

A study of the impact of DNS resolvers on CDN performance using a causal approach

Original

A study of the impact of DNS resolvers on CDN performance using a causal approach / Hours, Hadrien; Biersack, Ernst; Loiseau, Patrick; Finamore, Alessandro; Mellia, Marco. - In: COMPUTER NETWORKS. - ISSN 1389-1286. - STAMPA. - 109:(2016), pp. 200-210. [10.1016/j.comnet.2016.06.023]

Availability:

This version is available at: 11583/2655355 since: 2016-11-09T09:04:51Z

Publisher:

Elsevier

Published

DOI:10.1016/j.comnet.2016.06.023

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

A Study of the Impact of DNS Resolvers on CDN Performance Using a Causal Approach

Hadrien Hours^{a,b}, Ernst Biersack^c, Patrick Loiseau^a, Alessandro Finamore^{d,e},
Marco Mellia^d

^a*EURECOM, email: firstname.lastname@eurecom.fr*

^b*ENS Lyon, email: firstname.lastname@ens-lyon.fr*

^c*Caipy, email: erbi@e-biersack.eu*

^d*Telefonica, email: firstname.lastname@telefonica.com*

^e*Politecnico di Torino, email: firstname.lastname@polito.it*

Abstract

Resources such as Web pages or videos that are published in the Internet are referred to by their Uniform Resource Locator (URL). If a user accesses a resource via its URL, the host name part of the URL needs to be translated into a routable IP address. This translation is performed by the Domain Name System service (DNS). DNS also plays an important role when Content Distribution Networks (CDNs) are used to host replicas of popular objects on multiple servers that are located in geographically different areas. A CDN makes use of the DNS service to infer client location and direct the client request to the optimal server. While most Internet Service Providers (ISPs) offer a DNS service to their customers, clients may instead use a public DNS service. The choice of the DNS service can impact the performance of clients when retrieving a resource from a given CDN. In this paper we study the impact on download performance for clients using either the DNS service of their ISP or the public DNS service provided by Google DNS. We adopt a causal approach that exposes the structural dependencies of the different parameters impacted by the DNS service used and we show how to model these dependencies with a Bayesian network. The Bayesian network allows us to explain and quantify the performance benefits seen by clients when using the DNS service of their ISP. We also discuss how to further improve client performance.

Keywords: Causality, CDN, DNS, Traffic performance, Bayesian networks, Optimization, Monitoring, Knowledge inference, Computer networks.

1. Introduction

Each time an Internet user wants to access a resource, he uses a human readable name called Uniform Resource Locator (URL), containing the domain name of the administrative entity hosting this resource. However, a domain name is not routable and needs to be translated into the IP address of a server hosting the resource the client wants to access. This is taken care of by the DNS service. At the same time, many popular services such as YouTube, iTunes, Facebook or Twitter, rely on CDNs, where objects are replicated on different servers, and in different geographical locations to optimize the performance experienced by their users. When a client accesses an object hosted by a CDN, its default DNS server contacts the DNS server of the CDN that hosts the resource the client is requesting. Based on the origin of the request, the authoritative CDN DNS redirects the client to the optimal server. Most of the ISPs provide a DNS service, but it is now common to see customers using a public DNS service instead [Otto et al., 2012]. Clients using the DNS service of their ISP are served by a local DNS server that often provides a more accurate location information to the CDN compared to the information communicated by a public DNS service such as the Google DNS service. Indeed, public DNS servers are usually further away from the clients of a given ISP than the default ISP DNS server. There have been several studies suggesting that public DNS services do not perform as well as local DNS services provided by ISPs, mainly because of the impossibility of public DNS to correctly communicate the location of the clients originating the request [Huang et al., 2011; Ager et al., 2010]. This problem is addressed with ECS (edns-client-subnet) [Streibelt et al., 2013] but Akamai does not support it currently.

Studying the performance of the users accessing resources in the Internet is a complex task. Many parameters influence the end user experience and the relationships between these parameters is not always observable or intuitive. It is therefore necessary to use a simple, yet formal model that allows us to understand the role of a given parameter and its dependencies with other parameters. Bayesian networks offer a simple and concise way to represent complex systems [Darwiche, 2009]. In this paper, we use a Bayesian network to represent the causal model that captures the impact of the DNS service on the throughput performance experienced by clients accessing resources hosted by the Akamai CDN. Bayesian networks capture the dependencies between the different parameters impacting the throughput of the clients. One very interesting property of causal models is their *stability under intervention*. Causal models can be used to predict how the throughput of CDN users would evolve if we would *intervene*

on the different parameters influencing the performance of CDN users. Here an intervention consists in isolating a given parameter of the system being studied, removing all its direct and remote causes and fixing its variations to a pre defined value or distribution. Being able to predict the effect of interventions, we can use causal models to understand the observed performance of a given system and to design strategies to improve its performance. In this work, we infer and use the causal model of CDN performance to understand the impact of choosing one DNS service instead of another. From such a model we are able to explain why clients using the DNS service of their ISP experience better download performance than clients using the Google DNS service. We are also able to indicate how to further improve the performance of the clients using the DNS service of their ISP.

Our work differs from previous studies of DNS services in several important points:

- We use a causal approach that formally models the structural dependencies of the different parameters influencing the throughput obtained.
- Observing that the clients using the DNS service of their ISP (referred to as local DNS) experience higher throughput than the clients using the public DNS service (referred to as Google DNS), we can show that this performance difference is due to the fact that clients using the DNS service of their ISP are redirected to closer servers. We are also able to precisely quantify this performance improvement.
- The causal model of our system also reveals that the parameterization of TCP (initial congestion window) of the servers accessed by the users of the Google DNS plays a key role in their throughput performance. Besides fully explaining the observed performance, this result also indicates how to further improve the performance of the clients using the local DNS.

Overall, the main contribution of our work resides in the methodology adopted and in its use of counterfactuals to understand the causal dependencies of a complex system.

In Section 2, we introduce causal models and their use to predict interventions, summarizing some of the main concepts from [Pearl, 2009; Spirtes et al., 2001]. We then present, in Section 3, the environment of our study and the description of the parameters constituting our system. Section 4 presents our study of the DNS impact on the throughput. In particular we present the causal model of our system where we can observe the impact of the choice of the DNS service on the

throughput. Our approach also allows us to predict the improvement that could be achieved by modifying the parameterization of the servers accessed by the users of the local DNS service. Section 5 compares our approach to the related work and Section 6 summarizes our work and proposes directions to further extend our work.

Several methods mentioned in this paper were designed and validated with parallel studies that are presented in an Appendix. The Appendix is available with the online version of this paper.¹ We give references to these studies in the paper.

2. Causal model: Definitions and usage

To model a complex system such as a communication network and to organize the knowledge obtained from its passive observation is very challenging. Existing work typically looks for the presence of correlation between different events observed simultaneously (see [Mayer-Schönberger, 2013] and references therein). However, *correlation is not causation and the detection of correlation between two parameters A and B does not inform us on how they are related*. A can impact B , or the other way around, or an unobserved parameter can impact both A and B simultaneously. The difference between correlation and causation plays an important role if we want to find out how to improve our system by partly modifying its behavior. A causal approach uncovers the structural dependencies between the parameters of the system under study. The ability to predict the effects of a manipulation of the parameters of a system is a major strength of causal models as they are *stable under intervention*. Stability under intervention means that a causal model, inferred from the observations of a system in a given situation, is still valid if we manually change the system mechanisms, redefining the systems laws. The manual modification of the system parameters is called an *intervention*. Interventions consist in modifying the behavior of a component of the system, removing the influence of its direct and remote causes and manually setting its variations. The inference of a causal model and of a causal effect [Pearl, 2009; Spirtes et al., 2001] is made using passive observations only. The causal theory allows us to predict the behavior of the various parameters of the inferred model after an intervention *without the need of additional observations*.

In this section we present the PC algorithm [Spirtes and Glymour, 1991] that is used to infer the causal model of our system. We also describe the different properties of a causal model as described in [Pearl, 2009; Spirtes et al., 2001].

¹<http://linkincreation.com>

2.1. Causal model: Inference

For our work, we use the PC algorithm [Spirtes and Glymour, 1991] to build the Bayesian graph representing the causal model of our system. This algorithm takes as input the observations of the different parameters that characterize our system and infers the corresponding causal model. In our representation of a causal model as a Bayesian network, each node represents one parameter of our system and the presence of an edge from a node X to a node Y ($X \rightarrow Y$) represents the existence of a causal dependence of parameter Y on parameter X .

The PC algorithm starts with a fully connected and unoriented graph, called *skeleton*, where each parameter is represented by a node and connected to every other parameter. The PC algorithm then trims the skeleton by checking for independences between adjacent nodes:

- First, the unconditional independences ($X \perp\!\!\!\perp Y$) are tested for all pairs of parameters and the edges between two nodes whose corresponding parameters are found to be independent are removed.
- For the parameters whose nodes are still adjacent, the PC algorithm then checks if there exists a conditioning set of size one that makes two adjacent nodes independent. If this is the case, it removes the edge connecting the corresponding two nodes, otherwise the edge is kept.
- The previous step is repeated, increasing the conditioning set size by one at each step, until the size of the conditioning set reaches the maximum degree of the current skeleton (the maximum number of adjacent nodes for any node in the current graph), which means that no more independences can be found.

The final step of the PC algorithm consists in orienting the edges. First, the PC algorithm orients all the V-structures, i.e. subgraphs $X - Z - Y$ where X and Y are not adjacent, and then orients as many edges as possible without creating new colliders or cycles [Pearl, 2009]. A node Z is a *collider* if it is part of an oriented subgraph $X \rightarrow Z \leftarrow Y$ where X and Y are not adjacent. An illustration of the different steps of the PC algorithm is presented in the Appendix A.1.

2.2. Causal model: Properties and theorems

In this section we assume that we have the causal model of our system that is represented by a Bayesian network. We focus on two parameters X and Y , where Y represents the performance of our system and we are interested in the global effect

on Y when intervening on X , including the effects mediated by external parameters also impacted by this intervention. We call this causal effect the *total causal effect*. Details of the implementation of the methods presented in this section can be found in the Appendix B.

2.2.1. Atomic interventions

We denote by $do(X = x)$ (or $do(x)$) the intervention that consists in intervening on the parameter X by fixing its value to be x . An intervention that simply assigns to X a fixed value is called an *atomic intervention*. The difficulty of predicting the effect of an intervention comes from the possible presence of spurious associations between the intervention variable and the response variable. A *spurious association* between X and Y is an association between X and Y due to external parameters ($\notin \{X, Y\}$). To obtain an unbiased estimation of the effect of an intervention, we need to remove the effect of spurious associations. As an intervention is equivalent to isolating a given parameter from its direct and remote causes and to assigning it a fixed value, we need to remove the effects of direct and remote causes in our estimations. Such estimation is complex if one needs to consider all the possible inter dependencies between the different parameters influencing the performance of the system being studied. However, the use of a *graphical causal model*, where the different dependencies are present, makes it easy to estimate the outcome of interventions. Different criteria (c.f. [Pearl, 2009]) can be used to identify the minimum set of parameters that block the effects of direct and remote causes when estimating the effect of a given intervention.

If G denotes the Bayesian graph that represents the causal relationships between the parameters of our system, we use $G_{\overline{X}}$ to denote the sub-graph of G where all the edges entering X are removed and $G_{\underline{X}}$ the sub-graph of G where all the edges exiting X are removed. We can use the rules of *do-calculus* [Pearl, 2009] to estimate the distributions of the parameters of our system after an intervention based on their distributions prior to this intervention. Note that these rules do not make any assumption regarding the distributions or functional dependencies of the parameters.

We briefly recall the *Rules of do calculus* that will be used in Section 4.2 to predict the interventions we are interested in this work. Let P denote a (possibly multivariate) probability distribution specified by the probability mass function or probability density function, depending on the nature of the parameters.

Theorem 1 (3.4.1 from [Pearl, 2009]). (*Rules of do calculus*) *Let G be the directed acyclic graph associated with a causal model [...] and let $P(\cdot)$ stand for the*

probability distribution induced by that model. For any disjoint subsets of variables X , Y and Z we have the following rules.

