

Tactful Networking as a cornerstone for opportunistic human-aware D2D communication

Rafael L. Costa^{1,2}, Leobino N. Sampaio¹(Advisor), Aline C. Viana³(Advisor)

¹Dept. of Computer Science, Federal University of Bahia (UFBA), Salvador, Brazil

²Institut Polytechnique de Paris (IPP), Palaiseau, France

³Inria Saclay, 91120, Palaiseau, France

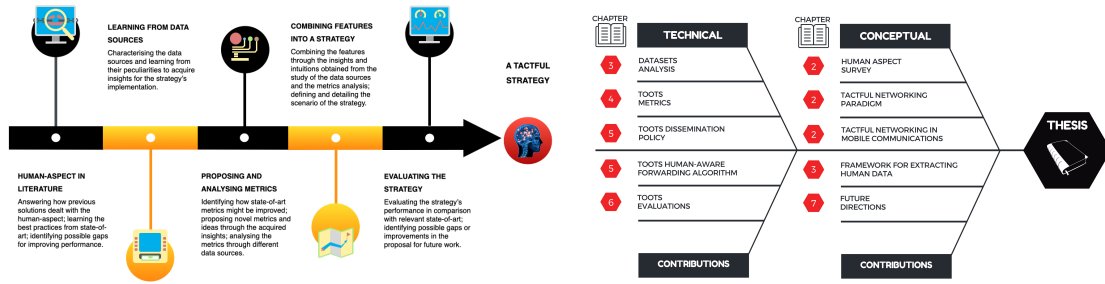
{rlimacosta, leobino}@ufba.br, aline.viana@inria.fr

Abstract. *In this thesis, we guided the reader through the whole process for building a novel Tactful Opportunistic COmmunicaTion Strategy (TOOTS). TOOTS leverages wireless encounter patterns, temporal, spatial, geographic, and direction awareness to improve cost-effectiveness content delivery in a mobile networks scenario. The proposal consists of: learning human-aspect state-of-art best practices; achieving insights to improve strategy's performance; using, proposing, and analyzing human-aware metrics; combining metrics and insights into the strategy targeting an improved performance in a more realistic mobile scenario; and finally, evaluating the strategy through its strengths and shortcomings. This thesis shows that TOOTS improved the performance of an opportunistic content delivery scenario in terms of overhead, delivery rate, and latency by following this process. The strategy can be applied in different scenarios to assist the operator in delivering content without not necessarily using their legacy network, e.g., by exploring the capillarity of their mobile users, for data offloading and other applications.*

1. Introduction

Previous work on D2D opportunistic forwarding algorithms tackled the cost-effective and timely delivery of data [Nunes et al. 2018], i.e., delivering as many contents as possible with less overhead and delay. In these scenarios, contents (or messages) are forwarded user-to-user, from source to destination. Most of such initiatives focused on proposing new algorithms that typically consider user encounters due to individual mobility, points of interest (PoIs), and time-evolving social ties between node pairs. Apart from that, not many initiatives approximated the evaluation metrics to broader inherent aspects of human mobility while targeting the Quality of Experience (QoE) of users and the Quality of Service (QoS) offered by the network.

The main goal of this thesis was to build an opportunistic communication strategy that leverages human-aware information to perform cost-effective content delivery in a mobile scenario. This strategy envisages scenarios such as unburdening legacy networks (i.e., data offloading), increasing the network's capillarity, or providing communication in challenging situations (e.g., crowded places). The thesis' main goal is linked to the following research questions: (i) Which lessons can we take from state-of-art opportunistic communication strategies when dealing with the human aspect?; (ii) What human characteristics peculiarities can we learn from mobility datasets to improve our strategy's



(a) Building process of the TOOTS proposal. (b) Summary of contributions per Thesis' chapter.

