Optimal percolation in correlated multilayer networks with overlap

Andrea Santoro [1,2] and Vincenzo Nicosia [1]

[1]*School of Mathematical Sciences, Queen Mary University of London, Mile End Road, London E1 4NS, United Kingdom*
[2]*The Alan Turing Institute, British Library, London NW1 2DB, United Kingdom*

Multilayer networks have been found to be prone to abrupt cascading failures under random and targeted attacks, but most of the targeting algorithms proposed so far have been mainly tested on uncorrelated systems. Here we show that the size of the critical percolation set of a multilayer network is substantially affected by the presence of interlayer degree correlations and edge overlap. We provide extensive numerical evidence which confirms that the state-of-the-art optimal percolation strategies consistently fail to identify minimal percolation sets in synthetic and real-world correlated multilayer networks, thus overestimating their robustness. We propose two targeting algorithms, based on the local estimation of path disruptions away from a given node, and a family of Pareto-efficient strategies that take into account both intralayer and interlayer heuristics and can be easily extended to multiplex networks with an arbitrary number of layers. We show that these strategies consistently outperform existing attacking algorithms, on both synthetic and real-world multiplex networks, and provide some interesting insights into the interplay of correlations and overlap in determining the hyperfragility of real-world multilayer networks. Overall, the results presented in the paper suggest that we are still far from having fully identified the salient ingredients determining the robustness of multiplex networks to targeted attacks.

I. INTRODUCTION

Network percolation theory has been shaken recently by the discovery that interdependencies and feedback loops between interacting networks change the character of the percolation transition and make it explosive [1–3]. These results have acquired even more relevance in the past few years, due to the increasing experimental evidence about real-world systems whose structures are naturally represented as multiplex [4,5] or multilayer networks [6]. Random percolation in multiplex networks is nowadays quite well understood [7–11] and the wide spectrum of possible percolation transitions, from abrupt to continuous ones [12–14], has been successfully related to some of the structural properties of these networks, such as the presence of interlayer degree correlations and edge overlap [10,15–18]. However, quite often targeted attacks can potentially drive a system to collapse by knocking down a much smaller fraction of nodes than required by random attacks [19–22]. Hence, optimal percolation, which is the problem of finding the minimal fraction of nodes whose removal would irreversibly fragment the system, has been extensively studied both in single-layer networks [23] and more recently in multilayer networks as well [24]. One of the reasons behind this renovated interest for optimal percolation is the fact that targeted attacks also play a central role in optimal strategies for influence maximization in opinion dynamics [25,26] and

for effective immunization in spreading processes [27–29]. The fact that most of the single-layer optimal attack strategies [29–32] cannot be easily extended to the multilayer case has resulted in an interesting and quite active line of research [24,33]. Although correlations and overlap are indeed a salient aspect of all real-world multiplex networks [6,15,16,34], the few strategies for optimal multiplex percolation proposed so far have been mainly tested on synthetic uncorrelated multilayer networks, thus neglecting interlayer degree correlations and edge overlap.

In this work we fill this gap by studying the problem of optimal percolation in multilayer networks with nontrivial interlayer degree correlations and non-negligible edge overlap. We find that the robustness of systems under targeted attacks is deeply affected by the presence of both interlayer degree correlations and edge overlap. In particular, all the current algorithms for optimal percolation systematically overestimate the size q of the minimal set of nodes to knock down in order to destroy the mutually connected giant component. Here we introduce two classes of algorithms based, respectively, on a generalization to duplex networks of the collective influence algorithm [30] and on the concept of Pareto efficiency [35,36], which allows us to combine layer-based and genuinely multilayer node properties. We show through extensive numerical simulations that all these algorithms provide consistently smaller critical sets in synthetic correlated multilayer networks and outperform other state-of-the-art algorithms in real-world systems.

II. TARGETED ATTACK STRATEGIES

Let us consider a multiplex network $\mathcal{M}$ with N nodes and two layers. The undirected and unweighted edges on

each layer are encoded in the adjacency matrices $A_{ij}^{[\alpha]}$, $\alpha = 1, 2$, whose generic element $A_{ij}^{[\alpha]} = 1$ if and only if nodes i and j are the end points of an edge at layer α and is zero otherwise. Two nodes of $\mathcal{M}$ belong to the same mutually connected component (MCC) if there exists at least one path on each layer that connects them and traverses only nodes belonging to the same MCC. The parameter of interest in percolation analysis is the relative size of the largest mutually connected component (LMCC), which is the largest maximal subgraph consisting of mutually connected nodes [1,6]. Note that the LMCC is a generalization of the giant connected component for single-layer graphs. The optimal percolation problem consists in finding the smallest set of nodes which, if removed, would reduce the size of the LMCC to $O(N^{1/2})$. We call this set a critical set or an attack set and we denote its relative size by q.

Optimal percolation is naturally a many-body problem. Indeed, interactions among nodes at all distances play an important role in the determination of the damage caused by the removal of a subset of nodes, which makes the problem NP-hard [25]. Although there are currently no studies about the computational complexity of optimal multiplex percolation, it is reasonable to assume that this problem is not easier than its classical single-layer counterpart, especially because the computation of the LMCC is based on the existence of paths connecting each pair of nodes on two graphs at the same time, hence the necessity to use heuristic algorithms to find approximate solutions. In most of the cases, heuristic algorithms proceed by assigning a score to each node, based on some structural indicator [23,37,38], and then iteratively removing nodes in decreasing order of their score. As confirmed by recent studies [24,33], single-layer attack strategies cannot be easily generalized to the case of multiplex networks, mainly because it is not immediate to combine node scores on different layers to obtain a meaningful ranking. The authors of Refs. [24,33] proposed several ways of integrating scores based on popular single-layer strategies, namely, (i) rankings based on the sum or product of the degrees in the two layers [high-degree adaptive (HDA)] [24], (ii) a generalization of the collective influence propagation algorithm [39], and (iii) a generalization of the so-called CoreHD algorithm (CoreHD) [31].

To give an idea of how hard it is to directly adapt a single-layer percolation strategy to a multilayer setup, let us consider the CoreHD algorithm, which is one of the most effective strategies to destroy the giant connected component of a single-layer graph. The algorithm proceeds by iteratively removing the nodes with the highest degrees from the 2-core of the graph (i.e., by effectively decycling the network). However, this idea cannot be directly applied to duplex networks, since the 2-core of a multiplex graph is not uniquely defined. As a consequence, there are several existing multiplex extensions of the CoreHD strategy, but none of them provides satisfactory results on duplex networks [33]. By contrast, the recently proposed effective multiplex degree (EMD) strategy [33] consistently improves over all the other existing methods. Indeed, the heuristic used by EMD takes into account multilayer adjacency at different distances and effectively exploits the degree heterogeneity between different layers.

A. Duplex collective influence

An efficient heuristic for optimal single-layer percolation was introduced in Ref. [30] by Morone and Makse. The authors mapped optimal percolation onto the minimization of energy of a many-body system, in which the interactions among units are expressed in terms of the nonbacktracking matrix of the graph, and proposed an efficient and scalable algorithm, called collective influence (CI), to identify the minimal set of influential nodes to remove. The CI algorithm iteratively removes nodes according to the highest values of CI scores, defined as

$$\text{CI}_\ell(i) = (k_i - 1) \sum_{j \in \partial \mathcal{B}(i, \ell)} (k_j - 1), \tag{1}$$

where k_i is the degree of node i, while $\partial \mathcal{B}(i, \ell)$ represents the frontier of the ball of radius ℓ containing all the nodes at distance smaller than or equal to ℓ from node i. This means that a node i is assigned a larger CI score if the set of nodes at distance ℓ from i has a large number of links. By removing a node with a large CI score, we are potentially removing a node that mediates a large number of walks. Remarkably, the attack strategy based on CI can be implemented by an algorithm with time complexity $O(N \log N)$, which is attained by using a max heap to keep and update the CI scores of nodes [39]. Some variations of the CI heuristic have managed to obtain relatively better performance (i.e., smaller attack sets) by including more structural information about the relevance of a given node for percolation [39] and at the cost of an increased time complexity. There has also been an attempt to extend the collective influence algorithm to the case of duplex networks by combining the bare CI scores of the nodes at the two layers [33], but the results are not competitive with other existing algorithms. The main reason is that the bare combination of the layer-based scores does not take into account the role played by edge overlap and interlayer degree correlations in triggering a cascade of node removals.