Rule 1(Insertion/deletion of observation):

$$P(y|do(x), z, w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z | X, W)_{G_{\bar{X}}} \quad (1)$$

Rule 2(Action/observation exchange):

$$P(y|do(x), do(z), w) = P(y|do(x), z, w) \text{ if } (Y \perp\!\!\!\perp Z | X, W)_{G_{\bar{XZ}}} \quad (2)$$

Rule 3(Insertion/deletion of intervention):

$$P(y|do(x), do(z), w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z | X, W)_{G_{\overline{XZ(W)}}}, \quad (3)$$

where $Z(W)$ is the set of Z -nodes that are not ancestor of any W -nodes in $G_{\bar{X}}$.

2.2.2. Enforcing intervention with a given probability

To study the impact of the DNS service on the performance seen by the clients (c.f. Section 4) we must estimate the effect of interventions on the parameters influenced by the DNS service and on the parameters influencing the throughput. To do so, we cannot use atomic interventions since we intervene on a given parameter by *changing its distribution*. If we want to predict how an intervention on X affects Y , where the intervention on X is enforced with the conditional probability distribution $f^*(X|Z)$, we obtain [Pearl, 2009, Section 4.2]:

$$f(y)_{|f^*(x|z)} = \int_{D_X} \int_{D_Z} f_{Y|do(X), Z}(y, x, z) f^*(x|z) f(z) dx dz. \quad (4)$$

2.3. D-separation

The d-separation criterion is a graphical criterion to decide, by looking at the graph, if two parameters, represented by their nodes, are independent. D-separation associates the notion of connectedness with dependence. If there exists a directed path between two nodes, the nodes are said to be d-connected and their corresponding parameters are dependent. On the other hand, if we condition on one of the nodes in the path from X to Y , then this node is said to block the path and X and Y are conditionally independent relative to this path. For X and Y to be independent, one must block *all* the paths d-connecting X and Y . When studying d-separation, an important notion is the one of collider. The presence of a collider on a *undirected* path blocks this path. While conditioning on a collider unblocks the path which can be explained by the fact that two independent causes become dependent if one conditions on their common consequence.

2.4. Density estimation

The theory of causality Pearl [2009] makes no assumption on the distribution of the parameters. We estimate the multidimensional probability density functions via Copulae [Jaworski et al., 2010], using the Sklar theorem.

The Sklar theorem stipulates that, if F is a multivariate cumulative distribution function with marginals $(F_1, \dots, F_i, \dots, F_n)$, there exists a copula C such that

$$F(x_1, \dots, x_i, \dots, x_n) = C(F_1(x_1), \dots, F_i(x_i), \dots, F_n(x_n)). \quad (5)$$

There are different types of copulae, in our work we focus on T-copulae [Demarta and McNeil, 2005] and G-copulae [Pitt et al., 2006]. T-copulae present the advantage that, by tuning the different parameters of the T-copula, one can better capture the tail dependencies between the different components of the multivariate distribution that is modeled. This is highly useful in our case where the performance (e.g. the throughput of a Web user) can be strongly affected by changes to the characteristics of the network such as packet loss or delay. Unfortunately, T-copulae are complex to parameterize, which implies that more data is needed to fit such model to our problem. In this paper, we are interested in counterfactuals such as “*How would the system behave under the condition C1 if one of its parameter was to behave as it has done under the condition C2, knowing that C1 and C2 are exclusive ?*”. Counterfactuals correspond to the predictions of complex interventions, each of which requires conditioning on several variables in order to block the different spurious associations.

We decided to use Gaussian copulae [Pitt et al., 2006], which are known to be less sensitive if the amount of data available is limited (see Appendix C).

In the bivariate case, the Gaussian copula is defined as:

$$C_\rho(u, v) = \Phi_\rho(\Phi^{-1}(u), \Phi^{-1}(v)), \quad (6)$$

where ρ represents the correlation matrix and Φ the CDF of the standard normal distribution. The marginals, $F_i(x_i)$, are estimated using normal kernels.

The choice of Gaussian copulae as well as the methods and their implementation to compute the conditional PDFs have been designed and validated based on studies made on artificial datasets that are presented in the Appendix C.

3. Experimental set up

In this section, we present how we do the data collection and how we extract the parameters of interest. We define our system as the set of parameters (see Section 2.1) and observe these parameters in different situations to capture their dependencies and infer the corresponding graphical causal model.

3.1. Experiment design

We collect IP packet traces at a Point of Presence (PoP) of a large European ISP and extract all the traffic directed to or coming from the Akamai CDN. To model the impact of the choice of DNS service on the client throughput, we make three choices: i) We only focus on the traffic carried by the TCP protocol. ii) To eliminate the impact of TCP slowstart, we only consider large TCP connections that carry at least 2MBytes of data. iii) As more than 90% of the observed connections use either Google DNS (*GDNS*) or the DNS of the local ISP (*LDNS*) we consider only these two DNS services.

The probe capturing the traffic is placed between the client and the server. We call internal network, denoted as *isp network*, the part of the network between the client and the probe. We call external network the part of the network between the probe and the server, assimilated to the Internet network and denoted as *inet network*. The traffic was captured on two different days, a Thursday and a Sunday, from 5.30 pm to 9.30 pm.

3.2. Parameters of our model

We use the Tstat software [Finamore et al., 2011] to extract from the packet traces relevant information on a per connection basis. We have about 7000 connections. We use domain knowledge to select a subset of the information obtained from Tstat that represents the parameters known to impact the throughput. In addition to the information obtained from Tstat we collect for each connection additional information such as the DNS service used, the number of hops between the client and the server and the server address.² The packet traces used for this study are confidential and cannot be shared publicly.

3.3. Summary of our data

Each connection is described by 19 parameters. In Table 1, we present the average (μ), minimum (*min*), maximum (*max*), standard deviation (σ) and coefficient of variation ($CoV = \frac{\sigma}{\mu}$) of each of the 19 parameters.

Since we are interested in comparing the performance of LDNS users and GDNS users, Table 2 presents the statistics for the connections where the LDNS is being used and for the connections where the GDNS is being used.

We use the following notations:

²Since the addresses were anonymized we represented the server address by the Autonomous System (AS) number of the AS the server is located in.

Table 1: Summary of the different parameters

Parameter	μ	<i>min</i>	<i>max</i>	σ	<i>CoV</i>
dstip	N.A.	1300	34000	N.A.	N.A.
dns	N.A.	1	3	N.A.	N.A.
dow	N.A.	4	7	N.A.	N.A.
tod (s)	7100	52000	78000	4400	0.1
isprttavg (ms)	76	0	19000	460	6.1
isprttstd (ms)	100	0	37000	960	9.2
ispnbhops	1.8	1	3	0.51	0.3
inetrtravg(ms)	26	0.48	660	27	1.0
inetrtrstd (ms)	8.2	0	4700	61	7.5
inetrnbhops	9.4	2	21	2.8	0.3
rwin0	0.83	0	360	11	13
rwinmin (kB)	31.3	0.004	65	22.5	0.9
rwinmax (kB)	213	17.5	2625	150	0.7
cwinmax (kB)	150	7.3	1625	103	0.7
cwinmin (kB)	0.9	0.001	1.5	0.6	0.7
retrscore	0.005	0	0.19	0.009	1.9
rto (bool)	0.11	0	1	0.32	2.8
nbbytes (MB)	23.8	2.1	3875	138	5.7
tput (Mbps)	3.2	0.006	35	2.6	0.8

Table 2: Summary of the different metrics for the two DNS: Local DNS (LD) and Google DNS (GD) (*dow* and *tod* are similar and provide no insight, so they were removed)

Par	μ		<i>min</i>		<i>max</i>		σ		<i>CoV</i>	
	LD	GD	LD	GD	LD	GD	LD	GD	LD	GD
isprttavg (ms)	80	61	0	0	19000	15000	470	440	5.9	7.2
isprttstd (ms)	1100	76	0	0	32000	37000	920	1100	8.3	14.0
ispnbhops	1.8	1.9	1	1	3	3	0.53	0.4	0.3	0.2
inetrtravg (ms)	20	48	0.48	11	510	660	20	38	1.0	0.8
inetrtrstd (ms)	8.6	6.5	0	0	4700	1400	65	44	7.6	6.8
inetrnbhops	8.7	12	2	5	17	21	2.4	2.7	0.3	0.22
rwin0	0.97	0.29	0	0	330	360	12	9.2	12.0	32.0
rwinmin (kB)	35	12	0.004	0.03	65	65	28	14	0.8	1.1
rwinmax (kB)	213	213	18	18	2625	2000	150	138	0.3	0.7
cwinmax (kB)	163	118	7.3	7.8	1625	738	108	72	0.7	0.6
cwinmin (kB)	0.9	1.2	0.001	0.001	1.5	1.5	0.6	0.5	0.7	0.4
retrscore	0.005	0.004	0	0	0.19	0.06	0.01	0.01	1.9	1.8
rto (bool)	0.11	0.11	0	0	1	1	0.32	0.31	2.8	2.9
nbbytes (MB)	29	7	2.1	2.1	3875	1375	150	44	5.3	6.5
tput (Mbps)	3.2	3	0.006	0.007	35	29	2.7	2	0.9	0.7

- Parameters with the prefix *isp* represent the *isp network* statistics, while the ones with the prefix *inet* represent the *inet network* statistics.
- The suffix *avg* represents the average value of a given parameter for a single connection (for example the average Round Trip Time between the client and the probe is denoted *isprttavg*).
- The suffix *std* represents the standard deviation of a given parameter for a single connection (for example the standard deviation of the Round Trip Time between the probe and a server is denoted *inetrttstd*).
- The *rto* parameter is set to true if there was at least one packet retransmission due to a time out and to false otherwise
- The *retrscore* parameter represents the fraction of retransmitted packets for a single connection ($= \frac{\text{retransmissions}}{\text{total transmissions}}$).
- The parameters *rwin** and *cwin** represent receiver window and congestion window metrics respectively.
- The day of the week and time of the day are captured by the variables *dow* and *tod* respectively.

Destination IP (*dstip*), DNS (*dns*) and days (*dow*) are categorical data for which the average value, standard deviation or coefficient of variation do not exist.

Without discussing in detail the values of the different parameters in Table 1, we would like to draw the attention to the difference in the RTT values observed inside the ISP network and the RTT values observed in the Internet: the average RTT value *isprttavg* is almost three times as high as the average RTT value *inetrttavg*. The use of an ADSL on the access link and the large buffers used in ADSL networks not only increase the RTT but also result in high variations of the RTT values observed that correspond to a standard deviation of the *isprttstd* being more than ten times bigger than the *inetrttstd*.

4. Causal study of the impact of the DNS service used on throughput

4.1. Modeling causal relationships

We use the PC algorithm [Spirtes and Glymour, 1991] and the kernel based independence test from [Zhang et al., 2012] to obtain the Bayesian network showing the causal model of our system (c.f. Figure 1). We briefly discuss some of the most interesting dependencies exhibited by this model.

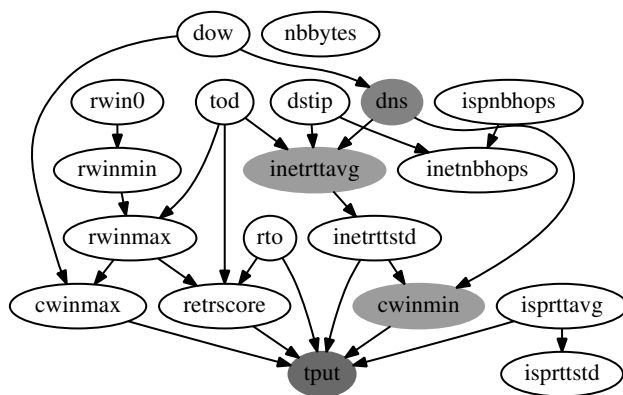


Figure 1: Bayesian network representing the causal model of Web performance using two different DNS: the public Google DNS and the DNS of the local ISP with the following parameters: Day of the Week (dow), Number of bytes exchanged during the connection ($Nbbytes$), first advertised receiver window ($rwin0$), minimum advertised receiver window ($rwinmin$), maximum advertised receiver window ($rwinmax$), minimum server congestion window ($cwinmin$), maximum server congestion window ($cwinmax$), time of the day (tod), retransmission score ($retrscore$), presence of time outs (rto), server IP address ($dstip$), number of hops between client and probe ($ispnbhops$), number of hops between probe and server ($inetnbhops$), average external delay ($inetrttav$), standard deviation external delay ($inetrttstd$), average internal delay ($isprttavg$) and standard deviation internal delay ($isprttstd$)

The day of the week (dow) and the time of the day (tod) are two nodes that have no parents, which is not surprising. The time of the day (tod) influences the RTT between the probe and the server ($inetrttav$), which captures the *peak hour effect* in the Internet.

