

1 Addressing the Node Discovery Problem in Fog 2 Computing

3 **Vasileios Karagiannis** 

4 Distributed Systems Group,
5 TU Wien, Austria
6 v.karagiannis@dsg.tuwien.ac.at

7 **Nitin Desai** 

8 Mälardalen University, Västerås, Sweden
9 nitin.desai@mdh.se

10 **Stefan Schulte** 

11 Distributed Systems Group,
12 TU Wien, Austria
13 s.schulte@dsg.tuwien.ac.at

14 **Sasikumar Punnekkat** 

15 Mälardalen University, Västerås, Sweden
16 sasikumar.punnekkat@mdh.se

17 — Abstract —

18 In recent years, the Internet of Things (IoT) has gained a lot of attention due to connecting
19 various sensor devices with the cloud, in order to enable smart applications such as: smart traffic
20 management, smart houses, and smart grids, among others. Due to the growing popularity of the
21 IoT, the number of Internet-connected devices has increased significantly. As a result, these devices
22 generate a huge amount of network traffic which may lead to bottlenecks, and eventually increase
23 the communication latency with the cloud. To cope with such issues, a new computing paradigm
24 has emerged, namely: fog computing. Fog computing enables computing that spans from the cloud
25 to the edge of the network in order to distribute the computations of the IoT data, and to reduce
26 the communication latency. However, fog computing is still in its infancy, and there are still related
27 open problems. In this paper, we focus on the node discovery problem, i.e., how to add new compute
28 nodes to a fog computing system. Moreover, we discuss how addressing this problem can have a
29 positive impact on various aspects of fog computing, such as fault tolerance, resource heterogeneity,
30 proximity awareness, and scalability. Finally, based on the experimental results that we produce by
31 simulating various distributed compute nodes, we show how addressing the node discovery problem
32 can improve the fault tolerance of a fog computing system.

33 **2012 ACM Subject Classification** Computer systems organization → Cloud computing; Computer
34 systems organization → Fault-tolerant network topologies

35 **Keywords and phrases** Fog computing, Edge computing, Internet of Things, Node discovery, Fault
36 tolerance

37 **Digital Object Identifier** 10.4230/OASICS.Fog-IoT.2020.8

38 **Acknowledgements** The research leading to this paper has received funding from the European
39 Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant
40 agreement No. 764785, FORA—Fog Computing for Robotics and Industrial Automation

41 **1** Introduction

42 The IoT paradigm envisions a world in which everyday objects (i.e., wearables, dumpsters,
43 phones, etc.) connect to the Internet [14]. Such objects may use this connectivity to exchange,
44 store, and process data in order to sense and to affect the surrounding environment [12].



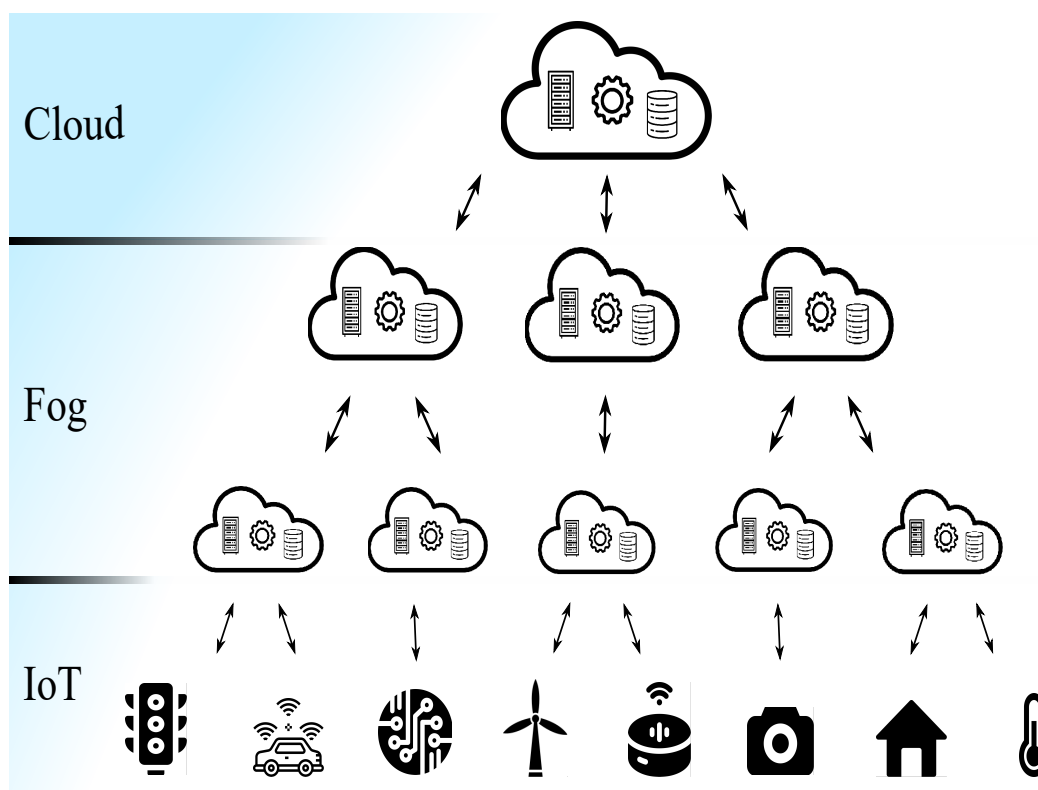
© Vasileios Karagiannis, Nitin Desai, Stefan Schulte, Sasikumar Punnekkat;
licensed under Creative Commons License CC-BY
2nd Workshop on Fog Computing and the IoT (Fog-IoT 2020).

