 2013. URL: <http://www.stevesouders.com/blog/2013/05/13/moving-beyond-window-onload/> (visited on 23/04/2018).
- [Spe] SpeedCurve. *Real User Monitoring: LUX*. URL: <https://speedcurve.com/features/lux/#overview> (visited on 08/05/2018).
- [Sta17] Stackify. *What is Real User Monitoring (RUM)? Examples and Tutorials*. July 2017. URL: <https://stackify.com/what-is-real-user-monitoring/> (visited on 23/04/2018).
- [Tra] TransferWise. *The Story of TransferWise*. URL: <https://transferwise.com/fb/our-story> (visited on 09/05/2018).
- [Tra17] TransferWise. *The TransferWise Mission*. Jan. 2017. URL: <https://transferwise.com/gb/blog/the-transferwise-mission> (visited on 09/05/2018).

- [Tra18] TransferWise. *From 2 founders to 1000 employees — how a small-scale startup grew into a global community*. Mar. 2018. URL: <https://medium.com/transferwise-ideas/from-2-founders-to-1000-employees-how-a-small-scale-startup-grew-into-a-global-community-9f26371a551b> (visited on 09/05/2018).
- [Wal] Philip Walton. *User-centric Performance Metrics*. URL: <https://developers.google.com/web/fundamentals/performance/user-centric-performance-metrics> (visited on 23/04/2018).
- [Wal17] Philip Walton. *Why Web Developers Need to Care about Interactivity*. Dec. 2017. URL: <https://philipwalton.com/articles/why-web-developers-need-to-care-about-interactivity/> (visited on 23/04/2018).
- [Web] WebPageTest. *Speed Index*. URL: <https://sites.google.com/a/webpagetest.org/docs/using-webpagetest/metrics/speed-index> (visited on 07/05/2018).
- [Zem] Mark Zeman. *Synthetic vs Real User Monitoring (RUM)*. URL: <http://support.speedcurve.com/testing-process/synthetic-vs-real-user-monitoring-rum> (visited on 23/04/2018).

Licence

Non-exclusive licence to reproduce thesis

I, **Oliver Viljamaa**,

1. herewith grant the University of Tartu a free permit (non-exclusive licence) to:

reproduce, for the purpose of preservation, including for the purpose of preservation in the DSpace digital archives until expiry of the term of validity of the copyright

of my thesis

Measuring the Effect of User-Perceived Load Metrics on Conversion Rate in the Context of TransferWise

supervised by Marlon Dumas and Mihkel Nõges

2. Making the thesis available to the public is not allowed.
3. I am aware of the fact that the author retains the right referred to in point 1.
4. This is to certify that granting the non-exclusive licence does not infringe the intellectual property rights or rights arising from the Personal Data Protection Act.

Tartu, 14.05.2018