We introduce here two generalizations of collective influence for duplex networks, which automatically take into account both interlayer degree correlations and edge overlap. The heuristics are based on two simple ideas: the first one is that nodes with high degrees and high edge overlap are more likely responsible for mediating many interdependent paths; the second one is that the removal of a given node i has a large impact on the size of the MCC if it triggers a larger cascade of node removals *away* from i. We define the duplex collective influence (DCI) as

$$\text{DCI}(i) = \frac{k_i^{[1]} k_i^{[2]} - k_i^{\text{int}}}{k_i^{\text{aggr}}} \left[\sum_j a_{ij}^{[1]} \left(k_j^{[2]} - 1 \right) + a_{ij}^{[2]} \left(k_j^{[1]} - 1 \right) \right], \tag{2}$$

where k_i^{int} is the degree of node i in the intersection graph (i.e., the graph containing only the links which appear in both layers) and k_i^{aggr} is the degree of node i in the binary aggregated graph (i.e., the union graph obtained by collapsing the two layers into one [40]). The DCI score of a given node i is indeed obtained as the product of two terms. The first contribution is due to the product of the degrees of node i at the two layers and to the local edge overlap of node i. It is easy to show that this term increases when k_i^{int}

increases, meaning that nodes with a high edge overlap and high degrees at the two layers will in general be ranked higher (see Appendix A for additional details). The term in square brackets instead takes into account potential cascades away from node i triggered by the removal of i. In particular, the term is larger if the neighbors of i on layer 1 have a high degree on layer 2 and vice versa. In this case, the removal of i (and of all its edges on both layers) will disrupt all the paths between the neighbors of i on layer 2 which are mediated by i, hence potentially disrupting the connected component to which i belongs at layer 2. This might in turn trigger further node removals in the neighborhoods of those nodes and let the cascade propagate away from node i. In the limiting case of a duplex network consisting of two identical layers (which is indeed equivalent to a single-layer network with respect to percolation), DCI yields the same node ranking as that induced by CI on the aggregated network when we set $\ell = 1$ in Eq. (1) (see Appendix A for details).

It is important to note that when nodes are iteratively removed from a duplex, the term $k_i^{[1]}k_i^{[2]}$ in Eq. (2) might become equal to zero, e.g., due to the removal of nodes around i which have left node i isolated in one of the two layers. However, node i might still have a relatively large degree on the other layer and its removal might trigger larger cascades away from i than a node which is still connected on both layers but has a small degree on each of them. This happens more frequently in duplex networks with heterogeneous degree distributions. To account for this inconvenience, we define a modified DCI score as

$$\mathrm{DCI}_z(i) = \frac{\left(k_i^{[1]} + 1\right)\left(k_i^{[2]} + 1\right) - 3k_i^{\mathrm{int}} - 1}{k_i^{\mathrm{aggr}}}$$
$$\times \left[\sum_j a_{ij}^{[1]}\left(k_j^{[2]} - 1\right) + a_{ij}^{[2]}\left(k_j^{[1]} - 1\right) \right], \quad (3)$$

which is obtained by replacing $k_i^{[\alpha]}$ with $k_i^{[\alpha]} + 1$ in Eq. (2) and enforcing that DCI_z induces the same node ranking as CI with $\ell = 1$ in the limiting case of duplex networks made of two identical layers (see Appendix A for details). The subscript z indicates that we are correcting for nodes with zero degree on at least one of the two layers.

We use DCI and DCI_z in an adaptive algorithm that iteratively removes nodes from the duplex according to their score recomputed on the remaining subgraph. This process is iterated until the size of the LMCC becomes nonextensive [i.e., $O(N^{1/2})$]. The time complexity of the direct implementation of this algorithm by means of simple data structures is $O(N^2 \log N)$, but a more efficient algorithm which uses a max heap to keep the list of scores sorted will have time complexity $O(N^{1.2})$ (see Appendix B for details).

B. Pareto efficiency for multiobjective optimization

A second class of attack strategies is based on the hypothesis that it should be possible to obtain smaller attack sets by combining layer-specific and genuinely multilayer information. We use here the concept of Pareto efficiency [35,36], which was originally devised to concurrently optimize multiple cost functions. The idea is illustrated in Fig. 1. We consider
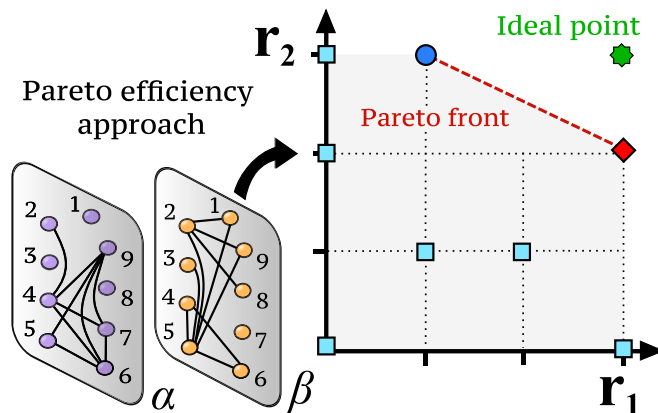

FIG. 1. Graphical representation of Pareto efficiency for two generic structural node descriptors r_1 and r_2. Each node of the multiplex is mapped onto a point in the (r_1, r_2) plane. The points for which no improvement can be achieved in one objective function without hindering the others are called Pareto efficient (blue circle and red diamond) and constitute a Pareto front. The Pareto-efficient points are iteratively ranked according to their Euclidean distance from the ideal point (green star), i.e., the point that maximizes all the objective functions. In this case, the node associated with the red diamond is ranked first.

a set of m node descriptors (also called objective functions), which we deem relevant for multilayer percolation, so that each node i is associated with the vector of ranks induced by each of the m scores $r^i = [r_1^i, r_2^i, \ldots, r_m^i]$ and is mapped onto a point of an m-dimensional space C. Assuming that optimal attack sets consist of nodes that are maximizing all the structural descriptors at the same time, we can employ the concept of dominance strict partial order [36] to identify Pareto-efficient nodes in the space C. A point is considered Pareto efficient if no single score associated with node i can be improved without hindering the other scores associated with node i. In general, for a given set of points there is more than one Pareto-efficient point, which constitutes the so-called Pareto front for that set (see Fig. 1).

At first glance, the Pareto-efficiency approach might appear similar to the hybrid methods presented in Ref. [41]; however, there are a few fundamental differences. In particular, the Pareto-efficiency approach (i) is agnostic with respect to the functions to be maximized (i.e., it is parameter-free), (ii) has a simple physical interpretation (i.e., multiobjective optimization arises naturally whenever a system is subject to at least two concurrent sets of constraints), and (iii) is known to have several advantages over scalarization methods [36,42].

Although multiobjective optimization is a quite appealing concept, the main drawback is that it proposes a set of equally viable "optimal" solutions at each iteration and such a set normally contains multiple solutions. This is indeed far from ideal, since comparing the performance of different cost functions (obtained from different ways of ranking nodes on the basis of their structural properties) can become somehow complicated. A common way to select only one of the Pareto-optimal solutions from a Pareto front, when no additional information is available about how preferable a certain solution is, consists in selecting the solution closest to the ideal point [35], i.e., the (possibly nonexistent) point
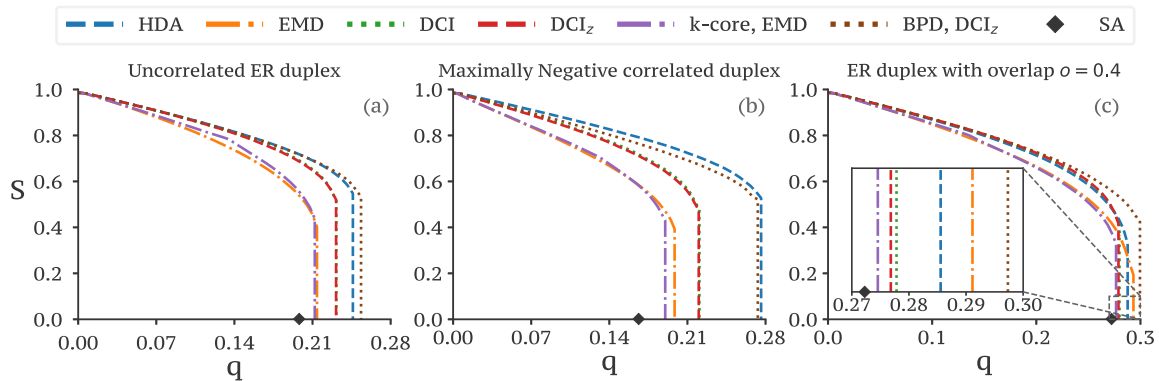
FIG. 2. (a) Relative size S of the LMCC of a multiplex network when a fraction q of nodes is removed by different targeted attack strategies. The multiplex consists of two Erdős-Rényi layers with $N = 10^4$ nodes, average degree $\langle k \rangle = 5$, no interlayer degree correlations ($\rho \approx 0$), and no edge overlap ($o \approx 0$). In this case, EMD provides a much smaller attack set than HDA, but the Pareto-efficient (k-core, EMD) strategy produces a smaller critical set. (b) If the two layers have maximally disassortative interlayer degree correlations ($\rho = -1$) but still no edge overlap ($o \approx 0$), HDA performs sensibly worse than in the uncorrelated case, while the Pareto-efficient (k-core, EMD) strategy finds a smaller attack set and again outperforms EMD. (c) If the duplex has a substantial edge overlap ($o = 0.4$) and no correlations ($\rho \approx 0$), the critical set is much larger than in the other two cases. The presence of edge overlap favors HDA, but the smallest attack set is still found by the (k-core, EMD) Pareto-efficient algorithm, immediately followed by DCI and DCI$_z$. For comparison, we also report in each plot the results obtained by simulated annealing (black diamond). All the curves are averaged over 20 realizations.

that simultaneously maximizes all the cost functions (see Fig. 1). Alternative ways to select Pareto-optimal solutions exist in the literature [36,42,43]; however, no consensus on the best approach has been reached yet. We adopted the ideal point method for the results shown in the following. In other words, for each Pareto strategy, we constructed the critical sets by iteratively removing the Pareto-efficient point having minimal Euclidean distance from the ideal point (potential ties are broken by selecting one of the points uniformly at random). We then recompute the set of Pareto-efficient points and iterate until the LMCC becomes nonextensive. Details about the time complexity of Pareto-efficient strategies are reported in Appendix B.