In Table 1 we saw that that the variance of the internal RTT ($isprttstd$) was much higher than the one of the Inet RTT ($inetrttstd$). This may lead one to expect that $isprttstd$ has a stronger impact on the throughput than the $inetrttstd$. However, the causal model shows something different: we have a direct dependence between ($inetrttstd$) and the throughput ($tput$) but not between the standard deviation of the internal RTT ($isprttstd$) and the throughput. This example illustrates the ability of causal model to exhibit non intuitive dependencies.

We observe that the day of the week (dow) influences the DNS service used by the clients (dns). As our observations are made on two days (a Thursday and a Sunday), our conclusions are a bit limited. However, it appears that on Thursday 72% of the connections use the LDNS service against 28% using the GDNS service, while on Sunday 93% of the connections use the LDNS service against 7% using the GDNS service. It would be interesting to identify the clients using one

DNS service and compare their locations with the ones of the clients using the other DNS service to better understand this dependence. The day of the week may capture the difference in the Internet usage and the devices used at home and at work. However, for privacy reasons, the IP addresses of the clients are obfuscated, which prevents us from investigating this hypothesis.

One of the most interesting dependencies, which motivated this work, is the one between the DNS service (*dns*) and the external RTT (*inetrttavg*). Our data show that most of the time, clients using the DNS of their ISP are redirected to an Akamai server located in the same AS. On the other hand, the clients using the Google DNS service are often redirected to servers located outside the client AS and even, in some cases, to a server outside of Europe.

It has been previously shown [Huang et al., 2011] that clients using the local DNS service benefit from a redirection to servers closer than the ones of the clients using a public DNS service. Our data (see Table 2) corroborate this observation since the average external RTT for the LDNS service users is of 20 ms, while the users of the GDNS service experience an average external RTT of 48 ms.

We can also see that congestion window metrics (*cwinmin*, *cwinmax*) have a direct impact on the throughput (*tput*). Additionally, the minimum congestion window (*cwinmin*) has the DNS (*dns*) as direct parent. Its average value for clients using GDNS is 1.2kB against 0.9kB for users served by the LDNS, see Table 2.

A parameter present in a causal model represents also the mechanisms captured by such parameter. This is the case of the *cwinmin* that also captures the tuning of the TCP parameters at the server side. Clients using the LDNS often access their objects from servers that are located *inside* the ISP network. These servers could have a configuration different from the servers accessed by the users of the GDNS. This hypothesis could also explain the fact that both DNS services result in a similar throughput performance despite the difference in the RTTs observed. Other reasons could be the impact of losses on the congestion window or the load of the servers being accessed by the clients. To capture the server load, we estimate the server processing time defined from the time at which a server sends the acknowledgment of the client HTTP/GET message and the time at which it sends the first data packet. However, the server processing time shows an expected value of 43 ms for the LDNS users against 64 ms for the GDNS users. A higher processing time for the servers accessed by the GDNS users suggests that they are more loaded. On the other hand, the congestion window is impacted by the loss. However in our data set, very few losses actually happen and no dependence is found between the loss (*retrscore*) and the DNS service (*dns*).

It is to be expected that the internal RTT (*isprttavg*) is a parent of the through-

put. Also, the absence of a dependence between the time of the day (*tod*) and the internal RTT can be explained by the fact that all the observed users are using the same “internal” path (the path from the users to the probe).

We see that the maximum receiver window advertised by the client (*rwinmax*) has the time of the day as one of its parents (*tod*). This could be due to the TCP buffer auto tuning mechanism [Ford et al., 2002] that adjusts the receiver window according to the quantity and frequency of data received by the client, which is influenced by the time of the day.

There is no edge between the DNS (*dns*) and the destination IP address (*dstip*) and the object size (*nbbbytes*) is not connected at all. This may be explained by the fact that the number of users of the LDNS service (80%) is much higher than the number of users of the GDNS service (20%). The same percentages are observed for the number of servers accessed by the users of the LDNS service (80%) and the number of servers accessed by the users of the GDNS service. The difference between these percentages weakens the dependence between *dns* and *dstip*. A solution to detect weaker dependences is to increase the acceptance rate in the independence tests. However, increasing the acceptance rate implies a higher risk of failing to reject weak independences and should be used with caution. The independence of the object size from other parameters influencing the throughput is not necessarily surprising as we consider long connections.

The two loss parameters (*retrscore* and *rto*) and the two RTT parameters (*in-etrtsstd* and *isprttavg*) are four of the six direct parents of the throughput, which is in line with our domain knowledge of TCP. The additional parents are the congestion window parameters of the server (*cwinmin* and *cwinmax*).

The fact that none of the receiver window metrics (*rwin**) is a direct parent of the throughput (*tput*) is not surprising. By comparing the throughput of a given connection with the minimum and the maximum quantity of information that the client can handle (see Figure 2), it appears that the receiver window advertised by the client is never limiting the throughput.

4.2. Asking what-if questions

We have seen that the Bayesian network reveals a rich set of causal relationships that indicate how the different parameters impact the throughput. We will now use this model to answer *what-if questions* using only the already collected data, i.e. without the need to collect more data or perform additional experiments.

This reasoning used to answer what-if questions is referred to as *counterfactual thinking*. By asking “What would be the performance of a user of the LDNS service if one of her parameter was to behave as it does when the GDNS is used,




Figure 2: Comparison of the throughput with the quantity of data a client can handle (rwin*)

knowing that the use of the LDNS and the GDNS are exclusive?”, we can estimate the impact of the choice of a DNS service on user performance. Such an approach allows to estimate the impact of choosing one DNS service instead of another and, even more interesting, allows us to estimate the impact of this choice on a given parameter that, in turn, impacts the user performance. In our work, we focus on the impact of the choice of a DNS service on the user throughput via the impact of the DNS service on the CDN server location (c.f. Section 4.2.1), and via the impact of the DNS service on the CDN server configuration (c.f. Section 4.2.2).

Since we deal with probabilities, we will compare the expected values of the throughput³ instead of its average values⁴ as we did in the previous section.

4.2.1. Distance and delay

To investigate the impact of the RTT on download performance we investigate the question: “What would have been the performance of a user served by the local DNS if it would have been redirected to a server whose inetrtt corresponds

³ $\mathbb{E}[TPUT] = \int_{D_{TPUT}} f_{TPUT}(tput) \cdot tput \cdot dtput$, with D_{TPUT} the throughput domain

⁴ $\mu_{TPUT} = \frac{1}{N} \sum_{i=1}^N throughput_i$, with N the total number of observations

to the one the Google DNS service would have redirected him to ?”.

To answer this question is equivalent to predicting the effect of an intervention where the external delay (RTT) experienced by clients served by the LDNS is modeled by the the distribution of the delay experienced by clients served by the GDNS; the distribution of the rest of the parameters being kept identical for the LDNS service users.

More formally, if RTT denotes the $inetrtaavg$ parameter, LD the local DNS and GD the Google DNS, we need to estimate the following distribution:

$$f(TPUT = tput|DNS = LD, do(RTT \sim f_{RTT|do(DNS)}(\cdot, GD))) \quad (7)$$

The causal graph in Figure 1 (cf the explanation of d-separation in Section 2.3) tells us that $(RTT \perp\!\!\!\perp DNS)_{G_{DNS}}$, which implies (Rule 2 from Theorem 1):

$$f_{RTT|do(DNS)}(rtt, GD) = f_{RTT|DNS}(rtt, GD). \quad (8)$$

To predict how an intervention on X affects Y , where the intervention on X is enforced with the conditional probability distribution $f^*(X|Z)$ we use Equation (4). The causal graph in Figure 1 (cf the explanation of d-separation in Section 2.3) tells us that $(RTT \perp\!\!\!\perp TPUT|DNS, TOD)_{G_{RTT}}$. It follows, from Rule 2 of Theorem 1 that

$$\begin{aligned} f_{TPUT|do(RTT), TOD, DNS}(tput, rtt, tod, dns) = \\ f_{TPUT|RTT, TOD, DNS}(tput, rtt, tod, dns). \end{aligned} \quad (9)$$

As a consequence, we can rewrite Equation (7) as:

$$\begin{aligned} f(tput|LD)_{f(rtt|do(GD))} = \\ \int_{D_{RTT}} \int_{D_{TOD}} f(tput|do(rtt), LD, tod) f(tod) f(rtt|do(GD)) P(GD) = \\ \int_{D_{RTT}} \int_{D_{TOD}} f(tput|rtt, LD, tod) f(tod) f(rtt|GD) P(GD) \end{aligned} \quad (10)$$

using Equation (8) and Equation (9).

The result of the intervention is presented in Figure 3. The CDF of the throughput for the LDNS before intervention is plotted as blue solid line and the CDF of the throughput for the LDNS service users after an intervention setting their external delays distribution to the delay distribution seen by the GDNS users is plotted as red dotted line. The throughput after invention is degraded due to the higher




Figure 3: Evolution of the throughput distribution before and after intervening on the external delay experienced by Local DNS (LDNS) clients

RTTs experienced by the clients’: The expected throughput for clients using the local DNS service prior to intervention is 3.5 Mbps and 3.0 Mbps after intervention (14% decrease). This result quantifies the gain in performance that the redirection to closer CDN servers, provided by the use of the local DNS service, represents.

This result also illustrates the use of *counterfactual thinking*. We can deduce the gain in performance for a user who chose the LDNS service by estimating the change in performance if the GDNS would have been chosen instead.

The results obtained cannot be validated in practice as this would require the modification of the behavior of the local DNS servers. In fact, this difficulty nicely illustrates the benefit of the causal approach: it offers the possibility to predict the effect of interventions that are impossible to perform experimentally. Our approach allows us to estimate what would have been the effect on a user performance if she would have chosen the GDNS service, knowing that in reality the LDNS was used.

Figure 4 shows the distribution of the external RTT for GDNS users and LDNS users. Both conditional probability distributions present a long tail and very few values are actually observed for a RTT > 200 ms. It is important to mention that RTT values are observed for the LDNS users for the range [0.5ms,200ms] and for GDNS users for the range [10ms,200ms]. This condition is necessary to perform the prediction performed in this section, which is a limitation of the method used: The prediction formulated in Equation (7) is only possible since the range of the




Figure 4: Histogram of the external RTT for the local DNS (LDNS) and Google DNS (GDNS)

external RTT values observed for GDNS users represents a subset of the range of values observed for the LDNS users.

If one wants to study the opposite intervention, where the users of the GDNS service would be given access to servers placed at the locations of the servers the LDNS service users are redirected to, the prediction would be more complex. We do not have samples to estimate $f(tput|rtt, GD, tod)$ for some of the smallest RTT values ($RTT < 10$ ms) for which we have $f(RTT|LDNS) > 0$. However, this limitation should not surprise us, since it is common to many machine learning problems where the amount of available information determines the predictions we can make. The reason why we cannot predict the opposite intervention is due to the use of kernels to estimate distributions, which requires the presence of samples in a given region to estimate the value of the distribution in this region. One possible way to overcome this problem would be to develop a parametric model that allows to extrapolate the different PDFs beyond the value range where the variables of our system are observed.

It is important to note that our model considers the impact of the change in the delay distribution but also the impact of the servers themselves, captured by the minimum congestion window and parameters such as the loss (*retrscore*) that are different between the two DNS services. In fact, the influence of these parameters may explain that the throughput experienced, in the original dataset, by the users of the GDNS service is only 7% smaller than for the users of the local DNS service. To evaluate the impact of the servers on download performance we focus on the impact of the minimum congestion window since *cwinmin* is a direct parent

of the throughput ($tput$) and is influenced by the DNS service choice (dns). Also, other parameters such as the loss parameters ($retrscore$ and rto), the delay parameters ($isprttavg$ and $isprttstd$) or the maximum congestion window ($cwinmax$) are not influenced by the choice of the DNS service (dns) (c.f. Figure 1).

4.2.2. Minimum congestion window

The minimum congestion window ($cwinmin$) is a direct parent of the throughput ($tput$), see Figure 1. Its average value is higher for the clients using the GDNS service than for the clients using the LDNS service (1.2kB and 0.9kB respectively). The difference in the expected value of the throughput of LDNS users (3.5 Mbps) and GDNS users (3.3 Mbps) is 6%, smaller than the gain for the LDNS users being redirected to closer server, that is estimated to be 14%. Our hypothesis is that the minimum congestion window represents a difference in the configuration of the servers accessed by the LDNS users and the configuration of the servers accessed by the GDNS users. To evaluate this hypothesis we estimate the causal effect of the minimum congestion window on the throughput, mediated by the choice of the DNS service. This is equivalent to asking the question: “*What would be the throughput for the clients using the local DNS if the servers they are redirected to would present the same minimum congestion window as the ones Google DNS users are redirected to ?*”.

