Heuristic Header Error Recovery for Corrupted Network Packets

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Diplom-Informatiker

Florian Schmidt

aus Köln

Berichter:

Prof. Dr.-Ing. Klaus Wehrle
Prof. Dr.-Ing. Wolfgang Kellerer

Tag der mündlichen Prüfung: 23. 10. 2015

Reports on Communications and Distributed Systems

edited by
Prof. Dr.-Ing. Klaus Wehrle
Communication and Distributed Systems,
RWTH Aachen University

Volume 12

Florian Schmidt

Heuristic Header Error Recovery for Corrupted Network Packets

Shaker Verlag
Aachen 2015

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Zugl.: D 82 (Diss. RWTH Aachen University, 2015)

Printed in Germany.

ISBN 978-3-8440-4174-3
ISSN 2191-0863

Shaker Verlag GmbH • P.O. BOX 101818 • D-52018 Aachen
Phone: 0049/2407/9596-0 • Telefax: 0049/2407/9596-9
Internet: www.shaker.de • e-mail: info@shaker.de

Abstract

Wireless communication provides many advantages over wired communication, such as easier deployment due the lack of cabling infrastructure and higher mobility for users. However, one of its most important downsides is the significantly higher error rate. This is exacerbated by the fact that traditional Internet communication enforces perfect bit-by-bit correctness of each packet. While this ensures high reliability in the received data, it is very inefficient: even a single bit error leads to a packet drop, leading to high packet loss even at comparatively low bit error rates. Such a behavior is especially wasteful when considering error-tolerant applications. For example, many media codecs have been designed to tolerate and mask bit errors in their data.

To better support these types of transmissions, suggestions in the past have aimed at tolerating errors in the payload portions of packets, the most well-known example being UDP-Lite. However, these solutions suffer from several drawbacks. First, they suffer from low acceptance unless they ensure they stay fully interoperable with standard protocols. Second, they focus on single protocols, without taking the layered nature of protocol combinations into account. This is problematic because error tolerance in one protocol can be rendered useless by combining it with a lower-layer protocol that drops all packets that contain errors. Third, focusing only on payload error tolerance means any errors in the header portions of packets still lead to drops. Especially in small packets, headers form a large part of the packet, limiting the effectiveness of payload error tolerance.

In this dissertation, we design and present solutions that address these shortcomings. First, by introducing error tolerance into existing standard protocols, we ensure interoperability. Second, by taking the whole stack into account, we ensure that packets are not dropped before error tolerance can recover them. Third, by allowing errors to also occur in the header portions of packets, we increase the effectiveness of error tolerance. This last contribution means that control information in packet headers is not reliable any more. Thus, packets for one application could be misattributed to another. Hence, we will present solutions to identify the correct application a packet belongs to, even under header errors, as well as ways to repair corrupted header information, to prevent this misattribution.

Our first contribution is a solution to introduce header error tolerance and repair into existing protocols at the examples of IPv4, UDP, and RTP. As a second contribution, we design a protocol-independent approach that can identify which connection a packet belongs to, as well as repair certain errors in protocol headers, without requiring any knowledge about the protocols it works on. Our final contribution focuses on the popular 802.11 wireless technology. To fully unlock the potential of header error recovery in 802.11, we design a novel rate adaptation algorithm that can adapt to changes in channel quality without relying on acknowledgments, which is not possible with state-of-the-art solutions.

Kurzfassung

Drahtlose Kommunikation bietet gegenüber drahtgebundener viele Vorteile, so zum Beispiel eine einfachere Einrichtung aufgrund des Verzichts auf Kabel-Infrastruktur sowie höhere Mobilität für Benutzer. Andererseits ist einer der größten Nachteile die im Vergleich deutlich höhere Übertragungsfehlerrate. Dieses Problem wird durch die in klassischer Internetkommunikation erzwungene exakte Bit-für-Bit-Korrektheit der übertragenen Pakete noch verschärft. Dies garantiert zwar eine hohes Maß an Zuverlässigkeit, ist aber sehr ineffizient: selbst ein einzelner Bitfehler führt zum Verwerfen des gesamten Paketes und damit zu hohem Paketverlust bereits bei niedrigen Bitfehlerraten. Ein solches Verhalten ist besonders verschwenderisch bei fehlertoleranten Anwendungen. So gibt es beispielsweise Media Codecs, die speziell im Hinblick auf Toleranz und Maskierung von Bitfehlern entwickelt wurden.