Figure 1. Thesis' proposal building process and contributions .

performance?; (iii) How can we translate the identified peculiarities and characteristics into mobility metrics, and how can we combine those to reach superior performance in our strategy? We answered those questions through *the full process for developing a tactful (i.e., human-aware) opportunistic communication strategy able to improve system performance in terms of cost-effective content delivery*. This process (Fig. 1(a)) relates to the main contributions of the thesis and is used for building and introducing the Tactful Opportunistic CommunicaTion Strategy (TOOTS). We summarize each thesis' chapter main technical and conceptual contributions and outcomes in Fig. 1(b).

The rest of this work is organized as follows: in Sec. 2 we detail the proposal while in Sec. 3 we present and discuss the experiments results. In Sec. we conclude this resume and finally, in Sec. we present the publications derived from the thesis.

2. The TOOTS proposal

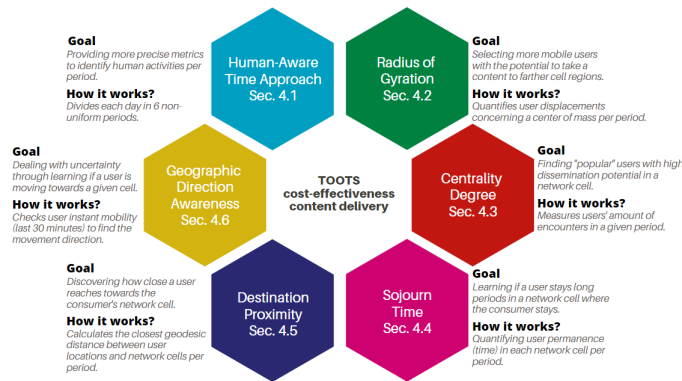


Figure 2. TOOTS' features and metrics for reaching cost-effectiveness.

TOOTS combines five mobility metrics and one temporal feature targeting cost-effective content delivery using opportunistic D2D communication among the nodes. Herein, the metrics of the strategy are briefly described in Fig 2. Further explanations appear on each thesis' section correspondent to the metrics. Fig. 3 describes TOOTS' evaluation scenario where we simulate a network with mobile nodes and a Mobile Edge Computing (MEC) [Mao et al. 2017] site. The MEC site stores popular content that the nodes commonly request.




Figure 3. TOOTS evaluation scenario.

The strategy relies on seven days for learning from users' mobility. Upon having a set of user requests for a given content, the proposed Tactful Dissemination Policy (TDP) runs to choose the disseminator nodes based on users' information. In Fig. 3, we illustrate the delivery of two distinct contents through the disseminators chosen by TDP. After that, the nodes carrying a content run locally another algorithm that takes a forwarding decision.

The TDP chooses disseminator nodes according to their social behavior (i.e., in terms of encounters) in each time window (through the human-aware time approach briefly described in Fig 2). The intention behind TDP is to start the content delivery task by taking the content closer to the consumer nodes for possibly reducing the overhead and delivery latency. We assume that the operator has a central entity that, at the end of a week k , receives the following information from each node u : (a) $\sigma(u)_p$ - a set of users v encountered by u for each period p ; (b) the average local improved centrality degree metric (C_{LID}) [Thilakarathna et al. 2017] of a user u for each period p . The policy uses the $\sigma(u)_p$ from the week k to select users with past direct contacts with the consumers since due to our human routines there is regularity in our human direct contacts [Oliveira et al. 2016]. Furthermore, the C_{LID} applies for identifying nodes' popularity and coverage through contacts duration and earliness. This metric is justified because selecting a contact that happens earlier in a time window and lasts enough to transmit content might decrease the delivery latency. Further, depending on the content's size, trying to transmit through short-duration contacts can waste resources without full content forwarding.

