

Increasing Network Resiliency by Optimally Assigning Diverse Variants to Routing Nodes

Andrew Newell, Daniel Obenshain, *Member, IEEE*, Thomas Tantillo, *Member, IEEE*, Cristina Nita-Rotaru, *Member, IEEE Computer Society*, and Yair Amir, *Member, IEEE Computer Society*

Abstract—Networks with homogeneous routing nodes are constantly at risk as any vulnerability found against a node could be used to compromise all nodes. Introducing diversity among nodes can be used to address this problem. With few variants, the choice of assignment of variants to nodes is critical to the overall network resiliency. We present the Diversity Assignment Problem (DAP), the assignment of variants to nodes in a network, and we show how to compute the optimal solution in medium-size networks. We also present a greedy approximation to DAP that scales well to large networks. Our solution shows that a high level of overall network resiliency can be obtained even from variants that are weak on their own. We provide a variation of our problem that matches the specific communication requirements of applications run over the network (e.g., Paxos and BFT). Also, we analyze the loss in resiliency when optimally assigning variants based on inaccurate information about compromises.

Index Terms—Overlay networks, computer security, integer linear programming

1 INTRODUCTION

NETWORKS with homogeneous routing nodes are constantly at risk as any vulnerability found against a single routing node could be used to compromise all nodes. Diversity can be employed at various levels on the routing nodes to address this problem by improving resiliency against different classes of attacks. In this work, we base resiliency on the number of surviving client-to-client connections offered by the network when under attack. Diversifying the operating system provides protection against common types of attacks that target operating system vulnerabilities [1]; utilizing multi-variant programming protects against programming vulnerabilities or logical programming errors [2], [3]; using different administrative personnel mitigates social engineering or insider attacks [4]. However, there are only a limited number of operating systems, software versions, and personnel to utilize as diverse variants. So then, how does one assign these limited number of diverse variants to the routing nodes in the network to achieve optimal resiliency?

Initially, we assumed that a random assignment of a few diverse variants would perform well. However, we were surprised to find that a random assignment performs rather poorly on a case study topology, in many cases providing less resiliency than using the best single variant at all routing nodes, and occasionally even less resiliency than using the worst single variant at all routing nodes. Clearly, a better approach is necessary to realize the benefits of diversity,

i.e., the use of diverse routing variants to limit the effects of correlated failures.

Our interest in this question arose from constructing a cloud service over a global network of data centers [5]. We needed to have an intrusion-tolerant infrastructure in order to monitor and control the cloud even in the case of sophisticated attacks. While designing intrusion-tolerant protocols for messaging and maintaining consistent state, we realized that without diversity all the nodes could be compromised by a single vulnerability. Inspired by [1], we were especially interested in diversifying the operating system (e.g., Linux, MacOS, and FreeBSD). The additional overhead of managing multiple operating systems within the cloud infrastructure led us to consider only a small number of variants to create diversity.

In this paper, we demonstrate that the way diverse variants are assigned across the network (i.e., which variant is assigned to which routing node) is of utmost importance to the overall network resiliency when the number of variants is smaller than the number of routing nodes in the network. To our knowledge, this work is the first to study the impact of variant assignment to routing nodes on overall network resiliency.

We present a novel problem, the Diversity Assignment Problem (DAP), which specifies how to optimize overall network resiliency when placing diverse variants that are compromised independently at routing nodes. While DAP is NP-hard, we show that it is feasible to solve it optimally on a variety of medium-size random network graphs. We also show an efficient algorithm that approximates DAP well for larger graphs, incurring a relatively small resiliency cost compared with the optimal solution.

To check the applicability of our approach in a real-world setting, we obtained a network graph representative of the global overlay topology used by the above cloud service. Even though this topology was constructed with high availability as the goal (rather than intrusion-tolerance), the

• A. Newell and C. Nita-Rotaru are with the Department of Computer Science, Purdue University, West Lafayette, IN 47907.
E-mail: {newella, crism}@cs.purdue.edu.

• D. Obenshain, T. Tantillo, and Y. Amir are with the Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218.
E-mail: {dano, tantillo, yairamir}@cs.jhu.edu.

Manuscript received 25 Oct. 2013; revised 2 Sept. 2014; accepted 3 Nov. 2014.
Date of publication 19 Nov. 2014; date of current version 13 Nov. 2015.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TDSC.2014.2373392

optimal variant assignment solution to the DAP ensures a system resiliency that is significantly higher than the resiliency achieved by any of the individual variants.

We initially choose an application agnostic metric for network resiliency that captures the expected client-to-client connectivity between all pairs. We investigate the advantages of considering the specific resiliency needs defined by the nature of a distributed application running at the clients. Specifically, we show how to find the optimal assignment for the underlying network supporting either the Paxos [6] or Byzantine Fault-Tolerant (BFT) [7] protocols. When applied to the mentioned global topology, we found that an assignment that is tailored to these application requirements can provide higher resiliency than an assignment that focuses on general network resiliency obtained by maximizing the expected client-to-client connectivity.

Our assignments are based on assumptions of accurate information about compromise probabilities of variants. Having inaccurate information results in a different assignment which can impact the resiliency of the system and the confidence of the network operator in the resulted assignment. We analyze and measure in a realistic scenario these two types of errors resulting from inaccurate information to understand the impact of inaccuracies in compromise probabilities on assignment resiliency and network operator confidence. Our results show that small inaccuracies in information only result in minor errors in assignment and confidence.

The contributions of this paper are as follows:

- We introduce the Diversity Assignment Problem. DAP describes how to assign diversity to routing nodes in order to maximize the probability of each client pair being connected.
- We formulate the DAP using mixed integer programming (MIP) [8] and find the optimal solution on random graphs constructed in a manner reminiscent of real overlay topologies. To support larger graphs, we extend this formulation to a fast greedy approximation and demonstrate results that are relatively close to the optimal solution in such larger graphs.
- We extend our approach to optimize network resiliency for a given application's demands, rather than for overall expected client-to-client connectivity, to maximize system resiliency.
- We analyze the loss in resiliency when optimally assigning variants based on inaccurate information about compromises.

This paper expands on an initial work first published in [9]. We add the three major components in this version. First, we demonstrate why random assignments are poor on our case study topology in Section 3.4. Second, we provide an alternative motivating experiment based on Paxos for our assignment that optimizes connected components in Section 5.3. Finally, we show the effects of inaccuracies in compromise information in Section 6. For space considerations we could not include all experiments that we performed after publishing the conference version. The interested reader can find further new material on additional examples of connected component optimization and a weighted traffic version of the problem in our accompanying technical report [10].

The rest of the paper is organized as follows. Section 2 describes our network and attacker models. Section 3 presents the general DAP along with an optimal solution. Section 4 describes and evaluates a greedy approximation algorithm to solve DAP in larger topologies. Section 5 shows the increased advantage of performing diversity assignment with client application knowledge. Section 6 analyzes how inaccurate compromise information affects assignment. Section 7 lists work related to ours. Section 8 concludes this work.

2 MODEL

We describe the model of the network and attacker which we consider in this work. These models are quite general as our approaches can be applied in various networking contexts with various of diversity techniques. Our motivation started with a scenario of cloud services being provided over a global network of datacenters while diversifying operating systems for improved resilience, but we noticed that the core problem is general to any network.

2.1 Network Model

We assume a network topology of routing *nodes* that provide communication to *clients*. We assume no control over the structure of the network topology as this is fixed based on the constraints of the networking context. In an overlay routing context, network links impose overhead to continuously monitor their latency and loss characteristics, thus the degree at each node must be limited while ensuring the entire network is still well connected. Alternatively, in a wireless context, network links are limited by the physical broadcast range of each node. We assume that we have a set of diverse variants and we can configure each routing node with a single variant. Our network goals are to maximize the number of client connections or an application-specific communication requirement of the clients.

2.2 Attacker Model

We assume that there is no way to configure a routing node that meets our network needs while being completely invulnerable to attacker attempts of compromise. Thus, we adopt a probabilistic attacker model where each variant is compromised with some probability. We capture the benefit of diversity by assuming any pair of variants are compromised independently. We assign a probability that an attacker is able to both find a vulnerability and create a successful exploit against a variant within a given time period, and then any routing node in the network with this variant will become compromised. As our probabilities are with respect to a certain time frame, a full long-term system would need mechanisms to detect and recover compromised variants. We consider such mechanisms as outside the scope of this work. Our probabilistic model of compromise offers a useful way to reason about an attacker's capabilities and measures a network's resilience. Even in realistic scenarios where an attacker is not modeled well probabilistically, we are still raising the bar for the attacker to ensure the attacker must find vulnerabilities and create exploits for different variants of routing nodes.

