A COMPREHENSIVE METHOD FOR ASSESSING THE RESILIENCE OF POWER NETWORKS IN THE FACE OF AN INTELLIGENT ADVERSARY

by

Sinan Tas

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Industrial Engineering)

at the

UNIVERSITY OF WISCONSIN-MADISON

2012

Date of final oral examination: 08/03/12

The dissertation is approved by the following members of the Final Oral Committee:

Vicki M. Bier Professor, Professor, Industrial and Systems Engineering
Ian Dobson Professor, Professor, Electrical and Computer Engineering, Iowa State University
Raj Veeramani Professor, Professor, Industrial and Systems Engineering
Jeffrey Linderoth Professor, Professor, Industrial and Systems Engineering
Oguzhan Alagoz Professor, Associate Professor, Industrial and Systems Engineering

To my parents, and my wonderful kids *Derin* and *Aiden*

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor, Professor Vicki M. Bier, for all of her guidance, support, and wisdom. Her critical thinking skills and insightfulness are remarkable, and I feel fortunate to have been under her guidance. This was an exceptional journey, and I am very grateful to have her next to me in every stage of my PhD studies. She is an amazing researcher, teacher, and mentor, and is my role model in my career.

I would also like to thank Dr. James Peerenboom from Argonne National Laboratory. His efforts and support were instrumental in my studies. My collaboration with him and researchers in his center helped me greatly in identifying an interesting and practical research question. I am looking forward to continuing my collaboration with everyone at Argonne.

I would like to thank my dissertation committee members, specifically professors Ian Dobson and Jeffrey Linderoth for their detailed guidance in my dissertation, professor Raj Veeramani for his insightful questions, and professor Oguzhan Alagoz for his valuable inputs in my job search.

Moreover, I would like to thank the wonderfully supportive family I have found at the University of Wisconsin-Madison. I feel truly honored to be a part of this greatness. In my time here, I have had the great pleasure of meeting many insightful people that have helped my greatly on my journey. I would especially like to thank (in alphabetical order): Onur Asan, Turgay Ayer, Mehmet Ayvaci, Nataliya Batina, Merve Bodur, Gizem Cavuslar, Mucahit Cevik, Weiwei Chen, James Codella, Vikas Dawar, Annie Duchek, Safa Erenay, Mehmet Ertem, Gregory Hammond,

Mahdi Hamzeei, Wen-Chieh Hu, Taher Jamshidi, Burcu Kabatas, Mustafa Rasim Kilinc, Janghoon Kim, Utkan Kolat, Fuat Kosanoglu, Shuji Liu, Wairimu Magua, Mahdi Namazifar, Mustafa Ozkaynak, Edmond Ramly, Debjit Roy, Lei Ruan, Yongjia Song, Lisa Tang, Rupa Valdez, Chen Wang and Jonathan Welburn. There have been so many people who have supported me along the way, and I would like to give my greatest gratitude to everyone that has helped me, although I have only listed a few names here.

Last, but not the least, I would like to express my deepest appreciation to my family. First of all, my mom and dad, Fatma and Yahya Tas, have showered such unconditional support for everything I have ever done. This thesis, in many ways, is the fruit of their labor. They have always been great parents to me, and I will strive to follow in their footsteps. Also, thank to my dearest sisters, Elif and Gulistan, I miss them so much. Along my journey, I found the greatest joy, and at times distraction, in my two beautiful children, Derin and Aiden. They are my biggest motivation, and the funnest two little people I know. More so, I would like to thank Laura Daly for all of her encouragement in this rather exhausting process. I am also grateful to my host parents, Peggy and Bruce Fealk for their unconditional love, support, and belief in me. They made me feel so welcome, and I will always be thankful for everything they have done. Lastly, special thanks to Pamela Moreno. She was a great supporter in my hardest times. She spent so many hours with me looking over my work, and giving me very insightful feedback. I owe a great deal to her to finishing it on time.

I am very honored and excited to be part of this distinct community of researchers and educators.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

After September 11, 2001, numerous methods of vulnerability analysis have been developed to help the owners and operators of infrastructure systems protect such systems against possible terrorist attacks. However, hardening a significant fraction of a large, complex infrastructure system is typically not cost-effective, and may even be infeasible. There is therefore an urgent need for methods of vulnerability analysis that take into account the ability of the system to respond to and recover from an attack (or, conversely, the vulnerability of the system to cascading failures and/or long restoration times).

In particular, with cascading failures, even small attacks can have a large impact. Cascading failures have historically been considered a major unsolved problem for complex networks such as electricity systems, but recent developments in probabilistic analysis of cascading failure are making it possible to take cascading failures into account in methods of vulnerability assessment.

Moreover, methods of vulnerability analysis can also be designed to highlight those vulnerabilities that are likely to lead to disproportionately long restoration times. This approach provides a more comprehensive method for vulnerability analysis of electric power systems and potentially other capacity-constrained complex networks, which enables us to analyze how the attacker can exploit different weaknesses within the system, and as a result, how we can protect the system against such weaknesses.

Our game-theoretic model can be used to measure the effectiveness of different investment types against intelligent attacks on power networks. Specifically, our model provides a tool to simulate power flows within the network, the impact of a greedy attack strategy, the possibility of cascading failure, and the restoration of the system back to its normal operating conditions. The model will enable us to analyze different attack scenarios, such as adversaries who might seek to cause cascading failures or long restoration times. Our model can also be used to compare different types of investments to make system more resilient, such as hardening components, increasing the capacities of critical transmission lines, adding new transmission lines, and decreasing the restoration times of specific component types.

1. Introduction

Infrastructure systems are critical to the economy, security, safety and quality of life of a society. Ensuring the security of such systems can be quite challenging, especially in the face of intentional threats.

September 11, 2001 was a turning point in the United States in recognition of the vulnerability of infrastructure systems, not only to natural disasters or technological failures, but more so to intelligent attacks, such as acts of terrorism, sabotage, or war. Infrastructure systems were generally not designed to resist malicious attacks (Apostolakis and Lemon, 2005), and protecting extensive infrastructure networks can be extremely costly or even impossible. Furthermore, an intelligent adversary may exploit the weaknesses of such systems (e.g., bottlenecks or capacity constraints in the network, aging components, etc.) so that the network will poorly respond to even one single attack.

Electric power networks are especially complex and important systems (Amin, 2003; Dueñas-Osorio and Vemuru, 2009), with large numbers of interacting components, complex responses to disturbances, and the possibility of high impact from even a small disruption. The complexity of these systems makes it nearly impossible to model all their interactions in detail; as a result, cascading failures have proven to be difficult to predict.

The objective of most past vulnerability methods on electric power networks was to identify the most critical components to prioritize for protection (i.e., hardening). However, protection of massive and geographically dispersed systems, such as electric power networks, may not be effective or cost-effective. Moreover, most methods of vulnerability analysis do not take into account the possibility of cascading failure, or the fact that some components (e.g., transformers) may have extremely long restoration times. Therefore, we develop a new vulnerability analysis method that makes it possible to study the resilience of a system, by addressing its ability to respond to, withstand, and recover from an attack. In particular, we develop a method for assessing the vulnerability of electric power systems to intentional threats, taking into account the vulnerability of the system to cascading failure in probabilistic manner, and the restoration times of components in the system. Moreover, our method is simple enough to be usable in practice, and capable of assessing the effectiveness and cost-effectiveness of possible defensive investments- not only of target hardening but also of other investments, such as increasing the capacity of the network (to reduce the potential for cascading failure), or decreasing restoration times (e.g., through stockpiling of spare components).

We also hypothesize that the approach outlined in this thesis can be applied not only to electric power networks, but also to any highly capacity-constrained networks, in which failure of a heavily loaded line or component may cause intolerable increases of load in other parts of the network, eventually leading to cascading failure. Examples include transportation systems, structures (which can be viewed as networks of structural members), water systems, etc. However, such applications are beyond the scope of this research, and would need to be explored in future work

In Chapter 2, we define resilience, and show its relationship with defensive measures. In Chapter 3, we review the literature on vulnerability-analysis methods for electric power networks, including models of cascading failure, and restoration times. In Chapter 4, we introduce the original model, and develop a methodology to include attacking any components and to model cascading failure and restoration times, and how we can use resilience to measure effectiveness of various defensive investments (including protection through hardening, increasing robustness through capacity investment or adding new transmission lines, and increasing the ability of the system to recover by decreasing restoration times of some critical components such as transformers). In Chapter 5, we show the result of extending the model to nodes as well as arcs (so that the attacker can attack components other than transmission lines). In Chapter 6, we show the results for modeling cascading failures. In Chapter 7, we show the results for modeling restoration model. Finally, in Chapter 8, we conclude with how we determine if our model is valid and useful, highlight some of the important findings, and discuss some possible future work.