We observe from the causal graph of Figure 1 (cf the explanation of d-separation in Section 2.3):

- $(CWINMIN \perp\!\!\!\perp DNS)_{G_{DNS}}$
- $(CWINMIN \perp\!\!\!\perp TPUT | DNS, INETRRTSTD)_{G_{CWINMIN}}$

For space reasons, and because the approach is the same as in section 4.2.1 for the external delay ($inetrttavg$), we only present the final equation.

Let denote $cmin$ the minimum congestion window (called $cwinmin$ in our model) and σ_{rtt} the standard deviation of the external rtt (called $inetrttstd$ in our model). As before LD refers to the local DNS and GD to the Google DNS. We obtain the following equation:

$$\begin{aligned}
 f(tput | LD)_{f(cmin|do(GD))} = & \\
 \int_{D_{CMIN}} \int_{D_{\sigma_{RTT}}} f(tput|do(cmin), LD, ts) f(\sigma_{rtt}) f(cmin|do(GD)) P(GD) = & \\
 \int_{D_{CMIN}} \int_{D_{\sigma_{RTT}}} f(tput|cmin, LD, \sigma_{rtt}) f(\sigma_{rtt}) f(cmin|GD) P(GD) & \quad (11)
 \end{aligned}$$




Figure 5: Evolution of the throughput distribution before and after intervening on the minimum congestion window of servers of the users of the local DNS

Equation (11) allows the prediction of the distribution of the throughput for the LDNS users after an intervention when we use for the minimum congestion window the distribution seen by GDNS users. The CDFs of the pre-intervention throughput (solid line) and post-intervention throughput (dotted line) are presented in Figure 5. We can see the gain in throughput due to the intervention on the minimum congestion windows of the LDNS servers. The expected throughput for LDNS service users after the intervention is 4.6 Mbps (compared to 3.5 Mbps prior to intervention), which represents an increase of more than 30%. This increase is due to the fact that the servers GDNS service users are redirected to use higher values for their minimum congestion window.

The study of the opposite intervention, where GDNS service users are redirected to servers with a minimum congestion window following the distribution of the minimum congestion window seen by the LDNS service users, in the original dataset, is not possible. The reason is the same as the one mentioned in Section 4.2.1. If we compare the distribution of the minimum congestion windows for LDNS service users and GDNS service users, Figure 6, we can notice the absence of $cmin$ values for GDNS users to estimate $f(tput|cmin, GD, \sigma_{rtt})$ for values of $cmin$ where $f(cmin|LD) > 0$.

If we summarize the findings of the last two sections, we can say that by using a causal model and its graphical representation we were able to quantify that it is not only the proximity of the server that has an important impact on the throughput but also the configuration of the server hosting the content a client wants to access.




Figure 6: Histogram of the minimum congestion window for the Local DNS (LDNS) and Google DNS (GDNS)

In a causal model such as the one presented in Figure 1, a given node X also represents the influence that external factors impacting *only this parameter* have on the rest of the system. This means that the difference in behavior of Akamai servers that the Google DNS redirects the clients to compared to the behavior of the servers the LDNS redirects the clients to may not be solely the effect of the minimum congestion window but may also be the effect of other un-observed parameters of TCP such as the *additive increase value* for each acknowledged packet. Unfortunately, we have no means to validate this hypothesis.

5. Related work

The two works closest to ours are WISE [Tariq et al., 2008] and Nano [Tariq et al., 2009].

WISE uses, as does our work, the PC algorithm [Spirtes and Glymour, 1991] to infer a graphical causal model from which interventions are then predicted. However, WISE requires a lot of domain knowledge in its feature selection and in the definition of external causes that guide the inference of the causal model. Also, WISE uses the Z-Fisher independence, which assumes linear dependencies. We have tested the Z-Fisher independence criterion in our work and obtained very poor results as the test fails to detect parameter independences resulting in incorrect models [Hours et al., 2015]. In addition, WISE considers much simpler scenarios of intervention and requires a much larger data set. Our approach takes

full advantage of the causal theory developed by Pearl [Pearl, 2009; Spirtes et al., 2001] to predict interventions and counterfactuals. Counterfactuals are very useful to understand the role of the different parameters of a system and, to our knowledge, scenarios such as the ones presented in Section 4.2 have not been treated so far.

Nano tries to detect network neutrality violation by assessing the direct causal effect between the quality of experience of a user from a given ISP and the type of content being accessed. A performance baseline is defined based on observations made for different ISPs sharing similar configurations and then compared to the one observed for a particular scenario. Again, this approach uses domain knowledge to define the possible confounders and to condition on these variables to remove spurious associations. Since Nano has not derived a formal causal model, its approach has serious limitations since one of the confounders could be a collider in the corresponding causal graphical model. Also, conditioning on a common effect induces a dependence between two independent causes whose influence tries to be canceled, questioning the obtained results.

Several papers study how the choice of the DNS service impacts client performance [Ager et al., 2010; Huang et al., 2011; Otto et al., 2012]. These works rely on active measurements and differ greatly in their approach and objectives from our work.

In our previous work [Hours et al., 2015] we presented solutions to the problem of causal model inference and to the prediction of atomic interventions for cases where the assumptions of normality and linearity do not hold. We also validated our approach and showed for simple systems and scenarios that it was possible to use a causal approach to study communication network performance.

The work presented in this paper goes much further. First, we study a more complex system with more parameters and diverse categories of data (including categorical data). We use the causal model obtained to explain non intuitive observations (namely a similar throughput for connections experiencing a different RTT). Second, the major contribution of the work presented in this paper is due to the use of *counterfactuals* and *counterfactual thinking*, Section 4.2. The use of *counterfactuals* gives us access to a deeper understanding of the causal mechanisms ruling the performance of the system and it allows us to quantify the impact of each of these mechanisms on the performance of this system.

6. Concluding remarks

The main contribution of our paper resides in the methodology based on the inference and in the usage of a causal model that allows us to estimate the causal effect of the DNS service on user performance. Using a causal approach and inferring the causal model, which is then represented as a Bayesian graph, we are able to study the causal effect of a DNS service on the TCP throughput. We compare the performance of clients using their ISP local DNS service to the performance of clients using the Google DNS service. The causal model allows to unveil dependencies that would be very difficult, if not impossible, to extract otherwise from the data. We showed that the choice of the DNS service has a strong impact on the location of the servers the clients are redirected to, which in turn impacts not only the distance from clients to servers but also the type of configuration of the servers. Distance and configuration are captured by the dependence between the DNS and the RTT and the dependence between the DNS and the server minimum congestion window.

A very interesting property of causal models is their “*stability under intervention*”. The model inferred from data following a given distribution is still valid when we predict the effect of modifying this distribution. When comparing the performance of the users of the local DNS and the users of the Google DNS, we can observe that the performance difference cannot simply be explained by the redirection of Google DNS users to more distant Akamai servers. Based on the causal graph obtained, we can formulate the hypothesis that the configurations of the Akamai servers Google DNS users are redirected to allow them to experience a performance close to the one of the local DNS service users. This hypothesis is confirmed by our prediction where we give to Akamai servers serving the local DNS users a minimum congestion window equivalent to the one of the Akamai servers serving Google DNS users. We estimate the gain in throughput corresponding to this intervention to be 32%. By comparison, the gain in terms of throughput corresponding to the redirection of the local DNS users to closer servers is estimated to be only 14%.

We demonstrated the potential of adopting a causal approach using counterfactuals. Counterfactuals are one of the possible way to approach Causality and we use this technique to evaluate the effect of a parameter on the system performance by predicting the effect that changing its parent would have with the rest of the system parameters left unchanged. We manage to answer questions such as “*How would the system behave under the condition C1 if one of its parameter was to behave as it has done under the condition C2, knowing that C1 and C2*”

are exclusive ?". The ability to make predictions for such scenarios illustrates the power of the inherent mechanisms underlying the development of Causality. Counterfactuals are relatively complex to study, explain and even more so to predict. However, thanks to the Bayesian network as a representation of the causal model of our system, using counterfactuals becomes easier.

Complex interventions, where many parameters are modified simultaneously, require important resources in terms of the amount of data and computational power. The results presented in this paper document a first successful attempt. Based on this work, we are confident that the underlying tools and methods can be improved to reduce the required resources and increase both, the accuracy of such predictions and the range and complexity of the interventions that one can consider.

We see the following directions for future work:

- By fitting a parametric model we could extend the prediction of counterfactuals for cases where the two conditional probabilities have only partial overlap .
- The weight of an edge, $X \rightarrow Y$, corresponds to what is known as the *direct effect* of one parameter, X on another, Y . However, in the absence of linearity, the estimation of the direct effect of X on Y is complex and requires predicting the effects of several interventions [Pearl, 2013] for each direct effect, which requires a lot of computational resources.
- Regarding a selection criterion, the absence of any assumption regarding the distribution of the parameters and the nature of their dependencies prevents us from using a classical selection criterion such as maximum likelihood. Two possibilities could be used instead: (i) A *Bootstrap* approach, where, by re-sampling the original dataset to create new datasets, we could infer one causal model for each dataset and, by comparison, derive a confidence level for our model. This approach is simple to implement. However, the inference of the causal model presented in this paper took up to one week running on a cluster of 30 machines. Therefore, a *bootstrap* approach requires important resources in terms of computation time. On the other hand, when creating sub-datasets, we work with smaller datasets, which has an impact on the accuracy of the results. (ii) We could use the *independence test p-values* to obtain a confidence in the presence or absence of any edge in the graph we obtain to give us a confidence in the model. This approach

becomes complex due to the number of tests to consider for a given pair of nodes and no general criterion has been designed as this stage of our work.

- We had to design several solutions to build a reliable framework for causal knowledge inference [Hours et al., 2015] that implied an increase in complexity and resource requirements. While we have used very small datasets to validate our approach and to show its benefits, there are many directions to explore to make Causal reasoning work more efficiently on large quantities of data thanks to the use of distributed computing. The parallelization of the independence testing for causal model inference [Hours et al., 2015] and the parallelization of the estimation of interventions (see Appendix B.2) fit very well a Big Data approach. Working with a bigger and partitioned dataset on which parallel computing could be done, would improve the performance of the Causal knowledge inference framework we presented in this work.

Acknowledgment

We would like to thank Elias Bareinboim from UCLA for his advice and support in developing the methodology to predict counterfactuals. We also would like to thank the reviewers for their constructive remarks. The research leading to these results has received fundings from the European Union under the FP7 Grant Agreement n. 318627 (Integrated Project "mPlane").

Ager, B., Mühlbauer, W., Smaragdakis, G., Uhlig, S., 2010. Comparing DNS resolvers in the wild. In: IMC. ACM, pp. 15–21.

Darwiche, A., 2009. Modeling and Reasoning with Bayesian Networks, 1st Edition. Cambridge University Press.

Demarta, S., McNeil, A., 2005. The T-copula and related copulas. *International Statistical Review* 73 (1), 111–129.

Finamore, A., et al., May 2011. Experiences of internet traffic monitoring with tstat. *IEEE Network Magazine* 25, 8–14.

Ford, P., Shelest, A., Srinivas, N., Aug. 15 2002. Method for automatic tuning of TCP receive window. US Patent App. 09/736,988.

Hours, H., Biersack, E., Loiseau, P., Dec. 2015. A causal approach to the study of TCP performance. *ACM Trans. Intell. Syst. Technol.* 7 (2), 25:1–25:25.

Huang, C., Maltz, D. A., Li, J., Greenberg, A. G., 2011. Public DNS system and global traffic management. In: INFOCOM. IEEE, pp. 2615–2623.