We assume a byzantine tolerant routing protocol is used for routing to ensure that communication can occur

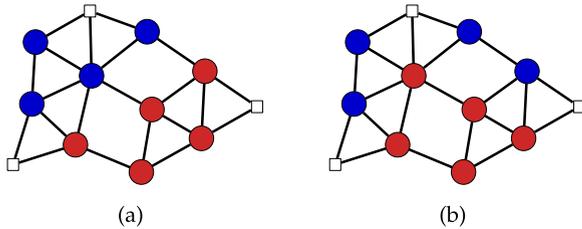


Fig. 1. Example of two assignments on the same topology where routing nodes are circles and clients are squares. We show two possibilities for diversity assignment to nodes where the two variants are red which has a 0.1 probability of being compromised and blue which has a 0.15 probability of compromise. (a) Diversity assignment with 0.838 expected client connectivity. Notice that only one client pair is connected if either red or blue is compromised. (b) Superior diversity assignment that has 0.957 expected client connectivity. Notice that three client pairs are connected if blue is compromised and two client pairs are connected if red is compromised.

between two clients as long as an honest path of routing nodes exists [11], [12].

3 DIVERSITY ASSIGNMENT

In this section we present the Diversity Assignment Problem. DAP describes how to assign variants to routing nodes in order to maximize the probability of each client pair being connected. We then describe existing mixed integer programming techniques and how these can be used to solve DAP. Lastly, we show the effectiveness of this technique on a realistic case study topology when compared with randomly assigning diversity.

3.1 Diversity Assignment Problem

We consider a network consisting of a set of nodes N and a set of clients M . A set of connections are defined among these nodes, so we can represent a network as a graph such as the one in Fig. 1. Each routing node is assigned a variant from the set of variants V , so there are $|V|^{|N|}$ possible assignments. We denote an assignment of one variant for each node as A . Note that $|V| < |N|$. Each variant $v_k \in V$ is associated with a compromise event e_k in the set of all compromise events E , so $|E| = |V|$. The probability of e_k occurring is $P(e_k)$. These events of compromise are independent,¹ so for any two compromise events $e_{k'}$ and $e_{k''}$ the following holds $P(e_{k'} \cap e_{k''}) = P(e_{k'}) * P(e_{k''})$.

We measure the goodness of an assignment of variants with the metric *expected client connectivity* (ECC). This metric is the expected value of the proportion of client pairs that are connected. To compute this value we consider the set of all possible combinations of compromise events C where $|C| = 2^{|E|}$ (C is the powerset [13] of E). An element $c \in C$ is a subset of the compromise events, E , and corresponds to those compromise events occurring while any other compromise events do not occur. We can compute the proportion of clients connected given that those variants are compromised. We consider two clients to be connected if a path of non-compromised nodes exists between them.

1. We make an assumption of independence among compromise events as it simplifies the presentation of the fundamental ideas in this work. We provide an analysis of what occurs when compromise events are not highly positively correlated in Section 6.

Our goal is to maximize the expected client connectivity of a graph by strategically assigning variants. We call this problem the Diversity Assignment Problem.

Definition 1. We define the Diversity Assignment Problem as finding the assignment of variants to nodes which maximizes the expected client connectivity. First, for a given assignment A and set of compromised variant events $c \in C$, we define a connectivity function $f_{A,c}(a, b)$ between two clients a and b as:

$$f_{A,c}(a, b) = \begin{cases} \binom{|M|}{2}^{-1} & \text{if clients } a \text{ and } b \text{ are connected} \\ & \text{by a set of non-compromised} \\ & \text{nodes,} \\ 0 & \text{otherwise.} \end{cases}$$

Then, the expected client connectivity is:

$$\begin{aligned} & E \left[\sum_{\{a,b \in M: a < b\}} f_{A,c}(a, b) \right] \\ &= \sum_{c \in C} \left(\prod_{e_k \in c} P(e_k) \prod_{e_k \notin c} (1 - P(e_k)) \right) \\ & \quad * \sum_{\{a,b \in M: a < b\}} f_{A,c}(a, b). \end{aligned}$$

The Diversity Assignment Problem is:

$$\operatorname{argmax}_A \left(E \left[\sum_{\{a,b \in M: a < b\}} f_{A,c}(a, b) \right] \right).$$

Theorem 1. The Diversity Assignment Problem is NP-hard with two or more variants (proof available in [10]).

We illustrate the meaning of DAP in Fig. 1 with an example topology graph. Figs. 1a and 1b show two ways to assign variants in this graph. Fig. 1b is the superior assignment as more client pairs are connected given that a single variant is compromised. The superiority of this assignment is also reflected by the expected client connectivity values.

3.2 MIP Approach to DAP

Despite DAP being NP-hard, many real-world network topologies are of limited size, so finding the optimal solution is of practical interest. To find the optimal solution, we chose to formulate the problem as a MIP and utilize an existing commercial solver, CPLEX [14]. A MIP is a linear program with the addition of integer constraints. The important implication of these integer constraints is that a MIP is not solvable in polynomial time (while a linear program can be), but these integer constraints allow for formulations of many difficult combinatorial problems. Problems from other domains have also resorted to MIP to find optimal solutions to practical problems in the area of operations research [15], [16], [17]. MIP formulations are good for problems where the optimal is desired and no efficient algorithm is known as many MIP solvers [14], [18], [19] employ a variety of techniques to avoid exhaustively searching the entire space of feasible solutions.

TABLE 1
Notation

Symbol	Description
N	Set of routing nodes. As our notation, these are x, y, z , etc. Depicted by circles in figures.
M	Set of client nodes. As our notation, these are a, b , etc. Depicted by squares in figures.
V	Set of variants. Depicted by colors of circles in figures.
E	Set of all compromise events. We index elements of E and V by k as their elements are related such that each e_k corresponds to the compromise event of the variant v_k .
C	Set of all possible compromise event sets, so $ C = 2^{ E }$. Each element $c \in C$ is a set of compromise events ($e \in E$) that are compromised.
$w_{i,j}$	Constants designating that edge $\{i,j\}$ exists. i and j can be either routing nodes or client nodes. Note that clients should not connect directly to other clients, so $i, j \in M \Rightarrow w_{i,j} = 0$. Depicted by lines between nodes in figures.
$f_{c,a,i,j}$	Measures the amount of flow that starts at client node a and travels on edge $\{i,j\}$ in compromise event set c . i and j can be either routing nodes or client nodes. Also, $c \in C$. This must be a non-negative value.
$s_{v,x}$	The variant assignment of routing node x . $s_{v,x}$ is 1 if x is variant v and 0 otherwise.

Our MIP formulation is set up as a collection of flow problems. We formulate a flow problem for every combination of source client and compromise scenario. Each flow problem aims to send as much flow from a source client to destination clients as possible. The constraints on a flow for a given source client and compromise scenario ensure that: (1) flow in and out of a router are equivalent, (2) source client accepts zero flow, (3) destination clients accept at most one unit of flow, and (4) no flow travels through a compromised router. Thus, maximizing the sum of flow out of a source client will result in a count of the number of connected clients, and this count is weighted by probabilities corresponding to the given compromise scenario. No flow is allowed through routers assigned a compromised variant for a particular compromise scenario. Variant assignment is given by integer variables to ensure each router is assigned exactly one variant. These variables must be integer to ensure a router is of a single variant type.

Table 1 describes each symbol that we use in our MIP formulation. We present the objective function (Equation (1)) followed by each constraint (Equations (2)-(10)).

DAP objective:

$$\max_{s,f} \frac{1}{2} * \binom{|M|}{2}^{-1} * \sum_{c \in C, a \in M, x \in N} \left(\prod_{e_i \in c} P(e_i) \prod_{e_i \notin c} 1 - P(e_i) \right) f_{c,a,a,x}. \quad (1)$$

We maximize the expected client connectivity of the graph, over all compromise events. The first term ($\frac{1}{2} * \binom{|M|}{2}$) ensures that the result will be out of 1, rather than out of the number of possible connections between clients. The two products ensure that each possible compromise event is weighted by the probability that it happens. The f term is a

measure of how much flow the given client a can push out onto the network (specifically, $f_{c,a,i,j}$ measures the amount of flow that started at source client a that travels on edge $\{i,j\}$ in compromise case c). Because of all the constraints below, this is exactly a measure of how many other clients a can connect to.

Variant constraints (I):

$$s_{v_i,x} = \{0, 1\}, v_i \in V, x \in N. \quad (2)$$

Routing nodes must be either entirely of a variant or entirely not of that variant. Fractional assignments are not allowed.

Variant constraints (II):

$$\sum_{v_i \in V} s_{v_i,x} = 1, x \in N. \quad (3)$$

Routing nodes must be exactly one variant.