2. Resilience in Electric Power Systems

2.1. Definition of Resilience

The U.S. Department of Homeland Security (DHS) defines resilience as the “ability to adapt to changing conditions and prepare for, withstand, and rapidly recover from disruption” (DHS, 2010). This comprehensive definition embraces all the efforts in a timeline from preparation to recovery as part of resilience– i.e., “to resist, tolerate, absorb, recover from, prepare for, or adapt to an adverse occurrence.” The National Infrastructure Advisory Council (NIAC) focuses on the outcome of a resilient system, and defines resilience as the ability to “reduce the magnitude and/or duration” of disruptive events. The effectiveness of a resilient infrastructure or enterprise depends upon its ability to *anticipate*, *absorb*, *adapt to*, and/or *rapidly recover from* a potentially disruptive event)” (NIAC, 2009). NIAC also distinguishes three key features of a resilient system: robustness; resourcefulness; and rapid recovery.

In the literature, however, resilience is often used with a narrower meaning. For example, Haimes et al. (1998) considers resilience one of four strategies for hardening a system, together with security, redundancy, and robustness. He defines security as measures that limit entrance to a system, such as surveillance, fences, and guards. Redundancy is defined as the ability of a system to replace the function of a failed component or mechanism. Robustness describes how a system can continue its designed function despite an error or failure. Finally, Haimes et al. (1998) defines resilience as the ability of system to return to its optimal condition in a short period of time, given that an adverse event has taken place. This focus on recovery time is narrower than the other definitions discussed above, and in particular ignores the magnitude or severity of the adverse event, as measured by cost, level of degradation, etc.

Holmgren (2007) likewise distinguishes robustness and resilience, using robustness to imply that the system will remain (nearly) unchanged even in the face of disruption, and resilience to describe the ability of the system to return to a stable condition after a disruption. Hansson and Helgesson (2003) note that robustness can be considered a special case of resilience, in which the recovery time is zero.

Rose and Liao (2005) define resilience as “the inherent ability and adaptive responses of systems that enable them to avoid potential losses.” Note that this definition does not explicitly include the ability of a system to recover after losses have already occurred, although many measures for avoiding or minimizing losses can still be taken after an adverse event has occurred.

Bruneau et al. (2003) define resilience in terms of three stages: the ability of a system to reduce the probability of an adverse event, to absorb the shock if the adverse event occurs, and to quickly re-establish normal operating conditions. According to Bruneau et al., resilience thus encompasses the four characteristics of robustness, redundancy, resourcefulness, and rapidity. This is parallel to the DHS definition. By comparison, note that NIAC does not consider redundancy at all, and Haimes et al. (1998) consider redundancy and robustness to be distinct from resilience. Bruneau et al. (2003) view the “desired ends” of robustness and rapidity as key aspects of resilience.

According to Holling (1973), resilience “determines the persistence of relationships within a system” and is “a measure of the ability of these systems to absorb changes of state variables, driving variables, and parameters, and still persist.” In other words, Holling considers resilience to be an inherent feature of the system that reflects its ability to manage changes. On the other hand, Wildavsky (1988) considers resilience to be an acquired feature rather than an inherent one, defining resilience as “the capacity to cope with unanticipated dangers after they have become manifest, learning to bounce back.”

Hollnagel (2006) defines resilience as “the intrinsic ability of an organisation (system) to *maintain or regain a dynamically stable state,* which allows it to continue operations after a major mishap and/or in the presence of a continuous stress.” Similarly, Woods (2006) defines resilience as “a system capability to create foresight, to recognize, to anticipate, and to defend against the changing shape of risk before adverse consequences occur.” According to Woods (2006), the concept of resilience should include some description of *how* a system adapts to disturbances, not just whether it adapts. Some aspects of this adaptation include *buffering capacity* (the sizes and types of disruptions that a system can absorb), *flexibility versus stiffness* (the ability of a system to restructure itself in response to external changes or pressures), *margin* (how close the system is operating relative to one or more performance boundaries), and *tolerance* (how a system behaves near a boundary – whether the system degrades gracefully as stress increases, or collapses quickly when pressures exceed its adaptive capacity).

Tierney and Bruneau (2007) define resilience as the ability of a system to “respond to” and “recover from” a major event. Similarly, Haimes (2009) defines resilience as “the ability of the system to withstand a major disruption within acceptable degradation parameters and to recover within an acceptable time and composite costs and risk." According to Haimes, vulnerability has to do with protection of a system, while resilience addresses the system's recovery; in other words, vulnerability measures a system's capability to resist a threat, while resilience represents the capability of the system to recover within an acceptable time. Hence, according to Haimes (2009), resilience is both threat- and time-dependent.

Finally, Vugrin et al. (2010) propose the following definition: "Given the occurrence of a particular disruptive event (or set of events), the resilience of a system to that event (or events) is the ability to efficiently reduce both the magnitude and duration of the deviation from targeted system performance levels." Moreover, the authors define the resilience of a system as depending on the system's "absorptive capacity", "adaptive capacity", and "restorative capacity." For some other definitions of resilience, see Appendix A.

One commonality in the literature is that resilience is considered meaningful only in the context of disruption or change. In other words, a resilient system is expected to perform well in the face of undesirable conditions. Using the comprehensive approach promoted by DHS, the term resilience can refer to any ability of the system to maintain or improve its performance in the face of disruption. Another commonality is that resilience generally refers to how the system responds to change over time. As noted above, some authors also consider robustness as a special case of resilience (Hansson and Helgesson, 2003).

Definition: Resilience is performance of a system over time in response to adverse change. Let R be the overall performance of the system to adverse change, T be the time period during which the system is needed, and $t = 0$ be the initiation time of the change, where $t \in T$. Then, $R(t)$ describes the resilience of the system at time t. A highly resilient system returns to desirable levels of performance quickly, so that $R(t)$ rapidly approaches 100% even for small values of t.

The disruption to which the system must respond can be a function of threat type, or the result of a specific scenario. Depending on the threat or scenario, the system may exhibit different reactions. Especially for complex dynamic systems, like electric power networks, with interdependencies to other systems, it may be quite challenging to estimate or predict the reaction prior to an event.

Note that for some disruptions, it may not even be possible for the system to go back to its normal operating conditions, in which case resilience would be only partial, even after a long period of time. Also, the resilience of the system is not necessarily proportional to the area under the function $R(t)$, because the importance of the system may vary over time. In other words, going without the functions provided by a system may become either more difficult over time, or easier (due to user adaptation).

2.2. Resilience and Time

A resilient system has the inherent ability to respond appropriately to every stage of an adverse event, including preparedness before the event, robustness as the event unfolds, restoration immediately after the event, and eventually complete recovery. Looking at the entire spectrum in this way is likely to be more practical, realistic, and cost-effective than focusing only on preventing electricity blackouts.

According to Talukdar et al. (2003), most policy discussions with regard to blackouts focus on preventing blackouts or reducing their frequency. Measures for doing so include "increasing transmission capacity, improving regulations and coordination, training for human operators, better automatic control systems, more data collection, more data processing, load management, and more programs to promote conservation." However, they argue that framing the problem only in terms of prevention is incomplete, and also difficult to solve.

Resilience of a system may take different forms before, during, and after an event; see Figure 2.1.

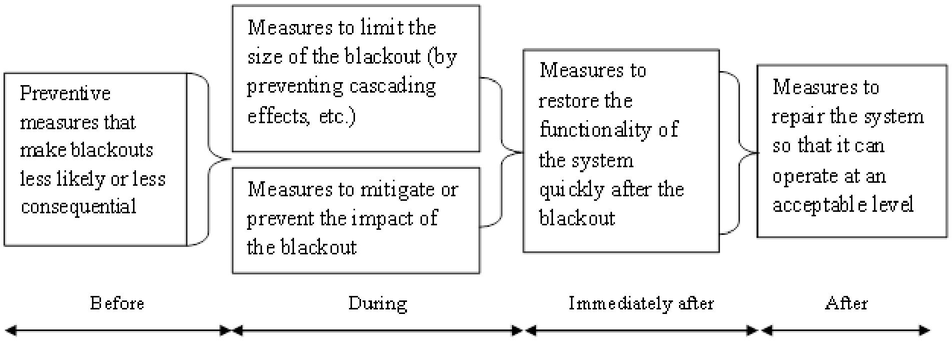
Preventive measures that make blackouts less likely or less consequential
Measures to limit the size of the blackout (by preventing cascading effects, etc.)
Measures to mitigate or prevent the impact of the blackout
Measures to restore the functionality of the system quickly after the blackout
Measures to repair the system so that it can operate at an acceptable level
Before
During
Immediately after
After

Figure 2.1-Resilience over time

Before an event, we can take measures to make the event less likely or less consequential. Resilience involves any inherent ability of system to prevent a disruption. Note that there are also measures that are independent of the system. During an event, the robustness of the system has two aspects: limiting the size of the blackout (such as by preventing cascading effects); and mitigating or preventing the impact of the blackout. The first measure can be simply simulated by modeling system's reaction including cascading failure. The second measure shows the importance of interdependencies among components. These interdependencies will be discussed in detail in the following chapters. After a blackout, some components of the system will typically have stopped working. Emergency restoration (energizing of substations, transformers, etc.) involves careful coordination of personnel and equipment, so restoration of the entire system may take a long time. Finally, some catastrophic events may damage components that would need to be replaced (e.g., by spare transformers) or repaired.