- Jaworski, P., Durante, F., Härdle, W., Rychlik, T., 2010. Copula theory and its applications. Lecture Notes in Statistics. Springer Berlin Heidelberg.
- Mayer-Schönberger, V., 2013. Big Data: A Revolution That Will Transform How We Live, Work and Think. John Murray Publishers.
- Otto, J. S., Sánchez, M. A., Rula, J. P., Bustamante, F. E., 2012. Content Delivery and the Natural Evolution of DNS: Remote DNS trends, performance issues and alternative solutions. In: IMC. ACM, pp. 523–536.
- Pearl, J., 2009. Causality: Models, Reasoning and Inference. Cambridge University Press.
- Pearl, J., 2013. Direct and indirect effects. CoRR abs/1301.2300.
- Pitt, M., Chan, D., Kohn, R., 2006. Efficient Bayesian inference for gaussian copula regression models. Biometrika 93 (3), 537–554.
- Spirtes, P., Glymour, C., 1991. An Algorithm for Fast Recovery of Sparse Causal Graphs. Social Science Computer Review 9, 62–72.
- Spirtes, P., Glymour, C., Scheines, R., Jan. 2001. Causation, Prediction, and Search, 2nd Edition. The MIT Press.
- Streibelt, F., Böttger, J., Chatzis, N., Smaragdakis, G., Feldmann, A., 2013. Exploring edns-client-subnet adopters in your free time. In: IMC. ACM, pp. 305–312.
- Tariq, M., Motiwala, M., Feamster, N., Ammar, M., 2009. Detecting network neutrality violations with causal inference. In: CoNEXT. ACM, pp. 289–300.
- Tariq, M., et al., 2008. Answering what-if deployment and configuration questions with WISE. In: SIGCOMM. ACM, pp. 99–110.
- Zhang, K., Peters, J., Janzing, D., B., S., 2012. Kernel-based conditional independence test and application in causal discovery. CoRR abs/1202.3775.

Appendix

A Study of the Impact of DNS Resolvers on CDN Performance Using a Causal Approach

Hadrien Hours^{a,b}, Ernst Biersack^c, Patrick Loiseau^a,
Alessandro Finamore^{d,e}, Marco Mellia^e

^a EURECOM, email: firstname.lastname@eurecom.fr

^b ENS Lyon, email: firstname.lastname@ens-lyon.fr

^c Caipy, email: erbi@e-biersack.eu

^d Telefonica, email: firstname.lastname@telefonica.com

^e Politecnico di Torino, email: firstname.lastname@polito.it

Appendix A. Causal model inference

Appendix A.1. PC algorithm

We have described the PC algorithm in Section 2.1. In Figure A.7 we now illustrate the different steps. In this example we try to infer the causal model of the system corresponding to four parameters, W, X, Y, Z . We could detect two independences: $\mathbf{I}_1 = (X \perp\!\!\!\perp Y|W)$ and $\mathbf{I}_2 = (W \perp\!\!\!\perp Z|\{X, Y\})$. In Figure A.7b and Figure A.7c the edge between X and Y and the edge between W and Z are removed for the independences detected with a conditioning set size of 1 and 2 respectively. Because of detected $X \perp\!\!\!\perp Y$ but $X \not\perp\!\!\!\perp Y|Z$ but we can orient $X - Z - Y$. However, the orientation of the second V-structure, $X - W - Y$, cannot be deduced from the set of detected independences. The three orientations presented in Figure A.7e, Figure A.7f and Figure A.7g verify the independences \mathbf{I}_1 and \mathbf{I}_2 .




Figure A.7: Different steps of the inference of causal model corresponding to a system of four parameters, $\{W, X, Y, Z\}$ with the detected independences $\mathbf{I}_1 = (X \perp\!\!\!\perp Y|W)$ and $\mathbf{I}_2 = (W \perp\!\!\!\perp Z|\{X, Y\})$

Appendix A.2. Independence test

The accuracy of the PC algorithm comes from the accuracy of the test used to test parameter independences. Compared to our previous works [Hours et al., 2015], one difference comes from the presence of a categorical variables, like the DNS service used by the clients observed in our study or the destination IP address. The test we use in our study is the KCI test [Zhang et al., 2012] combined with a bootstrap approach to solve numerical issues in its use of Cholesky factorization and to parallelize computations and decrease the algorithm completion time [Hours et al., 2015].

To validate the use of the KCI test in the presence of categorical variable, we generate two artificial datasets as follows:

- Dataset 1:
 - X_1 is a categorical variable with 4 levels: $X_1 \sim \mathbf{U}\{c_1, c_2, c_3, c_4\}$.
 - X_2 is a deterministic mapping of X_1 , adding 20% of Gaussian noise: $X_2 = f_2(X_1) + \varepsilon$.
 - X_3 is a deterministic mapping of X_1 , adding 20% of Gaussian noise, $X_3 = f_3(X_1) + \varepsilon$.
 - X_4 is a function of X_2 and X_3 , adding 20% of Gaussian noise: $X_4 = f_4(X_2, X_3) + \varepsilon$.
 - with
 - * f_2 and f_3 defined as $f_i(c_j) = c_{i,j}$ for $j \in \{1, 2, 3, 4\}$, with $c_{ij} \neq c_{i'j'}$ if $i \neq i'$ or $j \neq j'$
 - * $f_4(x, y) = \sqrt{x + y}$.
 - * ε an error terms following normal distribution with a mean equals to 0 and a variance equals to $0.2 \times \sigma_{f_i(X_i)}$.
- Dataset 2:
 - X_1 is a categorical variable with 4 levels: $X_1 \sim \mathbf{U}\{c'_1, c'_2, c'_3, c'_4\}$.
 - X_2 is a deterministic mapping of X_1 , adding 20% of Gaussian noise: $X_2 = f'_2(X_1) + \varepsilon$.
 - X_3 is a categorical variable with 4 levels, the probability of each level depends on X_1 : $X_3 = f'_3(X_1) + \varepsilon$.
 - X_4 is a function of X_2 and X_3 , adding 20% of Gaussian noise: $X_4 = f'_4(X_2, X_3) + \varepsilon$.
 - with
 - * f'_2 defined by, $f'_2(c'_j) = c_{2,j}$ for $j \in \{1, 2, 3, 4\}$,
 - * $f'_3(c'_j) \in \{c''_1, c''_2, c''_3, c''_4\}$, each value c''_k drawn from 4 different distributions, chosen base on the value of c'_j
 - * $f'_4(x, y) = c(x) + \sqrt{y}$, with $c(x)$ defined as a deterministic mapping of x , similarly to f'_2 .
 - * ε an error terms following normal distribution with a mean equals to 0 and a variance equals to $0.2 \times \sigma_{f_i(X_i)}$.

The graphical causal model corresponding to these dependencies is presented in Figure A.8. The definition of the two datasets leads to two independences $\mathbf{I}_1 = (X_2 \perp\!\!\!\perp X_3 | X_1)$ and $\mathbf{I}_2 = (X_1 \perp\!\!\!\perp X_4 | \{X_2, X_3\})$.

We test different independences for 20 artificial datasets of size 1000, generated according the definitions given above. The average p-values of the different tests for the two classes of artificial datasets are presented in Table A.3. We can see that the KCI performs correctly even in the presence of complex dependencies including the presence of categorical variables. The conclusion of this study is that the KCI performs as expected in the presence of categorical variable and can be used in our study.




Figure A.8: Graphical causal model illustrating the dependencies of the artificial dataset parameters

Table A.3: Results of the KCI test when testing independences with the presence of categorical parameters

Independence	p-value dataset 1	p-value dataset 2
$X_1 \perp\!\!\!\perp X_2$	0	0
$X_1 \perp\!\!\!\perp X_3$	0	0
$X_1 \perp\!\!\!\perp X_4$	0	0
$X_2 \perp\!\!\!\perp X_3$	0	0
$X_2 \perp\!\!\!\perp X_4$	0	0
$X_3 \perp\!\!\!\perp X_4$	0	0
$X_1 \perp\!\!\!\perp X_2 X_3$	0	0
$X_1 \perp\!\!\!\perp X_2 X_4$	0	0
$X_2 \perp\!\!\!\perp X_3 X_1$	0.4	0.6
$X_2 \perp\!\!\!\perp X_3 X_4$	0	0
$X_3 \perp\!\!\!\perp X_4 X_1$	0	0
$X_3 \perp\!\!\!\perp X_4 X_2$	0	0
$X_1 \perp\!\!\!\perp X_4 X_2$	2e-14	2e-2
$X_1 \perp\!\!\!\perp X_4 X_3$	0	3e-4
$X_1 \perp\!\!\!\perp X_4 \{X_2, X_3\}$	0.7	0.5

Appendix A.3. Markov Equivalence class

In our study of the impact of the DNS service choice on the Akamai CDN performance, we obtain the Bayesian network represented in Figure A.9. Using our understanding of CDN and the parameters present in our model, we orient the undirected edges and obtain the Bayesian network representing the causal model of our system represented in Figure A.10. Using the Tetrad software [Spirtes et al., 2001], it is easy to represent all the members of what is called the *Markov Equivalence Class*. There can be several graphs that represent the set of independences that were detected from the tests performed on a given series of observations. The set of all these graphs can be represented by a partially oriented graph, Figure A.9. In Figure A.11 we represent the eight members of the Markov Equivalence Class corresponding to the set of independences that were




Figure A.9: Output of the PC algorithm when no domain knowledge is used

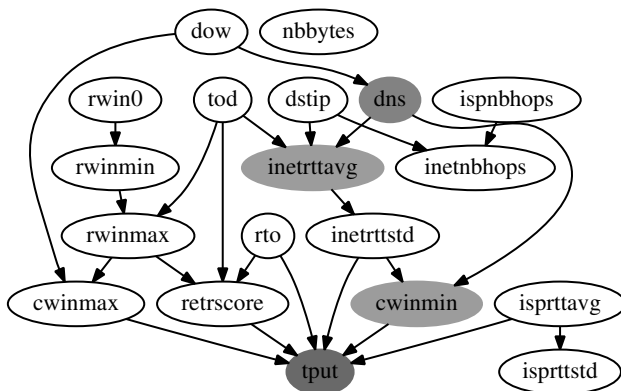


Figure A.10: Bayesian network representing the causal model of Web performance using two different DNS: the public Google DNS and the DNS of the local ISP

detected from the series of tests performed on the observations of the parameters of our system.

Appendix B. Predicting interventions

Appendix B.1. Theory

Appendix B.1.1. Atomic interventions

The description of atomic intervention was presented in Section 2.2.1. In this section, we only repeat the rules of Do-calculus that will be used in the following sections to present the details of our method.

If we use G to denote the Bayesian graph that represents the causal relationships between the parameters of our system, we use $G_{\bar{X}}$ to denote the sub-graph of G where all the edges entering X are removed and $G_{\underline{X}}$ the sub-graph of G where all the edges exiting X are removed. We can use the rules of *do-calculus* from [Pearl, 2009] to estimate the distributions of the parameters of

our system after an intervention based on their distributions prior to this intervention. Note that these rules do not rely on any assumption regarding the distributions or functional dependencies of the parameters. In particular, P represents the (possibly multivariate) probability distribution specified by the probability mass function or probability density function depending on the nature of the parameters.

Theorem 2 (3.4.1 from [Pearl, 2009]). *(Rules of do calculus) Let G be the directed acyclic graph associated with a causal model [...] and let $P(\cdot)$ stand for the probability distribution induced by that model. For any disjoint subsets of variables X , Y and Z we have the following rules.*

Rule 1 (Insertion/deletion of observation):

$$P(y|do(x), z, w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z \mid X, W)_{G_{\overline{X}}} \quad (\text{B.1})$$

Rule 2 (Action/observation exchange):

$$P(y|do(x), do(z), w) = P(y|do(x), z, w) \text{ if } (Y \perp\!\!\!\perp Z \mid X, W)_{G_{\overline{XZ}}} \quad (\text{B.2})$$

Rule 3 (Insertion/deletion of intervention):

$$P(y|do(x), do(z), w) = P(y|do(x), w) \text{ if } (Y \perp\!\!\!\perp Z \mid X, W)_{G_{\overline{XZ(W)}}}, \quad (\text{B.3})$$

where $Z(W)$ is the set of Z -nodes that are not ancestor of any W -nodes in $G_{\overline{X}}$.

Appendix B.1.2. Enforcing intervention with a given probability

In our study of the impact of the DNS service on CDN performance (throughput), we are interested in estimating the effect of interventions on parameters influenced by the DNS service and influencing the throughput. In such case, we do not limit ourselves to atomic interventions but we are interested in intervening on a given parameter to change its distribution.

From [Pearl, 2009, Section 4.2], if we want to predict how an intervention on X affects Y , where the intervention on X is enforced with the conditional probability distribution $f^*(X|Z)$, we obtain:

$$f(y)_{|f^*(x|z)} = \int_{D_X} \int_{D_Z} f_{Y|do(X), Z}(y, x, z) f^*(x|z) f(z) dx dz. \quad (\text{B.4})$$

From Equation (B.4), one should notice that we need to integrate on the intervention parameter, X . For performance reasons, the estimations of $f_{Y|do(X), Z}$ are made in parallel on different machines. Therefore, the estimation of $f_{Y|do(X), Z}(y, x, z)$ is done on a different machine for each x . As the data used in our study is not publicly available, we do not present the different parameterizations of the density estimation used for predicting $f(y)_{|f^*(x|z)}$.