Editors: Anton Cervin and Yang Yang; Article No. 8; pp. 8:1–8:10

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** A fog computing system consisting of various compute nodes that span from the cloud to the edge of the network.

45 Since the computational resources of the everyday objects alone may not be sufficient for
 46 handling the required computational efforts to achieve this, the IoT devices commonly make
 47 use of cloud-based computational resources [24].

48 However, despite the aid of the cloud, the traffic from a large number of Internet-connected
 49 devices can still lead to bottlenecks which increase the communication latency, and may even
 50 limit the expansion of the IoT [19]. Moreover, there are concerns related to preserving the
 51 privacy of the aggregated IoT data, and reducing the communication cost [20]. To cope with
 52 such issues, novel computing paradigms have emerged, two of the most popular being fog
 53 computing, and edge computing.

54 One distinguishing characteristic to separate fog computing from edge computing, is that
 55 the fog envisions a hierarchy of computational resources which span from the cloud to the
 56 edge of the network [3]. For example, Fig. 1 shows a fog computing system that includes
 57 various interconnected cloud and fog compute nodes which spread to the network edge where
 58 the IoT devices reside. Edge computing on the other hand, aims at pushing the computations
 59 towards the edge of the network wherever there are available computational resources (e.g.,
 60 cloudlets or fog nodes) without explicitly including interactions with the cloud [25].

61 The research efforts applied in the context of these two paradigms have resulted in
 62 architectures, models, and frameworks for performing computations in the proximity of
 63 the IoT devices. Due to such efforts, fog computing and edge computing systems have
 64 been observed to provide significant benefits for use cases like data stream processing [5],
 65 preserving privacy in the IoT [20], performing analytics of IoT data [1], online storage [21],
 66 and others [17].

67 To implement such architectures, compute nodes are provisioned at strategic positions
68 throughout the network in order to distribute the computations, avoid bottlenecks, and
69 reduce the communication latency [13]. A lot of research has been conducted in this context,
70 resulting in multiple computing systems which aim at leveraging the edge of the network in
71 order to satisfy the application requirements (e.g., regarding latency and bandwidth) and to
72 improve the user experience [15].

73 Despite the popularity of fog computing and edge computing in the distributed systems
74 research community, computing at the edge of the network is still a relatively recent research
75 topic. For this reason, there are still various important open research problems and challenges,
76 which require further investigation [22]. In this paper, we focus on the node discovery
77 problem [4, 23].

78 Typically, fog computing and edge computing research assumes that compute nodes are
79 already discovered and integrated in the system [11]. However, this can be a complicated
80 task because the current node discovery approaches usually used in cloud-based systems,
81 are not applicable to fog computing since the problem is very different when dealing with
82 compute nodes at the edge of the network [29]. For instance, fog computing systems are
83 expected to leverage on the proximity of the compute nodes while also considering compute
84 nodes with very diverse resource capacities. Such aspects which have not been considered
85 in the context of cloud computing, make novel node discovery techniques—tailored to fog
86 computing—necessary. For this reason, in this paper we analyze the node discovery problem
87 in fog computing, and we discuss the various related aspects that need to be taken into
88 account. Furthermore, we identify the related research questions which need to be addressed,
89 in order to tackle this problem efficiently.

