

Network Sharing for Reliable Networks: A Data-Driven Study

André Gomes*, Jacek Kibilda*, Arman Farhang[†],
Ronan Farrell[†], and Luiz A. DaSilva*

**CONNECT, Trinity College, The University of Dublin, Ireland, E-mail: {gomessaa,kibildj,dasilval}@tcd.ie*

[†]*Maynooth University, Ireland, E-mail: {arman.farhang,ronan.farrell}@mu.ie*

Abstract—The next generation of mobile networks will bring an appetite for reliable communication, which enables emerging critical-communication services. In this paper, we present network sharing between operators as a way to provide increased reliability while using the already existing mobile network infrastructure. Our results indicate that network sharing significantly improves network performance during periods of connectivity shortfalls, benefiting the design of reliable networks. Our conclusions are drawn from a real-world dataset of signal quality indicators for three mobile operators in Dublin, Ireland.

I. INTRODUCTION

Mobile networks are now so tightly integrated into our daily lives that many of us take them for granted, only noticing their existence in their absence, when outages occur. When communication is back on track, users are usually happy again and no major concerns remain. This approach does not cope well with emerging network use cases and services, for which reliable and resilient communication is a must [1].

Network reliability usually comes at a cost of additional network resources made available, such as spectrum, antennas, or denser network deployment, also bringing additional CAPEX and OPEX. Designing reliable networks (and developing a business case that supports their deployment) is one of the next grand challenges in mobile networking, being an enabler of critical-communication applications such as factory automation, public safety, cloud gaming, and many yet-to-be-defined use cases [2], [3].

A parallel trend is network sharing, where operators share resources in order to reduce investments while enhancing network performance. This may take different forms such as infrastructure sharing, where mobile operators share a common physical infrastructure [4], or multilateral agreements, where mobile operators provide network services to each other's users [5].

Studies have shown how multi-connectivity can help network reliability [6], [7], and in this paper we present a data-driven study of network sharing, in the form of multi-operator connectivity, as a way to achieve more reliable networks. Our study makes use of a real-world dataset of three different LTE networks deployed in the city of Dublin, for which we have collected signal strength indicators in a number of test locations (Fig. 1 shows an example of our traces). We base our analysis

on the channel capacity offered by different operators, as well as by multi-connectivity techniques. Unsurprisingly, multi-operator connectivity, where the user can simultaneously connect to multiple mobile operators, enhances the performance experienced by the user. However, our analysis indicates that these gains are particularly significant during period of connectivity shortfalls, benefiting services as critical-communication applications the most.

There are a few studies related to our work. In [8], we show how network capacity may be increased through network sharing between operators. This issue has also been studied in [6], [7], and [9]. The first of those works uses multiple radio technologies (such as WiFi, 2G, 3G, and 4G) to reduce latency, while the second achieves the same goal through multi-base station connectivity. The third is a multipath extension of the transmission control protocol (TCP) which specifies the use of multiple network interfaces. Our work differs from others in primarily two aspects. First, we base our analysis on a real-world dataset, which, to the best of our knowledge, is unique, for it comprises performance measurements from three mobile operators at the same times and locations. Second, we analyse network sharing in the form of multi-operator connectivity as a way to support the design of reliable networks, adopting a performance metric that captures the benefits to the user in periods of connectivity shortfalls.

The remainder of this paper is organised as follows. In Sec. II we present our multi-operator connectivity model. In Sec. III we describe our dataset. Sec. IV presents our data analysis, and the last section includes our final remarks and conclusions.

II. SYSTEM MODEL

We consider a user moving at a walking pace carrying a handset with support to N radio frequency (RF) front-ends, each of which allows connection to a different mobile network operator (MNO) in parallel. Operators occupy their own licensed spectrum and therefore their transmissions do not interfere with each other. In this work, we focus on the downlink communication.