Appendix B.2. Adapting the theory to our problem

Appendix B.2.1. Intervening on the DNS service

To understand how choosing one DNS service instead of another impacts CDN performance, we want to predict the throughput of a client who used a DNS service s_1 if the same client would have used a different DNS service, s_2 , instead. To do so, we are interested in the impact of the DNS service on a given parameter, X , that in turn influences the throughput. We use the Theorem 2 and

Equation (B.4) to estimate the distribution of the throughput of a client using a given DNS service, s_1 , if we intervene on the distribution of a parameter X , forcing its distribution to follow the one of X if the DNS service s_2 would have been used instead. This study is equivalent to study the impact of the parameter X for clients using the DNS service s_1 when intervening on the parameter X enforcing this intervention with the distribution $f_{X|DNS=s_2}$. If we denote Y the parameter capturing the performance of CDN users⁵, we want to estimate:

$$f(Y|DNS = s_1, do(X \sim f(X|DNS = s_2))). \quad (\text{B.5})$$

If we denote W the set of parameters blocking the spurious associations between X and Y , according to Theorem 2, we have

$$f(Y | DNS = s_1 | do(X \sim f(X|DNS = s_2))) = \int_X \int_W \overbrace{f(Y | X = x, DNS = s_1, W)}^{f(Y|DNS=s_1, do(X=x))} f(W) f(X = x | DNS = s_2) P(DNS = s_2) \quad (\text{B.6})$$

We can see from Equation (B.6) that the distribution of Y for users of the DNS s_1 , if we intervene on X and fix its distribution to follow the distribution of X seen by the users of DNS s_2 , is a weighted sum of the distribution of Y for DNS = s_1 after an atomic intervention on X ($do(X=x)$) with weights being the probability of observing $X = x$ under DNS = s_2 .

Such approach allows to (i) Capture the effect of the DNS on a given mechanism influencing the performance of CDN users; (ii) Divide our prediction in a set of predictions of atomic interventions that can be estimated from the results of Theorem 2. Finally, we can use Equation (B.4) to estimate the final distribution of Y for the intervention on X that modifies its distribution.

Appendix B.2.2. Interventions and conditional multivariate distributions

It should be noticed that, as X is a continuous variable, the probability of observing a given value is 0. Therefore, instead of selecting samples for which $X = x$ is observed, we define an interval I_x corresponding to $[x - \delta_X; x + \delta_X]$ and assume that the samples for which the X parameter falls into this interval can be approximated to take the value x .

From the samples where $X \in I_x$ and DNS = s_1 , we estimate the cumulative distribution function (CDF) of Y and W conditionally to DNS = s_1 . We then estimate $f(Y|X = x, W = w, DNS = s_1)$ using the Sklar theorem.

The Sklar theorem stipulates that, if F is a multivariate cumulative distribution function with marginals $(F_1, \dots, F_i, \dots, F_n)$, there exists a copula C such that

$$F(x_1, \dots, x_i, \dots, x_n) = C(F_1(x_1), \dots, F_i(x_i), \dots, F_n(x_n)). \quad (\text{B.7})$$

If we take the example of the bivariate distribution of two parameters X_1 and X_2 , which marginals are denote F_1 and F_2 and f_1 and f_2 for the CDFs and PDFs respectively, we obtain, taking the derivative of Equation (B.7) with respect to X_1 and X_2 :

$$f(x_1, x_2) = c(F_1(x_1), F_2(x_2))f_1(x_1)f_2(x_2), \quad (\text{B.8})$$

⁵In our study, we use the throughput to measure user performance

with f the bivariate PDF of X_1 and X_2 .

As we have:

$$f_{X_1|X_2}(x_1, x_2) = \frac{f(x_1, x_2)}{f_2(x_2)}, \quad (\text{B.9})$$

for values of X_2 for which $f_2(x_2) \neq 0$, we can deduce that:

$$f_{X_1|X_2}(x_1, x_2) = c(F_1(x_1), F_2(x_2))f_1(x_1). \quad (\text{B.10})$$

We can then estimate the conditional CDFs, $F(Y|X = x, DNS = s_1)$ and $F(W|X = x, DNS = s_1)$, using kernels and the samples where $X \in I_x$ and $DNS = s_1$, and use the previous formula to estimate $f(Y|X = x, W = w, DNS = s_1)$:

$$\begin{aligned} f(Y|X = x, W = w, DNS = s_1) &= c(F(Y|X = x, DNS = s_1), F(W|X = x, DNS = s_1)) \\ &f(Y|X = x, DNS = s_1), \end{aligned} \quad (\text{B.11})$$

Finally, by integrating Equation (B.6) on W we obtain the distribution $f(Y|do(X = x), DNS = s_1)$.

We then select the samples for which $DNS = s_2$ and use normal kernels to estimate $f(X|DNS = s_2)$ and frequencies to estimate $P(DNS = s_2)$.

After these steps, we have all the factors present in Equation (B.6) and we can integrate over X to obtain the distribution of Y post intervention.

Appendix B.2.3. Estimation of marginals

Some practical issues are silenced in the sequence of steps described in Appendix B.2.2:

1. How to define the intervals I_X ?
2. As we are working with continuous variables, the two distributions conditionally to different DNS values might not have the same support. How do we define the conditional probability so that we have common values to integrate on ?

The last point is solved by always defining the PDFs domains as equally spaced points between the minimum and maximum observed value of the corresponding parameter in the whole dataset ($DNS = s_1$ or $DNS = s_2$). We use normal kernels to estimate the different distributions.

The first point, however, is more complicate as many possibilities exist and there is the constraint of finding enough samples in each interval to estimate the CDFs from which the copula parameters will be estimated and the conditional distributions derived.

Several solutions have been tested

- Variable bin width histogram,
- Fixed bin width and interpolation,
- Fixed bin width, filtering, rescaling.

Variable bin width histograms. The first method consists in fixing an objective number of samples and, starting from a fixed width bin histogram, merging adjacent bins until obtaining bins of different size but with a minimum number of samples:

Pros This method ensures the maximum number of atomic predictions being successful.

Cons As many of the parameters we observe have a long tail distribution, and as it is often in the tail of the parameter distributions that we find the values for interesting predictions, we obtain very large bins for the extreme values of the intervention parameter. The approximation stating that this bin represents a single value is then too strong. Additionally, when multiplying the atomic intervention fixing X parameter to the value x by the value of the PDF conditionally to the other DNS for $X = x$ (Equation (B.6)) we need to take into account the actual range of X that the atomic intervention represents to have a consistent approximation.

Fixed bin width. The second method is going in the opposite direction. We use histograms with fixed bins and then try to make predictions of $f(Y|do(x))$ with the number of samples we found in a given bin corresponding to a X value. If the estimation of the prediction fails, then we use interpolation of the $f(Y|do(x))$ for the X values where the post intervention PDF of Y could be estimated.

Pros This method solves the issues of approximation inconsistency of the variable bin width method. We fix the bin width and decide on the approximation of assimilating an interval to a given value.

Cons It can often happen that we manage to predict $f(Y|do(x))$ even when there are few sample in the interval I_x corresponding to the x value. However, with very few samples, it is very likely that this estimation is not accurate. This lack of accuracy impacts in a very negative way the overall post intervention PDF accuracy, as the PDF of Y post intervention that could be computed with few values will be used in the interpolation to recover the PDF Y for x values where the estimation of the post intervention PDF failed.

Fixed bin width and high pass filter. The last method, the one eventually adopted, is based on the fact that, if there are very few samples in a given area of the distribution of $X|DNS = s_1$ then this value is very unlikely to be observed and, as the previous method will impact negatively the overall results by including these samples, we consider that the PDF of $X|DNS = s_1$ in the domains with few samples is null.

This method uses a fixed bin histogram and selects only the bins where a minimum number of samples is observed to predict atomic interventions. After predicting the distribution of Y after intervention we rescale it based on the following observation:

$$\int_Y f(Y|do(X = x), DNS = s_1) dY = 1 \quad (\text{B.12})$$

In practice, instead of varying the bin width and threshold we do the following: We select an objective number of samples, Δ_S , under which the atomic intervention $f(Y|X_2 = 0, do(X = x))$ is considered as null, and search for the optimal quantization leading to the maximum bins with a number of samples $\geq \Delta_S$. To do so we use this very simple algorithm:

Data: Vector X

Result: Set of bins (X_{final}) with fixed width (δ) defining the intervals around the values on which atomic interventions will be predicted:

```

nB = 10;
nVold = 0;
nVcur = 1;
Xold = [];
Xcur = [];
Hold = [];
Hprev = [];
while nVcur > nVold do
    Ht, Et = hist(X, nB);
    nVold = nVcur;
    nVcur = nvalues(Ht > ΔS);
    Xold = Xcur;
    Xcur = Et;
    Hold = Hcur;
    Hcur = Ht;
    nB = 2 * nB;
end
nVfinal = nVold;
Xfinal = Xold;
Hfinal = Hold;
δ = Xfinal(2) - Xfinal(1);

```

Algorithm 1: Dynamic quantization

This method solves the previous issues. There is a risk of not having any prediction for values in the tail of the distribution of $X|DNS = s_1$, that can represent the zone of overlap of the two distributions $X|DNS = s_1$ and $X|DNS = s_2$. However, limitations due to the lack of samples cannot be overcome by simple approximations as seen in the second method. The only way to solve this issue is to use parametric distributions (for example a mix of normal, gamma, beta, log-normal distributions). The studies made so far have shown that the approximation of the distributions of the parameters of our systems for values that are actually observed offers an acceptable accuracy. However, these models become inaccurate as soon as we try to use them to estimate the distribution of a parameter for values that have not been observed.

Appendix C. Parameterization of the method

In this section, we present a study that aims to answer the following questions:

1. How to choose the copula family that will model the dependencies between the different marginals ?
2. In which proportion the absence of values observed for both conditional distributions impacts the prediction accuracy and how to improve the accuracy in this case ?
3. As discussed previously, we estimate a complex intervention as the weighted average of atomic (simpler) interventions. Each atomic intervention is estimated on a sub domain

corresponding to an (small) interval around the value corresponding to our atomic intervention. On one hand, the bigger the interval is the more data we have to calibrate our model and the better prediction for the corresponding atomic intervention will be. On the other hand, if we manage to estimate many atomic interventions, we will have more inputs for estimating the global intervention we are interested in. The question boils down to finding the trade-off between the number of atomic interventions we estimate and the quality of each of these estimates.

To study these aspects we need a ground truth to estimate the accuracy of our prediction for different strategies and parameterization. Therefore, we generate a set of artificial datasets where the intersection between the two conditional distribution (DNS = s_1 and DNS = s_2) domains is modified and its impact on the accuracy of our method is studied.

Appendix C.1. Simulated dependencies

To simulate the same situation as the one met in the study of DNS service impact on the CDN performance, we randomly generate 4 parameters, X_1, X_2, X, Y , with dependencies illustrated by the graph presented in Figure C.13. To be closer to the situation observed in our study of the impact of DNS on CDN performance, we generate X_1 by randomly re-sampling the throughput observed in this study and we generate X_2 by randomly re-sampling the observed DNS from the same study. We eventually convert X_2 to binary value (0 or 1). The presence of a categorical data (the DNS in the corresponding study) is an important aspect that we want to keep in this study.