Node flow constraints:

$$\sum_{i \in N \cup (M - \{a\})} f_{c,a,x,i} - \sum_{i \in N \cup \{a\}} f_{c,a,i,x} = 0, c \in C, a \in M, x \in N. \quad (4)$$

The flow (originating at source client node a) entering routing node x must equal the flow (originating at source client node a) exiting routing node x . This is enforced for each of the $|M|$ clients and for each of the $|N|$ nodes, separately. In other words, flow cannot get stuck in the middle of the network; it has to end at client nodes.

Client flow constraints (I):

$$\sum_{x \in N} f_{c,a,x,b} \leq 1, c \in C, a, b \in M, a \neq b. \quad (5)$$

A client cannot accept more than one unit of flow from another client. This is so that we can count the total flow out of the source client to get the number of connected clients. Despite this constraint being ≤ 1 , it can only take a value of 0 or 1 due to the other constraints and the objective. For the CPLEX solver [14], it is more efficient to enforce fewer integer constraints whenever possible.

Client flow constraints (II):

$$f_{c,a,x,a} = 0, c \in C, a \in M, x \in N. \quad (6)$$

Traffic cannot start and end at the same client. In other words, a client cannot send to itself. Note that $\{x, a\}$ is any incoming edge into a .

Client flow constraints (III):

$$f_{c,a,b,x} = 0, c \in C, a, b \in M, x \in N, a \neq b. \quad (7)$$

A destination client cannot send out flow. So, flow cannot use a client to reach other clients.

Topology constraints:

$$f_{c,a,i,j} \leq (|M| - 1) * w_{i,j}, c \in C, a \in M, i, j \in (N \cup M). \quad (8)$$

Any pair of nodes with no edge between them (i.e., $w_{i,j} = 0$) cannot have any flow directly between them. It also

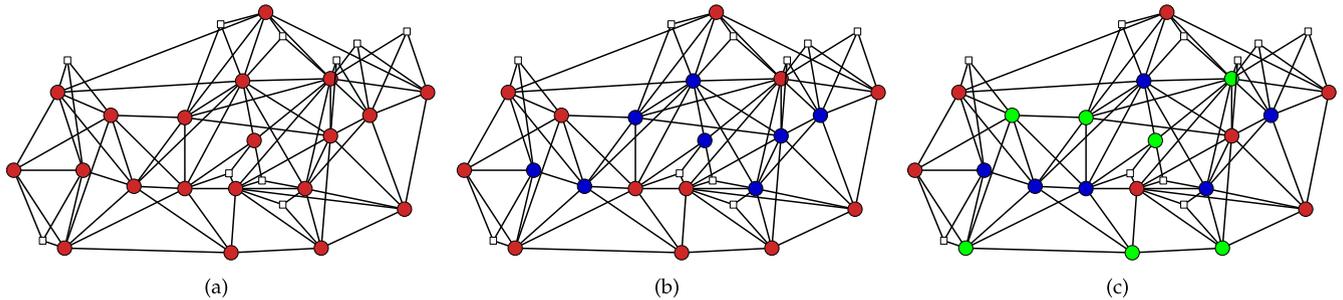


Fig. 2. Optimal assignments on case study topology: (a) one variant assignment achieves 0.9 expected client connectivity, (b) two variants assignment achieves 0.985 expected client connectivity, (c) three variants assignment achieves 0.997 expected client connectivity.

underlines the fact that up to $|M| - 1$ units of flow originating at the same client can share the same edge.

Variant flow constraints (I):

$$f_{c,a,i,x} \leq (|M| - 1) * \min_{e_i \in C} (1 - s_{v_i,x}), \quad (9)$$

$$c \in C, a \in M, x \in N, i \in N \cup M.$$

The amount of flow out of a routing node must be 0 if that node is compromised. It also underlines the fact that no edge can carry more than $|M| - 1$ units of flow from any source client node a .

Variant flow constraints (II):

$$f_{c,a,i,x} \leq (|M| - 1) * \min_{e_i \in C} (1 - s_{v_i,x}), \quad (10)$$

$$c \in C, a \in M, i \in N \cup M, x \in N.$$

The amount of flow into a node must be 0 if that node is compromised. It also underlines the fact that no edge can carry more than $|M| - 1$ units of flow from any source client node a .

3.3 DAP on the Case Study Topology

We investigate the benefit of optimal diversity assignment on a realistic overlay network topology. Then, various assignments of diversity are shown on the case study topology with their corresponding expected client connectivity. We show assignments for DAP with increasing number of variants being used, and we investigate random assignments as a comparison with the optimal solution.

For a case study topology, we took a connectivity graph from a cloud network provider [5]. The nodes of the graph represent data centers located around the globe. Each node is assigned a single variant which means that the overlay routing element at that data center will utilize the selected variant. The edges of the graph represent overlay connectivity used on that cloud to connect the different data centers. This connectivity is provided by a number of Internet Service Providers at each data center. The clients in the graph represent either clients external to the cloud or infrastructure components of the cloud. Each client has multiple connections to the cloud to avoid a single point of failure. In this example, we use three connections as that level of connectivity was quite prevalent in that network. This connectivity graph was designed with resiliency in mind, and without any consideration for diversity.

We assume some hypothetical scenario with three diverse variants represented by red, blue, and green having

a 0.1, 0.15, and 0.2 probability of being compromised over some arbitrary period of time, respectively. Note that this example, while simplistic, provides an interesting insight into the benefits and risks of diversity.²

Fig. 2a shows the optimal solution when only a single variant can be used. All the nodes are assigned with the least vulnerable variant. This corresponds to the situation where no diversity is used. The resulting network achieves an expected client connectivity of 0.9.

Fig. 2b shows the optimal solution when two variants can be used. Each node is assigned with either of the two least vulnerable variants. The resulting network achieves an expected client connectivity of 0.985. Note that this is better than either variant by itself.

Fig. 2c shows the optimal solution when three variants can be used. The resulting network achieves an expected client connectivity of 0.997. Notice that the optimal solution finds an assignment where any single variant is capable of connecting all clients. By adding a third, more vulnerable variant actually makes the system significantly more resilient.

As stated before, in this example, each client is connected to three routing nodes. If clients do not have at least three potential entry points into the network, then the availability of the connection is limited by the variants of the routing nodes that they are connected to. For example, if each client only connects to a single routing node, that connection would fail if either of the entry-point routing nodes is compromised. This is much more likely to occur than if there are three such entry-point routing nodes for each client, requiring at least three routing nodes to be compromised to cut the connection.

In this example, including variants *that have a higher but independent probability of being compromised* improves the overall system resiliency. This may be counterintuitive, as adding weaker components to a system usually makes it weaker, not stronger. The independence of the different variants and the overall robustness of the network mean that adding additional, more vulnerable variants makes a system more resilient.

As discussed earlier, random assignment could be used instead of the optimal MIP approach. One might expect this approach to do well, since randomness often helps in

2. The purpose of these values is to give preference to one variant over another and to quantify an estimate of the system resiliency with diversity. While we select numbers to illustrate the main concepts, the resulting assignment would not be significantly different if other values were selected.

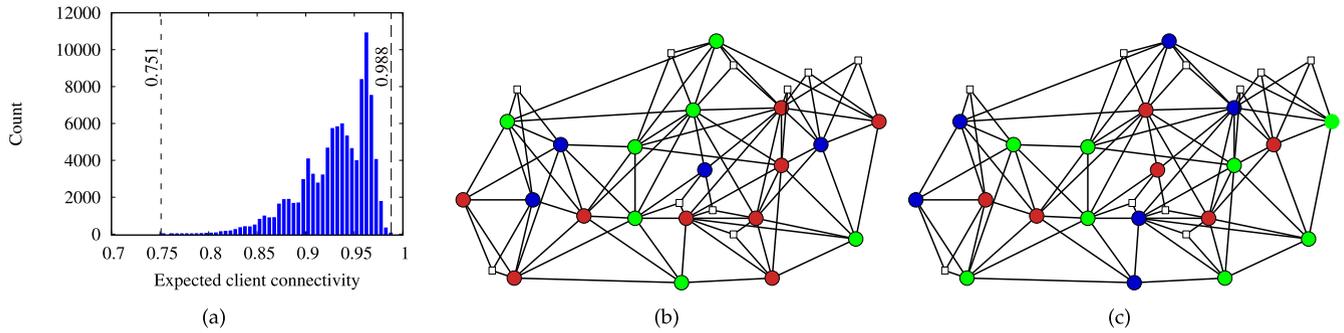


Fig. 3. Random and greedy assignments on case study topology: (a) histogram of expected client connectivity of 100,000 random assignments with vertical lines displaying the lower and upper bounds, (b) random assignment achieving 0.881 expected client connectivity, and (c) greedy assignment achieving 0.992 expected client connectivity.

adding diversity to systems. However, this does not necessarily lead to a good result. An example graph can be seen in Fig. 3b. This graph achieves an expected client connectivity of only 0.811, much worse than any of the other three graphs. In fact, it barely outperforms the worst of the three variants. This example graph comes from the bottom 1 percent of possible assignments and is given as an example of what could occur if the diversity assignment is not considered carefully.