90 The rest of this paper is organized as follows: Section 2 discusses related work from
91 the literature. Afterwards, in Section 3, we analyze the node discovery problem in fog
92 computing, and we identify related research questions. Subsequently in Section 4, we present
93 the preliminary evaluation results that we produce based on simulations which show some of
94 the benefits of addressing the proposed problem (regarding fault tolerance). Finally, Section 5
95 concludes this work, and describes our plans for further research on this topic.

96 **2 Related work**

97 The majority of related work, assumes that the various compute nodes of a fog computing
98 system are already discovered and integrated in the system [11]. Typically, these systems
99 follow a hierarchical architecture whereby the nodes are organized in layers [26]. For instance,
100 Bellavista et al. [2] discuss the execution of services on compute nodes at the edge of the
101 network using a three-layer architecture, and Deng et al. [6] discuss the provisioning of
102 services in distributed edge nodes. However, none of these approaches discuss how the
103 compute nodes are discovered and placed in appropriate positions in the hierarchy.

104 Kolcun et al. [18] present a distributed platform that allows IoT devices from wireless
105 sensor networks, to send data to cloud and local compute nodes. By shifting the computations
106 from the cloud to the local nodes, this approach reduces the network traffic. Furthermore,
107 the authors propose a node discovery algorithm which aids in finding an appropriate compute
108 node for each IoT device.

109 Similarly, Tomar and Matam [27] present a framework that allows the data from the
110 IoT devices to be processed in local compute nodes thereby lowering the dependency on
111 the cloud. This framework also includes a node discovery algorithm for finding appropriate
112 compute nodes for the IoT devices.

113 Finally, Venanzi et al. [31] address the same problem of node discovery for IoT devices
114 although, the focus of this approach is to prolong the lifespan of these devices by considering
115 energy efficiency aspects.

116 Notably, these approaches focus on the problem of selecting appropriate compute nodes
117 for processing the IoT data. In contrast, the work at hand focuses on the problem of
118 discovering new compute nodes that join a fog computing system. Even though these
119 problems seem similar, they require different solutions. The former problem relies on the
120 wireless communication of the IoT devices to discover potential compute nodes (i.e. the
121 compute nodes that reside within wireless range). In the latter problem, which is the problem
122 we address in our work, the compute nodes that span from the cloud to the edge of the
123 network may not integrate wireless communication. Therefore, the aforementioned solutions
124 that address the node discovery problem in the IoT, do not apply to the node discovery
125 problem in fog computing.

126 Further related work can be found in approaches that aim at creating fog computing
127 systems for handling applications related to safety. For instance, Dobrin et al. [8] discuss
128 safety-critical applications while focusing on the problem of having unexpected failures, and
129 Desai et al. [7] discuss various safety aspects (with a focus on safety-critical applications)
130 that need to be considered in fog computing systems.

131 In our work, we also address fault tolerance. However, these works consider fault tolerance
132 as an independent problem which makes it hard to cope with. In our work, we consider fault
133 tolerance at a very early stage, i.e., during the node discovery phase, which increases our
134 options regarding finding appropriate solutions, and based on this, we present promising
135 results.

136 Therefore, the papers discussed so far either briefly mention the node discovery problem
137 in fog computing, or assume that the compute nodes are already discovered and integrated
138 in the system. Thus, they do not provide an analysis of the problem, or any concrete ways
139 to solve it. On the contrary, in our work we analyze different aspects of this problem, we
140 propose related research questions, and we also present promising results towards addressing
141 the node discovery problem in fog computing efficiently.

142 **3 The Node Discovery Problem**

143 The node discovery problem refers to the way that new compute nodes are detected by
144 the system, as well as the process of integrating these nodes (this is also referred to as the
145 discovery phase). For instance, in Fig. 1 we show a fog computing system consisting of one
146 cloud compute node, and eight fog compute nodes (e.g., cloudlets, base stations, routers,
147 etc.), which are organized in three layers. If a new compute node becomes available, how is
148 this node detected by the system, and with which nodes should the new node communicate?
149 In other words, where should the new node be placed in the hierarchy. There are several
150 options because a new node can be placed in each one of the three layers, and connect to
151 different nodes from the adjacent layers. However, every option has a different impact on the
152 performance of the system. Since fog computing systems are expected to scale massively [9],
153 new compute nodes are likely to join the system frequently. Thus, node discovery is an
154 essential part of fog computing systems.