After running the TDP Algorithm, the content c is transmitted to each chosen disseminator node, which stores c in a local buffer. From this moment, when any node carrying c has an encounter, an algorithm runs locally for deciding if the content is forwarded. The algorithm input is: a destination node d , the content c , the present period p , the encountered node v , and the coordinates of d 's cell, given by l . All metrics coefficients are stored locally by each node in a table divided by period p . The content c is transmitted from u to v , only if an algorithm's condition is satisfied. The geographic space of both datasets was divided into network cells. If v is at the same cell reported by the destination (given by l), the algorithm forwards c if v has a higher forwarding potential (given by the Average Centrality Degree) inside the cell (i.e., v met more nodes during p). If the previous test is false, the algorithm checks if v 's average Radius of Gyration and average Sojourn Time are higher than u 's. If true, that means v has a higher avg. RG which gives it the potential to do larger displacements inside the cell given by l . Furthermore, v 's

routine makes it stay longer inside l (v has higher avg. ST). Due to these intuitions, c is forwarded to v .


When v is in a different cell than the one given by l , an instant mobility test towards l occurs. If v 's predominant direction in the last 30 minutes was towards l , c is forwarded, as v can potentially reach or get close to the cell domain given by l . Last, if the direction test is not true, the Destination Proximity (MP) metric is used. If v 's avg. MP is smaller than u 's, c is forwarded, as that means v got closer (or visited) the cell corresponding to l during the period p and due to human routines, this behavior is more likely to repeat. The algorithms and its full discussions and explanations are presented into the Thesis. The following section brings part of the experimentation results and analysis.

3. TOOTS' Evaluations

In order to assess TOOTS' effectiveness, we conducted experiments in four layers: (a) characterization of the datasets; (b) individual characterization and evaluations of each strategy's metrics; (c) Tactful Dissemination Policy evaluations; (d) TOOTS Forwarding evaluations. Our empirical evaluation is presented and performed through simulation analysis with the ONE simulator and the NCCU and GRM datasets. We use more realistic settings (e.g., content size, content time-to-live, and varying contact range) than most related work. Then, we analyze the results in terms of average delivery rate and latency from our TDP compared to state-of-art Store-wait-forward and Epidemic forwarding. Finally, we compare TOOTS with Bubble Rap (the most popular strategy in the context) enhanced with three dissemination policies: random-based (RDP), centrality-based (CDP), and tactful-based (TDP).

In our simulations, the contents range from 11-14 MB, an average size for 60 seconds YouTube 720P HD 30fps advertising videos. TOOTS uses seven days of users' mobility (extracted from NCCU and GRM datasets) during a learning phase. The contents are generated during one day, being one offloading task started every hour. The content delivery deadline is three days, which according to [Thilakarathna et al. 2017] is a reasonable time for delay-tolerant content. Each content has a set of random consumers with varying set sizes. The contents are injected into the network through the users. The first set of users carrying the content is called the Disseminator set. Upon having a contact (i.e., given by a specific communication range), they run an algorithm to make a forwarding decision. If the content is not fully transmitted during a contact, the forwarding is aborted. Every node in the simulation has an 802.11/11 Mbps network interface. We evaluated two communication ranges: 30 m (avg. for WiFi Direct) and 10 m (referred to as Bluetooth). Each experiment case ran 30 times, and the confidence intervals appear when necessary. Due to limited space, on this resume we focus on some TDP and TOOTS forwarding results and analysis. Full evaluation results and discussions are part of the Thesis.

In Fig. 4(a), the TDP-enhanced store-wait-forward show an improved average delivery rate of 82% and 75% respectively on NCCU and GRM, independent of the consumer set size evaluated (1%, 5%, and 10% of each dataset). This delivery rate improvement is explained by one of TDP insights, which is choosing nodes with direct encounters per period with the consumers. For instance, in comparison with the traditional store-wait-forward, the TDP-enhanced strategy's average delivery rate with 1%



(a) Avg. delivery rate on NCCU and GRM datasets with 30 m range
(b) Avg. delivery latency on NCCU and GRM datasets with 30 m range