Fig. 3a is a histogram created with data from 100,000 random assignments on the case study topology. For this data set, the minimum and maximum are 0.751 and 0.988 respectively. The mean is 0.931 and the median is 0.937. As can be seen, most of the random assignments perform better than if the best variant is used by itself ($0.937 > 0.9$). However, very few of the random assignments come close to performing as well as the optimal assignment found by MIP.

The optimal solution of 0.997 expected client connectivity exists while the best random solution out of the 100,000 random assignment shown in Fig. 3a was 0.988 expected client connectivity. Thus, even the best random solution out of numerous trials does not achieve the optimal solution. We define *expected client disconnectivity* to be the expected probability that communication between a client pair is broken, and this value is equivalent to $(\text{expected client disconnectivity}) = 1 - (\text{expected client connectivity})$. In terms of expected client disconnectivity the best random solution is 0.012 while the optimal solution is 0.003, so a client-to-client connection is broken four times less often with the optimal assignment.

Interestingly, the difference between what the optimal solution provides and the probability that at least one of the variants is non-compromised provides a metric for the quality of the connectivity resiliency of the graph.³ Ideally, we would want this distance to be zero, as in Figs. 2b and 2c of the provided example.

3.4 Near-Optimal Assignments on Case Study Topology

We aim to further understand why it is difficult to find an optimal solution, given that such an optimal solution is several factors better than random assignments from the perspective of the expected client disconnectivity. We compute the set of all assignments near the optimal solution in terms

of expected client disconnectivity. The number of assignments found compared to the size of the search space further supports our claim that random assignments are typically much worse than the optimal assignment. Thus, techniques to search for optimal assignments (like the ones we propose in this work) are important for any network aiming to achieve high resilience through diversity.

We search for solutions within a *disconnectivity factor* of the optimal solution. This value is computed from a given expected client disconnectivity as follows ($\text{disconnectivity factor} = (\text{expected client disconnectivity}) / (\text{OPT})$ where OPT is the optimal expected client disconnectivity). Intuitively, a disconnectivity factor of two for an assignment implies that clients on average are disconnected twice as much as the optimal assignment.

Exhaustive search of the entire search space is prohibitively expensive for the three variant case, and we could not use this strategy to find all near-optimal solutions. However, we were able to find all solutions within a factor of optimal by leveraging advanced features of MIP solvers. After finding an optimal solution, the solver can be set to continue searching for solutions. The solver avoids exhaustively searching the entire space by eliminating large portions of the search space through its branch and bound techniques. Given that the number of solutions found is small, this procedure is quite efficient.

Fig. 4 shows the number of solutions within a small factor of the optimal solution for the three variant scenario (note the log-scale of the y -axis). We show the proportion of the search space that these solutions represent on the right y -axis. The proportion of the search space indicates the probability that a random assignment has of

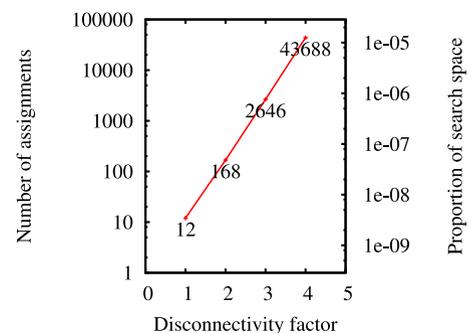


Fig. 4. Number of solutions within a given disconnectivity factor bound for the case study topology.

3. Thanks to Bob Balzer for this observation.

achieving an assignment within a small factor of the optimal solution. Thus, a random assignment has a probability of $3 \cdot 10^{-9}$ to achieve optimal, so that would require on the order of a billion topologies to be assigned and evaluated to find an optimal solution. The visually linear trend in this figure implies an exponential trend in the data due to the logscale of the y -axis. Thus, the number of solutions within a factor of optimal decreases exponentially with respected to decreasing factor, and this implies searching exponentially more assignments to expect to find such a solution.

4 SCALING DIVERSITY ASSIGNMENT

DAP is not tractable for large topologies since DAP is NP-hard (see Theorem 1). To scale to larger topologies, we sacrifice optimality in order to ensure the algorithm completes within a polynomially-bounded time. In this section we present the Approximate DAP (A-DAP), a greedy approach to A-DAP, an example on the case study topology, and an evaluation on random topologies.

4.1 Approximate DAP

A-DAP is similar to DAP, but A-DAP does not require that the problem be solved optimally. By relaxing this condition, we aim to find algorithms that run in polynomial time which are able to find large values of expected client connectivity. We do not formally define any restrictions on the goodness of the approximations as it is an open problem of whether a reasonable bound can be placed on the expected client connectivity achieved by a deterministic polynomial time algorithm. Instead, we used random topologies to validate the goodness of expected client connectivities achieved by a greedy approach to A-DAP when compared with the optimal.

4.2 Greedy Approach to A-DAP

Our greedy approach incrementally assigns nodes to variants. At each incremental assignment, the algorithm considers several candidate assignments and selects the one which provides the best immediate results. For a candidate set of incremental assignments we consider sets of nodes which can connect a client pair by a variant, so we consider at most $\binom{M}{2} * |V|$ candidate variant assignments. For a given client pair a and b and variant i , we compute the minimal number of unassigned nodes which must be assigned i to connect a and b by nodes assigned i . After this computation we have two values: the increase in expected client connectivity α and the number of newly assigned nodes β .

Given a set of candidate assignments that each have an α and β value, we select the one which maximizes $\frac{\alpha}{\beta}$. It is obvious why we want to find large α values, but it is equally important to ensure the β value is small as well. Smaller values of β allow for more nodes to remain unassigned and to be used to connect more client pairs by other variants in future assignments. This approach is analogous to the greedy choice in bin packing, as we select items with the highest payoff versus weight ratio to ensure that items are selected that increase overall payoff while allowing for more items to be picked in the future. Note, that $\beta = 0$ is a trivial case where the candidate is simply removed from consideration as the client pair

is already connected via the considered variant. We provide more details, including pseudo-code, in the technical report version of this work [10].

4.3 A-DAP on the Case Study Topology

We consider the same scenario as in Section 3.3 with three variants. Fig. 3c shows the assignment found by our greedy solution which achieves 0.992 expected client connectivity. Notice that all clients are connected via just the blue or green variants. However, two clients remain disconnected from the rest if only the red variant is uncompromised. The optimal solution found with the MIP formulation finds an assignment which connects all clients as long as any single variant is uncompromised. This loss of expected client connectivity is due to the greedy algorithm making choices in the early steps of the algorithm to connect clients via blue and green variants (the more resilient variants) which leaves fewer choices to connect clients via the red variants. The greedy approach for the A-DAP took 0.38 seconds to complete while the MIP approach for the DAP took 396.13 seconds to complete. With far less computational requirements, the greedy algorithm does outperform the best of the 100,000 random assignments (0.988 client connectivity) and comes close to the optimal solution.

4.4 A-DAP on Random Topologies

We answer the following three questions through simulation.

- 1) How does the goodness of the assignment of the greedy algorithm compare to other algorithms (random assignment and optimal) for the DAP on typical topologies?
- 2) How does the running time of the greedy algorithm for the A-DAP and the MIP approach for the DAP vary with typical topologies created with different parameters?
- 3) What are trends in the expected client connectivity over all the assignment algorithms when varying topology parameters?

Simulation methodology. We use expected client connectivity and running time to evaluate each algorithm. Expected client connectivity is a measure of how well the algorithm performs. Running time is a measure of how quickly the algorithm will terminate with an expected client connectivity.

We generate random topologies for given parameters of number of nodes and density. Density is the average number of neighbors each node has. We place the desired number of nodes and five clients uniformly at random in a two-dimensional square. Any client and node within a calculated communication radius have an edge between them (except pairs of clients, which do not have an edge). The communication radius is selected to ensure the desired average density. Topologies constructed in this way are obviously representative of wireless contexts, but they are also quite similar to overlay topologies, because overlay topologies include many short, well-behaved links.