155 To address this problem, we analyze the different aspects of a fog computing system that
156 are affected by the manner whereby nodes are discovered and integrated in the system. To this
157 end, the following sections discuss the reason that the node discovery problem affects different
158 aspects of fog computing, and why these aspects are important. Specifically, Section 3.1

159 discusses fault tolerance, Section 3.2 addresses the potential resource heterogeneity of the
160 nodes, Section 3.3 discusses the importance of proximity awareness, and Section 3.4 addresses
161 scalability. Finally, Section 3.5 presents the research questions that need to be answered in
162 order to address the node discovery problem efficiently.

163 3.1 Fault Tolerance

164 In fog computing, some of the participating compute nodes may be unreliable, and might
165 fail unexpectedly at any moment, which can divide a fog computing system into disjoint
166 parts [16], and affect the system's reliability [32]. For this reason, mechanisms for handling
167 node failure become essential. However, this can be especially challenging in fog computing
168 because when a node fails, moving the computations to neighbor nodes or to the cloud, may
169 affect the performance of the system (e.g., might increase the communication latency) [30].

170 Nevertheless, it is possible to cope with this problem by integrating efficient mechanisms
171 for handling potential future node failures, at the discovery phase, i.e., when a new node
172 joins the system. This can be achieved by having each new node store additional nodes
173 which may not reside in proximity, and are not necessarily used for processing the IoT data,
174 but can be used for maintaining connectivity in case the neighbors fail (cf. Section 4).

175 3.2 Resource Heterogeneity

176 Fog computing systems consist of various resource-heterogeneous compute nodes [28]. This
177 means that the participating compute nodes may have very different resource capacities, e.g.,
178 regarding CPU and memory, but they may also have different capabilities, e.g., regarding
179 hosted services and applications. This diversity should be taken into account during the
180 discovery phase, because different nodes need to be treated differently. For example, upon
181 discovery, a cloud compute node which is able to provide a huge amount of computational
182 resources should go to the top of the hierarchy. This way, the nodes of lower layers will be
183 able to send the IoT data to that node (for processing) by forwarding the data upwards the
184 hierarchy (cf. Fig 1). On the contrary, a compute node at the edge of the network should be
185 placed close to the IoT devices (cf. Fig 1) in order to leverage on the low communication
186 latency. Therefore, the resource heterogeneity of the compute nodes needs to be considered
187 during the discovery phase in order to ensure the efficient operation of a fog computing
188 system.

189 3.3 Proximity Awareness

190 Since processing data in nearby compute nodes improves the communication efficiency [9], fog
191 computing systems leverage on the proximity among the various compute nodes, and the IoT
192 devices, in order to process the IoT data with low communication latency. Most approaches
193 assume that the participating compute node are already discovered and integrated in the
194 system based on proximity (as discussed in Section 1). However, in order to take into account
195 the proximity among the nodes, new nodes need to take proximity measurements (e.g., using
196 round-trip time or hop count), and then connect to the neighbors of the closest proximity.

197 Taking into account the proximity among the nodes during the discovery phase is a
198 challenging task in fog computing, because proximity measurements may have conflicts
199 with other aspects, e.g., with the resource heterogeneity aspect (cf. Section 3.2). This can
200 happen for instance, upon discovery of a new compute node which integrates a big amount
201 of computational resources, and should be placed in a high layer so that many nodes of lower

202 layers can use these resources. At the same time, this new node may be in the proximity
203 of nodes in lower layers. This means that according to proximity, the new node should be
204 placed in a low layer. Thus, during the discovery phase, there may be conflicts based on the
205 different goals of the discovery problem.

206 3.4 Scalability

207 As discussed in Section 1, fog computing systems can include compute nodes that span from
208 the cloud to the edge of the network and thus, they may need to scale to a large degree [9].
209 This means that during the discovery phase, there can be a huge number of possible positions
210 for a new node. Examining all the possible options means taking proximity measurements
211 for a very large number of potential neighbors. However, this may not be possible since this
212 process generates a considerable amount of network traffic which is part of the overhead of the
213 discovery phase. Furthermore, more messages need to be exchanged in order to discover and
214 store additional nodes for fault tolerance, and in order to examine the resource heterogeneity
215 of the other nodes, as discussed in Sections 3.1 and 3.2. Since generating a significant amount
216 of overhead can compromise the scalability of the system, the overhead of the discovery
217 phase needs to be considered, especially because in fog computing new compute nodes may
218 be discovered at any time [16].