Given the nature and objectives of the system, performance can be measured in different ways, such as production capacity, satisfied customer demand, etc., or combinations of these. By showing a system's response to change as a function of time, we can illustrate how long the system resists before the system degrades (robustness), how much the system performance degrades initially, how well the system recovers, and how long it takes to return to normal. For example, see Figure 2.2.

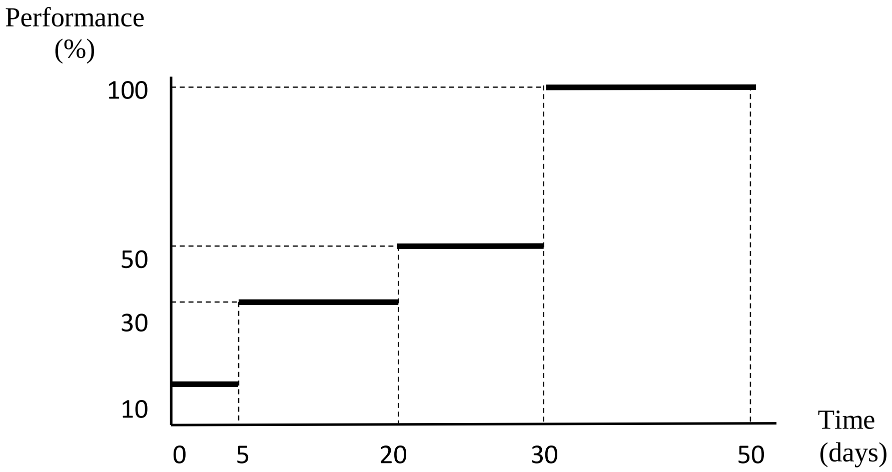

Figure 2.2-Resilience as performance of a system

This figure shows a system's performance (as a percentage of normal or intended performance) as a function of given time. This system experiences an immediate degradation, followed by several stages of improvement. Note, however, that performance may not necessarily be monotonically increasing after an adverse event. Some possible reasons for decreased performance could be cascading effects, human error, etc.

In some systems, performance can also increase continuously rather than discretely over time. However, it may be desirable to simplify such continuous functions; for example, in Figure 2.3, Tierney and Bruneau (2007) describe the "resilience triangle" as a simple way to capture the effectiveness of restoration and recovery efforts. The triangle shows both how much functionality is lost as a result of a catastrophic event after the initial impact and cascading failure, and how long it takes for the system to go back to pre-catastrophe levels of performance.

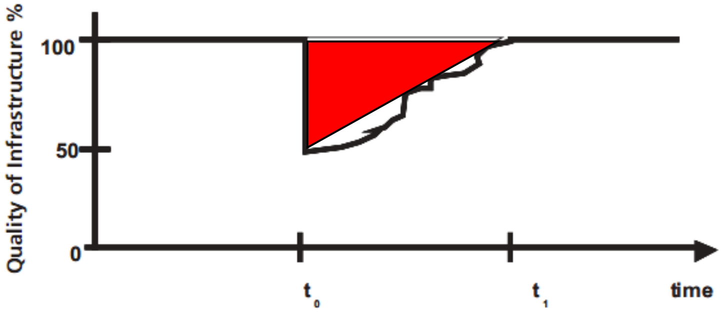

Figure 2.3-Resilience triangle (Tierney and Bruneau, 2007)

2.3. Relationship between Resilience and Protection, Robustness, and Recovery

We define three types of defensive investments: protection; robustness; and recovery. Protection is defined as reducing the likelihood of damage, robustness reduces the extent of damage, and recovery reduces the duration of damage. Resilience encompasses the ability of the system to protect itself from, respond robustly to, and recover from a disruption. In this study, we view protection, robustness, and rapid recovery as alternative means for achieving a more resilient system. Figure 2.4 illustrates the three types of defensive investments.

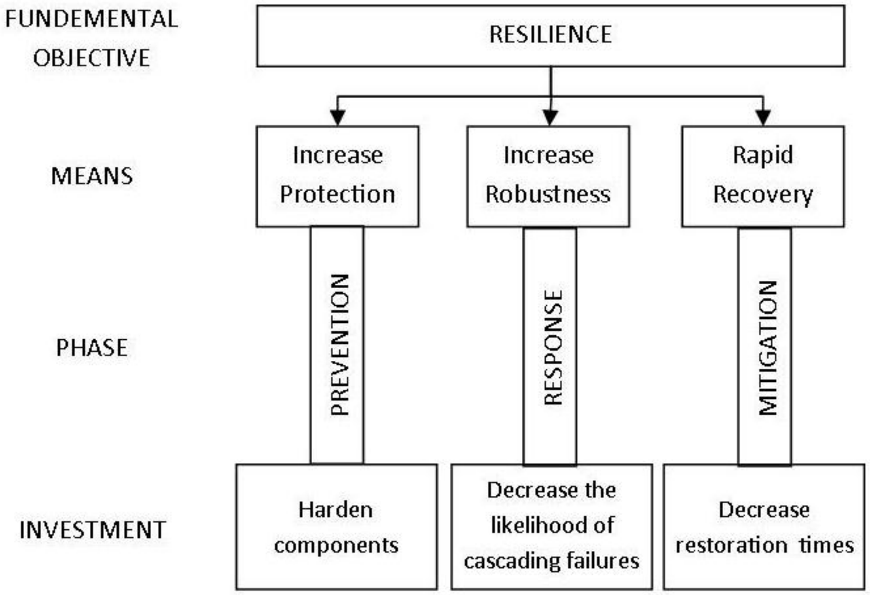

Figure 2.4-Categorization of defensive investments

Specifically, for protection, we consider hardening components. For robustness, we consider increasing the capacities of some components (to decrease the likelihood of cascading failure). Finally, for rapid recovery, we consider decreasing the restoration times of some components.

Figure 2.5 shows the relationship between resilience and defensive investments to make the system more resilient before, during, and after an attack. (Note that some investments may contribute to multiple goals.)

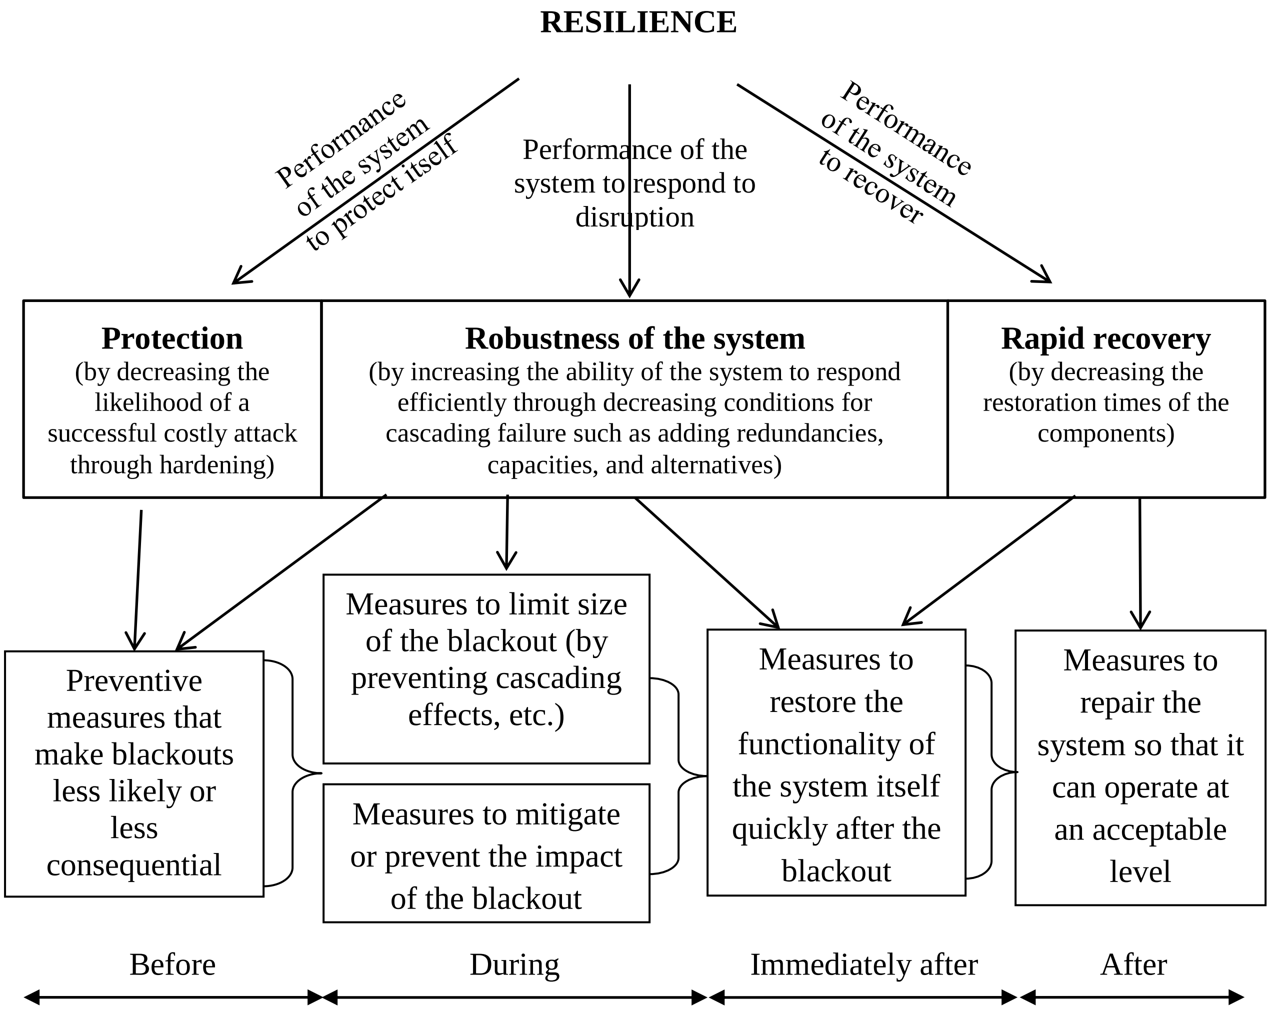
RESILIENCE
Performance of the system to protect itself
Performance of the system to respond to disruption
Performance of the system to recover
Protection
(by decreasing the likelihood of a successful costly attack through hardening)
Robustness of the system
(by increasing the ability of the system to respond efficiently through decreasing conditions for cascading failure such as adding redundancies, capacities, and alternatives)
Rapid recovery
(by decreasing the restoration times of the components)
Preventive measures that make blackouts less likely or less consequential
Measures to limit size of the blackout (by preventing cascading effects, etc.)
Measures to mitigate or prevent the impact of the blackout
Measures to restore the functionality of the system itself quickly after the blackout
Measures to repair the system so that it can operate at an acceptable level
Before
During
Immediately after
After