Given topology parameters, we create 100 random topologies and run the optimal, greedy, and random algorithms on these topologies. We average the expected client connectivity and running times obtained for each algorithm over

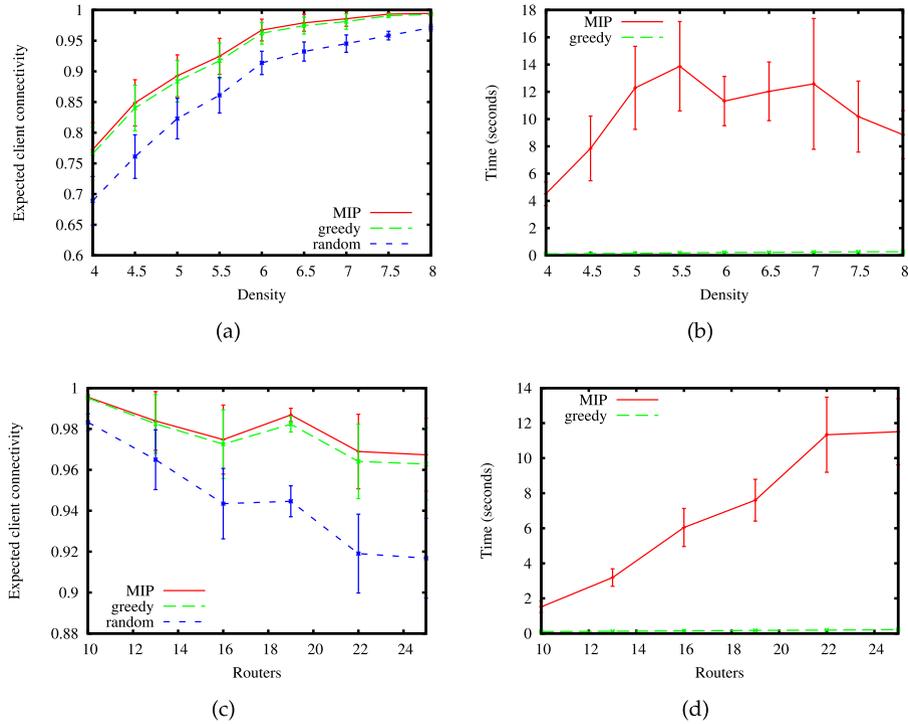


Fig. 5. Experiments for the random, optimal (MIP approach), and greedy algorithms. (a) and (b) show results of random, optimal, and greedy algorithms on random topologies with 25 nodes and varied density. (c) and (d) show the random, optimal, and greedy algorithms on random topologies with six density and varied nodes.

the 100 runs and show 95 percent confidence intervals for those averages. For the running time values of the MIP formulation, it is important to note that we use the software package CPLEX with a quad-core 3.4 Ghz Intel processor which does leverage all cores.

The results are shown in Fig. 5. We describe how they answer each of the initial questions that we proposed.

Question 1. The goodness of an algorithm's assignment is the expected client connectivity. This is upper-bounded by the optimal value (which the MIP approach always achieves). The greedy algorithm outperformed the random assignment and was quite close to the optimal value, independent of varying either density (Fig. 5a) or the number of nodes (Fig. 5c).

Question 2. The running time of the greedy algorithm is on the order of milliseconds, which is barely visible when compared to the running time of the MIP-based approach. Fig. 5b shows the MIP approach running time for varying density values. The running time is low for small density values since most variant assignments result in poor expected client connectivity, allowing the branch-and-bound algorithm of CPLEX to avoid searching the majority of variant assignments. The running time is also low for high density values since a dense graph has many possible optimal assignments and the branch-and-bound algorithm can terminate early after finding any of them. Thus, the problem is the most difficult for networks with moderate density values. The running time of both algorithms when varying the size of the network is shown in Fig. 5d. The MIP approach running time grows nearly linearly over these input parameters, but this relationship is potentially exponential according to Theorem 1. The MIP approach running time is still significantly greater than the greedy approach.

For many networking scenarios the running time of the MIP is reasonable, days or weeks. However, as the problem is NP-hard, on much larger topologies, the MIP approach could take years to produce optimal results, making such a technique prohibitive. The greedy algorithm running time increases polynomially with the topology size resulting in acceptable performance on larger topologies.

Question 3. The trend of expected client connectivity is similar among all three algorithms. The expected client connectivity increases as density increases (Fig. 5a), which is expected since more edges allow more possibilities for clients to become connected. The expected client connectivity decreases as the number of nodes increases (Fig. 5c). By keeping the density constant and increasing the number of nodes, the graph becomes less connected and therefore less resilient.

From these results we see that the greedy algorithm outperforms the random algorithm while being quite close to the optimal solution, and the greedy algorithm is far more efficient in terms of running time and is polynomially-bounded while the MIP formulation is not. Hence, on larger topologies where the MIP formulation cannot be computed, the greedy algorithm is a decent substitute. Another interesting result is that the expected client connectivity decreases with more nodes when keeping the density constant. So, the density or node degree must increase to retain high levels of expected client connectivity when the number of nodes increases in the topology.

5 DIVERSITY ASSIGNMENT FOR SPECIFIC APPLICATIONS

Certain distributed systems that maintain consistent state pride themselves on their ability to tolerate part of the system

failing. State machine replication protocols with this property include Paxos [6], Byzantine Fault Tolerance (BFT) [7], Prime [20], and Aardvark [21], where Prime and Aardvark give additional performance guarantees even while the system is under attack. These protocols explicitly state their assumptions about the proportion of replicas that must be correct for safety and liveness properties to hold. However, an equally important consideration is that a sufficient number of correct replicas must be able to communicate with each other via the underlying network. If we view the state machine replicas as clients of the underlying network, then applying diversity to the network improves the resiliency of the overall system.

We use these state machine replication protocols as an example of how to customize DAP for a specific client application. State machine replication protocols have specific connectivity needs among replicas that must be satisfied to ensure safety and liveness. We show how DAP is customized to better ensure the network meets these requirements, and we show how such customization can be helpful in a realistic scenario. The steps we take here to customize DAP can be followed to create other versions that meet the specific connectivity needs of other distributed systems.

The expected client connectivity from DAP maximizes the expected value of the proportion of client pairs that are connected. This is a reasonable metric for resiliency of many applications, and it could even work well for state machine replication in certain scenarios. However, an approach that takes into account the connectivity requirements of the specific application (in this case, state machine replication) may result in higher overall resiliency. We refine DAP to exactly match the needs of a replicated state machine protocol by maximizing the probability that a specific sized connected component exists among the replicas.

5.1 Connected Component DAP (CC-DAP)

The goal of this algorithm is to optimize the probability that g clients can communicate with each other. The connected component size g can be derived from the specific state machine replication protocol. We denote this problem as the Connected Component Diversity Assignment Problem with formal details in Definition 2 (we use the notation from Table 1). Unsurprisingly, this problem is also NP-hard as stated in Theorem 2.

Definition 2. *The Connected Component Diversity Assignment Problem is to find the assignment of variants to nodes which maximizes the probability of a component of clients being connected. First, we define the random variable X_A which is the size of the largest connected component of clients given a variant assignment A . This variable is random as it depends on the random events E . Then, the Connected Component Diversity Assignment Problem is:*

$$\operatorname{argmax}_A (P(X_A \geq g)).$$

Theorem 2. *The Connected Component Diversity Assignment Problem is NP-hard with two or more variants (proven in [10]).*

5.2 MIP Approach to CC-DAP

For the MIP formulation we keep the constraints in Equations (2)-(10) from Section 3.2, reformulate the objective

function, and add new constraints. Our new objective and constraints include new variables which are used to keep track of which subset of clients are used for a connected component $\beta_{c,a}$ as well as variables to check if the connected component is large enough α_c . We describe the purpose of the new objective and each new constraint in detail to show how it captures the CC-DAP problem.

CC-DAP objective:

$$\max_{s,f,\alpha,\beta} \sum_{c \in C} \left(\prod_{e_i \in c} P(e_i) \prod_{e_i \notin c} 1 - P(e_i) \right) \alpha_c. \quad (11)$$

We maximize the probability that a g -sized connected component exists, over all compromise events. The two products ensure that each possible compromise event is weighted by the probability that it happens. α_c is 1 if a connected component of size g is present under compromise event c and 0 otherwise.

Component constraint (I):

$$\alpha_c = \{0, 1\}, c \in C. \quad (12)$$

A g -sized connected component either exists under compromise event c , or it does not.

Component constraint (II):

$$\beta_{c,a} = \{0, 1\}, c \in C, a \in M. \quad (13)$$

$\beta_{c,a}$ is 1 if client a is in the g -sized connected component under compromise event c , and 0 otherwise.

