Towards Automation & Augmentation of the Design of Schedulers for Cellular Communications Networks

Michael Fenton, David Lynch, David Fagan, Stepan Kucera, Holger Claussen, Michael O'Neill

ABSTRACT

Evolutionary Computation is used to automatically evolve small cell schedulers on a realistic simulation of a 4G-LTE heterogeneous cellular network. Evolved schedulers are then further augmented by human design to improve robustness. Extensive analysis of evolved solutions and their performance across a wide range of metrics reveals evolution has uncovered a new human-competitive scheduling technique which generalises well across cells of varying sizes. Furthermore, evolved methods are shown to conform to accepted scheduling frameworks without the evolutionary process being explicitly told the form of the desired solution. Evolved solutions are shown to out-perform a human-engineered state-of-the-art benchmark by up to 50%. Finally, the approach is shown to be flexible in that tailored algorithms can be evolved for specific scenarios and corner cases, allowing network operators to create unique algorithms for different deployments, and to postpone the need for costly hardware upgrades. This work appears in full in Fenton et al., "Towards Automation & Augmentation of the Design of Schedulers for Cellular Communications Networks", Evolutionary Computation, 2018. DOI 10.1162/evco_a_00221.

KEYWORDS

Augmentation, Scheduling, Genetic Programming, Grammatical Evolution, Heterogeneous Networks

ACM Reference Format:

Michael Fenton, David Lynch, David Fagan, Stepan Kucera, Holger Claussen, Michael O'Neill. 2018. Towards Automation & Augmentation of the Design of Schedulers for Cellular Communications Networks. In *GECCO '18 Companion: Genetic and Evolutionary Computation Conference Companion, July 15–19, 2018, Kyoto, Japan,* Jennifer B. Sartor, Theo D'Hondt, and Wolfgang De Meuter (Eds.). ACM, New York, NY, USA, Article 4, 2 pages. https://doi.org/10.1145/3205651.3208211

1 BACKGROUND

Wireless communications networks are a global trillion dollar industry. The GSM Association reports the mobile industry comprised 4.4% of global GDP in 2016, amounting to $3.3 trillion [4]. In order to remain relevant in a vast and increasingly competitive market, network operators value any performance improvements that yield an edge over competitors. Globally, network operators are forecast to spend upwards of $1.4 trillion upgrading their systems through to 2020 [4]. As such, small performance improvements can scale to deliver significant cost savings in such a large domain.

Until recently, the main focus for the optimisation of wireless communications networks has observed conflicting goals of maximising coverage and network performance whilst minimising power consumption [5, 8]. However, with the exponential increase in mobile traffic [3] arising from both rapid growth in the mobile devices market and the onset of the internet of things, this focus has shifted to pure capacity maximisation as network operators struggle to meet demand [2].

As part of capacity maximisation problem faced by network operators, it is now common for these operators to densify their networks through the deployment of small cells [2]. Effectively, existing high-powered Macro Cell (MC) deployments are supplemented by lower-powered Small Cells (SCs) in a Heterogeneous Network, or HetNet. These SCs can be deployed *Ad hoc* within the operational range of the MC in order to offload User Equipments[1] (UEs) from the MC tier. As bandwidth is scarce and expensive, MCs and SCs typically operate on a co-channel deployment, using the same bandwidth.

Optimisation of HetNets can occur on a number of fronts, including SC transmit power optimisation and packet transmission scheduling in the time domain. Intelligent timeframe scheduling at the SC level is attractive to network operators as it represents a relatively cheap software solution, and does not require re-configuration of the network. As such, it is the focus of this study.

2 OVERVIEW

In this study we bring evolutionary computation closer to producing solutions which can be deployed in real networks. Real-world network deployments are extremely limited in the quality/granularity of measurement reports. Not only are reports highly constrained and limited, but reported data is quantized and averaged from its true form. Furthermore, such inaccurate reporting can have a significantly detrimental effect on end-user performance, to the extent where data transmissions can be permanently dropped if actual end-user signal differs too greatly from reported signal [1]. This information paucity adds an extra layer of complexity to the problem, and as such presents a far greater challenge to optimisation methods.

An important aspect of algorithmic controllers for self-optimising networks is that network engineers should be able to fully understand and implement the devised techniques, i.e. they should not be black-box solutions. While this may be a constraint for other techniques, this is a core aspect of many EC approaches.

Michael Fenton is with the Data Science and Machine Learing Group in Corvil Ltd., Dublin, Ireland (email: michaelfenton1@gmail.com).
David Lynch, David Fagan, and Michael O'Neill are with the Natural Computing Research and Applications group (NCRA) in UCD, Ireland (e-mail: david.lynch.1@ucdconnect.ie, david.fagan@ucd.ie, m.oneill@ucd.ie).
Stepan Kucera and Holger Claussen are with Bell Laboratories, Nokia, Dublin, Ireland (e-mail: stepan.kucera@nokia-bell-labs.com, holger.claussen@nokia-bell-labs.com).

GECCO '18 Companion, July 15–19, 2018, Kyoto, Japan

ACM ISBN 978-1-4503-5764-7/18/07…$15.00
https://doi.org/10.1145/3205651.3208211

[1] Any network-connected devices, such as smartphones, tablets, laptop computers, etc.

