



# LUND UNIVERSITY

## The confluence of Cloud computing, 5G, and IoT in the Fog

Tärneberg, William

2019

*Document Version:*

Publisher's PDF, also known as Version of record

[Link to publication](#)

*Citation for published version (APA):*

Tärneberg, W. (2019). *The confluence of Cloud computing, 5G, and IoT in the Fog*. Department of Electrical and Information Technology, Lund University.

*Total number of authors:*

1

### General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

### Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117  
221 00 Lund  
+46 46-222 00 00

# The confluence of Cloud computing, 5G, and IoT in the Fog

*William Tärneberg*



LUND UNIVERSITY

Doctoral Dissertation  
Electrical Engineering  
Lund, March 2019

William Tärneberg  
Department of Electrical and Information Technology  
Electrical Engineering  
Lund University  
P.O. Box 118, 221 00 Lund, Sweden

Series of licentiate and doctoral dissertations  
ISSN 1654-790X; No. 120  
ISBN 978-91-7895-010-2 (Print)  
ISBN 978-91-7895-011-9 (PDF)

© 2019 William Tärneberg  
Typeset in Palatino and Helvetica using  $\text{\LaTeX}$  2 $\epsilon$ .  
Printed in Sweden by Tryckeriet i E-huset, Lund University, Lund.

Cover designed by William Tärneberg.

No part of this dissertation may be reproduced or transmitted in any form or by any means, electronically or mechanical, including photocopy, recording, or any information storage and retrieval system, without written permission from the author.

# Abstract

**I**n the wake of the arrival of cloud computing, future applications are poised to become more resilient and adaptive by embracing elasticity in an osmotic manner. Although cloud computing is a strong attractor for application developers, there are still unconquered performance frontiers. Latency-sensitive and mission-critical applications make up a significant portion of all software systems, and their owners are eager to reap the benefits of cloud computing. However, they are hindered by significant delay, jitter in the delay, and relatively low resilience when operating on traditional, distant, cloud data centres.

Fog computing is emerging as a remedy. Fog computing is a heterogeneous hyper-distributed cloud infrastructure paradigm, ranging from small compute nodes close to the end-users to traditional distant data centres. With greater proximity to the end-users, delay and jitter in the delay can be reduced, and intermediate network reliability improved. Additionally, with increased heterogeneity of resources, applications have a richer tapestry of resources to take advantage of for their objectives. However, managing and taking advantage of this heterogeneity in resources and objectives is a challenge for both the infrastructure providers and application owners alike. Only where to place and scale application components and how to manage system resources to meet the objectives of both parties, is non-trivial. Application placement implies elaborate optimisation objectives, hard-to-find solutions, and operational conflicts.

The objective of this thesis is to investigate the performance-related properties of fog computing, how such an infrastructure can be managed while applications can osmotically take advantage of the infrastructure, and what Fog computing's potential practical performance gains are. These are fundamental topics that need to be answered for providers and application owners alike to be able to invest in fog computing. In general terms, the work in this thesis seeks the trade-offs between infrastructure, applications, and software platform in contrast to the traditional cloud offering.

The thesis provides modelling and simulation tools for evaluating the performance and feasibility of Fog computing. Based on which, the thesis goes on to propose holistic infrastructure management algorithms. The requirements of latency-sensitive and mission-critical applications and use cases are discussed for a fog computing paradigm. These requirements are then translated to Fifth Generation Wireless Specifications (5G) Massive Multiple Input Multiple Output (MIMO) specifications. An original 5G-based fog computing test-bed for time-sensitive and mission-critical applications is implemented. The test-bed is used to evaluate the potential application performance gains of fog computing and to what extent the applications can practically take advantage of a fog infrastructure. The thesis also investigates the architecture of the applications that are proposed to benefit from fog computing and how they perform in traditional cloud offerings.

The included works show that fog computing indeed has a performance advantage over the traditional distant cloud, not only in latency but also in robustness. The benefits of 5G on a time-sensitive application deployed in a fog computing infrastructure are shown to be significant. It is also shown that a fog computing infrastructure with a high degree of heterogeneity and with multiple objectives can be successfully managed scalably. Additionally, the thesis sheds some light on the challenges of implementing latency-sensitive and mission-critical applications with traditional cloud service offerings.