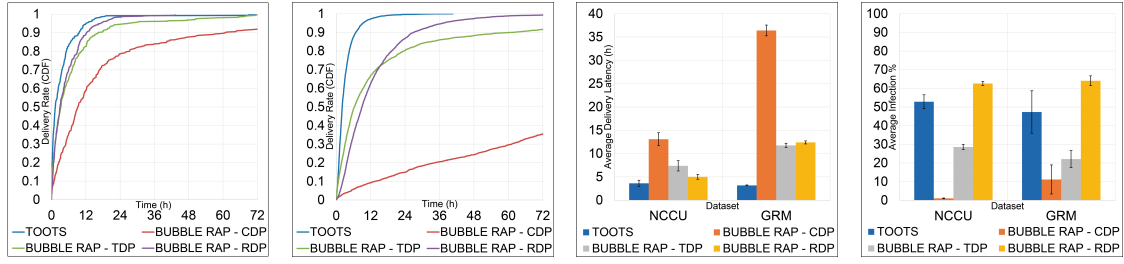
Figure 4. TDP average delivery rate and delivery latency performance evaluation. TDP is combined with epidemic and store-and-wait forwarding in scenarios with varying amount of consumers and different communication range.

of consumer set size was 18% and 60% higher respectively on NCCU and GRM. In Fig. 4(a), the TDP-enhanced epidemic forwarding, as expected, reaches 100% average delivery rate regardless of the consumer set size and dataset. In comparison with the traditional Epidemic forwarding algorithm, the TDP-enhanced epidemic had decreased average hop count, which impacts also in the delivery latency (analysis to follow).

Following in the evaluations, in Fig. 4(b), we plot the average delivery latency. The TDP-enhanced store-wait-forward, regardless of the size of consumers set, takes 13-15h to deliver the contents on NCCU, and 18h-19h on GRM. With that said, the TDP policy makes possible to deliver at least 75% of the delay-tolerant contents in an acceptable time [Thilakarathna et al. 2017], regardless of the dataset, and with zero overhead. In comparison with the traditional store-wait-forward, the TDP-enhanced strategy’s average delivery latency with 1% of consumers was reduced 23% on NCCU and 48% on GRM. As previously stated, thanks to the TDP policy, the number of hops in epidemic forwarding decreased, and so the delivery delay (18.6% and 64.5% smaller respectively on NCCU and GRM with 1% of consumers). Still, the networking infection of epidemic forwarding makes this algorithm very costly or unfeasible in real scenarios.

Finally, we present results from the evaluation of the full strategy in terms of delivery rate, delivery latency, and overhead within NCCU and GRM datasets. As with the state-of-art epidemic forwarding evaluation, the scenario has random consumers with the set size equal to 1% of dataset total nodes and varying communication range (10 m and 30 m). Due to space limitations, this resume features the results with 30 m range only. We compare TOOTS’ performance with Bubble Rap.

First, in Fig. 5(a) TOOTS and Bubble Rap are evaluated in terms of delivery rate through the NCCU dataset with 30 m communication range. TOOTS reached 100% delivery ratio and was also the fastest strategy, with approx. 90% of the generated contents delivered in up to 9 hours. Even in the NCCU dataset which has a smaller number of nodes, those with high centrality degree, Bubble Rap-CDP failed to deliver approx. 10% of the contents. Adopting a CDP is not feasible in real-world, as this strategy originates a “bottleneck” tending to train the device resources of the most “popular” users. Despite reaching 100% delivery ratio, a higher overhead is expected in Bubble Rap-RDP (analysis to follow). Finally, Bubble Rap-TDP also delivered 100% of the contents, but in such



(a) Delivery rate on NCCU with 30 m range (b) Delivery rate on GRM with 30 m range (c) Avg. delivery latency on NCCU and GRM with 30 m range (d) Avg. infection % on NCCU and GRM with 30 m range

Figure 5. Delivery rate, delivery latency, and overhead performance comparison of TOOTS, Bubble Rap-CDP, Bubble Rap-TDP, and Bubble Rap-RDP on NCCU and GRM datasets with 30 m communication range.

experiments, this strategy took slightly more time than TOOTS to deliver all contents.