219 3.5 Research Questions

220 There are many aspects of fog computing that can be improved by considering the node
221 discovery problem (cf. Sections 3.1 - 3.4). For this reason, and in order to be able to solve
222 this problem efficiently, we identify the following research questions (RQ):

223 **RQ1** To what degree can fog computing systems be fault-tolerant, by storing additional
224 nodes during the discovery phase, which are used in case of node failures?

225 **RQ2** How should the proximity and the resource heterogeneity of the compute nodes, affect
226 the position of a new node that joins a fog computing system?

227 **RQ3** How to make sure that the overhead from new compute nodes joining, does not
228 compromise the scalability of a fog computing system?

229 When we are able to answer these research questions, then we will be in the position to
230 design efficient discovery mechanisms that aid in improving various aspects of fog computing.

231 4 Evaluation

232 In this section, we report the preliminary results of our efforts to tackle the node discovery
233 problem in fog computing. The setup we use in order to produce these results is described in
234 Section 4.1. Afterwards in Section 4.2, we perform various experiments which focus on the
235 fault tolerance aspect of the node discovery problem, and we present our results.

236 4.1 Evaluation Setup

237 In order to perform experiments, and examine the fault tolerance of a fog computing system,
238 we have built a simulator using Java. The reason we do not use a simulator developed in the
239 scope of related work from the literature (e.g., iFogSim [10]), is that alternative simulators
240 lack the necessary functionality to address the proposed problem (e.g., compute nodes that
241 fail or become unavailable temporarily).

242 By using our simulator, we are able to simulate hierarchical fog computing systems
243 consisting of compute nodes that span from the cloud to the edge of the network. The
244 number of the participating compute nodes in these systems is configurable, but the layout
245 is always hierarchical. In the hierarchy, every parent node selects as neighbors up to three
246 children nodes, as shown in Fig. 1. For this evaluation, we perform 50 experiments with
247 100 nodes. The reason we have selected these specific numbers, is that after experimenting
248 extensively with this simulator, we found these numbers to produce results which can be
249 considered representative of the general case.

250 In each one of the 50 experiments, we select various percentages of the participating
251 compute nodes to become unresponsive, and then we examine the percentage of the responsive
252 nodes that remain connected. Since node failure can divide a fog computing system into
253 disjoint parts (as discussed in Section 3.1), with this experiment we aim at measuring the fault
254 tolerance of the system. The specific nodes that fail are chosen randomly using the uniform
255 distribution. Using this evaluation setup, we examine two node discovery mechanisms.

256 In the first mechanism, each new node requests to join from a preexisting node of the
257 system (i.e., a contact node), and stores only nearby neighbors which are found through the
258 contact node. In the second, the new node requests to join through the contact node again,
259 but apart from storing the nearby neighbors, it also stores the neighbors of the contact node.
260 The neighbors of the contact node may not reside nearby so they might not be suitable for
261 processing data with low communication latency. However, these nodes are used in case the
262 other neighbors fail.