# Contents

<b>Contents</b>	<b>v</b>
<b>Preface</b>	<b>xi</b>
<b>Acknowledgments</b>	<b>xv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Cloud computing . . . . .	3
1.1.1 What makes a cloud . . . . .	5
1.1.2 Elasticity . . . . .	8
1.1.3 High-level concerns . . . . .	11
1.1.4 Who is the cloud for today? . . . . .	11
1.2 Tomorrow's applications and the cloud frontier . . . . .	12
1.2.1 Emerging application types . . . . .	12
1.2.2 Latency and uncertainty challenges . . . . .	14
1.3 Fog computing . . . . .	19
1.3.1 Infrastructure convergence and Fog computing attractors . . . . .	21
1.3.2 Elasticity in the fog and applications . . . . .	24
1.3.3 Fog computing detractors . . . . .	25
 <b>I Modelling and managing a Fog computing infrastructure</b>	 <b>27</b>
<b>2 Mobility</b>	<b>29</b>

2.1	Targeted system . . . . .	30
2.2	Targeted scenario . . . . .	30
2.3	Simulation model . . . . .	31
2.3.1	Application model . . . . .	31
2.3.2	Network model and topology . . . . .	31
2.3.3	Mobility model . . . . .	32
2.3.4	Data Center (DC) model . . . . .	32
2.4	Experiments . . . . .	32
2.5	Results and discussions . . . . .	34
2.5.1	Waiting time degradation . . . . .	34
2.5.2	Session and Virtual Machine (VM) migration . . . . .	35
2.5.3	VM migration time . . . . .	35
2.5.4	Request migration . . . . .	36
2.5.5	Session migration versus node residency time . . . . .	37
2.6	Conclusions . . . . .	37
<b>3</b>	<b>Modelling and system architecture</b>	<b>39</b>
3.1	Existing Fog computing models . . . . .	39
3.1.1	Workload Models . . . . .	40
3.1.2	Set-up Models . . . . .	41
3.1.3	Costs Models . . . . .	43
3.2	Fog computing Meta-model . . . . .	44
3.2.1	Workload Model . . . . .	45
3.2.2	Model parameters . . . . .	47
3.2.3	Objectives Model . . . . .	49
3.2.4	Limitations . . . . .	49
3.3	Simulation showcase . . . . .	50
3.3.1	Experiments . . . . .	50
3.3.2	Results . . . . .	51
<b>4</b>	<b>Centralised Fog computing resource management</b>	<b>53</b>
4.1	Resource Management Challenges . . . . .	54
4.1.1	Service paradigm . . . . .	54
4.1.2	Resource management objectives . . . . .	55
4.1.3	Challenges . . . . .	55

4.2	Extended Fog model . . . . .	56
4.2.1	Data centre Model . . . . .	57
4.2.2	Network Model . . . . .	58
4.2.3	Application Model . . . . .	59
4.2.4	User model . . . . .	59
4.3	Optimisation Formulation . . . . .	60
4.3.1	Resource utilisation metrics and constraints . . . . .	60
4.3.2	Optimisation problem . . . . .	62
4.4	Proposed Application Placement Method . . . . .	63
4.4.1	Exhaustive search . . . . .	64
4.4.2	Iterative local search . . . . .	65
4.4.3	Re-evaluation interval . . . . .	65
4.5	Evaluation model . . . . .	66
4.5.1	Evaluation method . . . . .	66
4.5.2	Application Demand . . . . .	67
4.5.3	Infrastructure and topology . . . . .	68
4.5.4	Application types . . . . .	68
4.5.5	Placement algorithm parametrisation . . . . .	69
4.5.6	Simulator . . . . .	69
4.6	Experiments . . . . .	70
4.6.1	Workload scenarios . . . . .	70
4.6.2	Infrastructure . . . . .	71
4.6.3	Application types . . . . .	72
4.6.4	Placement algorithms . . . . .	72
4.7	Results . . . . .	73
4.7.1	Cost . . . . .	73
4.7.2	Round Trip Time (RTT) . . . . .	75
4.7.3	Resource utilisation . . . . .	75
4.8	Related work . . . . .	77
4.8.1	Replica placement . . . . .	78
4.8.2	CDN and caching . . . . .	79
4.8.3	Inter-and-Intra data centre VM-placement . . . . .	79
4.9	Conclusions . . . . .	80