Figure 2.5-Relationship between resilience and defensive investments

We now demonstrate how we can represent our three types of defensive investments with an example. Figure 2.6 is a simple illustration of the impact of various investment types against intelligent attacks.

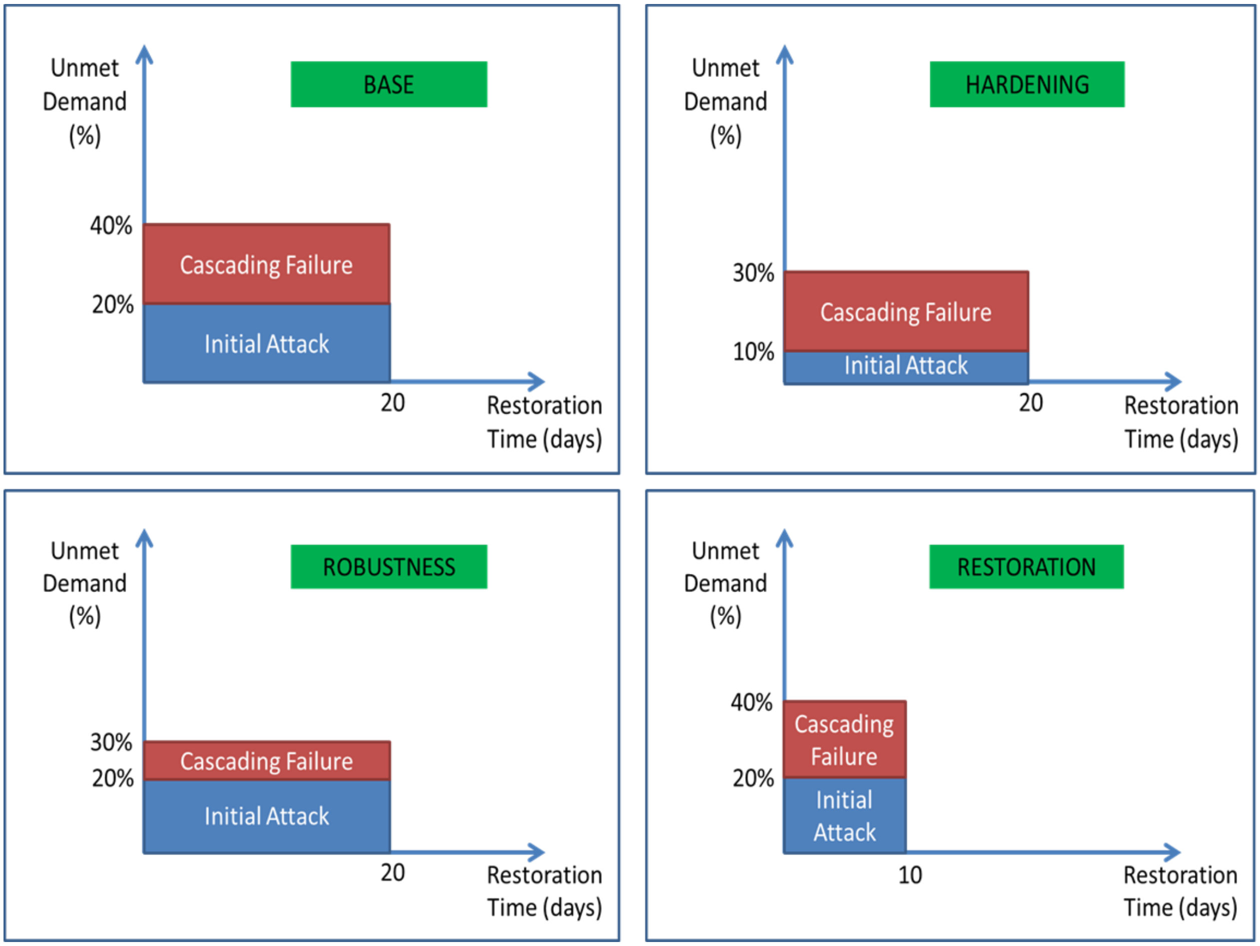

Figure 2.6-Illustration of the impact of various investment types

In the base case (the graph on the upper left corner), attacks cause 20% unmet power demand as an initial impact. Moreover, because of cascading failure, unmet demand doubles. Then, it takes 20 days to fully recover the system (for simplicity, we show the restoration process as one step). In this scenario, the total energy loss is the area under the curve (which is 40% times 20 days, or 800% of a day’s energy use, or 8 days of energy lost in overall). When we invest in hardening of the system (as shown in the upper right corner), the initial impact of an attack is expected to be less (10% instead of 20%). In this illustration, the impact of cascading failure stays the same as the base case (20%); however, because of the complex nature of cascading failure, this impact could be lower (or conceivably even higher). Since the electric power networks are highly capacity-constrained, we can also increase the robustness of the system by making it less likely to cascade. If we make the system less vulnerable to cascading failure by increasing line capacities or adding new lines (as seen in the lower left corner), the initial impact of an attack remains unchanged in our illustration (although it may in fact be different, since the network will have changed as a result of the investment); however, the cascading impact drops by half (from 20% to 10%), resulting in 6 days of energy lost in overall. Finally, we may want to help the system recover faster (as seen in the graph in the lower right corner), which would not change the initial impact or the impact of cascading failure, but would lead to less energy loss; for example, if the system recovers in half the time as in the base case, there would be 4 days of energy lost in overall. Thus, in our simple illustration, investing in the recovery process leads to the greatest saving in energy loss.

2.4. Resilience in Electric Power Systems

2.4.1. Why is resilience important in electric power systems?

It is virtually impossible to protect an entire electricity system against all possible natural and intelligent threats. As a result, it may be more reasonable to increase the system’s capacity to recover quickly from catastrophic events, rather than trying to protect the system from such events (Farrell et al., 2002).

Electric power systems are resistant to the failure of one or two random components. However, even a few overloaded components can sometimes cascade into larger failures. In a typical blackout, it may take days to fix some transmission lines; replacing a transformer or repairing a failed generating station can take months. Generating stations are generally protected by fences and guards, as are main substations located within generating facilities. On the other hand, transformers or substations near load centers are often open targets for intelligent adversaries. According to Crane (1990), destroying a few well-selected substations can cause a serious blackout. Even though some power would be restored almost immediately, region could suffer rolling blackouts for many months, especially during peak demand periods such as hot summers or cold winters.

The North American Electric Reliability Corporation (NERC) aims to ensure that the electric power system in North America is adequate, reliable, and secure. According to Blume (2007), some of the indicators of a stable and reliable interconnected electricity network are the “inertia” of the power grid, balanced electricity generation, and the grid’s ability to handle disturbances or faults. However, interconnectedness brings challenges, such as managerial complexity, increased reliance on supervisory control and data acquisition (SCADA) systems, scheduling problems, the possibility of cascading effects from other grids, etc. Amin (2003) notes that electric power systems are not only aging, but also increasingly complex and stressed, and remarks that it is unclear how long current systems will be able to support the ever-increasing demand for electricity. All of these issues make resilience in electric power networks an important goal.