Zur besseren Unterstützung solcher Anwendungsarten wurden in der Vergangenheit Ansätze vorgeschlagen, die Fehler in den Nutzdaten von Paketen tolerieren; das bekannteste Beispiel ist hier UDP-Lite. Allerdings leiden diese Verfahren unter mehreren Nachteilen. Erstens leiden Sie unter geringer Akzeptanz, wenn sie keine vollständige Interoperabilität mit bestehenden Standardprotokollen sicherstellen. Zweitens konzentrieren sie sich auf einzelne Protokolle, ohne dabei das Schichtenmodell und die daraus resultierenden Kombinationen von Protokollen in Betracht zu ziehen. Dies führt zu Problemen, wenn Protokolle niederer Schichten fehlerbehafte Pakete verwerfen. Drittens werden Pakete mit Headerfehlern weiterhin verworfen. Vor allem bei kleinen Paketen bilden Header einen großen Teil des Pakets, so dass eine Fehlertoleranz lediglich für Nutzdaten von beschränkter Wirkung ist.

In dieser Dissertation werden Lösungen für diese Nachteile entworfen und präsentiert. Erstens wird durch das Einbinden von Fehlertoleranz in bestehende Protokolle eine Interoperabilität mit anderen Systemen sichergestellt. Zweitens wird durch das Betrachten des gesamten Netzwerkstapels ein vorzeitiges Verwerfen von Paketen verhindert, bevor Fehlertolernzmechanismen das Paket behandeln können. Drittens wird durch das Zulassen von Fehlern im Headerbereich die Effektivität von Fehlertoleranz erhöht. Eine Herausforderung ist in diesem Zusammenhang die Tatsache, dass fehlerhafte Headerinformationen zu einer Fehlzuordnung des Pakets führen können. In diesem Fall werden Pakete einer Anwendung fälschlicherweise einer anderen zugeordnet. Die Vermeidung solcher Fehlzuordnungen, selbst bei Datenfehlern in den Identifikationsdaten der Paketheader, stellt einen der Hauptbeiträge dieser Dissertation dar.

Im ersten Teil dieser Dissertation wird ein Konzept für Fehlertoleranz anhand einer Implementierung in IPv4, UDP und RTP beispielhaft vorgestellt und untersucht. Im zweiten Teil wird ein protokollunabhängiger Ansatz vorgestellt, der die Zugehörigkeit von (potentiell fehlerbehafteten) Paketen zu einer Verbindung erkennt, ohne über Wissen über die verwendeten Protokolle zu verfügen. Schließlich wird ein neuartiger Ansatz für Ratenadaption in 802.11-Netzwerken vorgestellt, der aufgrund seiner im Gegensatz zu Standardverfahren von ACK-Paketen unabhängigen Adaptation die effektive Verwendung von Fehlertoleranz in WLAN-Netzen ermöglicht.

Acknowledgments

This dissertation concludes a long chapter of my life that I spent in Aachen, from my days as a young and confused undergraduate student to a “finished” PhD. This transformation, and the resulting work you are reading, would not have possible without the support and input by many people. Even though I am sure I will unjustly forget some of them in the following, I want to thank at least some of those that had the greatest influence on both me and my work:

My first thanks rightfully belong to my advisor, Klaus Wehrle. Many years ago, he put me on track as a young PhD student on a project that eventually would eventually branch out in many directions and form this dissertation. Through all this time, he gave me the support I needed, and fostered a creative research group with immense freedom to pursue your goals. I also want to thank Wolfgang Kellerer not only for agreeing to act as a second opponent, but also for valuable feedback during the final parts of this work.

While Klaus made sure that I found a great environment to flourish in, it is Elias Weingärtner who I am most grateful to for paving the way to even enter this environment. Elias picked me up as an insecure Diploma student and saw more scientific capability in me than I did myself. Without his encouragement, I would have never pursued a PhD.

During my work on this dissertation, I had the chance to collaborate with and advise many brilliant students, which helped me shape my work the way it stands now. Thank you, Caj-Julian, David, Dominik, Erwin, Mario, Matthias, Niklas, René, and Tobias, for our many discussions and your dedication to your work. I especially want to thank Anwar and Martin for their scientific rigor and high motivation to push their thesis topics to the limit. Both of them decided to pursue a PhD themselves, and I am sure that soon, I will see them defend brilliant dissertations.

Over the years, I met many great people at COMSYS who I wish to thank. Ismet and Raimondas, for being the best office mates to wish for, both for scientific discussions and non-scientific small talk. Jan and Torsten, for the same reasons, and for enduring my advice. Matteo and Oliver for, both in their own and different ways, opening up new insights for me into how academia works. Tobi, for making sure I would always go the extra mile to make every idea and result that little critical bit better and more well-rounded. Stefan, for teaching me more about scientific writing than anybody else, and for finding the most incredible movies for after-work movie nights. Petra and Ulrike, for always working around organizational issues and solving them for us. Kai, Rainer, and Dirk, for keeping my own and the whole chair’s work–life balance in order. Additionally, all the others I did not mention here specifically: you made (and make) COMSYS a great place.

Finally, I want to thank my brother and my parents for their unconditional love and support in all the years. Without encouraging my curiosity from an early age, I might not have ended up a scientist; without the computer we finally bought after my nagging for a long time (and whose operating system I killed several times before I learned to fix it—sorry for all the lost files!), I might not have ended up a computer scientist.

Contents