In this paper we set out to ascertain:

(1) whether it is possible for the evolutionary process to successfully produce viable solutions given sparse and inaccurate information about the true state of the network,
(2) how easily and successfully these solutions can be understood and potentially augmented by human experts, and
(3) whether these evolved and augmented solutions can outperform a state-of-the-art human-designed benchmark across a range of scenarios.

We report the successful application of evolutionary computation, in particular a grammar-based form of Genetic Programming [7], to this pressing real-world communications network problem, which achieves beyond human-competitive performance, significantly outperforming human-designed state-of-the-art solutions reported in the communications networks literature. An additional advantage of the adopted encoding leaves the evolved solutions transparent to the network engineers, making them amenable to human understanding and augmentation. Furthermore, we demonstrate how an in-depth examination of both the evolved solutions and their semantic performance can yield an intuitive understanding of how human-competitiveness has been achieved.

3 RESULTS & CONCLUSIONS

Extensive analysis of the evolved solutions reveal that EC has uncovered a new technique for scheduling SC-attached UEs which is not only generalisable but is both intuitive and easy to implement. Furthermore, evolution is shown to have *twice* produced a solution which conforms to accepted scheduling frameworks which match the literature, despite evolution being given no information about this form of solution and despite being trained on different datasets. These solutions do not break down as a result of changes to the deployment environment, and address corner cases through their use of gradient.

These presented methods are human-competitive in the traditional Koza sense [6], as:

(1) they are equal to or better than a result that was accepted as a new scientific result at the time when it was published in a peer-reviewed scientific journal, and
(2) they are equal to or better than the most recent human-created solution to a long-standing problem for which there has been a succession of increasingly better human-created solutions.

Specifically, the evolved solutions manage to significantly increase cell throughput with respect to proportional fairness over a state-of-the-art human-designed benchmark. Figure 1 shows the percentage changes in downlink rates over the baseline scheduling method. In terms of outright downlink rates expressed as a percentage improvement over the baseline rate, the single best improved UE in the network (with respect to percentage improvement) sees on average around a 200% increase in downlink rates. Up to the 5^{th} percentile, all UEs see greater than a 100% increase in downlink rate performance over the baseline scheduling method. This translates to the bottom 30% of SC attached users achieving a ~1 MB/s performance improvement under the evolved scheme. Furthermore, the top 60% of UEs in the network see an average downlink rate improvement of 15% over both baseline and benchmark. Unlike the benchmark scheduling method, no UEs under the evolved scheme see worse performance than the baseline.

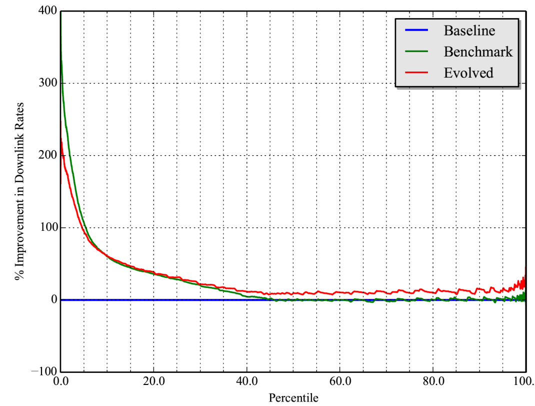

Figure 1: Percentage improvement in downlink rates over baseline for SC-attached UEs.

As network demand rises, SC densification is seemingly the most cost-effective method for operators to increase capacity within their networks. However, evolutionary computation provides a means to not only automatically generate tailored algorithms for specific scenarios, but for human experts to further augment and enhance these solutions. Targeted solutions can be evolved for different network deployments that are capable of handling highly congested/overloaded SCs. This presents a low-cost software alternative to hardware upgrades, thus postponing the need for network operators to supplement their networks with additional SCs. Moreover, higher attachment numbers allow for more fine-grained performance trade-offs, enabling increased fairness.

ACKNOWLEDGEMENTS

This research is based upon works supported by Science Foundation Ireland under grant 13/IA/1850.

REFERENCES

[1] 3GPP. 2014. 3rd Generation Partnership Project (3GPP). http://www.3gpp.org/. (December 2014).
[2] Yan Q. Bian and Deepak Rao. 2014. Small Cells Big Opportunities. Global Business Consulting, Huawei Technologies Co. Ltd.
[3] Cisco. 2015. *Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014 – 2019.* Technical Report. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html.
[4] GSMA. 2017. *The Mobile Economy.* Technical Report. https://www.gsmaintelligence.com/research/?file=9e927fd6896724e7b26f33f61db5b9d5&download.
[5] Erik Hemberg, Lester Ho, Michael O'Neill, and Holger Claussen. 2011. A symbolic regression approach to manage femtocell coverage using grammatical genetic programming. In *Proceedings of the 13th annual conference companion on Genetic and evolutionary computation.* ACM, 639–646.
[6] John R Koza. 2010. Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines* 11, 3-4 (2010), 251–284.
[7] R.I. McKay, X.H. Nguyen, P.A. Whigham, Y. Shan, and M. O'Neill. 2010. Grammar-based Genetic Programming - A Survey. *Genetic Programming and Evolvable Machines* 11, 3-4 (2010), 365–396.
[8] Jie Tang, Daniel KC So, Emad Alsusa, Khairi Ashour Hamdi, and Arman Shojaeifard. 2015. Resource allocation for energy efficiency optimization in heterogeneous networks. *IEEE Journal on Selected Areas in Communications* 33, 10 (2015), 2104–2117.