Interruption of power delivery can have severe impact on the society. Holmgren (2007) underlines the importance of disruption by its duration, size of power loss, and number of people affected. Note that resilience triangle as a performance measure of electric power systems demonstrates the size of power loss as a function of time (duration), but lacks the ability to show the impact of the disruption to the society, such as area/magnitude of impact or number of people affected. Moreover, magnitude of an outage may be small in size and short in duration, but its impact can still be significant, such as computer problems and equipment jamming in commercial facilities after a few seconds of power disruption (Crane, 1990). See Appendix B for some major blackouts in North America.

2.4.2. The multidimensional nature of resilience

Bruneau et al. (2003) consider four types of resilience: technical; organizational; economic; and social. They note that different measures of resilience are needed to adequately address these different dimensions. For illustrative purposes, they present measures of seismic performance for electric power systems as shown in Table 2.1.

Technical resilience concerns the ability of the system to function. Some measures of technical resilience for electric power systems are the percentage of demand met, the ratio of supply to demand, time to restoration, time to full recovery, etc. *Organizational resilience* concerns the ability of the organization(s) to manage the system. For example, measures of organizational resilience could include how well emergency units function, how quickly spare parts are replaced, how quickly repair crews are able to reach the affected components of a system, etc. *Social resilience* concerns how well society copes with the loss of services as a result of a blackout. For severe blackouts, social resilience can be the most critical dimension of resilience. Finally, e*conomic resilience* concerns the ability to reduce direct and indirect economic losses. Rose and Liao (2005) note that direct costs manifest themselves in four ways: lost sales;

equipment damage/restart costs; spoilage of variable inputs; and idle labor costs (in addition to the costs of measures to reduce potential losses, such as backup generators and capacity expansion). Indirect costs are multipliers that ripple through the economy, such as impacts on the customers and suppliers of a disrupted firm, decreased consumer spending, decreased investments in the disrupted firm, public-health problems (such as dysfunctional sewage treatment), and economic disorder (looting, etc.). They also note that some Crane (1990) argues that, while direct losses can be avoided by backup systems, indirect costs may be partially mitigated through contingency planning, improved communications, customer education, social programs, and other planning approaches.

Table 2.1–Seismic performance measures for electric power systems (modified from Bruneau et al., 2003)

Dimension/ Property	**Robustness**	**Redundancy**	**Resourcefulness**	**Rapidity**
Technical	Maximize availability of operational power supply (units) after EQ (e.g., % of pre-earthquake level following small earthquake)	Replacement inventories (e.g., % available for small earthquake)	Models to assess network vulnerability and damage (e.g., EPRI model)	Maximize provision target power supply level (e.g., restoration to 95% of pre-earthquake level within 1 day)
Organizational	Emergency organization and infrastructure in place; critical functions identified	Replacement inventories for critical equipment (e.g., transformers, bushings)	Plans for mobilizing supplies and personnel (e.g., mutual aid agreements); identification of emergency work-around strategies	Maximum restoration of power supply
Social	% of all households with power immediately after EQ	Alternative power supplies for all critical emergency facilities (e.g., hospitals)	No form of rationing needed to meet minimum power needs	Partial power restored to all households within 1 hour
Economical	% of all businesses with power immediately after EQ	Alternative power supplies (backup power) for all key businesses	Voluntary power conservation program implemented	Pre-EQ economic activities re-established within 1 day

Similarly, Bush et al. (2005) define six consequence categories, including sector-specific consequences, human health and safety, economic, environmental, socio-political, and national-security consequences. However, some of these categories of consequences, such as social-political effects, may be difficult to predict and measure.

2.4.3. Threat and resilience

As shown in Figure 2.3, a resilient electric system demonstrates the necessary responses over time. However, this reaction may depend on the nature of the threat. Threats may be intelligent, natural, or operational in nature. The effect of natural threats may vary; for example, earthquakes may affect any system equipment, whereas hurricanes mainly affect transmission and distribution systems (Crane, 1990). However, for natural threats (such as earthquakes, hurricanes, severe weather, fires, solar flares, etc.), we can often predict resilience ahead of time, based on previous experience with similar events. By contrast, an intelligent threat may observe our protective actions and develop unexpected and effective attack strategies. For such threats, the nature of the risk may be far from static over time. Moreover, the measures needed to reduce damage to a transformer from an earthquake may be completely different from the measures needed to deter trespassing (in order to prevent sabotage). As a result, we may need to talk about the ability of the system to respond to a specific type of event.

Farrell et al. (2002) analyze 58 incidents of interruptions, unusual occurrences, demand or voltage reductions, and public appeals that took place in 2000. Nearly half of the events were due to weather (mostly thunderstorms), one fifth due to operator and maintenance error, another one fifth due to faulty equipment, one tenth due to insufficient power to meet demand, and finally a couple of events due to forest fires. Natural and operational threats accounted for roughly equal numbers of events, whereas none of the events were due to intelligent threats. However, the circumstances and public perceptions in the United States with respect to terrorist threat changed dramatically after September 2001. Moreover, an intelligent threat can evolve more sophisticated technologies over time. For example, Lave et al. (2007) note that smaller, non-nuclear electromagnetic-pulse weapons might be used for local attacks on critical power-system components or networked computers and telecommunication systems, without actually penetrating a facility.

In 2002, NERC developed *Security Guidelines for the Electricity Sector* (NERC, 2002). This document identifies a “spectrum of threats” ranging from simple trespassing to vandalism, civil disturbances, or dedicated acts of terror and sabotage, as well as natural disasters such as earthquakes, hurricanes, major floods, and ice storms. The security guidelines recommend that companies understand how to respond to the full spectrum of threats. However, these guidelines are generic, rather than tailored for different types of threats; for example, the guidelines simply consider a facility or combination of facilities to be critical if its damage or destruction would have significant impact for an extended period of time.

For a specific threat scenario, we can estimate the resilience of each component in the electricity system; however, as noted above, these components may have differing levels of resilience for different failure types. Moreover, failures may have cascading effects, or be initiated by a common cause. Finally, an intelligent adversary with knowledge of the system may seek optimal or near-optimal attack strategies, in which case the threat may become a function of the observed vulnerabilities and defenses of the system.

Figure 2.7 illustrates interdependencies of an electric power system. The system requires inputs in order to function properly, such as fuel, staff, spare parts, etc. In return, electricity is generated, transmitted, and distributed to users. Other infrastructure systems are among the users of electricity, but deserve special attention due to the interdependent nature of infrastructure systems. In the following sections, we will address what resilience means for each part of the system, some of the resilience measures that are commonly used, and options to improve resilience.

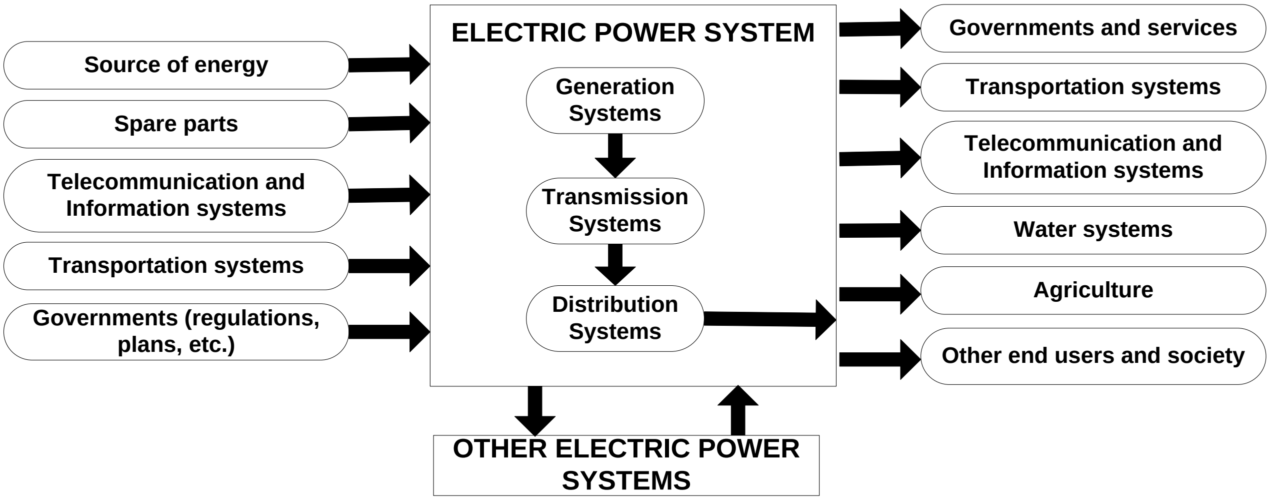

Figure 2.7-Interdependencies of an electric-power system

2.5. Resilience within Electric Power Systems

Electric power systems primarily consist of generation, transmission, and distribution systems; for a simple diagram of a generic electric power system, see Figure 2.8. A resilient electric power system has a strong capability to respond to any disturbances that may take place. According to Blume (2007), limited generation capacity, coupled with inadequate transmission lines, can force utilities to work at or near full capacity, and lower their reserve margins over time. This can be especially problematic in conditions of high demand, such as cold winters.