III. COMPARISON OF TARGETED ATTACK STRATEGIES

Here we compare the two state-of-the-art algorithms for optimal multiplex percolation proposed so far, namely, HDA [24] and EMD [33], with a variety of multiplex targeted attack strategies from three classes, namely, (i) alternative genuinely multiplex strategies, (ii) Pareto-efficient strategies based on the combination of the scores of single-layer targeted attack strategies on the two layers, and (iii) Pareto-efficient strategies obtained by combining single-layer descriptors with one genuinely multiplex algorithm. In the following we will discuss in detail the performance obtained by six strategies, namely, HDA, EMD, DCI, DCI$_z$, and the two Pareto-efficient strategies obtained by combining the k-core ranking on the two layers with the ranking induced by EMD, which we call k-core and EMD, and the score assigned on each layer by believe propagation decimation and the ranking induced by DCI$_z$. Note that when considering HDA, we iteratively remove nodes from the duplex that have the highest product of the degrees in the two layers, as done in [24]. As a reference, we also report the results obtained by simulated annealing (SA) as described in Ref. [13], which is able to find very small targeted attack sets at the expense of heavier computations.

The results obtained with all the other methods we tested are reported in Appendix C.

In Fig. 2(a) we report the percolation diagrams of duplex networks with uncorrelated Erdős-Rényi layers. Note that the duplex network in Fig. 2(a) is consistent with that used in Ref. [33], where the authors showed that the critical set found by EMD is usually much smaller than that found using HDA. Interestingly, the combination of EMD and k-core provides a smaller critical set than either EMD or HDA alone. This is because by targeting nodes which have high EMD scores and at the same time belong to the inner k-core on each layer, we have a higher probability of simultaneously damaging the LMCC of the multiplex and the giant connected component on each layer. Even more interesting results are reported in Fig. 2(b) for a duplex with maximally disassortative interlayer degree correlations (and no edge overlap) and in Fig. 2(c) for a duplex with high edge overlap. It is evident from the figures that the relative performance of each targeting algorithm depends quite substantially on the structure of the multiplex and in particular on the presence of interlayer degree correlation and edge overlap. For instance, EMD still outperforms HDA by a large margin when the multiplex has no edge overlap and disassortative degree-degree correlations [Fig. 2(b)], while EMD is the worst-performing strategy when edge overlap is not negligible. In general, the algorithms based on Pareto-efficiency perform better than either EMD and HDA, while both DCI and DCI$_z$ find relatively smaller critical sets in the case of networks with non-negligible overlap. This is a confirmation of our intuition that heuristics that perform better in one specific condition (e.g., where the two layers are uncorrelated and edge overlap is negligible) do not always achieve the same performance under other conditions.

A. Dependence on edge overlap

The edge overlap of a two-layer multiplex measures the fraction of edges that are present on both layers [6,16,40]. It

can be measured as

$$o_s = \frac{\sum_{i,j}^N o_{ij}}{2\sum_{i,j}^N \Theta(o_{ij})}, \qquad (4)$$

where $o_{ij} = \sum_{\alpha=1}^2 A_{ij}^{[\alpha]}$ and $\Theta(\cdot)$ is the Heaviside step function. In particular, $o_s = 1/2$ when the two layers do not share any edge in common [44]. By contrast, the maximum value $o_s = 1$ is obtained when the two layers are identical. For the sake of convenience, we consider the linear transformation $o = 2(o_s - 1/2)$ that maps the edge overlap o_s into the interval [0, 1]. In general, real-world multiplex networks exhibit relatively large values of edge overlap [16,40], indicating the presence of nontrivial correlations between the two layers. Nevertheless, targeted attack strategies have been compared mainly (if not exclusively) on duplex networks with Erdős-Rényi layers having a negligible edge overlap.

We investigated the impact of edge overlap on the performance of different targeted attack strategies by considering a class of synthetic duplex networks with tunable edge overlap o. In particular, we employed an approach similar to the one presented in [45]. That is, starting from two identical layers ($o = 1$ and maximal interlayer degree correlation $\rho = 1$), we iteratively rewire the edges of one of the two layers to reduce the edge overlap until we get to $o = 0$ while maintaining untouched the degree sequence of each layer (see Appendix D for details). In Fig. 3(a) we plot the relative size of the critical set q obtained by each of the six algorithms as a function of the edge overlap in a duplex with Erdős-Rényi layers. We notice that in general q is an increasing function of o. This fact is somehow expected, since the existence of an extensive MCC imposes more stringent constraints on the graph than the existence of a giant connected component in a single-layer graph. Indeed, a duplex with $o = 1$ is indistinguishable from the single-layer graph obtained by combining the two (identical) layers; hence the optimal attack set in that case corresponds to that of each layer.

However, each attack strategy behaves slightly differently as o increases. For instance, for $o > 0.3$ the critical set found by EMD is always larger than that obtained by all the other strategies. In the limit of $o = 1$, however, the EMD and HDA heuristics coincide, since the EMD weight of each node i becomes proportional to the degree k_i. By contrast, DCI, DCI_z, and the two Pareto-efficient strategies perform relatively poorly in networks with small overlap, but they generally outperform both EMD and HDA as the amount of overlap increases. This is because targeted methods that indirectly disrupt a large number of interdependent paths are more likely to trigger cascades in the system. Note that some Pareto-efficient strategies already outperform the results of the simulated annealing achievable in a reasonable computing time (same implementation as the one presented in [33] with temperature steps equal to 10^{-7}). A similar qualitative behavior is observed when considering duplex systems having a heterogeneous degree distribution on each layer [Fig. 3(b)], although the typical values of q are overall smaller. This indicates that the heterogeneity of the degree distribution of each layer has some impact on the efficiency of each attack strategy, but the presence of edge overlap effectively determines the relative performance of different strategies. Interestingly, for

FIG. 3. Size of the critical attack set q as a function of edge overlap for different attack strategies in duplex networks with $N = 10^4$ nodes, $\langle k \rangle = 5$, and whose layers have (a) Erdős-Rényi or (b) scale-free degree distributions ($\gamma = 2.6$). The plots are obtained by starting from two identical layers ($o = 1$ and $\rho = 1$) and then iteratively rewiring the edges of one of the two layers to reduce the edge overlap until we get to $o = 0$ [45]. Again, strategies based on Pareto efficiency yield the best results. Results are averaged over 20 realizations. Error bars are smaller than the marker size.

both the topologies, the best (smallest) critical set is always obtained by one of the methods proposed in this paper, that is, methods that combine layer-based and genuinely multilayer node properties through Pareto efficiency, with DCI and DCI_z following closely when $o > 0.4$ (see Appendix C for the corresponding figure with a comparison of all the targeted attack strategies considered).