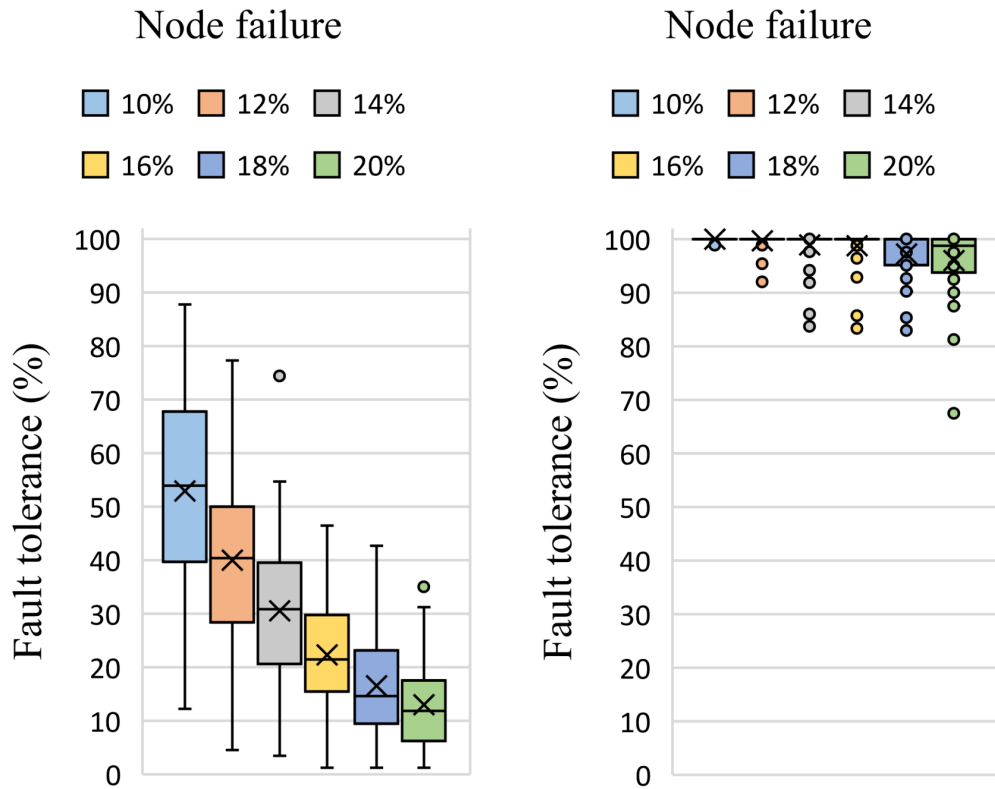
263 4.2 Evaluation Results

264 In Fig. 2, we show the results of our experiments. For Fig. 2a, the nodes store only neighbors,
265 i.e., using the first node discovery mechanism (cf. Section 4.1). In this experiment, we
266 induce node failure of 10%, 12%, 14%, 16%, 18%, and 20% of the nodes, and we measure
267 the corresponding percentages of the responsive nodes that remain connected. Each box
268 plot includes 50 values from the 50 experiments we have conducted. Notably, the average
269 percentage of responsive compute nodes that remain connected is approximately 53% with
270 10% node failure, and the fault tolerance of the system decreases, while the percentage of
271 node failures increases.

272 For Fig. 2b, we repeat the same experiment, but we change the node discovery mechanism.
273 Instead of storing only neighbors (as done for Fig. 2a), in this experiment every node stores
274 additional nodes to be used in case of failures, i.e., the second node discovery mechanism (cf.
275 Section 4.1). Thus, when a responsive node detects (e.g., using heartbeat messages) that the
276 neighbors have failed, this node tries to connect to the system using the additional nodes.
277 Notably, the average percentage of responsive compute nodes that remain connected in this
278 experiment, is approximately 99% with 10% of node failure. Again, the fault tolerance of
279 the system decreases, while the node failures increase although, until the node failures reach
280 20%, the average fault tolerance remains always above 90%.

281 Based on Fig. 2a, we note that creating a fog computing system whereby each node stores
282 only its neighbors, is not an efficient approach with regard to fault tolerance. This is claimed
283 because, when various nodes fail, the percentage of remaining responsive nodes which remain
284 connected decreases radically.

285 However, according to Fig. 2b, we note that the fault tolerance of a fog computing system
286 can be increased significantly, by storing additional nodes during the node discovery phase.
287 Similarly, we believe that addressing the node discovery problem can aid in improving various
288 aspects of fog computing systems, as discussed in Section 3.



(a) When each compute node stores only nearby neighbors. (b) When each compute node stores neighbors and additional nodes to be used in case of failures.

■ **Figure 2** Fault tolerance of a fog computing system.

289 **5 Conclusion**

290 In this paper, we present the node discovery problem in fog computing systems. To this end,
 291 we analyze various aspects of fog computing that can be affected from the way new nodes
 292 are discovered and integrated in the system, such as: fault tolerance, resource heterogeneity,
 293 proximity awareness, and scalability. Furthermore, we identify related research questions
 294 which need to be addressed in order to tackle the proposed problem efficiently. Finally, we
 295 simulate fog computing systems, and we perform experiments with various compute nodes
 296 which integrate a node discovery mechanism that focuses on improving the fault tolerance of
 297 the system. By analyzing the results, we show that when each new node that joins, stores
 298 additional nodes during the discovery phase, the fault tolerance of a fog computing system
 299 improves significantly.

300 Due to the promising results, in the future we plan to focus on node discovery mechanisms
 301 that improve fog computing systems. Specifically, we plan to design node discovery mechan-
 302 isms tailored to fog computing systems by considering not only the fault tolerance of the
 303 system, but also aspects related to proximity awareness, resource heterogeneity, scalability,
 304 and others.