We assess the system's performance based on a figure of merit $Z = f(\Gamma, t)$, which we define as a measure of the capacity of the effective communication channel(s) in use by the mobile. The figure of merit is shown in Eq. (1). It is a function of

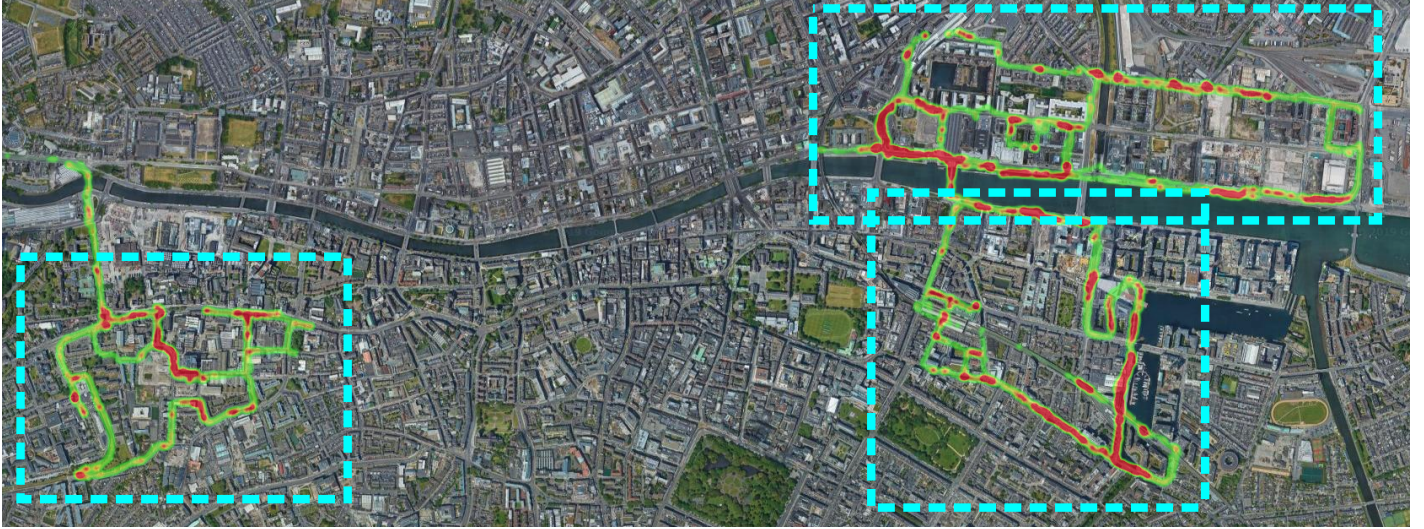


Fig. 1. Signal-to-noise ratio traces of an operator for locations (from left-to-right, top-to-bottom): (a) North Dublin, (b) Guinness Storehouse, and (c) South Dublin. Color legend: ● $\leq -10dB$, ● $-10dB \leq \leq 0dB$, and ● $0dB \leq$.

$\gamma_i = (\gamma_i^0, \gamma_i^1, \dots, \gamma_i^t)$, a time series vector of signal-to-noise ratio (SNR) measurements from mobile operator i . Uppercase $\Gamma = (\gamma_1, \gamma_2, \dots, \gamma_N)$ is a tuple of vectors, each corresponding to an RF interface.

$$f(\Gamma, t) = B \times \log_2(1 + G(\Gamma, t)) \text{ (bps)} \quad (1)$$

where $G(\cdot)$ is the effective SNR either from a connection to a single operator or after combining the N wireless interfaces together (we compare several different alternatives), and B is the channel bandwidth in Hz. We assume all operators use channels of the same bandwidth.




Fig. 2. Example of multi-operator connectivity with three operators: MNO 0, 1, and 2.

We analyse three different multi-operator connectivity techniques in the form of diversity combining schemes [10]. An overall example is depicted in Fig. 2. In all cases, the same information is transmitted by all networks. Each combining algorithm merges the redundant information from multiple operators into a single output data stream as follows:

- 1) *Selection Combining (SC)*: In SC, the output is selected from the wireless interface of highest SNR over the N in-

terfaces available, while others are discarded. The effective SNR $G(\cdot)$ is given by:

$$G(\Gamma, t) = \max\{\gamma_i^t, \forall i \in N\}. \quad (2)$$

- 2) *Equal-Gain Combining (EGC)*: In EGC, the output signal is a coherent summation of all input signals. If we assume that all N communication channels are subject to the same white Gaussian noise, we can write $G(\cdot)$ as follows [11]:

$$G(\Gamma, t) = \frac{1}{N} \left(\sum_i^N \sqrt{\gamma_i^t} \right)^2. \quad (3)$$

- 3) *Maximal-Ratio Combining (MRC)*: In MRC, the output signal is a weighted coherent summation of all input signals according to their respective SNR. Also assuming the same white Gaussian noise for all interfaces, we obtain [12]:

$$G(\Gamma, t) = \frac{1}{N} \sum_i^N \gamma_i^t. \quad (4)$$

For all cases above, γ_i^t is the SNR of operator i at time t , and N represents the number of RF interfaces available.