Component constraint (III):

$$g = \sum_{a \in M} \beta_{c,a}, c \in C. \quad (14)$$

A valid connected component under compromise event c must be of size g . In any other case, this constraint will not be met. Note, if a larger connected component could exist, this constraint ensures that only g clients are considered, which is required for other constraints.

Component flow constraint (I):

$$f_{c,a,x,b} \leq \beta_{c,b}, c \in C, a, b \in M, x \in N, a \neq b. \quad (15)$$

A client b , in the connected component under compromise event c , cannot accept more than one unit of flow from another client a . If b is not in the connected component, it will not accept any flow.

Component flow constraint (II):

$$f_{c,a,a,x} \leq (g-1) * \beta_{c,a}, c \in C, a \in M, x \in N. \quad (16)$$

A client a , in the connected component under compromise event c , cannot send more than $g-1$ units of flow, enough for every other client in the connected component. If a is not in the connected component, it will not send any flow.

Component satisfaction constraints:

$$g * (g-1) * \alpha_c = \sum_{a \in M, x \in N} f_{c,a,a,x}, c \in C. \quad (17)$$

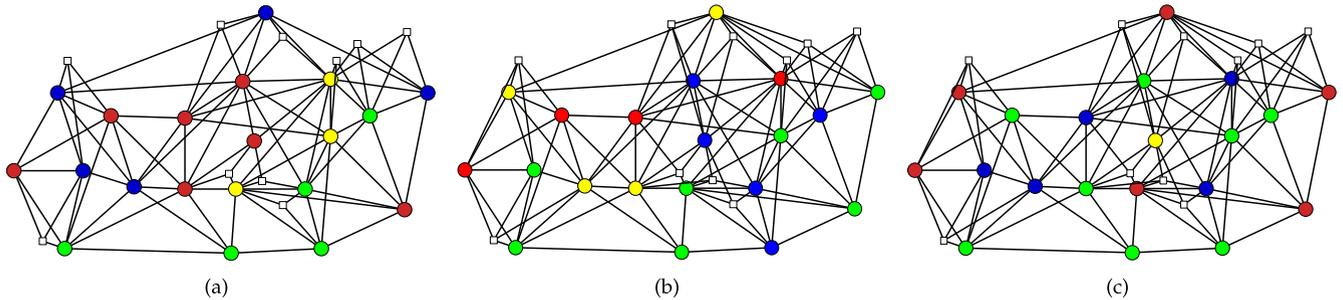


Fig. 6. Assignments illustrating effectiveness of application-specific assignments with CC-DAP: (a) assignment with CC-DAP where Paxos has a probability of 0.99925 to make progress and expected client connectivity is 0.9675, (b) assignment with CC-DAP where BFT has a probability of 0.99925 to make progress and expected client connectivity is 0.9806, and (c) assignment with DAP where BFT has a probability of 0.997 to make progress and expected client connectivity is 0.9975. Note the differences in topologies as (a) has 3 connections from each client to servers while (b) and (c) have 4.

If there exists a g -sized connected component under compromise event c , then there are a total of $g * (g - 1)$ units of flow in the network. If no such connected component exists, the total flow is 0.

5.3 CC-DAP for Paxos

Scenarios where DAP connects all client pairs by every variant individually are trivial for CC-DAP, since an optimal DAP assignment is also an optimal CC-DAP assignment. Thus, we slightly change the setup from Section 3.3 to ensure a non-trivial comparison between DAP and CC-DAP. In the topology we use for Paxos, we add a new variant v_4 where $P(e_4) = 0.25$ represented in the figures by the color yellow. In the topology we use for BFT, we start with the topology used for Paxos and add new connections between clients and routing nodes. BFT requires this extra modification of including new connections since the nature of BFT requires larger connected components.

Paxos maintains consistent state given that there are at most f_s fail-stop failures when using a total of $n = 2f_s + 1$ replicas. In the Paxos scenario, we assume replicas may be partitioned from each other due to attacks on the routing nodes. A client being partitioned from the others is equivalent to a fail-stop failure. For the purposes of this example, we do not consider any other forms of failure, that is, the network may fail but the replicas themselves do not fail. Given that we have 10 replicas in total, this implies that $f_s = 4$. As a result, the required connected component size is $g = n - f_s = 6$.

Fig. 6a shows the assignment when using the MIP approach for CC-DAP while Fig. 2c from before shows the assignment when using the MIP approach for DAP. In Fig. 6a, the probability that six of the clients will be able to communicate is 0.99925 with an expected client connectivity of 0.9675. In contrast, in Fig. 2c, the probability that six of the clients will be able to communicate is only 0.997 while having an expected client connectivity of 0.997 as well. In essence, CC-DAP is able to sacrifice some of the expected client connectivity to increase the probability that a connected component of the desired size will be present.

5.4 CC-DAP for BFT

BFT tolerates up to f Byzantine failures when using a total of $n = 3f + 1$ replicas. We will view these f failures as a combination of f_b , Byzantine replicas, and f_s , fail-stop replicas (indistinguishable from replicas that have been partitioned

away). The choice of values for f_b and f_s are left to the system designer. There is trade-off between f_b and f_s , governed by the trustworthiness of the replicas versus the trustworthiness of the network routing nodes, but further details are beyond the scope of this paper. For our example, we choose $f_b = 1$. Given that we have 10 replicas in total, implying that $f = 3$, the system can tolerate two replicas being partitioned away ($f_s = 2$) and still tolerate one Byzantine fault. As a result, the required connected component size is $g = n - f_s = 8$.

For the results of assignments for BFT, we observe a similar trend to the results of the Paxos scenario. Fig. 6b shows the assignment when using the MIP approach for CC-DAP that achieves a probability of 0.99925 that eight of the clients communicate. Fig. 6c shows the assignment when using the MIP approach for DAP which has only a probability of 0.997 that eight of the clients communicate.

6 ERRORS IN COMPROMISE INFORMATION

Up to now, we have assumed the true assignment compromise values are known and independent with each other. In a realistic scenario, these assignment values could be selected based on expert opinion or extracted from real-world statistics. Both techniques cannot be perfectly accurate. In this section we investigate what occurs when assignment is based on imperfect information.

6.1 Methodology to Investigate Erroneous Information

We establish certain parameters and values that we use to investigate the effects of errors in information.

We define three scenarios for obtaining an ECC from solving DAP:

- `A_ECC_A_INFO` is the ECC value based on available information for an assignment solved with the available information. This is the connectivity that a network operator expects when using an assignment based on solving DAP with available information.
- `R_ECC_A_INFO` is the ECC value based on real information for an assignment solved with available information. This is the realistic connectivity that a network operator will actually achieve when using an assignment based on solving DAP with available information.
- `R_ECC_R_INFO` is the ECC value based on real information for an assignment solved with real

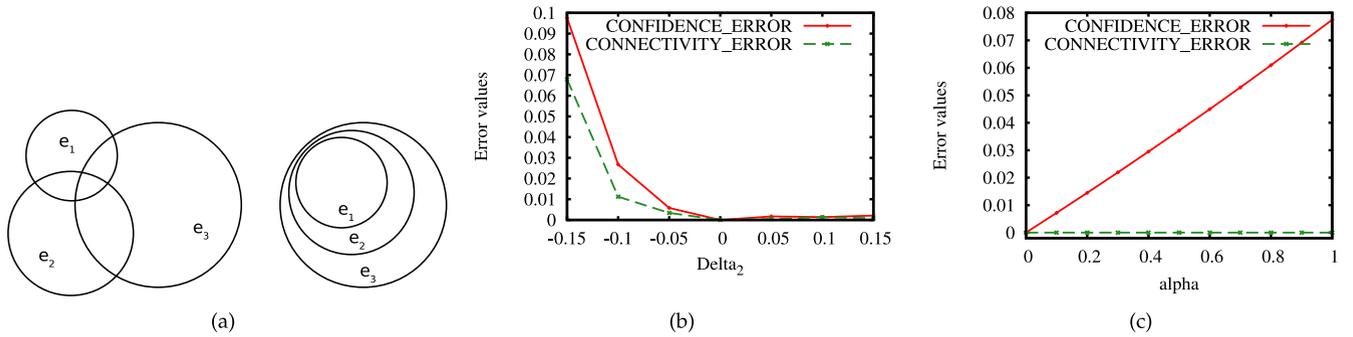


Fig. 7. (a) Depiction of the difference between independence on the left ($\alpha = 0$) and full dependence on the right ($\alpha = 1$), (b) error values with a discrepancy between real and available information in Δ_2 , and (c) error values with a discrepancy between real and available information in α .

information. This is the connectivity that could have been achieved if the network operator had perfect information.