Appendix C.2. Intervention prediction

First, from the causal model, G , represented by the Bayesian network of Figure C.13, we can use the d-separation criterion to deduce the following independences:

- $(X \perp\!\!\!\perp Y | X_1, X_2)_{G_X}$
- $(X_2 \perp\!\!\!\perp X)_{G_{X_2}}$

From the do-calculus rules from [Pearl, 2009] we can deduce that the distribution of Y under the condition $X_2 = 0$ intervening on X to fix its distribution to $X \sim X | do(X_2 = 1)$ is given by

$$f(Y | X_2 = 0, do(X \sim X | do(X_2 = 1))) = \int_X \int_{X_1} f_{Y|X, X_1, X_2}(x, x_1, 0) f_{X_1}(x_1) f_{X|X_2}(x, 1) Pr(X_2 = 1) dx_1 dx \quad (C.1)$$

Appendix C.3. Ideal situation

In this case we generate our first artificial dataset as follows:

$\mathbf{X}_1 = \text{random_re-sampling}(\text{Throughput})$

$\mathbf{X}_2 = \text{random_re-sampling}(\text{DNS})$

$$\mathbf{X} \sim \begin{cases} \Gamma(k_1, \theta_1) + \sqrt{X_1} * \frac{\mu_1}{2} + \varepsilon : X_2 = 0 \\ \Gamma(k_2, \theta_2) + \sqrt{X_1} * \frac{\mu_2}{2} + \varepsilon : X_2 \neq 0 \end{cases}$$

$$\mathbf{Y} \sim \begin{cases} 10. \sqrt{5.X + 10.X_1} + \varepsilon : X_2 = 0 \\ 25. \sqrt{3.X + 6.X_1} + \varepsilon : X_2 \neq 0 \end{cases}$$

with ε representing an error term.

In this first case we chose the following values $\{k_1 = 5, \theta_1 = 1.0, k_2 = 2.0, \theta_2 = 2.0, \mu_1 = 5, \mu_2 = 8\}$. The resulting distributions are presented in the Figure C.14. To make sure that the parameters are correctly generated, we infer the corresponding causal model using the PC algorithm [Spirites and Glymour, 1991] with the independence test from [Zhang et al., 2012].

Notice that, for testing our method under the same main constraint we limit our sample size to 10000 (against 7500 in the real case scenario) and we keep the ratio between the number of samples where $X_2 = 0$ is observed and the number of samples where $X_2 = 1$ is observed equal to the ratio between the number of connections where the ISP DNS service was observed (80%) and the number of connections where the Google DNS service was observed (20%) in the real case scenario.

Appendix C.3.1. Prediction of Y given $X_2 = 0$ after intervention on X giving it the distribution of X given $X_2 = 1$

Using the Equation (C.1) we are able to compute the PDF $f(Y|X_2 = 0, do(X \sim X|X_2 = 1))$ and obtain the expected value $\mathbb{E}[Y|X_2 = 0, do(X \sim X|X_2 = 1)]$.

As mentioned in Appendix B.2.3, the choice of the number of bins and the threshold to decide or not to estimate the post atomic intervention PDF, should have an impact on the prediction accuracy. Consequently, in this ideal scenario where the two distributions of $X|X_2 = 0$ and $X|X_2 = 1$ have very similar domains, we vary these two parameters and study their impact on the prediction accuracy.

To obtain the distribution of $f_{Y|X_2=0, do(X \sim X|X_2=1)}$ we generate X and Y as following:

$$\begin{aligned} \mathbf{X} &\sim \Gamma(k_2, \theta_2) + \sqrt{X_1} * \frac{\mu_2}{2} + \varepsilon : X_2 = 0 \\ \mathbf{Y} &\sim 10. \sqrt{5.X + 10.X_1} + \varepsilon \end{aligned} \quad (\text{C.2})$$

We obtain an expected value of $\mathbb{E}[Y|X_2 = 0, do(X \sim X|X_2 = 1)]$ of 101.5.

We summarize the different results we obtained in Table C.4.

We can see that the use of a T-copula (*T-cop*) in the modeling of the multi dimensional PDF gives slightly better results, in terms of accuracy, than the modeling of multidimensional PDF with a Gaussian copula (*G-cop*). An interesting advantage of the T-copula comes from its ability to capture tail dependencies between the different components of the multivariate distribution.

If X is a d -dimensional random vector following a multivariate t -distribution with ν degrees of freedom, mean vector μ and a positive-definite dispersion Σ , denoted $X \sim t_d(\nu, \mu, \Sigma)$, [Demarta and McNeil, 2005] showed that the tail dependency coefficient is given by:

$$\lambda = 2t_{\nu+1}(\sqrt{\nu+1} \sqrt{1\rho} / \sqrt{1+\rho}) \quad (\text{C.3})$$

where ρ is the off-diagonal element of the correlation matrix implied by the normalization of the scatter matrix Σ .

This result shows that, by tuning the different parameters of the T-copula we can better capture the tail dependencies between the different components of the multi-variate distribution we want to

Table C.4: Effect of varying Δ_S on the prediction accuracy for the first artificial dataset for two different multi dimensional PDF modeling, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{\mathbf{E}}[Y X_2 = 0, do(X \sim X X_2 = 1)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	93.80	96.22	7.8%	5.5%	195	0 (0%)	5 (2.6%)
30	93.28	95.90	8.6%	6.0%	118	0 (0%)	3 (2.5%)
50	95.40	96.79	6.5%	5.2%	76	0 (0%)	2 (2.6%)
70	94.41	96.94	7.5%	5.0%	54	0 (0%)	1 (1.9%)
100	94.24	96.74	7.6%	5.2%	39	0 (0%)	0 (0%)

model. This property is really interesting when modeling communication networks performance, namely the throughput, as it is often the case that we find dependencies in the tail of parameter distributions, such as the one of delay or loss, with the throughput. As stated previously, the counter part of the T-copulae in practice comes from the higher sensitivity to data shortage.

Very likely due to the functions used for generating the artificial dataset and the use of a Gamma distribution, the T-copula is not always able to model the PDFs $f_{Y|X,X_1,X_2}(y, x, x_1, 0)$ that are necessary to compute the post intervention PDF of Y in Equation (C.1).

It is also important to notice that the choice of using a Gaussian copula was motivated by the prediction of an intervention in the opposite case ($f_{Y|do(X \sim X|X_2=0),X_2=1}(y)$), see next section. In addition, we generated the artificial dataset in order to have the same domains for $X|X_2 = 0$ and $X|X_2 = 1$ and do not observe the same tail dependence as we would have for the throughput in the real case. From this perspective, the usage of T-copula is not fully justified and the usage of a Gaussian copula, for this particular dataset, should still be preferred.

Appendix C.3.2. Prediction of Y given $X_2 = 1$ after an intervention on X giving it the distribution of X given $X_2 = 0$

Without repeating the explanations given in the previous section, we use Equation (C.1) to predict the expected value of $\mathbf{E}[Y|X_2 = 1, do(X \sim X|X_2 = 0)]$.

The results are presented in Table C.5 in which we compared the expected value obtained using Equation (C.1) and G-copulae or T-copulae with the value obtained from the definitions of the parameters, Equation (C.2). We also study the impact of the number of samples used for estimating the atomic post-intervention distribution, Δ_S . Note that by increasing the minimum number of samples required to estimate a given atomic intervention we decrease the number of atomic interventions used for predicting the distribution of $f(Y|X_2 = B, do(X \sim X|X_2 = \neg B))$ (with $B \in 0, 1$), represented by the parameter #AI.

Several important remarks can be made from the results we obtain:

- The usage of a T-copula for modeling the PDF of $f_{Y|X,X_1,X_2}(y, x, x_1, 0)$ for different values of X fails more than 50% of the times.
- The utilization of wider bins, gathering more data to approximate $X = x$, does not improve the the success rate of the predictions of the PDFs post-intervention.

Table C.5: Effect of varying Δ_S on the prediction accuracy for the first artificial dataset for two different multi dimensional PDF models, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{E}[Y X_2 = 1, do(X \sim X X_2 = 0)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	173.9	N.A.	2.5%	N.A.	43	0 (0%)	23 (53%)
30	173.9	N.A.	2.5%	N.A.	32	0 (0%)	20 (63%)
50	175.7	N.A.	1.4%	N.A.	17	0 (0%)	13 (76%)
70	175.7	N.A.	1.6%	N.A.	13	0 (0%)	10 (77%)
100	174.0	N.A.	2.3%	N.A.	9	0 (0%)	8 (89%)

- The usage of a G-copula gives better results (in terms of prediction accuracy) than the previous predictions of interventions.

Remarks. It should be noticed that, despite the apparent symmetry between the prediction of Y conditionally to $X_2 = 1$ if we perform an intervention on X where we fix its distribution to the one of $X|X_2 = 0$ and the one of the prediction of the value of Y conditionally to $X_2 = 0$ if we perform an intervention on X where we fix its distribution to the one of $X|X_2 = 1$, the problematic is different. From an external point of view, looking at the completion time and success rate, Gaussian copulae are less data demanding. We generated X_2 by randomly re sampling the DNS parameters from our real dataset. Doing so, we have 80% of the samples where $X_2 = 0$ is observed against 20% where $X_2 = 1$ is observed.

This second scenario shows the sensitivity of our approach to resource limitation. Even in this “optimal scenario” where both conditional PDF of $X|X_2 = 0$ and $X|X_2 = 1$ have the same domains, the shortage of data for the second conditional PDF prevents us from using a model which could be more accurate (T-copula). This can be seen when comparing the predictions made in this section with the ones made when we had more data (Appendix C.3.1) where the T-copula model gave more accurate predictions.

Appendix C.3.3. Concluding remarks

In this section we presented the optimal case where we have both conditional PDFs having almost perfectly overlapping domains. We tried to predict interventions where we condition on one value of the categorical parameter, X_2 , and intervene on another parameter, X , and fix its distribution to follow the conditional distribution corresponding to the complementary value of the categorical parameter. The results of this study represent two important findings:

- Our method works and makes an accurate prediction (error < 3%) when enough data is present
- The use of T-copulae better captures the multi dimensional PDFs dependences but fails as soon as the data becomes scarcer, while G-copulae still give an accurate prediction.

Again, these conclusion are made using an artificial dataset (see definition in Appendix C.3) where we could not faithfully mimic the tail dependencies of the parents of the throughput that

Table C.6: Effect of varying Δ_S on the prediction accuracy for the second artificial dataset for two different multi dimensional PDF modeling, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{E}[Y X_2 = 0, do(X \sim X X_2 = 1)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	96.73	97.12	4.7%	4.3%	195	0 (0%)	5 (2.6%)
40	95.97	96.13	5.5%	5.3%	97	0 (0%)	2 (2.1%)
60	96.39	96.69	5.1%	4.8%	60	0 (0%)	1 (1.7%)
80	95.16	95.87	6.3%	5.6%	48	0 (0%)	1 (2.1%)
100	94.24	95.05	7.2%	6.4%	39	0 (0%)	0 (0.0%)

we observed in our real case scenario nor the variability of the observed values. These choices are inherent to the design of an artificial dataset and should not be seen as a limitation of the presented method.

Appendix C.4. Removing samples for $X | X_2 = 1$ outside of the zone of the concentration of the distribution $f_{X|X_2=0}$

To estimate the impact of the absence of some values in both domains of the conditional PDFs of $X|X_2$, we remove some samples from the dataset where $X_2 = 1$. For this second artificial dataset, we do not remove samples in the domain where the distribution $f_{X|X_2=0}$ takes its biggest values. Figure C.16 represents the original distributions of $X|X_2$ (our first artificial dataset) and Figure C.17 presents the resulting distributions of $X|X_2$ after removing samples from the distribution of $X|X_2 = 1$ (we remove the samples for $X_2 = 1$ where $X \in [0, 5] \cup [15, 20]$).

As in Appendix C.3.1 and Appendix C.3.2, we estimate the expected value of Y conditionally to $X_2 = 0$ when we intervene on X and we fix its distribution to the one of $X \sim X|X_2 = 1$. We also estimate the expected value of Y conditionally to $X_2 = 1$ when we intervene on X and we fix its distribution to the one $X \sim X|X_2 = 0$

We first estimate the effect of intervening on X to fix its distribution to the one of $X \sim X|X_2 = 1$. We can see from Table C.6 that the conclusions drawn previously are still valid when we remove samples from the distribution of $X|X_2 = 1$. We can make the following remarks:

- The precision decreases with the increase of Δ_S and the decrease of the number of atomic interventions,
- The number of failures of the multidimensional PDF modeling using a T-copula is sensibly similar to the number of failures that were observed for the first dataset.

When looking at the opposite case (conditioning on $X_2 = 1$ and giving to X the distribution of $X \sim X|X_2 = 0$) we can observe that the modeling of the conditional PDFs using a T-copula fails often and prevents the estimation of the post-intervention PDF, see Table C.7. The estimation of conditional PDF when we use G-copulae gives very good results.