## 5 Distributed Fog computing resource management 83



5.1	Extended Fog computing model . . . . .	83
5.1.1	Topology . . . . .	84
5.1.2	Data centre model . . . . .	85
5.1.3	Network model . . . . .	85
5.1.4	Application model . . . . .	85
5.2	Distributed resource management algorithm . . . . .	86
5.2.1	Common objective function . . . . .	87
5.2.2	Data centre agent . . . . .	89
5.2.3	Application agent . . . . .	91
5.3	Experiments . . . . .	91
5.3.1	Infrastructure . . . . .	92
5.3.2	Data Centers . . . . .	92
5.3.3	Links . . . . .	92
5.3.4	Topology . . . . .	92
5.3.5	Workload and applications . . . . .	94
5.3.6	Comparison methods . . . . .	95
5.3.7	Evaluation metrics . . . . .	96
5.4	Results . . . . .	97
5.4.1	Convergence . . . . .	97
5.4.2	Step response . . . . .	100
5.4.3	Allocation distribution . . . . .	101
5.5	Conclusions . . . . .	102

## **II Smart cities & Internet of Things 103**

### **6 Realising smart city services with Internet of Things (IoT) and Function-as-a-Service (FaaS) 105**

6.1	Research gap . . . . .	106
6.2	Targeted system . . . . .	107
6.2.1	System components . . . . .	108
6.2.2	System properties . . . . .	109
6.3	Implementation . . . . .	110
6.3.1	Amazon Web Services (AWS) Components . . . . .	110
6.3.2	Testbed Architecture . . . . .	112
6.3.3	Simulated testbed architecture . . . . .	114

6.4	Evaluation . . . . .	115
6.4.1	Representative Scenario . . . . .	115
6.4.2	Performance . . . . .	116
6.5	Conclusions . . . . .	117
<b>7</b>	<b>Bounding shared state inconsistency in distributed IoT systems</b>	<b>119</b>
7.1	System model . . . . .	121
7.2	Cross-Layer Controller . . . . .	123
7.2.1	Objective . . . . .	123
7.2.2	Queuing dynamics . . . . .	124
7.2.3	Lyapunov drift . . . . .	124
7.2.4	Controller design . . . . .	125
7.2.5	Parameter estimation . . . . .	127
7.3	Evaluation . . . . .	127
7.3.1	Comparison policies . . . . .	127
7.3.2	Metrics . . . . .	127
7.3.3	System parameter values . . . . .	128
7.3.4	Input values . . . . .	128
7.4	Results . . . . .	129
7.4.1	Expected deferred state traffic . . . . .	130
7.4.2	Stability and system utility . . . . .	131
7.4.3	Choice of $V$ . . . . .	131
7.4.4	Quantifying the trade-off . . . . .	132
7.5	Conclusions . . . . .	133
<b>III</b>	<b>5G and IoT</b>	<b>135</b>
<b>8</b>	<b>Ultra-Reliable and Low-Latency Communication for the mission-critical applications</b>	<b>137</b>
8.1	Bilateral tele-operation . . . . .	139
8.2	Reliability . . . . .	140
8.2.1	The role of massive MIMO . . . . .	141
8.2.2	Performance of massive MIMO . . . . .	141
8.3	Latency . . . . .	144
8.3.1	System view . . . . .	145

8.3.2	Latency and reliability . . . . .	146
8.3.3	Precoding design . . . . .	147
8.4	Conclusions . . . . .	147
<b>IV</b>	<b>A Fog computing test-bed</b>	<b>149</b>
<b>9</b>	<b>A 5G edge cloud test-bed</b>	<b>151</b>
9.1	Related work . . . . .	153
9.2	Research test-bed . . . . .	154
9.2.1	5G . . . . .	154
9.2.2	Fog computing and network . . . . .	155
9.2.3	Cloud native application framework . . . . .	156
9.3	Evaluation . . . . .	157
9.3.1	Control application . . . . .	157
9.3.2	System characteristics . . . . .	160
9.3.3	System adaptability . . . . .	161
9.3.4	Tightened constraints . . . . .	162
9.4	Conclusions . . . . .	163
	<b>Bibliography</b>	<b>167</b>