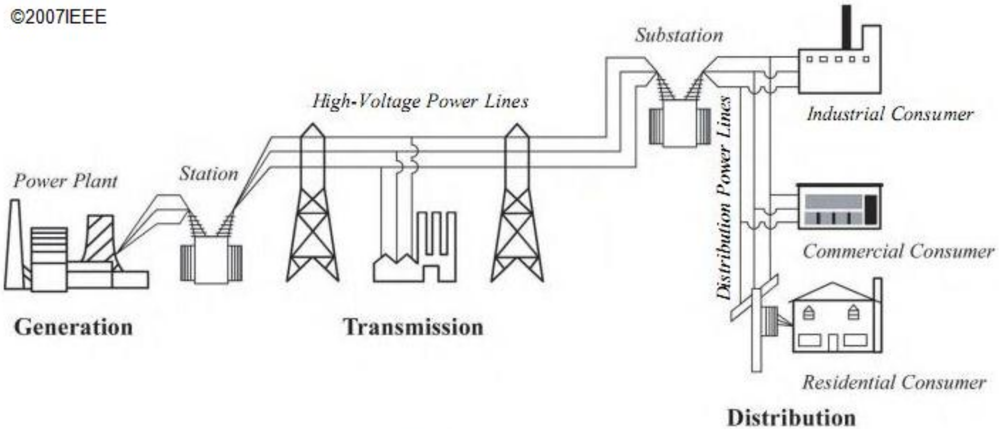

Figure 2.8-A generic electric power system (Blume, 2007)

Interconnectedness stems from the fact that electric grids are connected to each other, and use each other's transmission systems to transfer electricity. The interconnectedness of the grid can reduce the cost of providing electricity, avoid voltage collapse, and reduce the chance of undesirable load shedding (Blume, 2007). However, this interconnectedness can also lead to longer outages, particularly when transmission is scarce. For example, in some states, electricity deregulation has meant less investment in new transmission capacity. The potential for cascading effects and regional blackouts also means that problems in one state or region can affect other parts of the country. To reduce this dependence, Patterson (2007) recommends loosely interconnected and largely independent sub-networks, in which electricity may not need to be delivered over long distances, but of course, this would reduce or eliminate the economic benefits of interconnectedness.

Cascading failures happen when an initial failure triggers other failures in a short period of time. Major blackouts often occur as a result of cascading failures. One way to avoid this type of failure is to design networks that are resilient to cascading failures, as proposed by Ash and Newth (2007). Lovins and Lovins (1982) list some desirable characteristics of resilient electric power networks, such as modular structure, early fault detection capability, redundancy and substitutability, selective coupling, diversity (to avoid common cause failures), diverse or remote location of key components, standardization (making it possible to plug in common replacement components), hierarchical embedding (to isolate faults on the level at which they occur), stability (to gain time to decouple a faulty component before it affects others), simplicity (to aid in identifying and tracking failure modes), and accessibility (to improve ease of maintenance).

Ideally, one would like to minimize the negative impacts of interconnectedness, while still obtaining its benefits. One way to avoid negative impacts is to develop online alert systems that can provide short-term predictions or notifications of disturbances in other electric grids (Capodieci et al., 2010). Since utility companies are affected by cost considerations, a long-term solution is to develop incentives so that utilities would find it more beneficial to invest in transmission capacity. Large blackouts due to cascading failures may be particularly likely when transmission systems are operated near a critical point (Carreras et al., 2002). Zimmerman et al.

(2007) note that transmission systems are the most frequent disabled systems (60% of international outages due to terrorism and 90% of US-based non-terrorist events). They also note that the worst-case scenarios are those where transmission lines run along few corridors, with only a few power sources in the transmission system, and no urban power generation is available.

In general, capacity limitations and low reliability are major obstacles to a resilient transmission system. According to Apt and Lave (2003), greater resilience can be achieved both by building more transmission lines, and also by improving the capacity and controllability of existing lines. These investments are in general quite costly, and under current regulatory conditions usually have a low rate of return. Despite their cost, however, some improvements in data acquisition and control can pay for themselves by decreasing the costs of operations, scheduling, and maintenance. Crane (1990) also advocates making transmission towers more resilient to natural disasters, such as high wind and earthquakes.

NERC recognizes the need for reliability metrics for generation and transmission systems. As of today, there are seven approved metrics to measure reliability:

1. Planning Reserve Margin
2. Transmission Related Events Resulting in Loss of Load
3. Average Percent Non-Recovery of Disturbance Control Standard (DCS) Events
4. Disturbance Control Events Greater Than Most Severe Single Contingency

5. Percent of Automatic Outages caused by Failed Protection System Equipment

6. Energy Emergency Alert 3

7. Energy Emergency Alert 2

NERC's metrics are intended to measures a system's ability to respond to unexpected changes. However, they assume access to the required inputs (personnel, fuel, etc.). Moreover, these metrics are intended for use primarily within a single region, rather than for comparison between regions.

Planning Reserve Margin is the difference between the forecasted generation capacity and the projected peak demand, normalized by peak demand. Forecasts are based on median weather (such that the actual weather is equally likely to be either warmer or cooler). For example, Figure 2.9 shows that in the winter peak in Canada in 2011, the planning reserve margin is expected to be 20%, which is higher than the goal of 10%. This measure shows how much flexibility the generation system has during its peak period to respond to adverse events such as extreme weather or unexpected outages. However, since this measure is based on median weather, it is not directly comparable from one region to another. For example, worst-case summer weather in Hawaii may be only a few degrees warmer than median summer weather, while worst-case winter weather in Wisconsin or Canada could conceivably be thirty degrees colder than median winter weather.

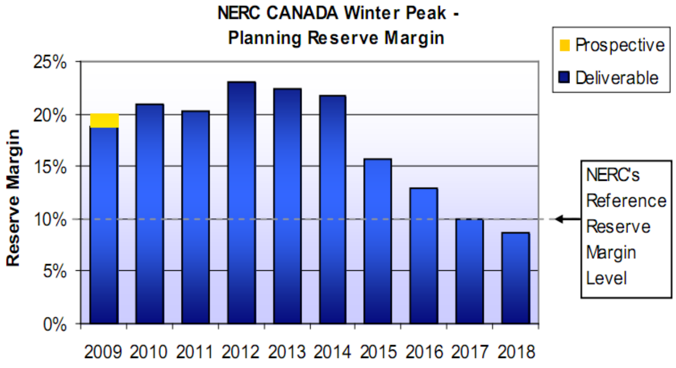

Figure 2.9–NERC's planning reserve margin for winter in Canada (NERC, 2010)*

Transmission Related Events Resulting in Loss of Load are defined as transmission failures due to equipment failure or misuse, causing significant loss of load greater than 200MW and/or 50% of demand (NERC, 2010). See Figure 2.10 for the number of significant transmission-related events between 2002 and 2009. NERC is also considering developing different metrics for different types of transmission outages, such as failed protection systems, human error, failed circuits, failed station or substation equipment, etc.

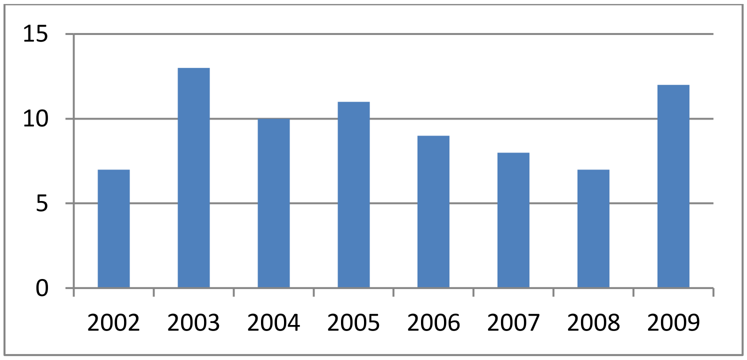

Figure 2.10–Transmission-related significant events between 2002-2009 (NERC, 2010)*

* These images from the North American Electric Reliability Corporation's website are the property of the North American Electric Reliability Corporation and are available at http://www.nerc.com/docs/pc/rmwg/RMWG_AnnualReport6.1.pdf. This content may not be reproduced in whole or any part without the prior express written permission of the North American Electric Reliability Corporation.

Average Percent Non-Recovery of Disturbance Control Standard (DCS) Events measures the ability of load-balancing authorities or reserve-sharing groups to compensate for any significant loss of supply with contingency reserves within a specified recovery period (e.g., 15 minutes). The goal is to return to the pre-disturbance level of electric generation (or at least to a situation where generation equals demand, if pre-disturbance generation was greater than demand) within a predetermined amount of time after a significant (reportable) disturbance. This metric is the arithmetic average of the calculated non-recovery percentages of all reportable disturbances in a given year, even for different sized outages. (Note that these numbers are again not directly comparable between regions, since factors such as the definition of a reportable disturbance and the length of the specified recovery period may vary.)