Within the GRM dataset and 30 m communication range (Fig. 5(b)), TOOTS reaches 100% delivery rate, with approx. 97% of the contents in up to 12 hours. Bubble Rap-CDP delivery rate is even worst than in NCCU. This is explained by the fact GRM nodes' has much lower centrality degrees, and, by a communication bottleneck generated by higher node density, where nodes' interfaces are able to transmit only one content at once. As in the previous analysis, Bubble Rap-RDP reached 100% delivery ratio, but with a higher overhead expected (analysis to follow). Bubble Rap-TDP delivered approx. 94% of the contents. This can be explained by a characteristic of GRM: nodes with low centrality degrees, which makes Bubble Rap-TDP need more time to reach 100% delivery rate.

In Fig. 5(c), the strategies are evaluated in terms of delivery latency, considering a 30 m communication range. On both datasets, TOOTS had the lower average delivery latency, followed by Bubble Rap-RDP, and Bubble Rap-TDP respectively on NCCU and GRM. Reminding that only TOOTS and Bubble Rap-RDP delivered 100% of the contents in up to 72h. On GRM, Bubble Rap-CDP had a much higher average delivery latency, as this strategy struggles with the bottlenecks on fewer nodes with higher centrality, and by a characteristic from this dataset: from the time the contents are generated till their time-to-live (i.e., end of the simulation), over 90% of the nodes has less than 0.4 centrality degree, which are low coefficients. In Fig. 5(d) we evaluate the overhead of each strategy. Bubble Rap-CDP appears with the smaller overhead both on NCCU and GRM datasets with 30 m communication range. On NCCU, this is explained by the existence of higher centrality degree nodes, which makes possible for the disseminator nodes to find the destination directly for most of the generated contents. Still, as previously discussed, there is a bottleneck, lowering Bubble Rap-CDP's delivery rate in comparison with other strategies. As expected, Bubble Rap-RDP had the worst-case in terms of overhead. In comparison with Bubble Rap-TDP, the latter performs with 53% less overhead on NCCU, and 65% less on GRM. Nevertheless, we remark that the TDP-enhanced Bubble Rap needs more time to reach 100% delivery rate on GRM, when there is a 72 h deadline. Applying other larger real-world datasets will be interesting when evaluating this strategy. TOOTS was the fastest strategy, it reached 100% delivery rate, and had 10% and 17% less overhead

than Bubble Rap-RDP respectively on NCCU and GRM. The following section resumes the achievements of this thesis.

4. Thesis' Contributions and Outcomes

The contributions of this thesis are 10-fold (Fig 1(b)) and link to the main goal: presenting TOOTS, a novel strategy to improve cost-effectiveness content delivery in a mobile opportunistic networks scenario. The thesis started as a doctoral project at the Post-graduation Program in Computer Science (PGCOMP) from the Federal University of Bahia (UFBA) in collaboration with the Brazilian National Scientific Computing Laboratory (LNCC). After one year, it became the first international joint-supervision thesis of the program. The signed "cotutelle international" agreement involved the Polytechnic Institute of Paris (IPP) and France's National Institute for Research in Computer Science and Automation (INRIA). The research starting point was the investigation of human aspects in networking. We began with a Survey on the human aspect in networking research covering over a decade of initiatives through a timeline. We found that, throughout the years, there has been an evolution in how the human perspective deals with computer networking challenges. This broader human-aware vision culminates into the Tactful Networking perspective, a novel concept presented in the thesis and later published in a high impact factor journal. We discussed several application examples of tactful networking into mobile network concepts, technologies, and communication models. Among the examples are 5G, IoT, SDN, and NFV. After that, we focused on reviewing state-of-art opportunistic forwarding strategies and algorithms.