B. Role of interlayer degree correlations

Interlayer degree correlations are known to have a substantial role in determining the robustness of many real-world systems [34] and are responsible for consistent shifts in the position of the random percolation threshold [12,46]. Several studies have found that maximally disassortative interlayer degree correlations improve the robustness of multiplex systems to both random [15,34] and targeted attacks based on the selection of nodes with the largest degrees [15]. However, most of those studies have only considered the case of duplex
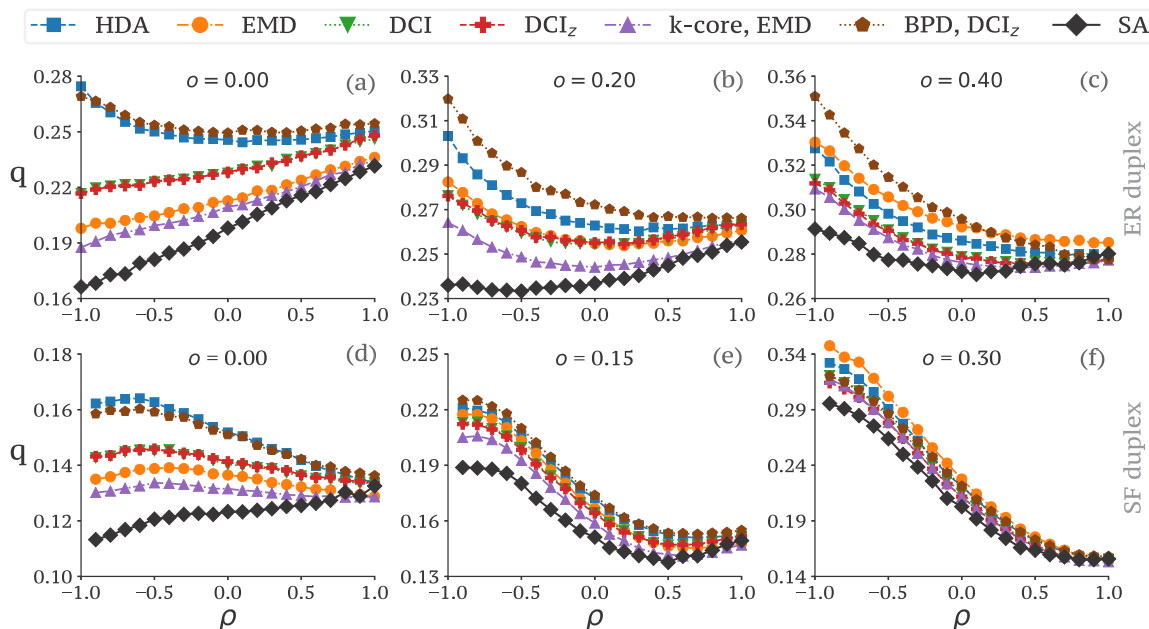
FIG. 4. Size of the critical set q as a function of the interlayer degree correlations coefficient ρ [16] for duplex networks with $N = 10^4$ nodes, whose layers are (a)–(c) Erdős-Rényi graphs with $\langle k \rangle = 5$ and (d)–(f) scale-free networks with $\gamma = 2.6$. For each topology, we report three different overlap conditions: (a) and (d) no edge overlap, (b) and (e) low edge overlap, and (c) and (f) moderate edge overlap. The concurrent presence of interlayer degree correlations and edge overlap strongly affects the robustness of a system against targeted attacks. This is particularly evident when considering duplex networks with heterogeneous degree distributions on each layer [(d)–(f)]. When the edge overlap is non-negligible, duplex networks with maximally negative degree correlations are extremely robust under targeted attacks with respect to their maximally positive counterparts. Conversely, when the overlap is negligible, disassortatively correlated duplex networks are more fragile. Overall, DCI, DCI_z, and the Pareto-efficient strategies that simultaneously combine single- and multiple-layer attacks consistently detect smaller critical sets. The results obtained by simulated annealing are reported for comparison. Results are averaged over 20 realizations. Error bars are smaller than the marker size.

systems with identical degree distributions on the layers and either maximally positive or maximally negative interlayer degree correlations. Here we use the procedure explained in Ref. [16] to tune interlayer degree correlations between the maximally disassortative case (also called maximally negative) and the maximally assortative one (maximally positive). In order to isolate the effect of interlayer degree correlations, we study the performance of the six targeted attack strategies as a function of the interlayer degree correlation coefficient ρ [16], imposing that each realization of the multiplex has $o \approx 0$. To simultaneously account for the joint effect of overlap and interlayer degree correlations, we also consider the sequences of multiplex networks obtained by increasing ρ while keeping the edge overlap fixed at a given value. To obtain those sequences, we first increase the value of interlayer degree correlation [16] and then we set the desired value of edge overlap through biased edge rewiring [44,45] (see Appendix D for a more detailed description of the method).

In Fig. 4 we plot the size of the critical set q identified by the six targeted attack strategies as a function of the interlayer degree correlations ρ and for different values of edge overlap o. We report the results obtained on duplex networks with Erdős-Rényi layers [Figs. 4(a)–4(c)] and with scale-free layers [Figs. 4(d)–4(f)]. Interestingly, in all the scenarios considered the state-of-the-art EMD and HDA are outperformed by one or more of the heuristics proposed in this paper. In particular, the smallest critical set is often obtained by the (k-core, EMD) Pareto-efficient strategy. However,

depending on the interplay between edge overlap and interlayer degree correlations, profound differences among the six methods emerge. For instance, when considering a duplex with Erdős-Rényi layers and negligible edge overlap [Fig. 4(a)], the discrepancy between the overall best strategy (k-core, EMD) and the worst one (HDA) is maximal when $\rho \simeq -1$. In particular, the critical set found by the k-core, EMD strategy when $\rho \simeq -1$ is around 19% (smaller than the one found for $\rho \simeq 1$, i.e., around 23%), while HDA finds a much larger critical set (28%), which is even larger than the one it finds for $\rho \simeq 1$ (25%). As a consequence, the presumed increased robustness of multiplex networks with disassortatively correlated degrees is probably just an artifact of the algorithm used to determine the critical set [15]. By looking at the size of the critical set found by simulated annealing in Fig. 4(a), it seems clear that negatively correlated multiplex systems without overlap are generally hyperfragile compared to positively correlated ones. However, some of the attack strategies considered, including HDA and especially in uncorrelated systems, provide a diametrically opposite picture and suggest that in the absence of edge overlap positively correlated degree sequences are more fragile. The results shown in Figs. 4(b) and 4(c) shed light on the interplay between edge overlap and interlayer correlations. In both cases, the sizes of the critical sets found by the six algorithms are higher than those shown in Fig. 4(a) (i.e., when the edge overlap is negligible). In particular, the (sub)optimal critical set q found by simulated annealing reveals that both edge overlap
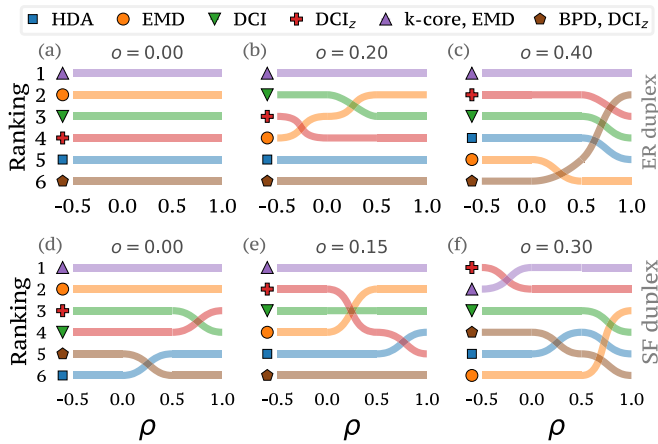
FIG. 5. Rankings of the six targeted attack strategies presented in Fig. 4 for different values of edge overlap and different network topology, namely, (a)–(c) Erdős-Rényi and (d)–(f) scale-free, as a function of four values of interlayer degree correlations. Interestingly, the (k-core, EMD) Pareto-efficient strategy has a considerably better performance in most of the conditions considered. By contrast, the EMD strategy appears to have a good performance only in duplex networks with negligible or small overlap.

FIG. 6. (a) Relative performance of the six targeted attack strategies for real-world duplex networks. Notice that strategies that perform well in synthetic duplex networks are not the best ones when it comes to real-world systems. (b) Overall pair performance rate, defined as the relative number of times either of a pair of algorithms identifies the smallest critical set. In this case, the combination of DCI and DCI$_z$ with other strategies results in the best performance rates. However, the fact that the highest value of pair performance is 58% suggests that we are still far from having fully understood the robustness of real-world systems to targeted attacks. Results are averaged over ten realizations.

and interlayer degree correlations contribute to determine the robustness of a duplex system.

Similar conclusions can be drawn by examining duplex networks with a scale-free degree distribution [see Figs. 4(d)–4(f)]. Also in this case both edge overlap and interlayer degree correlations have a substantial impact on the performance of each algorithm. However, the difference between maximally negative and maximally positive correlated duplex networks becomes more relevant when edge overlap increases, mainly due to the fact that degree heterogeneity on each layer has a stronger impact on the percolation of the MCC. It is interesting to notice here that, since the relative performance of the algorithms considered clearly depends on both interlayer degree correlations and edge overlap, there is no algorithm that clearly outperforms all the others. This is made evident in Fig. 5, where we highlight the behavior of the ranking of the six heuristics based on increasing size of the critical set q (i.e., the algorithm ranked first is the one providing the smallest critical set). Although the (k-core, EMD) Pareto-efficient strategy seems to perform consistently well across the board, being ranked first or second more often that the other five strategies, there are several combinations of layer structure, edge overlap, and interlayer degree correlations for which other algorithms perform much better. An overview of the critical sets found by all the attack strategies we have considered as a function of overlap and correlations is reported in Appendix C.