III. OUR DATA

We have conducted walk tests and collected measurements from three mobile operators in Dublin, Ireland. The resulting data includes performance metrics such as reference signal received power (RSRP), reference signal received quality (RSRQ), and SNR. These measurements were collected using G-MoN, a freeware passive observation application [13]. For a fair comparison, we used three LTE handsets (one per operator) of the same model and brand. The use of G-MoN involved activating the application to record data to a script file. It is worth mentioning that there were no alterations to G-MoN's default settings. During the walk tests, the phones were encased in a frame in a backpack to keep them at a similar orientation

to each other so as to limit bias across the different operators. The phones recorded the data once every second. In open areas of heavy traffic, the walkers would walk to a small number of separate points and pause for two to five minutes to gather a larger number of data samples. There was no downtime, i.e. the phones were always connected to a base station. Fig. 1 depicts our traces over three different geographical areas of the city, namely, North Dublin, the region around the Guinness Storehouse, and South Dublin.

IV. DATA ANALYSIS

In this section, we explore how network sharing, in the form of multi-operator connectivity, impacts mobile communication. We evaluate network performance by accessing the downlink transmission channel capacity of mobile network operators in isolation (namely MNO_0 , MNO_1 , and MNO_2) and the multi-operator connectivity techniques discussed in Sec. II. To this end, we base our analysis and evaluation on a real-world signal-to-noise dataset, assessing the channel(s) capacity as the figure of merit.

We set the bandwidth of each MNO to 5MHz in accordance with licensing information related to our data, so that we can estimate the channel capacity in Eq. (1). For SC, we report multiple results, corresponding to applying the technique over the 1, 2, or 3 operators to select the one of highest average SNR. For instance, $SC_{N=3}$ corresponds to applying selection combining, Eq. (2), over all three mobile operators, while $SC_{N=2}$ only combines signals from the two operators of highest average SNR over a given region. Thus, $SC_{N=1}$ represents the operator of maximal average SNR, and for that reason we use $SC_{N=1}$ as a natural baseline while assessing the gains of other techniques (i.e., $SC_{N>1}$, EGC, and MRC).

A. The Complementarity of Mobile Network Operators

Intuitively, the main benefit of redundancy comes from complementarity. By combining signals from multiple operators, most gains in reliability arise when a given operator is capable of providing connectivity and another is not.

In order to visualise that, Fig. 3 shows, for each mobile operator, the points corresponding to the lowest 5-percentile of SNR values in the region of the Guinness Storehouse, illustrating the geographical coordinates where subscribers might experience lower capacity. On the one hand, the white dashed circle highlights a region where all three mobile operators seem to underperform, which intuitively restricts the gains of multi-operator connectivity. On the other hand, the remainder of the data points indicate the opposite, where operators exhibit different patterns from each other. For instance, the white arrow in Fig. 3 points to a region where only MNO_0 underperforms, suggesting that both MNO_1 and MNO_2 could complement MNO_0 by providing greater channel capacity in that specific area.

We go a step further in that analysis in Fig. 4, which shows the estimated probability density function of the capacity under single- and multi-operator connectivity. It is noticeable that multi-operator connectivity ($SC_{N=3}$, EGC, and MRC)




Fig. 3. The lowest 5-percentile of SNR data points in the Guinness Storehouse region per mobile network operator: MNO_0 ●, MNO_1 ●, and MNO_2 ●.