Table C.7: Effect of varying Δ_S on the prediction accuracy for the second artificial dataset for two different multi dimensional PDF modeling, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{E}[Y X_2 = 1, do(X \sim X X_2 = 0)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	176.63	N.A.	0.6%	N.A.	40	0 (0%)	23 (57.5%)
40	176.95	N.A.	0.4%	N.A.	22	0 (0%)	14 (63.6%)
60	178.21	N.A.	0.3%	N.A.	12	0 (0%)	10 (83.3%)
80	179.73	N.A.	1.2%	N.A.	10	0 (0%)	9 (90.0%)
100	177.43	N.A.	0.1%	N.A.	9	0 (0%)	8 (88.9%)

Appendix C.4.1. Concluding remarks

In this scenario we removed samples from our initial dataset to study the impact of data shortage on our method precision. In this case, we removed samples from the conditional distribution $X_2 = 1$ corresponding to x values where the conditional probability distribution $f_{X|X_2=0}$ takes small values. We predicted the effect on $Y|X_2 = b$ when intervening on $X|X_2 = b$ and enforcing this intervention with probability $f_{X|X_2=\bar{b}}$, where b can take the value 0 or 1 and \bar{b} is b complementary.

The prediction of this intervention consists in *i*) Predicting the effect, on Y , of the atomic intervention $f_{Y|do(X=x), X_2=b}$ *ii*) Assign to each atomic intervention a probability being $f_{X|X_2=\bar{b}}$.

$$f(Y|do(X \sim X|do(X_2 = \bar{b}), X_2 = b) = \int_X f(Y|do(X = x), X_2 = b)f(X = x|X_2 = \bar{b})dx \quad (C.4)$$

Therefore, the prediction is more complex in the case where we give to $f_{Y|do(X=x), X_2=b}(y, x)$ a probability $f_{X|X_2=\bar{b}}(x) > 0$ but no value is actually observed for $X = x|X_2 = b$.

This appears in the second case where we want to predict the distribution $f_{Y|X_2=1, do(X \sim X|X_2=0)}$. As there are x values for which $X|X_2 = 1$ is not observed, we cannot enforce the conditional distribution $f_{Y|do(X=x)|X_2=1}(y, x)$ with the distribution $f_{X|X_2=0}(x)$.

Nevertheless, because we removed observation corresponding to x values where $f_{X|X_2=0}(x)$ takes small values, the impossibility to compute the result of Equation (C.4) for some x values has a relatively small impact on the estimation of the overall post-intervention distribution and our method keeps performing well.

The next section presents a more complex scenario where we remove samples in intervals where the PDF of $X|X_2 = 0$ takes important values (values of x where $X|X_2 = 0$ has a high probability).

Appendix C.5. Removing samples $X | X_2 = 1$ at the zone of the concentration of the distribution $f_{X|X_2=0}$

We generate a third artificial dataset with a distribution $X|X_2$ as represented in Figure C.18. For this new dataset, we remove samples from the domain where the distribution $f_{X|X_2=0}$ takes its biggest values. The results for the two interventions are presented in Tables C.8- C.9 with:

- The prediction of expected value of Y after intervention corresponding to PDF of Y : $f_{Y|X_2=0, do(X \sim X|X_2=1)}(y)$ in Table C.8

Table C.8: Effect of varying Δ_S on the prediction accuracy for the third artificial dataset for two different multi dimensional PDF modeling, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{E}[Y X_2 = 0, do(X \sim X X_2 = 1)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	66.79	80.81	34.2%	20.4%	195	0 (0%)	5 (2.6%)
40	76.15	85.80	25.0%	15.5%	97	0 (0%)	2 (2.1%)
60	67.07	82.14	33.9%	19.1%	60	0 (0%)	1 (1.7%)
80	69.55	83.08	31.5%	18.2%	48	0 (0%)	1 (2.1%)
100	N.A.	N.A.	N.A.	N.A.	39	0 (0%)	0 (0.0%)

Table C.9: Effect of varying Δ_S on the prediction accuracy for the third artificial dataset for two different multi dimensional PDF modeling, using a T-copula (*T-cop*) and Gaussian copula (*G-cop*). #AI stands for the number of Atomic Interventions

Δ_S	$\hat{E}[Y X_2 = 1, do(X \sim X X_2 = 0)]$		Error		#A.I.	#Failures	
	G-cop	T-cop	G-cop	T-cop		G-cop	T-cop
20	153.68	N.A.	13.5%	N.A.	12	0 (0%)	7 (58.3%)
40	154.50	140.18.	13.0%	21.1%	6	0 (0%)	3 (50.0%)
60	157.90	N.A.	11.1%	N.A.	5	0 (0%)	3 (60.0%)
80	163.52	N.A.	8.0%	N.A.	3	0 (0%)	2 (66.7%)
100	157.25	N.A.	11.5%	N.A.	2	0 (0%)	1 (50%)

- Prediction of expected value of Y after intervention corresponding to PDF of Y: $f_{Y|X_2=1, do(X \sim X|X_2=0)}(y)$ in Table C.9

Appendix C.5.1. Concluding remarks

We can observe the accuracy of the prediction of the expected value of Y conditionally to X_2 when setting X distribution to the one of $X|X_2 = 1$ is highly impacted by the absence of samples in the zone where the PDF $f_{X|X_2=0}$ takes important values. The usage of a T-copula for conditional PDF gives slightly better results than the predictions based on the usage of a Gaussian copula. However with 20% error rate, we cannot use our method any more.

For the prediction of the expected value of Y conditionally to $X_2 = 1$ when intervening on X and fixing its distribution to $f_{X|X_2=0}(x)$, we penalize the usage of T-copula for modeling conditional PDFs more than the G-copula. Gaussian copulae seem to require less data to estimate a parameterization offering an acceptable modeling of the dependencies between the marginals of the multivariate distribution we need to estimate our model.

Appendix C.6. Conclusion

The questions we wanted to answer were:

- How to choose the copula family that will model the dependencies between the different marginals ?
- In which proportion the absence of values observed for both conditional distributions impacts the prediction accuracy and how to improve the accuracy in this case ?
- Should we favor the number of atomic predictions from which the final distribution is estimated or the quality of the atomic intervention predictions by increasing the number of samples from which these atomic predictions are computed ?

The first question can be answered by “*data dictates the choice*”. In our case, as we are working with a limited amount of data, Gaussian copulae are used to capture the dependencies between the marginals of the multivariate distributions we need to estimate to predict the effect of interventions on parameters impacted by the DNS parameter.

The absence of values observed for both conditional distributions can be overcome using Gaussian copulae if we observe enough values in the zones corresponding to high probabilities for the conditional distributions. T-copulae suffer more from data shortage than G-copula but if the observations of the conditional distributions are too sparse then both models become inaccurate and cannot be used.

The number of atomic interventions should be preferred to the number of samples used for estimating a given intervention but a minimum number of samples should be present (≥ 30).

Given these conclusions, we have defined and parameterized the methods that can be used to study the impact of DNS on CDN performance.

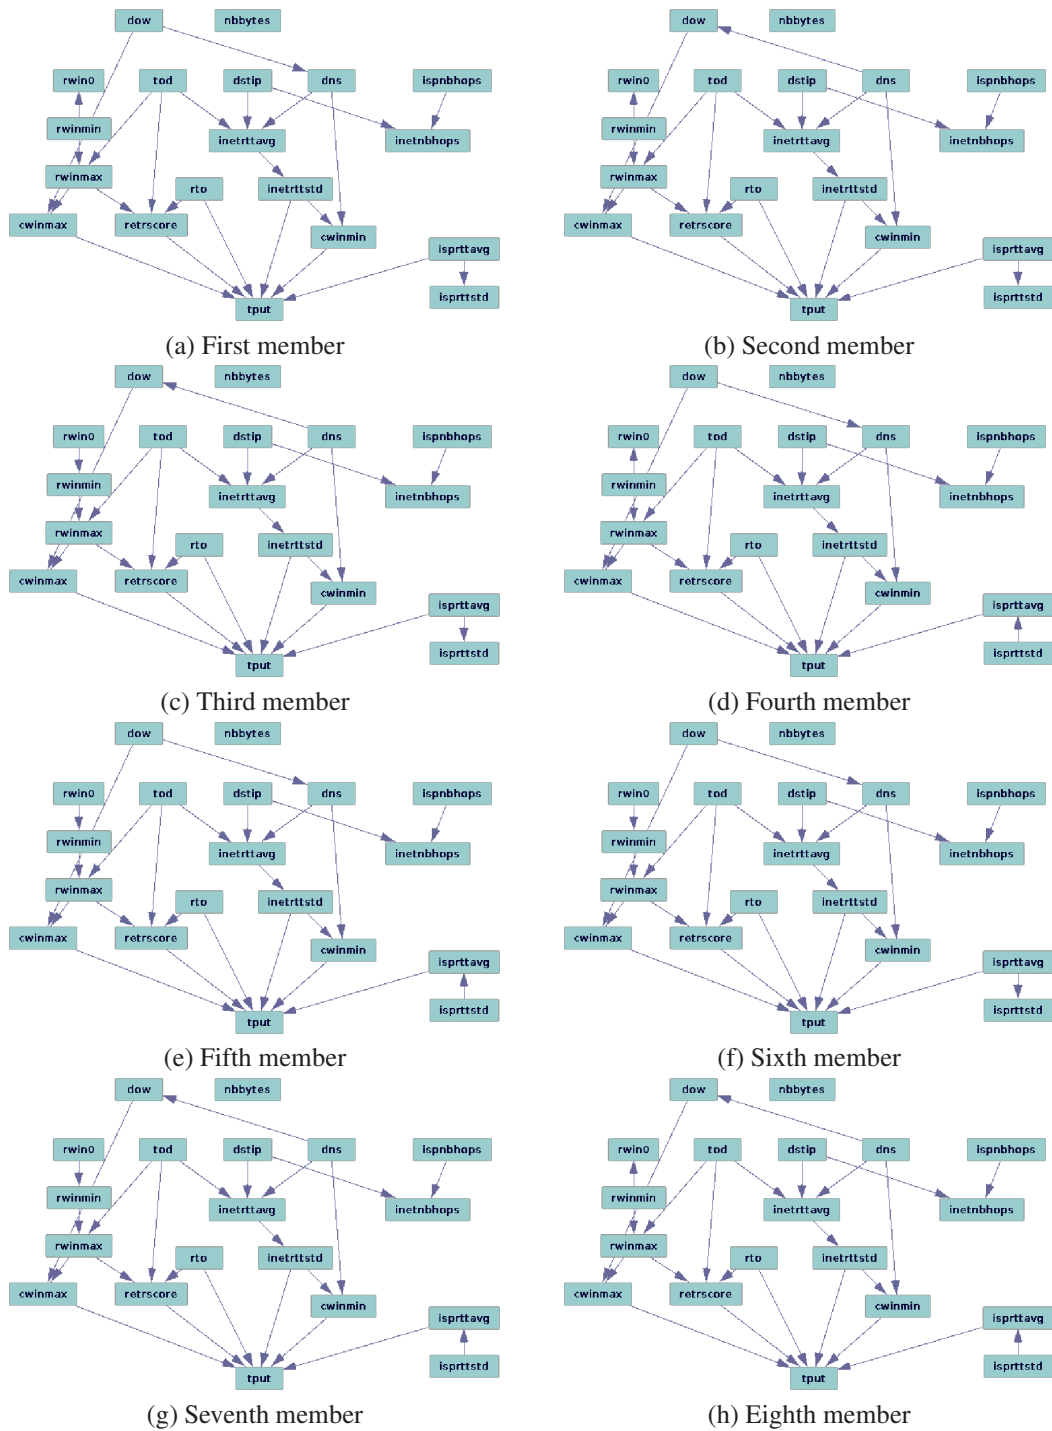


Figure A.11: The eight members of the Markov Equivalence Class corresponding to the set of independences that were detected from our observations. These graphs were obtained using the Tetrad software [Spirtes et al., 2001]




Figure B.12: Illustration of the different subgraphs $G_{\underline{X}}$ and $G_{\overline{X}}$ for a Bayesian network representing the causal model of a four parameter system $\{W, X, Y, Z\}$




Figure C.13: Artificial dataset dependencies




Figure C.14: Distribution of the different parameters for both values of X_2




Figure C.15: Causal Model of the first dataset




Figure C.16: Conditional probability density function of X conditional on X_2 in the original dataset




Figure C.17: Conditional probability density function of X conditional on X_2 after removing samples from $X_2 = 1$ in the domain of $X|X_2 = 0$ where the distribution $f_{X|X_2=0}$ is not taking high values, second artificial dataset




Figure C.18: Conditional probability density function of X conditional on X_2 after removing samples from $X_2 = 1$ in the domain of $X|X_2 = 0$ where the distribution $f_{X|X_2=0}$ is concentrated, third artificial dataset