IV. OPTIMAL PERCOLATION IN REAL-WORLD MULTIPLEX NETWORKS

One of the main aims behind the study of targeted attacks is to try to find efficient ways to mitigate the fragility of real-world infrastructures, which are normally characterized by layer heterogeneity, non-negligible edge overlap, and

interlayer degree correlations. For this reason, we tested the targeted attack strategies presented in this paper in 26 real-world multiplex networks [47]. The size of the critical set found by each of the algorithms is reported in Table I. The systems considered in the table range in size from a few dozens to thousands of nodes, with different values of edge overlap and interlayer degree correlations. Since many of those multiplex networks have more than two layers, for each system we consider the duplex subnetworks corresponding to the pairs of layers yielding the largest MCC, as already done, for instance, in the main text of Ref. [34]. Unsurprisingly, there is no single strategy that works better than all the others in all the cases. What is surprising instead is that those strategies yielding the best performance when considering synthetic duplex systems, e.g., the (k-core, EMD) Pareto-efficient algorithm, do not perform as well in real-world systems. By contrast, DCI and DCI$_z$ quite often find the smallest critical set. This can be easily visualized in Fig. 6(a), where we plot the relative amount of times (i.e., performance rate) that a certain strategy identifies the smallest critical set in all the 26 real-world duplex networks considered. The best-performing strategy here is DCI, with a rate of 42%, followed by DCI$_z$ (35%) and EMD (31%). We also considered the pair performance, which is defined as the relative number of times that at least one of two algorithms identifies the smallest critical set. The results are reported in Fig. 6(b). Remarkably, combinations of targeted attacks including DCI and DCI$_z$ yield the best pair performance rate, where the DCI-EMD pair is able to find the smallest critical set in 58% of the cases. Overall, these results warn against the quest to find a single targeted attack strategy that performs well whatever the multiplex network it is applied to. In particular, the generalization to real-world networks of targeting strategies that perform well in specific classes of synthetic graphs can result in the gross overestimation of the robustness of a system.

TABLE I. Size of the attack sets for all the different strategies considered in this paper for 26 different real-world duplex networks [44,47]. For each data set we report the size of the initial MCC, the value of interlayer degree correlations within the MCC (ρ), and the normalized edge overlap [$o_{norm} = o/o_{max}$, where o_{max} is the maximum overlap in the corresponding configuration model ensemble [44,45] (see Appendix D for details)]. The performance of each strategy heavily depends on the presence of edge overlap and interlayer degree correlations. Interestingly, DCI, DCI$_z$, and Pareto-efficient strategies based on them perform better than all the other strategies. For methods that rely on random tiebreaking, the numbers reported correspond to the minimum value found by the method over ten independent realizations. The size of the minimal attack set found on each duplex set is underlined.

Data set	MCC	ρ	o_{norm}	HDA	k-core EMD	k-core DCI	k-core$_z$ DCI$_z$	CoreHD EMD	CoreHD DCI	CoreHD$_z$ DCI$_z$	CI$_{\ell=2}$ EMD	CI$_{\ell=2}$ DCI	CI$_{\ell=2,z}$ DCI$_z$	BPD EMD	BPD DCI	BPD$_z$ DCI$_z$	CoreHD DCI	CoreHD$_z$ DCI$_z$	EMD	DCI	DCI$_z$	CoreHD	CI$_{\ell=2}$	BPD
Air. FR-U2	28	0.25	0.14	4	5	4	4	4	4	4	4	4	4	4	4	4	4	4	5	5	4	4	4	5
Air. AA-DL	191	0.72	0.60	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	9	8	9
Air. AA-UA	204	0.77	0.58	9	7	7	7	8	8	8	9	8	8	8	8	8	8	8	9	9	9	9	9	8
UK Train L 26-41	59	0.38	0.19	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
UK Train L 30-41	43	0.19	0.18	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
ArXiv L 2-6	916	0.85	0.77	91	119	106	106	139	108	108	101	104	100	115	103	108	121	115	137	104	103	117	104	119
ArXiv L 3-6	790	0.92	0.81	85	84	98	98	115	103	103	85	97	99	96	99	97	100	96	106	97	97	106	87	107
CS Aarhus L 1-5	58	0.32	0.63	14	15	15	12	13	15	15	12	14	15	12	15	14	15	13	15	14	14	18	14	14
FAO L 3-24	193	0.94	0.70	83	85	87	85	86	90	90	82	82	88	82	88	82	86	88	82	82	82	88	82	85
IMDb Com.-Dra.	181	0.97	0.82	88	90	87	89	88	90	90	88	90	90	91	90	89	89	89	90	90	89	89	88	89
Terr. L Tru.-Op.	61	0.22	0.50	18	16	15	15	19	19	19	19	19	17	17	17	19	15	15	15	19	19	17	21	21
Arabid. L 1-2	442	0.65	0.40	34	32	29	30	34	31	31	27	31	31	31	31	31	39	34	35	31	31	33	35	21
Drosoph. L 1-2	299	0.18	0.07	10	9	14	12	10	14	14	9	8	12	11	8	10	10	10	8	8	10	10	12	9
Drosoph. L 1-3	202	0.27	0.06	4	4	3	3	5	4	4	3	3	3	3	3	3	4	4	3	3	3	3	3	3
Drosoph. L 1-4	1024	0.13	0.09	26	26	27	26	35	39	39	24	27	31	25	37	30	42	38	38	27	30	40	33	42
Drosoph. L 2-3	449	0.68	0.35	49	48	50	53	52	50	50	55	51	49	47	48	51	54	51	53	51	51	54	54	52
Homo L 1-2	9312	0.57	0.32	1052	1058	1026	1021	1017	1045	1045	1088	1060	1069	1042	1051	1095	1060	1092	1088	1060	1095	1087	1101	1172
Homo L 1-5	3886	0.31	0.16	312	299	271	297	301	315	315	344	332	319	307	305	347	347	358	329	332	347	341	363	432
Homo L 2-5	4944	0.44	0.17	420	366	406	387	389	413	413	412	434	395	386	410	422	439	435	410	434	422	481	473	534
Hum. HIV L 1-2	144	0.54	0.41	4	4	3	3	4	3	3	4	4	2	4	2	4	3	3	4	4	4	2	4	3
Mus L 1-3	1059	0.56	0.37	60	60	52	50	57	52	52	58	57	65	61	69	60	68	56	61	57	60	60	58	63
S. Cerev. L 1-2	4531	0.36	0.10	785	757	768	769	743	785	785	853	823	787	795	782	831	811	825	811	823	831	859	885	956
S. Cerev. L 1-7	4720	0.28	0.07	982	940	898	905	911	930	930	1108	1000	942	994	958	1047	953	973	1025	1000	1047	1009	1025	1115
S. Pombe L 3-4	1112	0.20	0.14	56	41	50	52	54	59	59	54	54	52	44	54	53	62	60	57	54	53	72	79	75
S. Pombe L 3-6	956	0.14	0.06	39	32	40	35	42	38	38	38	38	35	35	40	37	41	39	45	38	37	50	49	58
S. Pombe L 4-6	2292	0.61	0.01	370	361	353	360	358	366	366	385	371	366	377	369	376	377	372	383	371	376	382	380	404

V. CONCLUSION

Optimal multiplex percolation is characterized by a variety of subtleties and we are probably far from understanding it completely. The massive comparison of many different attack strategies that we presented here has revealed that the performance of all the state-of-the-art attack strategies on uncorrelated multiplex networks does not guarantee their ability to identify sufficiently small critical sets in networks having non-negligible edge overlap and nontrivial interlayer degree correlations. In particular, extensive simulations on synthetic networks have shown that both edge overlap and interlayer degree correlations play an important role in determining the robustness of a system and that their combination can be effectively exploited to modulate the robustness of a system against targeted attacks. One of the main ingredients to identify critical nodes using only local information seems to be their potential to disrupt many paths among their second neighbors, since this would indirectly contribute to the disruption of the LMCC. The duplex collective influence algorithm presented here is based on this assumption and indeed shows a better performance than all the existing state-of-the-art algorithms, especially in duplex networks with non-negligible edge overlap. It seems clear to us that further improvements might possibly be obtained by considering extensions of DCI that take into account the impact on farther-away neighborhoods. Another important ingredient seems to be the possibility to combine structural descriptors on each of the two layers with more genuinely multiplex information. In this respect, the family of Pareto-efficient strategies that we presented here represents a quite promising approach. We find it remarkable that these strategies consistently outperform all the existing algorithms in both synthetic and real-world duplex networks.