References

- 305 —
- 306 1 Hamid Reza Arkian, Abolfazl Diyanat, and Atefe Pourkhalili. Mist: Fog-based data analytics
307 scheme with cost-efficient resource provisioning for iot crowdsensing applications. *Journal of*
308 *Network and Computer Applications*, 82:152–165, 2017.
- 309 2 Paolo Bellavista, Alessandro Zanni, and Michele Solimando. A migration-enhanced edge
310 computing support for mobile devices in hostile environments. In *International Wireless*
311 *Communications and Mobile Computing Conference (IWCMC)*, pages 957–962. IEEE, 2017.
- 312 3 Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. Fog computing and its
313 role in the internet of things. In *Workshop on Mobile Cloud Computing (MCC)*, pages 13–16.
314 ACM, 2012.
- 315 4 Rustem Dautov, Salvatore Distefano, Dario Bruneo, Francesco Longo, Giovanni Merlino,
316 Antonio Puliafito, and Rajkumar Buyya. Metropolitan intelligent surveillance systems for
317 urban areas by harnessing iot and edge computing paradigms. *Software: Practice and*
318 *Experience*, 48(8):1475–1492, 2018.
- 319 5 Marcos Dias de Assuncao, Alexandre da Silva Veith, and Rajkumar Buyya. Distributed data
320 stream processing and edge computing: A survey on resource elasticity and future directions.
321 *Journal of Network and Computer Applications*, 103:1–17, 2018.
- 322 6 Shuiguang Deng, Zhengzhe Xiang, Jianwei Yin, Javid Taheri, and Albert Y Zomaya.
323 Composition-driven iot service provisioning in distributed edges. *IEEE Access*, 6:54258–54269,
324 2018.
- 325 7 Nitin Desai and Sasikumar Punnekkat. Safety of fog-based industrial automation systems. In
326 *Workshop on Fog Computing and the IoT (IoT-Fog)*, pages 6–10. ACM, 2019.
- 327 8 Radu Dobrin, Nitin Desai, and Sasikumar Punnekkat. On fault-tolerant scheduling of time
328 sensitive networks. In *Workshop on Security and Dependability of Critical Embedded Real-Time*
329 *Systems (CERTS)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 330 9 Pedro Garcia Lopez, Alberto Montresor, Dick Epema, Anwitaman Datta, Teruo Higashino,
331 Adriana Iamnitchi, Marinho Barcellos, Pascal Felber, and Etienne Riviere. Edge-centric
332 computing: Vision and challenges. *ACM SIGCOMM Computer Communication Review*,
333 45(5):37–42, 2015.
- 334 10 Harshit Gupta, Amir Vahid Dastjerdi, Soumya K Ghosh, and Rajkumar Buyya. ifogsim: A
335 toolkit for modeling and simulation of resource management techniques in the internet of things,
336 edge and fog computing environments. *Software: Practice and Experience*, 47(9):1275–1296,
337 2017.
- 338 11 Cheol-Ho Hong and Blesson Varghese. Resource management in fog/edge computing: a survey
339 on architectures, infrastructure, and algorithms. *ACM Computing Surveys (CSUR)*, 52(5):1–37,
340 2019.
- 341 12 Vasileios Karagiannis. Building a Testbed for the Internet of Things. *Alexander Technological*
342 *Educational Institute of Thessaloniki*, pages 1–92, 2014.
- 343 13 Vasileios Karagiannis. Compute node communication in the fog: Survey and research challenges.
344 In *Workshop on Fog Computing and the IoT (IoT-Fog)*, pages 1–5. ACM, 2019.
- 345 14 Vasileios Karagiannis, Periklis Chatzimisios, Francisco Vazquez-Gallego, and Jesus Alonso-
346 Zarate. A survey on application layer protocols for the internet of things. *ICAS Transaction*
347 *on IoT and Cloud Computing*, 3(1):11–17, 2015.
- 348 15 Vasileios Karagiannis and Apostolos Papageorgiou. Network-integrated edge computing
349 orchestrator for application placement. In *International Conference on Network and Service*
350 *Management (CNSM)*, pages 1–5. IEEE, 2017.
- 351 16 Vasileios Karagiannis, Stefan Schulte, Joao Leitao, et al. Enabling fog computing using
352 self-organizing compute nodes. In *International Conference on Fog and Edge Computing*
353 *(ICFEC)*, pages 1–10. IEEE, 2019.
- 354 17 Vasileios Karagiannis, Alexandre Venito, Rodrigo Coelho, Michael Borkowski, and Gerhard
355 Fohler. Edge computing with peer to peer interactions: Use cases and impact. In *Workshop*
356 *on Fog Computing and the IoT (IoT-Fog)*, pages 1–5. ACM, 2019.