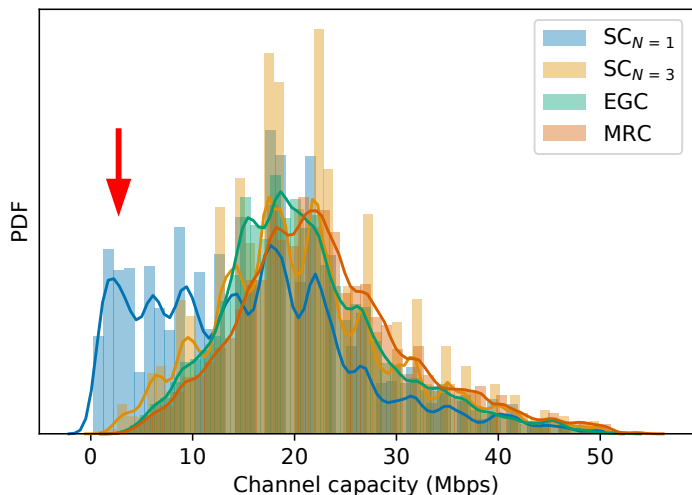


Fig. 4. Probability density function (pdf) of channel capacity in the Guinness Storehouse region.

mostly suppresses the lower-tail of the density function (red arrow) of the channel capacity, in comparison with single-operator connectivity ($SC_{N=1}$). This result reinforces the idea of complementarity in Fig. 3. It also indicates that the gains are concentrated in areas of capacity shortfalls rather than the average or well-performing (upper-tail) cases.

B. The Gains of Network Sharing

Our goal now is to quantify the gains that multi-operator connectivity can bring in terms of dependability of channel capacity. As discussed previously, these gains mostly affect areas of capacity shortfalls, and hence average analysis is unlikely to effectively capture them.

A conventional approach while evaluating the performance of the tail of a distribution is to rely on α -quantiles. This corresponds to specifying a probability $\alpha \in [0, 1]$ and determining the least channel capacity corresponding to that probability, i.e.:

$$q_\alpha(Z) = \sup\{z \in \mathbb{R} : F_Z \leq 1 - \alpha\} \quad (5)$$

where Z represents the figure of merit (channel capacity, in our case) and F_Z its cumulative distribution function.

The superquantile is a similar measure of the tail of a distribution and was proposed by [14], [15] as a replacement of quantiles because of its greater mathematical properties. For instance, superquantiles account the magnitude that exceeds the quantile, being also used as a measure of risk. This is useful for heavy tail distributions, where quantiles do not account for the magnitude of low-probability measures. Also, superquantiles offer mathematical tractability when incorporated into optimisation problems, being easier to handle and better-behaved than quantiles. Superquantiles represent the expectation over the lower-tail of the density function and are formally defined as follows:

$$\bar{q}_\alpha(Z) = \frac{1}{1 - \alpha} \int_0^{1-\alpha} q_\beta(Z) d\beta. \quad (6)$$

Another good property of the α -superquantile is that it approaches the expected value of Z when $\alpha \rightarrow 0$ and approaches the α -quantile when $\alpha \rightarrow 1$, depicting both average and quantile in the extreme cases while also quantifying the tail performance in between. In this paper, we deploy superquantiles as a measure of the network reliability so that we can quantify the gains of network sharing in terms of channel capacity.

Let us first consider Fig. 5. The blue bars represent the estimated average channel capacity for single- (MNO₀, MNO₁, MNO₂, and SC_{N=1}) and multi-operator connectivity (SC_{N=2}, SC_{N=3}, EGC, and MRC). There is a small improvement in performance as we move from single- to multi-operator connectivity. The horizontal blue lines contrast the performance gap between the best result from single-operator connectivity (SOC) baseline (dashed line) and the best-performing multi-operator connectivity (MOC) technique (continuous line).

In counterpart, the orange bars and lines depict the 0.95-superquantile of single- and multi-connectivity techniques. Different from the average performance, Fig. 5 shows a significant improvement in terms of \bar{q}_α . We can clearly see the discrepancy between average and 0.95-superquantile gains by comparing the horizontal lines corresponding to each case. This is a corroborating result to Fig. 4, as we can now quantify how much multi-operator techniques suppress the lower-tail distribution in comparison to the average case. While the blue lines are mostly restricted to an incremental gain of 10%, the superquantile gains approach one order of magnitude.

This result also relates to Fig. 3. Recall that when $\alpha = 0.95$, the superquantile is a measure of the (lower) 5%-tail distribution, whose data point coordinates are depicted in Fig. 3 for the Guinness Storehouse area of Dublin. This suggests that operators tend to complement each other in the worst case scenario (e.g., the 5% least performing data points) but this complementarity does not bring significant gains to the average case.