We consider two types of discrepancies between available and real information. First, some compromise events have inaccurate values, that is, $P'(e_i) = P(e_i) + \Delta_i$ where P' is the available probability distribution, P is the actual probability distribution, and Δ_i is the error for a particular compromise event. Second, the compromise events are not fully independent, that is, $P'(E) = (1 - \alpha) * P(E) + \alpha * D(E)$ where E is a set of compromise events, $D(\cdot)$ is the probability distribution if there is complete dependence among the events, and α is a parameter determining how correlated the variants actually are ($\alpha = 0$ is complete independence while $\alpha = 1$ is the most extreme dependence). We illustrate the independent and full dependence scenarios with Venn diagrams in Fig. 7a.

With a discrepancy between the available and real information and letting $x = A_ECC_A_INFO$, $y = R_ECC_A_INFO$, and $z = R_ECC_R_INFO$ we observe the following two types of errors.

- $CONFIDENCE_ERROR = \frac{|x-y|}{y}$ is the error in how confident a network operator is with the created assignment.
- $CONNECTIVITY_ERROR = \frac{|y-z|}{z}$ is the error in how much worse an assignment based on available information is versus an assignment based on the real information.

6.2 Error Analysis on Random Topologies

We show the effect of a discrepancy in the compromise probability of a single variant. Then, we show the effect of discrepancy in the assumption of complete independence among variants. We use random topologies with similar settings to the random topologies in Section 4.4. Each topology had five clients and three variants with compromise probabilities $P(e_1) = 0.1, P(e_2) = 0.15, P(e_3) = 0.2$. In that section we showed results when varying density and number of nodes. For varying density we fixed the number of nodes at 25, and for varying the number of nodes we fixed the density at 6. These values were chosen as these parameters produced interesting topologies, that is, the topologies were connected but not too connected that assignment was trivial. Thus, in this section we fix the number of nodes to 25 and density to 6 for interesting topologies to investigate the effects on assignment when there are discrepancies between available and real information.

Fig. 7b shows the $CONFIDENCE_ERROR$ and $CONNECTIVITY_ERROR$ when the $P(e_2)$ used for assignment is different from the real $P(e_2)$. We show the errors when the available information has a compromise probability greater than the actual compromise probability ($\Delta_2 < 0$) and less than the actual compromise probability ($\Delta_2 > 0$). We see the greatest errors (for both types) when $P(e_2)$ is believed to be a weaker variant than it truly is, that is, $\Delta_2 < 0$ this is due to the assignment algorithm preferring to select v_3 over v_2 when forced to make a choice between these two. We observe little errors when $\Delta_2 > 0$ which is the case that the available information indicates v_2 is a stronger variant than it actually is. This is due to the random topologies having many client pairs that can be connected by two paths, so it is not so detrimental for the assignment to prefer v_1 over v_2 . There is some error which indicates that the preference of v_2 over v_1 is slightly detrimental.

Fig. 7c shows the errors when the assignment selected is based on the assumption of complete independence. We note in this case that $R_ECC_A_INFO = R_ECC_R_INFO$, since the assignments are actually the same despite the change in independence information, and thus $CONNECTIVITY_ERROR$ is always equal to zero in this case. However, $CONFIDENCE_ERROR$ is nonzero since any connectivity believed to be found by a network operator is less than the realistic connectivity since dependence among variants is detrimental to diversity. We see that this error increases linearly with α , the parameter controlling dependence.

7 RELATED WORK

Diversity assignment. The work most similar to ours considers diversity assignment over nodes of a distributed system [22], but the goal of that work is to prevent the spread of malware. In contrast, we assume that if a node of some variant is compromised, then all nodes of that variant are also compromised, as the attacker is not restricted to only using links within the network. When using diversity to prevent the spread of malware, the computation problem in [22] is different from ours as they intend to minimize the number of links which contain two nodes of the same variant. Thus, their underlying optimization problem for variant assignment is a version of the classic graph coloring algorithm. This problem is NP-hard, so their work also explores a heuristic solution which can scale to large networks.

Fault-tolerant topology construction. Existing work has introduced the concept of the *fault-diameter* of a graph,

which is a metric that bounds the diameter of a graph given that a bounded number of nodes may fail [23], [24], [25], [26]. For a network topology, this means that if the number of failures is bounded, then the maximum number of hops between any two correct nodes will not exceed the fault diameter. This translates to acceptable latency and overhead even in the worst case. Work in this area has considered various ways to create graphs with good fault-diameters, but these methods only consider unweighted graphs where edges are possible between any pair of nodes. In our work, we assume the topology is chosen ahead of time and fixed to ensure good link quality, and we do not need to add edges for our technique.

In wireless contexts, work has studied the allocation of energy among nodes in a wireless adhoc network to ensure high connectivity even when some bounded number of nodes fail [27], [28], [29]. The work assumes that node positions are fixed and an amount of energy can be assigned to each node. Higher energy at a node implies a larger transmission range and more possible connections for that node. The optimization problem is to find a power assignment to nodes which minimizes the global power consumption while ensuring connectivity among correct nodes given a bounded number of nodes can fail. This optimization problem is studied in detail, providing a MIP and exploring various approximation techniques.

Wireless Sensor Network (WSN) key distribution. Wireless sensor networks consist of resource constrained devices which sense physical phenomena and deliver this information over a wireless network to a base station. In this context, PKI and full pair-wise key initialization are prohibitive due to the limitations of sensors. Thus, various work proposes special key distributions, where secret information is shared among more than a single pair of nodes [30], [31], [32], [33], [34]. This has similarities to diversity assignment as the physical capture of a single node allows an attacker to utilize the secret information on that node to attack links of other nodes which share similar secret information. Our work does fundamentally differ as we perform diversity assignment with the complete topology information to maximize a resiliency metric while WSN key distribution work focuses on assigning initial secret information to nodes to maximize the potential of many links are secure. With the potential for many secure links, a random wireless topology can be created and have certain resiliency properties.

Path diversity. Other work has studied the possible geographically diverse paths of real-world topologies [35]. The assumptions of this work are that problems on today's Internet are correlated geographically, so having multiple paths which contain nodes that are geographically diverse will result in higher reliability, i.e., reduced probability of lost packets. The main contributions of this work are defining the metric of geographic diversity for a graph and analyzing this value for realistic graphs. No assignment problem exists in this context as diversity is fixed by geographic location.

8 CONCLUSION

This work illustrates the resiliency benefits gained when shifting from homogeneous networks with potential vulnerabilities shared across all routing nodes to networks that

leverage optimally-assigned diversity. We summarize our key findings. First, randomly assigning diversity to a realistic network has surprisingly poor results, which motivated the need to formulate and solve the Diversity Assignment Problem (DAP). Second, we propose an algorithm that solves DAP optimally, and show the results on medium-sized random networks as well as a realistic network. Third, we propose an algorithm that approximates the optimal solution, scaling well to large networks, and show that on random networks, the resulting resiliency is close to that of the optimal solution. Fourth, we show how to optimize for the specific resiliency needs of an application running on the network. We applied this to Paxos and BFT, finding that the probability of making progress can be significantly increased. Lastly, as it is difficult to exactly estimate compromise probabilities we showed how discrepancies between compromise probabilities used for assignment and the real compromise probabilities affect assignment and resilience.

ACKNOWLEDGMENTS

This work was supported in part by US Defense Advanced Research Projects Agency (DARPA) grant N660001-1-2-4014. Its contents are solely the responsibility of the authors and do not represent the official view of DARPA or the Department of Defense. Andrew Newell is the corresponding author of this paper.