Similarly, the other NERC metrics are also either numbers of events, or the percentage of events with a given characteristic. This makes the metrics straightforward to compute, but may hide some key characteristics, such as the cause, severity, and duration of events, which are likely to be highly related to the resilience of the system. For example, a system could have excellent performance on the NERC metric for a period of several years, but still suffer catastrophic failure due to a natural disaster or intentional attack that exceeds the system's response capacity.

Farrell et al. (2002) claim that large central generators and long transmission lines are inherently vulnerable, and propose a dispersed system with many small generators situated near large population centers. Similarly, Patterson (2007) claims that reliability and control are key drivers behind the move toward greater on-site generation. He believes that as small-scale generating technologies (e.g., microcogeneration, microturbines, fuel cells) become more mature, even facilities with relatively small loads (such as office buildings, hospitals, airports, hotels, and schools) will become candidates for on-site generation, to the extent that facilities without access to on-site generation may be at a “severe disadvantage.” Makansi (2007) also advocates distributed power systems.

Generation and transmission systems cause only a few outages every year, mainly due to weather and natural disasters, human errors, and equipment problems. Farrell et al. (2002) note that distribution systems are responsible for the majority of electricity outages, despite the limited impact of most distribution failures. In particular, distribution systems are more vulnerable to hurricanes, floods, and flying debris, since they tend to use low lines. However, Crane (1990) nevertheless recommends greater generating and transmission reserve margins, to reduce the impact of generation and transmission outages on customers. See Table 2 for some general investment strategies to prevent damage, limit consequences, speed recovery, and generally reduce the vulnerability of an electric-power network. Note, however, that the prevention measures in this table are generally focused on intelligent threats.

Table 2.2–Investments to increase resilience of electric power systems (Crane, 1990)

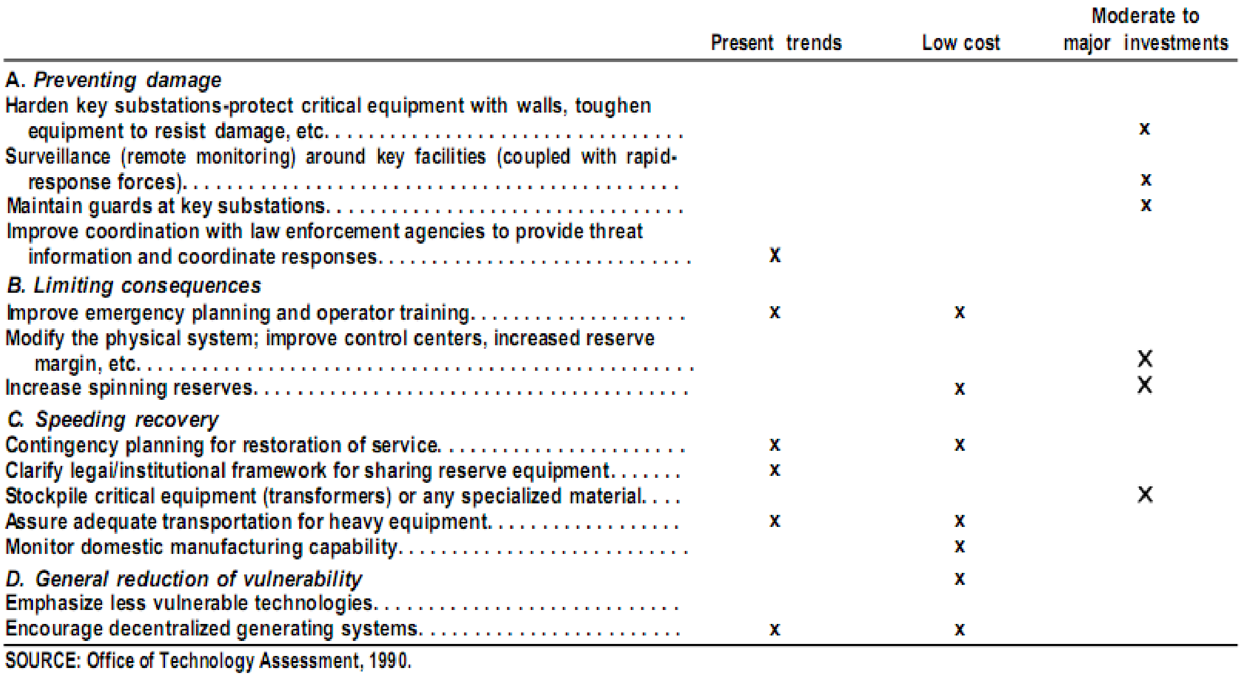

	Present trends	Low cost	Moderate to major investments
A. Preventing damage			
Harden key substations-protect critical equipment with walls, toughen equipment to resist damage, etc.			x
Surveillance (remote monitoring) around key facilities (coupled with rapid-response forces).			x
Maintain guards at key substations.			x
Improve coordination with law enforcement agencies to provide threat information and coordinate responses.	x		
B. Limiting consequences			
Improve emergency planning and operator training.	x	x	
Modify the physical system; improve control centers, increased reserve margin, etc.			X
Increase spinning reserves.		x	X
C. Speeding recovery			
Contingency planning for restoration of service.	x	x	
Clarify legai/institutional framework for sharing reserve equipment.	x		
Stockpile critical equipment (transformers) or any specialized material.			X
Assure adequate transportation for heavy equipment.	x	x	
Monitor domestic manufacturing capability.		x	
D. General reduction of vulnerability		x	
Emphasize less vulnerable technologies.			
Encourage decentralized generating systems.	x	x	

SOURCE: Office of Technology Assessment, 1990.

2.6. Resilient Inputs to Electric Power Systems

Bush et al. (2005) highlight the dependence of electric power systems on the information and communication sector for operations of its supervisory control and data acquisition (SCADA) systems, on the transportation sector for the movement of spare parts and repair personnel, and on the government for the institution of electricity-conservation alerts or the event of shortages. To this list, we can also add fuel and other resources needed to generate and transmit electricity, and crew and spare parts to maintain and repair them.

The most crucial input to an electric system is the fuel or energy source to run the generators; e.g., coal, natural gas, and uranium. Nuclear power plants require refueling only every year or two, while other technologies require frequent or even continuous fuel delivery. Utility companies may be able to increase resilience by ensuring adequacy of supply or alternate supplies. For example, Lave et al. (2007) note that increasing use of natural gas increases the dependence of electric-power systems on gas transmission systems; however, gas transmission systems are relatively invulnerable because they are underground. Utility companies may want to look into availability of alternate fuel suppliers and/or alternate modes of transport (e.g., road rather than rail for delivery of coal in emergencies).

It may also be critical to have sustainable energy resources that do not jeopardize the ecosystem. Thus, investing in renewable energy, such as solar and wind power sources, may make electricity systems more resilient to potential scarcity of current energy sources.

The components of the electric-power network must also be maintained efficiently. While transmission lines are relatively easy to replace, generating plants often have redundant systems that enable them to maintain their operations in case of component failures. Crane (1990) recommends the use of spare transformers to speed recovery from transformer failures. Spares are often purchased for other key equipment, but transformers are expensive, so are often shared between multiple plants. Spare transformer programs can reduce the time to recover from a transformer failure from over a year to several months (Crane, 1990). One concern is delays in the manufacturing of replacement transformers, since they are manufactured in only a few countries. Farrell et al. (2004) also note that most utilities keep their spares at substations, so an attack on a facility may also cause damage to the spare transformer.

According to Amin (2002), existing control systems were initially designed as stand-alone systems, and over time were connected to each other; however, the changes needed to make these systems more secure were not immediately incorporated. In recent years, there have been increasing efforts to make the internet more secure, by limiting personnel authorization, training IT personnel, using data encryption, and implementing advanced intrusion-detection systems through firewalls, multiple check points, and security procedures. There has also been an industry trend towards more decentralized and dedicated communication sources rather than the networked systems (Amin, 2002), which make events less likely to occur, and also less likely to cascade across multiple regions.

SCADA systems are widely used today to monitor facility operations in real time. These software programs increase efficiency by making inspection, adjustment, and data collection easier and more convenient. On the other hand, these systems suffer from externalities common to the software industry. Anderson (2001) notes that, because of the high fixed costs and low marginal costs of IT products, and the difficulty of customers switching from one product or service to another, “time-to-market” is critical for survival in the software industry. As a result, the software industry is dominated by the few firms that manage to obtain first-mover (or early-mover) advantage. One side effect of such rapid product development may be less concern for security in the initial development of IT products. Another problem is difficulty in getting competitors to coordinate their efforts, which may result in excessive vulnerability of SCADA systems.

Utilities also invest heavily in continuity-of-operations planning to ensure that they can continue to function effectively in the event of natural disasters that damage their facilities, and/or pandemics that limit the availability of their work force. In order to achieve this end goal, utilities invest in protection of essential facilities, training and backup for critical personnel, measures to decrease the probability of disruptions, data protection, management information systems, etc.