Let us broaden our analysis by looking at the superquantile as




Fig. 5. Channel capacity when $\alpha = 0.95$. From left to right: (a) North Dublin, (b) Guinness Storehouse, and (c) South Dublin.




Fig. 6. Channel capacity superquantile as a function of α . From left to right: (a) North Dublin, (b) Guinness Storehouse, and (c) South Dublin.




Fig. 7. Spectral efficiency as a function of α . From left to right: (a) North Dublin, (b) Guinness Storehouse, and (c) South Dublin.

a function of α . In Fig. 6, we see that the discrepancy between single- and multi-operator connectivity in fact increases as a function of α , being the smallest in terms of the average capacity ($\alpha = 0$) and the largest when $\alpha = 1$. Thus, the gains of multi-operator connectivity mostly lie in where each operator alone under-performs and tend to complement each other the most.

C. The Cost of Network Sharing

The gains of multi-operator connectivity come at a cost, where extra bandwidth (from multiple connections) is used to transmit information. We study that cost by analysing the spectral efficiency of single- and multi-operator connectivity as a function of superquantiles, Eq. (7). The total amount of bandwidth used depends on how many connections are deployed, where SC_{N=1} has a cost factor of one ($n = 1$ in Eq. (7)), SC_{N=2} has a cost factor of two ($n = 2$), and others (SC_{N=3}, EGC, and MRC) three ($n = 3$).

$$\text{Spectral efficiency} = \frac{\bar{q}_\alpha}{\sum_{i=1}^n b_i} \text{ (bps/Hz)} \quad (7)$$

We plot the spectral efficiency as a function of α in Fig. 7. Conversely to Fig. 6, the discrepancy between single- and multi-operator connectivity decreases as $\alpha \rightarrow 1$. Here, combining operators leads to poorer spectral efficiency in the average case ($\alpha \rightarrow 0$) but approaches single-operator efficiency as α increases, being even slightly better when $\alpha \rightarrow 1$.

D. The Effectiveness of Network Sharing

So far, we have discussed different aspects of network sharing, in the form of multi-operator connectivity, individually: the complementarity, the gains, and the cost. Here, we reason about the effectiveness of network sharing by analysing them together.

From Sec. IV-A and Sec. IV-B, we conclude that network sharing brings the most significant benefits when multi-connectivity can be used to suppress channel capacity shortfalls. The gains are greater as we focus on outages, for the operators tend to complement each other where single-operators alone underperform, being less likely that multiple (or even all) degrade in performance simultaneously. Further, the spectral efficiency also improves in those occasions, resulting in higher gains than single-operators alone while keeping the same cost.

Those are relevant results specially for critical-communication applications, which are sensitive to connectivity shortfalls and have higher demands for resilient and reliable networks.

Finally, we devote our attention to the feasibility of the different multi-operator connectivity techniques we have analysed in this paper. It is important to notice that those techniques require coherent signal decoding. This approach has been implemented by [7] in a small-scale LTE testbed deployment where the mobile simultaneously connects to multiple base stations. A large-scale deployment involving multiple operators may still be challenging in practice. As an alternative, one can also extend selection combining, which also provided significant gains in our analysis, to upper network layers (apart from signal processing at physical layer) in the form of timing combining, where the earliest frame/packet is selected rather than the strongest signal [6], [16]. This diminishes the implementation cost by not requiring coherent signal decoding.

V. CONCLUSIONS AND FUTURE WORK

It is expected that the next generation of mobile communication will face different challenges from the current networks, including both higher demand for reliability and new business models. In this paper, we addressed the problem of offering reliability to cope with critical-communication services through network sharing, in the form of multi-operator connectivity.

To this end, we evaluated the channel capacity of single- and multi-operator connectivity using a real-world dataset of three LTE operators in Dublin, Ireland. We started by analysing the impact of combining the networks of multiple operators to offer more reliable connectivity. Our analysis showed that the gains mostly lie in periods of connectivity shortfalls. To quantify that, we adopted a superquantile performance metric that captures the gains of multi-operator connectivity. We also discussed the cost, in terms of spectral efficiency, of single- and multi-connectivity approaches.