REFERENCES

- [1] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Obelheiro, "OS diversity for intrusion tolerance: Myth or reality?" in *Proc. IEEE/IFIP 41st Int. Conf. Dependable Syst. Netw.*, 2011, pp. 383–394.
- [2] B. Cox, D. Evans, A. Filipi, J. Rowanhill, W. Hu, J. Davidson, J. Knight, A. Nguyen-Tuong, and J. Hiser, *N-Variant systems: A Secretless Framework for Security Through Diversity*. Defense Technical Information Center, Ft. Belvoir, VA, USA, 2006.
- [3] I. Gashi, P. Popov, and L. Strigini, "Fault tolerance via diversity for off-the-shelf products: A study with sql database servers," *IEEE Trans. Dependable Secure Comput.*, vol. 4, no. 4, pp. 280–294, Sep./Oct. 2007.
- [4] Y. Deswarte, K. Kanoun, and J. Laprie, "Diversity against accidental and deliberate faults," in *Proc. Comput. Secur., Dependability Assur.: From Needs to Solut.*, 1998, pp. 171–181.
- [5] (2012). LTN global communications [Online]. Available: <http://www.ltnglobal.com/>
- [6] L. Lamport, "The part-time parliament," *ACM Trans. Comput. Syst.*, vol. 16, no. 2, pp. 133–169, 1998.
- [7] M. Castro and B. Liskov, "Practical byzantine fault tolerance," presented at the 3rd Symp. Operating Syst. Des. Implementation, New Orleans, LA, USA, 1999.
- [8] A. Schrijver, *Theory of Linear and Integer Programming*. New York, NY, USA: Wiley, 1998.
- [9] A. Newell, D. Obenshain, T. Tantillo, C. Nita-Rotaru, and Y. Amir, "Increasing network resiliency by optimally assigning diverse variants to routing nodes," in *Proc. 43rd Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2013, pp. 1–12.
- [10] A. Newell, D. Obenshain, T. Tantillo, C. Nita-Rotaru, and Y. Amir. (2013). Increasing network resiliency by optimally assigning diverse variants to routing nodes, Johns Hopkins Univ., Tech. Rep. CNDS-2013-1 [Online]. Available: <http://dsn.jhu.edu/pub/papers/cnds-2013-1.pdf>
- [11] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Médard, "Resilient network coding in the presence of byzantine adversaries," in *Proc. IEEE 26th Int. Conf. Comput. Commun.*, 2007, pp. 616–624.
- [12] B. Awerbuch, R. Curtmola, D. Holmer, C. Nita-Rotaru, and H. Rubens, "ODSBR: An on-demand secure byzantine resilient routing protocol for wireless ad hoc networks," *ACM Trans. Inform. Syst. Secur.*, vol. 10, no. 4, p. 6, 2008.

- [13] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms, Third Edition*. Cambridge, MA, USA: MIT Press, 2009, pp. 1161–1161.
- [14] (2012). High-performance software for mathematical programming and optimization [Online]. Available: <http://www-01.ibm.com/software/integration/optimization/cplex-optimizat ion-studio/>
- [15] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Proc. Eur. Control Conf.*, 2001, pp. 2603–2608.
- [16] L. Pallottino, E. Feron, and A. Bicchi, "Conflict resolution problems for air traffic management systems solved with mixed integer programming," *Trans. Intell. Transportation Syst.*, vol. 3, no. 1, pp. 3–11, 2002.
- [17] G. Huang, B. Baetz, and G. Patry, "Grey integer programming: an application to waste management planning under uncertainty," *Eur. J. Oper. Res.*, vol. 83, no. 3, pp. 594–620, 1995.
- [18] Coin-or [Online]. Available: <http://www.coin-or.org/>, 2015.
- [19] Scip [Online]. Available: <http://scip.zib.de/>, 2015.
- [20] Y. Amir, B. Coan, J. Kirsch, and J. Lane, "Prime: Byzantine replication under attack," *Dependable Secure Comput.*, vol. 8, no. 4, pp. 564–577, 2011.
- [21] A. Clement, E. Wong, L. Alvisi, M. Dahlin, and M. Marchetti, "Making byzantine fault tolerant systems tolerate byzantine faults," in *Proc. 6th USENIX Symp. Netw. Syst. Des. Implementation*, 2009, pp. 153–168.
- [22] A. O'Donnell and H. Sethu, "On achieving software diversity for improved network security using distributed coloring algorithms," in *Proc. 11th ACM Conf. Comput. Commun. Secur.*, 2004, pp. 121–131.
- [23] M. Krishnamoorthy and B. Krishnamurthy, "Fault diameter of interconnection networks," *Comput. Math. Appl.*, vol. 13, no. 5, pp. 577–582, 1987.
- [24] S. Latifi, "On the fault-diameter of the star graph," *Inform. Proc. Lett.*, vol. 46, no. 3, pp. 143–150, 1993.
- [25] S. Latifi, "Combinatorial analysis of the fault-diameter of the n-cube," *Trans. Comput.*, vol. 42, no. 1, pp. 27–33, 1993.
- [26] K. Day and A. Al-Ayyoub, "Fault diameter of k-ary n-cube networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 8, no. 9, pp. 903–907, Sep. 1997.
- [27] M. Hajjaghay, N. Immorlica, and V. Mirrokni, "Power optimization in fault-tolerant topology control algorithms for wireless multi-hop networks," in *Proc. MobiCom*, 2003, pp. 300–312.
- [28] X. Jia, D. Kim, S. Makki, P. Wan, and C. Yi, "Power assignment for k-connectivity in wireless ad hoc networks," *J. Combinatorial Optim.*, vol. 9, no. 2, pp. 213–222, 2005.
- [29] D. Panigrahi, P. Duttat, S. Jaiswal, K. Naidu, and R. Rastogi, "Minimum cost topology construction for rural wireless mesh networks," in *Proc. IEEE INFOCOM*, 2008, pp. 771–779.
- [30] S. Çamtepe and B. Yener, "Combinatorial design of key distribution mechanisms for wireless sensor networks," *IEEE Trans. Netw.*, vol. 15, no. 2, pp. 346–358, 2007.
- [31] L. Oliveira, H. Wong, M. Bern, R. Dahab, and A. Loureiro, "SecLEACH—A random key distribution solution for securing clustered sensor networks," in *Proc. 5th IEEE Netw. Comput. Appl.*, 2006, pp. 145–154.
- [32] H. Chan, A. Perrig, and D. Song, "Random key predistribution schemes for sensor networks," in *Proc. IEEE Symp. Secur. Privacy*, 2003, pp. 197–213.
- [33] H. Chan and A. Perrig, "Pike: Peer intermediaries for key establishment in sensor networks," in *Proc. IEEE INFOCOM*, 2005, pp. 524–535.
- [34] W. Du, J. Deng, Y. Han, P. Varshney, J. Katz, and A. Khalili, "A pairwise key predistribution scheme for wireless sensor networks," *ACM Trans. Inform. Syst. Secur.*, vol. 8, no. 2, pp. 228–258, 2005.
- [35] J. Rohrer, A. Jabbar, and J. Sterbenz, "Path diversification for future internet end-to-end resilience and survivability," *Springer Telecommun. Syst.*, vol. 56, pp. 49–67, May 2014.



Andrew Newell received the BS degree in computer science and mathematics from Southern Illinois University at Carbondale in 2008. He is currently working toward the PhD degree in computer science at Purdue University. He is a member of the Dependable and Secure Distributed Systems Laboratory. His research interests include resilient network design, wireless networks, network coding, and machine learning.



Daniel Obenshain received the BS degree from the California Institute of Technology in 2011 and the MSE degree from the Johns Hopkins University in 2013. He is currently working toward the PhD degree at the Johns Hopkins University, where he is a Beauchamp fellow. His research interests include distributed systems and intrusion tolerant systems. He is a member of the ACM and the IEEE.



Thomas Tantillo received the BS degree in computer engineering in 2010 and the MSE degree in computer science in 2013 from the Johns Hopkins University. He is currently working toward the PhD degree in computer science from the Johns Hopkins University. He is a member of the Distributed Systems and Networks laboratory and his research interests include security and intrusion tolerance for networks and distributed systems. He is a member of the IEEE.



Cristina Nita-Rotaru received the BS and MS degrees from Politehnica University of Bucharest, Romania, in 1995 and 1996, and the PhD degree in computer science from Johns Hopkins University in 2003. She is an associate professor in the Department of Computer Science at Purdue University. She leads the Dependable and Secure Distributed Systems Laboratory. She served on the technical program committee of more than 40 conference in networking, distributed systems, and security. She received the National science Foundation (NSF) CAREER award. She served as an associate editor for the *ACM Transactions on Information Security* and she is currently an associate editor for the *IEEE Transactions on Dependable and Secure Computing*, and the *IEEE Transactions on Mobile Computing*. Her research interests include security and fault-tolerance for distributed systems and networks. She is a member of the ACM and IEEE Computer Society.



Yair Amir received the BS and MS degrees in 1985 and 1990, respectively, from the Technion, Israel Institute of Technology, and the PhD degree in 1995 from the Hebrew University of Jerusalem, Israel. He serves as a professor of computer science, The Johns Hopkins University since 1995. Prior to the PhD degree, he gained extensive experience building C3I systems. He is a creator of the Spread and Secure Spread group communication toolkits, the Backhand and Wackamole clustering projects, the Spines overlay network messaging system, and the SMesh wireless mesh network. He has been a member of various program committees including the IEEE International Conference on Distributed Computing Systems, the ACM Conference on Principles of Distributed Computing, and the IEEE/IFIP International Conference on Dependable Systems and Networks. He currently serves as an Associate Editor for the *IEEE Transactions on Dependable and Secure Computing*. He co-founded Spread Concepts LLC (2000) and LTN Global Communications Inc (2008), and is a member of the ACM and the IEEE Computer Society.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.