One of the main motivations behind studying percolation is to improve our ability to assess the robustness and to mitigate the fragility of real-world networks, i.e., of concrete systems presenting non-negligible edge overlap and nontrivial interlayer degree correlations. The most surprising results indeed came from the analysis of real-world multiplex systems and provide a clear warning against hasty generalizations. On the one hand, the heuristics that are good at finding small critical sets in uncorrelated multiplex networks often perform rather poorly in real-world systems, thus tending to overestimate their robustness. On the other hand, the variability in performance shown by almost all the algorithms we have considered confirms that a fair assessment of the robustness of a multilayer system must be based on the usage and comparison of multiple attack strategies. We believe that these results constitute a solid springboard for a more in-depth investigation of optimal percolation in multilayer systems.

Implementations of the 20 targeted attack strategies used in this paper, and of the algorithm to tune overlap and interlayer degree correlations in synthetic duplex systems, are available [48].

APPENDIX A: ADDITIONAL DETAILS ON DCI and DCI$_z$

1. Dependence of duplex collective influence score on edge overlap

Here we study the character of the DCI score of a node i as a function of k_i^{int}, which is the degree of node i in the intersection graph, obtained by considering all and only the links that exist on both layers. Notice that k_i^{int} is intimately connected to the edge overlap around node i. Indeed, the fraction of edges attached to node i that exist in both layers can be expressed as $o_i = k_i^{\text{int}}/k_i^{\text{aggr}}$. Since we have $k_i^{\text{aggr}} = k_i^{[1]} + k_i^{[2]} - k_i^{\text{int}}$, the DCI score of node i can be rewritten as

$$\text{DCI}(i) = \frac{k_i^{[1]}k_i^{[2]} - k_i^{\text{int}}}{k_i^{[1]} + k_i^{[2]} - k_i^{\text{int}}} \left[\sum_j a_{ij}^{[1]}\left(k_j^{[2]} - 1\right) + a_{ij}^{[2]}\left(k_j^{[1]} - 1\right) \right].$$

The term inside the large square brackets does not depend on k_i^{int}, so we can just focus on the ratio outside, which can be conveniently rewritten as

$$\frac{a - k_i^{\text{int}}}{b - k_i^{\text{int}}}, \tag{A1}$$

where we have set $a = k_i^{[1]}k_i^{[2]}$ and $b = k_i^{[1]} + k_i^{[2]}$. Note that $k_i^{\text{int}} \in [0, \min(k_i^{[1]}, k_i^{[2]})]$ and in particular $k_i^{\text{int}} = 0$ if the neighborhoods of node i at the two layers are disjoint, while $k_i^{\text{int}} = \min(k_i^{[1]}, k_i^{[2]})$ if the intersection between those two neighborhoods is maximal, where the case $k_i^{[1]} = k_i^{[2]}$ corresponds to identical neighborhoods on the two layers. It is easy to show that Eq. (A1) is an increasing function of k_i^{int} for $a > b$, which holds whenever $\min(k_i^{[1]}, k_i^{[2]}) > 1$. This means that, all other things being equal, a node having degree larger than one on both layers will have a larger DCI score if it has a larger edge overlap. Similar reasoning holds for DCI$_z$.

2. DCI in multiplex networks with identical layers

It is easy to show that in a duplex network with identical layers the ranking of nodes induced by the DCI score defined in Eq. (2) coincides with that induced by the CI score on the corresponding aggregated graph when $\ell = 1$. In fact, if the two layers are identical, we have $a_{ij}^{[1]} = a_{ij}^{[2]} = a_{ij} \, \forall \, i, j = 1, \dots, N$ and also $k_i^{[1]} = k_i^{[2]} = k_i^{\text{int}} = k_i^{\text{aggr}} = k_i \, \forall \, i = 1, \dots, N$, so we get $\text{DCI}(i) = 2(k_i - 1) \sum_j a_{ij}(k_j - 1) = 2\text{CI}_{\ell=1}(i)$, which means that the two rankings are identical.

APPENDIX B: TIME COMPLEXITY

1. Time complexity of DCI and DCI$_z$

The adaptive targeted strategies based on DCI and DCI$_z$ require us to recompute the DCI scores of all the remaining nodes after each node is removed. An implementation with simple data structures (basically, the list of neighbors of each node) guarantees a worst-case time complexity $O(N^2 \log N)$, where N is the number of nodes of the graph. Indeed, the
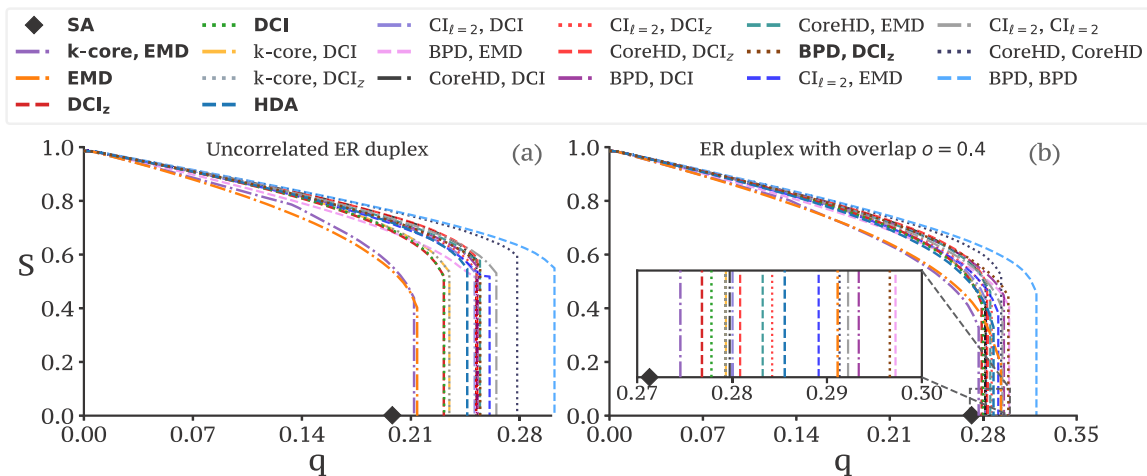
FIG. 7. Percolation diagram for the 20 different attack strategies examined in this paper, for the same duplex networks reported in (a) Fig. 2(a) and (b) Fig. 2(c). Notice that Pareto-efficient strategies combining only single-layer metrics do not perform well in the two cases reported. By contrast, methods which effectively combine single- and multiple-layer information yield the best performance. Labels are sorted in ascending order of size of the critical set from (a). We highlighted in bold the six strategies presented in Fig. 2. Results are averaged over 20 realizations.

initial DCI (or DCI$_z$) score of all the nodes can be computed in $O(K)$ (where K is the total number of edges of the multiplex) and sorted in $O(N \log N)$. The removal of the ith node from the network will modify the DCI scores of all its neighbors on the two layers, which are at most $N - i - 1$. Since we need to keep the list of DCI scores ordered, the usage of simple structures requires us to sort again the scores, which has time complexity $O(N \log N)$ at each step. As a consequence, updating DCI scores throughout the percolation procedure has time complexity $O(N^2 \log N)$. A direct computation and update of the size of the LMCC would run in N^3, but its efficiency can be improved to $O(N^{1.2})$ by using the algorithm explained in Refs. [49,50]. So overall the DCI (DCI$_z$) algorithm for DCI and DCI$_z$ has time complexity $O(N^2 \log N)$. However, the usage of a max heap to store and update the list of DCI scores would guarantee a worst-case time complexity of $O(N^{1.2} + K \log N)$, which is dominated by $O(N^{1.2})$ in sparse graphs.

2. Time complexity of Pareto-efficient strategies

The time complexity of Pareto-efficient strategies can be expressed as $O(F + S)$, where $O(S)$ is the time complexity of computing and updating the scores used for multiobjective optimization, while $O(F)$ is the time complexity of computing and updating the Pareto front throughout the percolation procedure. Identifying the Pareto front at each step has time complexity $O(N \log N)$ when the number of objective functions m is at most $m = 3$, which is the case for all the Pareto-efficient strategies considered in the present paper. If the number of functions to optimize is $m > 3$, then the time complexity becomes $O(N(\log N)^{m-2})$ [51]. As a consequence, $O(F) = O(N^2 \log N)$ in the worst case. The time complexity of computing and updating the scores depends on the details of the functions used, but all the functions we used in this paper are dominated by $O(N^2 \log N)$.