The trade-off between gains and cost benefits multi-operator connectivity as we focus on outages. On the other hand, there is a small improvement to the average case, which also comes at a higher cost than single-operator connectivity.

Finally, our findings motivate the investigation of network sharing in the form of multi-operator connectivity as a way to

design reliable networks. They also motivate the development of more efficient sharing schemes. As we have shown, there is a small benefit of sharing resources in the average case in opposition to outages. In our future works, we intend to investigate dynamic resource allocation as a way to exploit that trade-off, so that network operators may leverage network sharing to enhance network reliability where and when it is needed but avoid the cost of sharing where no significant benefit is provided.

ACKNOWLEDGEMENTS

This work received funding support from the Science Foundation Ireland under Grants 17/NSFC/5224 and 13/RC/2077 (CONNECT).

REFERENCES

- [1] "Setting the scene for 5G: opportunities & challenges," ITU, Tech. Rep., 2018. [Online]. Available: https://www.itu.int/en/ITU-D/Documents/ITU_5G_REPORT-2018.pdf
- [2] S. Chandrashekar, A. Maeder, C. Sartori, T. Höhne, B. Vejlgard, and D. Chandramouli, "5G multi-RAT multi-connectivity architecture," in *IEEE International Conference on Communications workshops (ICC)*, 2016, pp. 180–186.
- [3] P. Popovski, J. J. Nielsen, C. Stefanovic, E. De Carvalho, E. Strom, K. F. Trillingsgaard, A.-S. Bana, D. M. Kim, R. Kotaba, J. Park *et al.*, "Wireless access for ultra-reliable low-latency communication: principles and building blocks," *IEEE Network*, vol. 32, no. 2, pp. 16–23, 2018.
- [4] T. Li and L. Bai, "Model of wireless telecommunications network infrastructure sharing & benefit-cost analysis," in *2011 International conference on information management, innovation management and industrial engineering*, vol. 2. IEEE, 2011, pp. 102–105.
- [5] T. Fiatal, "Mobile virtual network operator," US Patent 8 107 921, January, 2012.
- [6] J. J. Nielsen, R. Liu, and P. Popovski, "Ultra-reliable low latency communication using interface diversity," *IEEE Transactions on Communications*, vol. 66, no. 3, pp. 1322–1334, 2018.
- [7] A. Wolf, P. Schulz, M. Dörpinghaus, J. C. S. Santos Filho, and G. Fettweis, "How reliable and capable is multi-connectivity?" *IEEE Transactions on Communications*, vol. 67, no. 2, pp. 1506–1520, 2018.
- [8] P. Di Francesco, F. Malandrino, and L. A. DaSilva, "Mobile network sharing between operators: a demand trace-driven study," in *Proceedings of the 2014 ACM SIGCOMM workshop on Capacity sharing workshop*, 2014, pp. 39–44.
- [9] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP extensions for multipath operation with multiple addresses," IETF, RFC 6824, January 2013, <http://www.rfc-editor.org/rfc/rfc6824.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc6824.txt>
- [10] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [11] D. A. Zogas, N. C. Sagias, G. S. Tombras, and G. K. Karagiannidis, "Average output SNR of equal-gain diversity receivers over correlative Weibull fading channels," *European Transactions on Telecommunications*, vol. 16, no. 6, pp. 521–525, 2005.
- [12] B. Holter and G. E. Øien, "The optimal weights of a maximum ratio combiner using an eigenfilter approach," in *Proceedings of the fifth IEEE Nordic Signal Processing Symposium*, 2004.
- [13] Knuetter, C., "G-MoN." [Online]. Available: <https://play.google.com/store/apps/details?id=de.carknue.gmon2>
- [14] R. T. Rockafellar and S. Uryasev, "Conditional value-at-risk for general loss distributions," *Journal of Banking & Finance*, p. 29, 2002.
- [15] R. T. Rockafellar and J. O. Royset, "Engineering decisions under risk averseness," *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, vol. 1, no. 2, June 2015.
- [16] M. Rentschler and P. Laukemann, "Performance analysis of parallel redundant WLAN," in *Proceedings of 2012 IEEE 17th International Conference on Emerging Technologies & Factory Automation (ETFA 2012)*, 2012, pp. 1–8.