Following this, we started to work with three datasets: MACACO, NCCU, and GRM. Throughout this task, we presented a framework for enhancing raw human data hidden into datasets. Data Management, analytics, and privacy tasks were discussed through best practices to deal with challenges in this context, data source examples, and other aspects. We applied those best practices to characterize and analyze the datasets and later validate the strategy. After acquiring the insights, we translated the knowledge obtained into the proposal of metrics and features of our strategy (TOOTS). This knowledge consists of a deeper understanding of the human aspect acquired through: (a) surveying human-behavior state-of-art in networking solutions and the Tactful Networking paradigm proposal. We learned that for trying to reach superior performance in our strategy, it was necessary: (a) identifying and using inherent aspects of the human users hidden into the datasets; (b) learning the contributions and shortcomings from the state-of-art forwarding strategies; (c) acquiring the insights and characterization results from datasets' analysis. Those gave us intuitions for proposing and combining the metrics into a complete strategy for reaching cost-effective content delivery in a mobile scenario. We explained each metric's motivations, insights, and goals individually and how they work into the strategy. Then, we described all metrics, including their mathematical formulations and characterization results in the three datasets. TOOTS uses a probabilistic approach to learn from past users' mobility, identifying node behavior through five different metrics and a time approach feature.

Finally, we introduced TOOTS, its evaluation scenario, and the formal definition for the choosing disseminator nodes problem (TOOTS' 1st phase). As a motivation to develop the strategy, we analyzed the overhead and latency metrics within the state-of-art Epidemic forwarding. Finally, we introduce the two phases of TOOTS: a TDP (Tactful

Dissemination Policy) and a Human-aware forwarding algorithm. Strategy evaluations occurred in each one of its steps. We started detailing the evaluation setup, followed by the 1st phase of TOOTS (TDP policy) evaluations combined with state-of-art Store & Wait and Epidemic forwarding algorithms. Results show that the TDP Policy was able to improve performance. Finally, we compared TOOTS (i.e., the entire proposal) with traditional Bubble Rap and enhanced versions of the latter with TDP, CDP (Centrality-based Dissemination Policy), and RDP (Random-based Dissemination Policy). Results show that TOOTS was the fastest strategy and delivered 100% of the contents on both datasets. The above-mentioned works generated the following publications from the authors R. L. Costa, L. N. Sampaio, A. Ziviani, and A. C. Viana:

Tutorial & Book Chapter: “Humanos no ciclo de comunicação: facilitadores das redes de próxima geração” (Humans in the Communication Loop: enablers of next-generation networks), SBRC, May 2018.

Invited Paper: “Towards Human-Aware D2D Communication” in 2nd International Workshop on Urban Computing, part of the proceedings of the 16th International Conference on Distributed Computing in Sensor Systems – DCOSS’2020, IEEE, May 2020.

Main Track - Conference Paper: “Extração e Análise de Dados Como Suporte a Estratégias de Comunicação D2D Cientes do Humano” (Data Extraction and Analysis for Supporting Human-Aware D2D Communication Strategies), SBRC, December 2020.

Survey - Journal Paper: “Tactful Networking: Humans in the Communication Loop”, Transactions on Emerging Topics in Computational Intelligence – TETCI Volume: 5, Issue: 1, IEEE, February 2021.

Book Chapter & Main Track - Conference Paper: “Tactful opportunistic forwarding: What human routines and cooperation can improve?”, International Conference on Advanced Information Networking and Applications – AINA, Springer, Advanced Information Networking and Applications, May 2021.

Open Call - Journal Paper (Under Minor Review): “On Building Human-aware Opportunistic Communication Strategies for Cost-effective Content Delivery”, in Computer Communications, Elsevier.

References

- Mao, Y., You, C., Zhang, J., Huang, K., and Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Comm. Surveys & Tutorials*, 19(4):2322–2358.
- Nunes, I. O., Celes, C., Nunes, I., Vaz de Melo, P. O. S., and Loureiro, A. A. F. (2018). Combining spatial and social awareness in d2d opportunistic routing. *IEEE Communications Magazine*, 56(1):128–135.
- Oliveira, E. M. R., Viana, A. C., Sarraute, C., Brea, J., and Alvarez-Hamelin, I. (2016). On the regularity of human mobility. *Pervasive and Mobile Computing*, 33:73 – 90.
- Thilakarathna, K., Viana, A. C., Seneviratne, A., and Petander, H. (2017). Design and analysis of an efficient friend-to-friend content dissemination system. *IEEE Transactions on Mobile Computing*, 16(3):702–715.