- 357 18 Roman Kolcun, David Boyle, and Julie A McCann. Optimal processing node discovery
358 algorithm for distributed computing in iot. In *2015 5th International Conference on the*
359 *Internet of Things (IOT)*, pages 72–79. IEEE, 2015.
- 360 19 Yang Liu, Jonathan E Fieldsend, and Geyong Min. A framework of fog computing: Architecture,
361 challenges, and optimization. *IEEE Access*, 5:25445–25454, 2017.
- 362 20 Rongxing Lu, Kevin Heung, Arash Habibi Lashkari, and Ali Akbar Ghorbani. A lightweight
363 privacy-preserving data aggregation scheme for fog computing-enhanced iot. *IEEE Access*,
364 5:3302–3312, 2017.
- 365 21 Ivan Lujic, Vincenzo De Maio, and Ivona Brandic. Efficient edge storage management based
366 on near real-time forecasts. In *International Conference on Fog and Edge Computing (ICFEC)*,
367 pages 21–30. IEEE, 2017.
- 368 22 Carla Mouradian, Diala Naboulsi, Sami Yangui, Roch H. Glitho, Monique J. Morrow, and
369 Paul A. Polakos. A comprehensive survey on fog computing: State-of-the-art and research
370 challenges. *IEEE Communications Surveys & Tutorials*, 20(1):416–464, 2017.
- 371 23 Ilir Murturi, Cosmin Avasalcui, Christos Tsigkanos, and Schahram Dustdar. Edge-to-edge
372 resource discovery using metadata replication. In *International Conference on Fog and Edge*
373 *Computing (ICFEC)*, pages 1–6. IEEE, 2019.
- 374 24 Carlo Puliafito, Enzo Mingozzi, Francesco Longo, Antonio Puliafito, and Omer Rana. Fog
375 computing for the internet of things: A survey. *ACM Transactions on Internet Technology*,
376 19(2):18, 2019.
- 377 25 Mahadev Satyanarayanan. The emergence of edge computing. *Computer*, 50(1):30–39, 2017.
- 378 26 Vitor Barbosa Souza, Xavi Masip-Bruin, Eva Marín-Tordera, Sergi Sánchez-López, Jordi
379 Garcia, Guang-Jie Ren, Admela Jukan, and Ana Juan Ferrer. Towards a proper service
380 placement in combined fog-to-cloud (F2C) architectures. *Future Generation Computer Systems*,
381 87:1–15, 2018.
- 382 27 Nitendra Tomar and Rakesh Matam. Optimal query-processing-node discovery in iot-fog
383 computing environment. In *2018 International Conference on Advances in Computing, Com-*
384 *munications and Informatics (ICACCI)*, pages 237–241. IEEE, 2018.
- 385 28 Luis M Vaquero and Luis Roderó-Merino. Finding your way in the fog: Towards a compre-
386 hensive definition of fog computing. *ACM SIGCOMM Computer Communication Review*,
387 44(5):27–32, 2014.
- 388 29 Blesson Varghese, Nan Wang, Sakil Barbhuiya, Peter Kilpatrick, and Dimitrios S Nikolopoulos.
389 Challenges and opportunities in edge computing. In *International Conference on Smart Cloud*
390 *(SmartCloud)*, pages 20–26. IEEE, 2016.
- 391 30 Prateeksha Varshney and Yogesh Simmhan. Demystifying fog computing: Characterizing
392 architectures, applications and abstractions. In *International Conference on Fog and Edge*
393 *Computing (ICFEC)*, pages 115–124. IEEE, 2017.
- 394 31 Riccardo Venanzi, Burak Kantarci, Luca Foschini, and Paolo Bellavista. MQTT-driven node
395 discovery for integrated IoT-fog settings revisited: The impact of advertiser dynamicity. In
396 *Symposium on Service-Oriented System Engineering (SOSE)*, pages 31–39. IEEE, 2018.
- 397 32 Zhenyu Wen, Renyu Yang, Peter Garraghan, Tao Lin, Jie Xu, and Michael Rovatsos. Fog
398 orchestration for internet of things services. *IEEE Internet Computing*, 21(2):16–24, 2017.