APPENDIX C: ADDITIONAL RESULTS ON SYNTHETIC NETWORKS

In this Appendix we report the results obtained by the multiplex targeted strategies constructed by considering all Pareto-efficient combinations of different of single- and multiple-layer methods. In Fig. 7 we show the percolation
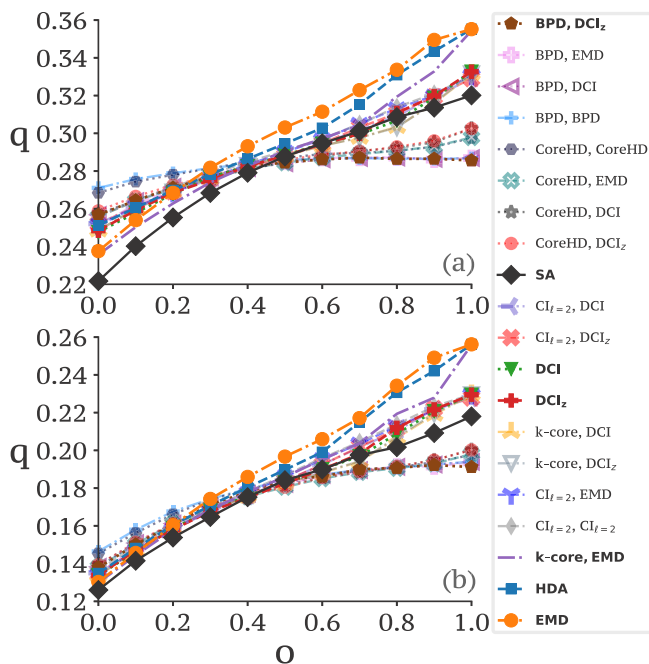

FIG. 8. Size of the critical attack set q as a function of edge overlap for the 20 different attack strategies in the same duplex networks as in Fig. 3. Labels are sorted in ascending order of size of the critical set from (a) when $o = 1$ (i.e., it is analogous to the single-layer percolation). Results are averaged over 20 realizations. Error bars are smaller than the marker size.
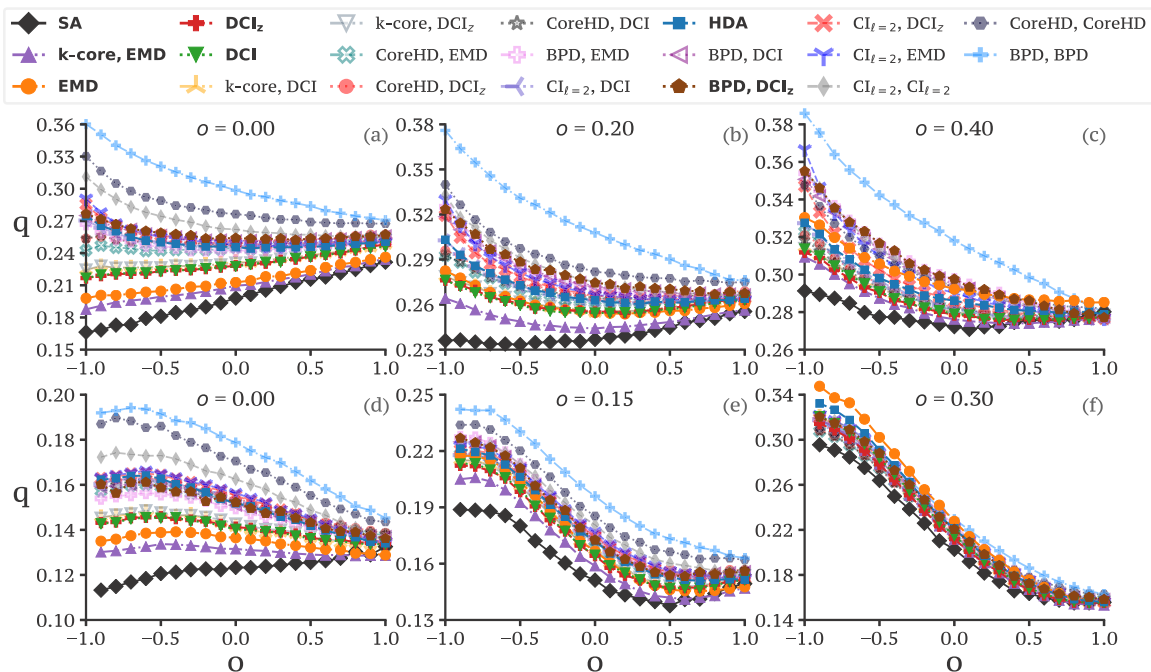
FIG. 9. Size of the critical set q found by each of the 20 attack strategies as a function of the interlayer degree correlations coefficient for the same duplex networks as in Fig. 4. Notice that also in this case the attack strategies which combine only single-layer metrics perform quite poorly compared to the other Pareto-efficient methods. Labels are sorted in ascending order of q when considering $\rho = -1$ of (a). We highlighted in bold the methods presented in Fig. 4. Results are averaged over 20 realizations. Error bars are smaller than the marker size.

diagrams (including also the results already presented in Fig. 2) in duplex networks with no overlap [Fig. 2(a)] and with high overlap [Fig. 2(b)]. It is clear that the strategies that incorporate multilayer information (i.e., HDA, EMD, DCI, and DCI$_z$), as well as a Pareto-efficient strategy that takes one of them into account, perform consistently better than those based exclusively on single-layer metrics [i.e., (CI$_{\ell=2}$, CI$_{\ell=2}$), (CoreHD, CoreHD), and (BPD, BPD)]. This is because, as noted in Refs. [24,33,46], the presence of interdependencies in the multiplex structure deeply affects the overall robustness of duplex networks against random and targeted attacks and this information is not present in either of the layers considered separately.

In Fig. 8 we report the size of the critical attack set q as a function of structural edge overlap. As expected, the best performing targeted strategies for $o = 1$ are those based on BPD, which is the best-performing strategy on single-layer graphs [23,31]. It is interesting to notice that some Pareto-efficient strategies outperform simulated annealing for large values of overlap (same implementation as the one presented in [33] with temperature steps equal to 10^{-7}).

Finally, in Fig. 9 we show the size of the critical set found by each of the 20 strategies for different combinations of interlayer degree correlations and edge overlap (same conditions presented in Fig. 4). It is clear that Pareto-efficient strategies combining multiple- and single-layer information perform better than the others also in this case and especially much better than methods relying only on single-layer metrics. This is even more evident when the duplex has $o \approx 0$ while the gap becomes smaller as the overlap increases, as expected.

APPENDIX D: TUNING INTERLAYER DEGREE CORRELATION AND EDGE OVERLAP

The algorithm to tune interlayer degree correlations and edge overlap in a duplex network used in the paper is based on biased edge rewirings. The procedure works along the same lines of the two procedures to separately tune edge overlap and interlayer degree correlations originally considered in Refs. [16,44,45]. More precisely, to decrease edge overlap we start from two (possibly different) layers, we iteratively select at random two edges on a randomly chosen layer, and we rewire at random the end points of the two links only if such rewiring results in a reduction of the edge overlap. The procedure is iterated until we reach the desired value of edge overlap $o^\star$. In this way, the degree sequence on each layer is preserved throughout the process. Notice that the actual range of edge overlap attainable with this method actually depends on the degree sequences at the two layers.

Increases in edge overlap are obtained with a similar procedure, where a rewiring is accepted only if it results in the increase of edge overlap of at least one of the two edges involved in the rewiring. As a consequence, also this procedure does not modify the degree sequence of each layer. In general, the actual range of edge overlap obtained by this procedure depends on the actual degree sequences of the two layers. For instance, a value of $o = 1$ is attainable only if the degree sequences of the two layers are identical. Since the maximum value $o_{\max}$ of edge overlap for a generic pair of layers is not known *a priori*, in our simulations we compute an approximation of $o_{\max}$ by iteratively increasing the overlap of the system until no further increase is attainable (i.e., the

termination criterion is such that the edge overlap does not increase after 5×10^7 random rewirings).

The procedure for tuning the interlayer degree correlation ρ is identical to the one presented in [16,52]. Briefly, starting from a generic duplex network, we consider R to be the $N \times N$ matrix that accounts for the coupling between the nodes of the two layers. Here the generic entry $r_{ij} = 1$ if node i in layer α corresponds to node j in layer β. Since we are dealing with a duplex network, there is a one-to-one correspondence between the nodes in the two layers, so we have to impose $\sum_j r_{ij} = 1 \, \forall i$. The main idea is that the coupling R can be realized in many ways, and among all these possibilities we choose one that corresponds to a given level of interlayer degree correlation $\rho^\star$. We define the cost function $F(R) = |\rho - \rho^\star|$ and we iteratively modify the structure of the matrix assignment in order to minimize $F(R)$. The minimization procedure is obtained by using a simulated annealing algorithm. In particular, starting from a certain matrix assignment R, we rewire at random two edges of such a matrix in order to obtain a new assignment R'. We then accept the new assignment with probability

$$p = \begin{cases} 1 & \text{if } F(R') < F(R) \\ e^{-\frac{F(R')-F(R)}{\beta}} & \text{otherwise,} \end{cases}$$

where β has the role of an inverse temperature. The algorithm stops when $F(R) < \varepsilon$, where ε is a threshold set by the user. In our simulations, we consider $\beta = 10^{-7}$ and $\varepsilon = 0.005$. The two algorithms to tune edge overlap and interlayer degree correlations can be combined to obtain a duplex network with prescribed values of o and ρ. We start by tuning the value of interlayer degree correlation ρ and we then iteratively increase edge overlap o through biased edge rewiring. Notice that the combination of these two procedures (in this order) does not alter the original degree distribution on each layer of the duplex network.

[1] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, Catastrophic cascade of failures in interdependent networks, Nature (London) **464**, 1025 (2010).

[2] S. Boccaletti, J. A. Almendral, S. Guan, I. Leyva, Z. Liu, I. Sendiña-Nadal, Z. Wang, and Y. Zou, Explosive transitions in complex networks structure and dynamics: Percolation and synchronization, Phys. Rep. **660**, 1 (2016).

[3] R. M. D'Souza, J. Gómez-Gardeñes, J. Nagler, and A. Arenas, Explosive phenomena in complex networks, Adv. Phys. **68**, 123 (2019).

[4] M. De Domenico, A. Solé-Ribalta, E. Cozzo, M. Kivelä, Y. Moreno, M. A. Porter, S. Gómez, and A. Arenas, Mathematical Formulation of Multilayer Networks, Phys. Rev. X **3**, 041022 (2013).

[5] S. Boccaletti, G. Bianconi, R. Criado, C. I. del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña Nadal, Z. Wang, and M. Zanin, The structure and dynamics of multilayer networks, Phys. Rep. **544**, 1 (2014).

[6] G. Bianconi, *Multilayer Networks: Structure and Function* (Oxford University Press, Oxford, 2018).

[7] J. Gao, S. V. Buldyrev, H. E. Stanley, and S. Havlin, Networks formed from interdependent networks, Nat. Phys. **8**, 40 (2012).

[8] D. Cellai, E. López, J. Zhou, J. P. Gleeson, and G. Bianconi, Percolation in multiplex networks with overlap, Phys. Rev. E **88**, 052811 (2013).

[9] F. Radicchi, Percolation in real interdependent networks, Nat. Phys. **11**, 597 (2015).

[10] G. Bianconi and F. Radicchi, Percolation in real multiplex networks, Phys. Rev. E **94**, 060301 (2016).

[11] F. Coghi, F. Radicchi, and G. Bianconi, Controlling the uncertain response of real multiplex networks to random damage, Phys. Rev. E **98**, 062317 (2018).

[12] G. J. Baxter, S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, Avalanche Collapse of Interdependent Networks, Phys. Rev. Lett. **109**, 248701 (2012).

[13] G. J. Baxter, G. Bianconi, R. A. da Costa, S. N. Dorogovtsev, and J. F. F. Mendes, Correlated edge overlaps in multiplex networks, Phys. Rev. E **94**, 012303 (2016).

[14] S. D. S. Reis, Y. Hu, A. Babino, J. S. Andrade, Jr., S. Canals, M. Sigman, and H. A. Makse, Avoiding catastrophic failure in correlated networks of networks, Nat. Phys. **10**, 762 (2014).

[15] B. Min, S. D. Yi, K.-M. Lee, and K.-I. Goh, Network robustness of multiplex networks with interlayer degree correlations, Phys. Rev. E **89**, 042811 (2014).

[16] V. Nicosia and V. Latora, Measuring and modeling correlations in multiplex networks, Phys. Rev. E **92**, 032805 (2015).

[17] D. Cellai, S. N. Dorogovtsev, and G. Bianconi, Message passing theory for percolation models on multiplex networks with link overlap, Phys. Rev. E **94**, 032301 (2016).

[18] I. Kryven and G. Bianconi, Enhancing the robustness of a multiplex network leads to multiple discontinuous percolation transitions, Phys. Rev. E **100**, 020301 (2019).

[19] R. Albert, H. Jeong, and A.-L. Barabási, Error and attack tolerance of complex networks, Nature (London) **406**, 378 (2000).

[20] D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts, Network Robustness and Fragility: Percolation on Random Graphs, Phys. Rev. Lett. **85**, 5468 (2000).

[21] G. Dong, J. Gao, R. Du, L. Tian, H E. Stanley, and S. Havlin, Robustness of network of networks under targeted attack, Phys. Rev. E **87**, 052804 (2013).

[22] D.-w. Zhao, L.-h. Wang, Y.-f. Zhi, J. Zhang, and Z. Wang, The robustness of multiplex networks under layer node-based attack, Sci. Rep. **6**, 24304 (2016).

[23] S. Pei, J. Wang, F. Morone, and H. A. Makse, Influencer identification in dynamical complex systems, J. Complex Netw. cnz029 (2019).

[24] S. Osat, A. Faqeeh, and F. Radicchi, Optimal percolation on multiplex networks, Nat. Commun. **8**, 1540 (2017).

[25] D. Kempe, J. Kleinberg, and É. Tardos, *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, 2003), pp. 137–146.

[26] M. Kitsak, Lazaros K Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse, Identification of influential spreaders in complex networks, Nat. Phys. **6**, 888 (2010).

[27] R. Cohen, S. Havlin, and D. Ben-Avraham, Efficient Immunization Strategies for Computer Networks and Populations, Phys. Rev. Lett. **91**, 247901 (2003).

[28] R. Pastor-Satorras and A. Vespignani, Immunization of complex networks, Phys. Rev. E **65**, 036104 (2002).

[29] P. Clusella, P. Grassberger, F. J. Pérez-Reche, and A. Politi, Immunization and Targeted Destruction of Networks Using Explosive Percolation, Phys. Rev. Lett. **117**, 208301 (2016).

[30] F. Morone and H. A. Makse, Influence maximization in complex networks through optimal percolation, Nature (London) **524**, 65 (2015).

[31] L. Zdeborová, P. Zhang, and H.-J. Zhou, Fast and simple decycling and dismantling of networks, Sci. Rep. **6**, 37954 (2016).

[32] S. Mugisha and H.-J. Zhou, Identifying optimal targets of network attack by belief propagation, Phys. Rev. E **94**, 012305 (2016).

[33] G. J. Baxter, G. Timár, and J. F. F. Mendes, Targeted damage to interdependent networks, Phys. Rev. E **98**, 032307 (2018).

[34] K.-K. Kleineberg, L. Buzna, F. Papadopoulos, M. Boguñá, and M. Á. Serrano, Geometric Correlations Mitigate the Extreme Vulnerability of Multiplex Networks Against Targeted Attacks, Phys. Rev. Lett. **118**, 218301 (2017).

[35] K. Deb, *Multi-Objective Optimization using Evolutionary Algorithms* (Wiley, New York, 2001), Vol. 16.

[36] K. Miettinen, *Nonlinear Multiobjective Optimization* (Springer Science + Business Media, New York, 2012), Vol. 12.

[37] M. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).

[38] V. Latora, V. Nicosia, and G. Russo, *Complex Networks: Principles, Methods and Applications* (Cambridge University Press, Cambridge, 2017).

[39] F. Morone, B. Min, L. Bo, R. Mari, and H. A. Makse, Collective influence algorithm to find influencers via optimal percolation in massively large social media, Sci. Rep. **6**, 30062 (2016).

[40] F. Battiston, V. Nicosia, and V. Latora, Structural measures for multiplex networks, Phys. Rev. E **89**, 032804 (2014).

[41] Ş. Erkol, C. Castellano, and F. Radicchi, Systematic comparison between methods for the detection of influential spreaders in complex networks, Sci. Rep. **9**, 15095 (2019).

[42] M. Ehrgott, *Multicriteria Optimization* (Springer Science + Business Media, New York, 2005), Vol. 491.

[43] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Trade-Offs* (Cambridge University Press, Cambridge, 1993).

[44] A. Santoro and V. Nicosia, Algorithmic complexity of multiplex networks, Phys. Rev. X **10**, 021069 (2020).

[45] M. Diakonova, V. Nicosia, V. Latora, and M. San Miguel, Irreducibility of multilayer network dynamics: The case of the voter model, New J. Phys. **18**, 023010 (2016).

[46] S. V. Buldyrev, N. W. Shere, and G. A. Cwilich, Interdependent networks with identical degrees of mutually dependent nodes, Phys. Rev. E **83**, 016112 (2011).

[47] M. De Domenico, V. Nicosia, A. Arenas, and V. Latora, Structural reducibility of multilayer networks, Nat. Commun. **6**, 6864 (2015).

[48] https://github.com/andresantoro/Multiplex_optimal_percolation.

[49] S. Hwang, S. Choi, D. Lee, and B. Kahng, Efficient algorithm to compute mutually connected components in interdependent networks, Phys. Rev. E **91**, 022814 (2015).

[50] D. Lee, S. Hwang, S. Choi, and B. Kahng, Decremental dynamic algorithm to trace mutually connected clusters, SoftwareX **7**, 273 (2018).

[51] H. T. Kung, F. Luccio, and F. P. Preparata, On finding the maxima of a set of vectors, J. ACM **22**, 469 (1975).

[52] MAMMULT, https://github.com/KatolaZ/mammult.