2.7. Resilience to Loss of Electric Power

Society is highly dependent on electricity. Fresh food, air conditioning, heating, telecommunications (cell phones, landlines, internet, and the media), security systems, hot water, and emergency energy (such as traffic lights and hospital equipment) are basic requirements of modern life (Makansi, 2007). In the absence of electricity, commercial facilities may not be able to maintain some necessary functions or services (e.g., air conditioning, refrigeration, secure money transfer). Moreover, industrial organizations may need to cut production levels, which can result in cancellations, scheduling problems, etc. Outages can also cause spoilage of both raw material and finished and unfinished goods. Emergency-management organizations, law enforcement, firefighters, and the National Guard also depend on electricity to varying degrees. Talukdar et al. (2003) provide some examples of critical services that depend on electricity, such as pumps for water and sewer systems, urban mass transit, emergency-service systems, escalators and elevators, navigation aids for air traffic, and financial services.

Crane (1990) lists some of the direct and indirect costs of the 1977 blackout in New York City. Some of the affected groups included businesses (food spoilage, banking, the need for emergency aid to the private sector), government (federal and state assistance programs), utility companies (restoration cost, overtime payments, new capital equipment), the insurance sector (federal crime insurance, fire insurance, private property insurance), public health (overtime in emergency rooms), transportation (transportation losses, transportation-overtime costs, transportation vandalism, new transportation equipment), and other public services (fire-department overtime, police-department overtime, state court overtime, etc.).

Many infrastructure systems are heavily dependent on electric power, mainly due to functional dependencies and cascading effects. In a major blackout, within a short amount of time, cell phones and landlines typically go out of service, banking services (such as money-transfer facilities and ATMs) stop functioning, roads become congested (since traffic lights may stop working), transit systems and airports may be forced to discontinue or delay services, computer servers may shut down (unless they have adequate backup power), gasoline pumps will stop working, transportation of other energy resources (natural gas, oil, etc.) may be delayed. Yet, it is a major challenge to measure how dependent all of these systems are on electric power.

Federal and local governments and agencies are responsible for coordinating efforts in the case of power outages, despite the fact that they are also vulnerable to these outages. The Federal Emergency Management Agency (FEMA) is charged with coordinating responses to disasters, including major power outages. Emergency-response efforts become more challenging in power outages, because of the high dependency of other infrastructure systems on electricity. It is difficult to even simulate realistic blackout conditions, or to ensure well-coordinated emergency-management efforts under all possible disaster scenarios, because of uncertainties about post-event conditions. Power outages and their impacts may make coordination, transportation, and emergency operations more challenging. Therefore, FEMA attempts to increase citizen awareness and knowledge, so that emergency-response efforts can be maintained and function as efficiently as possible.

After any large blackout, federal and state governments initiate assistance programs. Hospitals, fire departments, and police departments may accrue additional emergency and overtime charges, and may have limited capacity. For example, backup generators in hospitals may provide limited support to emergency units, operating rooms, intensive care units, x-ray devices, air conditioning, refrigeration, elevators, etc. (Crane, 1990). In the 1998 ice storm in Canada, fire departments were busy dealing with downed live power lines (Scanlon, 1999). If a blackout also leads to social disturbances, then state courts and prosecution/correction authorities may need to work overtime.

The banking system is also heavily dependent on electricity. After the northeast blackout in August 2003, the U.S. Treasury Department reported that the U.S. financial system had been extremely resilient, but this was in part because the power outage started after the markets had closed, and lasted only a few hours. In a longer blackout, the need for bank machines, credit cards, and electronic funds transfer will become more significant, since people will need money to pay their bills and buy essentials for survival. Chang et al. (2007) assessed the impact of the 1998 ice storm in Canada, and documented problems with usability of ATMs, banks, and credit cards due to lack of power.

Reliance on electricity for transportation (by car, truck, train, and especially transit) enhances the dependency of the transportation sector on electricity (Ibáñez, 2010). Transit systems use electricity to energize the main system, signaling devices, and essential telecommunication systems, and to provide ingress and egress for transportation facilities (Bush et al., 2005). In addition, the economic impact of blackouts can cause lost revenue due to disruption in transit services, and overtime costs after system functionality is restored. Emergency backup systems, emergency preparedness and evacuation efforts, and the availability of alternative energy sources may increase the resilience of transit systems to blackouts. However, in the northeast blackout of 2003, people leaving work shortly after the blackout created massive congestion in all available means of transportation within approximately 10 minutes (DOT, 2004).

Lack of electricity may eventually affect other means of transportation as well, through dysfunctional pumps in gasoline stations, failed traffic signals, dysfunctional automated tollbooths, electrical tools needed for repairs, dysfunctional pumps to control flooding in tunnels and other low areas, lack of lighting, lack of emergency support (due to communication failures and staffing shortages), etc. For example, in the 2003 blackout, eight oil refineries in the U.S. and Canada were shut down, which threatened gasoline shortages (especially in the Detroit metropolitan area), and created a potential energy emergency; as a result, certain air-quality regulations were suspended (Electricity Consumers Resource Council, 2004). For more examples, see DOT (2004).

It is also important to note that the magnitude of the impact may depend on the concentration of traffic in the affected location at the time of the incident. For example, the New York region (with a large number of industrial, commercial, and residential users, and easily congested roadways) may be more affected by lack of electricity than some other areas (Zimmerman et al., 2007).

However, Farrell et al. (2002) claim that simple solutions to maintain or restore essential services may sometimes be more cost-effective than trying to prevent blackouts. For example, they note that traffic signals could use low-power LED lights with uninterruptible power supplies, so that traffic flow could be maintained even after blackouts. Some other simple solutions that increased resilience after the northeast blackout of 2003 were the use of portable stop signals on congested roads, requests for people to stay home after they got there, etc. (DOT, 2004). DOT also notes additional measures that could decrease the impact of future blackouts, such as telephones with separate power sources, voice-messaging systems with dial-up capability, redundant communication systems (for example, push-to-talk services or dialup network connections), fuel-storage tanks for backup generators, and mutual-aid agreements with partner agencies and neighboring communities.

Telecommunication systems are highly dependent on electricity. Businesses and public services (such as police, fire, etc.) all rely on timely flow of information. An extended outage may have a significant impact on telecommunication networks, since emergency-backup systems are generally both short-term and costly. As noted by Rabkin et al. (2004), communication systems are heavily used in emergencies, and as a result, the batteries for cell and satellite phones may be discharged quickly. One way to avoid this problem is to have multiple communication technologies, especially systems that use little or no electricity (for example, radios that rely on solar or mechanically generated power). Having spare batteries on hand will also extend the usability of cell and satellite phones.

Water systems depend on electricity for pumping, treating, and distributing water (Bush et al., 2005). Hydraulic and solar pumps may decrease dependence on electricity, as can the availability of local water storage for drinking, irrigation, fire suppression, farming, food preparation, chemical manufacturing, etc. In particular, water towers can maintain the distribution of water even during power outages. Another option is to use solar water heating systems in warm climates. Electricity is also necessary for treatment and pumping of sewage. In an extended outage, backup systems may be exhausted; as a result, untreated sewage could flow directly into the environment, creating a potential public-health hazard (Crane, 1990).

The food sector deserves special attention, because extended and large-scale outages could create severe social impacts if food becomes scarce. Electricity is needed for incubation, milking, refrigeration, heating, and air conditioning (Crane, 1990). Due to previous large blackouts, the food and agriculture industries increasingly rely on backup generators. Residential customers can also decrease the impact of extended outages by stocking food and drinking water. In the 1998 ice storm in Canada, three million people lost electric power from periods of several hours to as much as a month in some areas. A generator group was organized within an hour after declaration of emergency to find and allocate limited generators to the locations where power is urgently needed, such as nursing homes, fire departments, and individuals with serious health conditions. Similarly, a firewood team was established to alleviate heating problems (Scanlon, 1999). Dairy farmers were particularly affected by the power loss, because of heavy dependence on milking machines. Farms could not milk their cows, which made them vulnerable to mastitis.

Moreover, with no ventilation and low temperatures in the barns, many cows developed infectious diseases such as pneumonia (Kerry et al., 1999).

Social resilience to the interruption of electricity can also be increased by providing practical knowledge of what people should do in case of emergency, reducing reliance on electric systems, recommending the use of manually charged flashlights in emergency toolkits, etc. Clarke (2002) notes that widespread fear, panic, and disorder are rare following major disasters. Rapid information flow on the nature of the outage may also help to maintain order and increase social resilience (Lave et al., 2007), but attention may need to be paid to achieving rapid information flow at times when many communications technologies may not be working.

2.8. Conclusion

Resilience is performance of a system against an undesired event over time. This performance is negatively correlated to the consequences of that event. We can use resilience to gauge the effectiveness of our defensive measures, i.e. protection, robustness, and rapid recovery.

Electric-power systems are critical to society, and closely linked to many other critical infrastructure systems and essential services. Electric-power systems require inputs and services from other infrastructure systems, interact with each other through the grid, and provide power to critical infrastructure systems and other end users. A systematic approach to analyzing and enhancing the resilience of electric power systems requires considering not only the resilience of the electricity system itself (e.g., ensuring adequate reserve margin and minimizing equipment outages), but also ensuring the resilience of needed inputs (fuel transportation, etc.) and reducing the dependence of key sectors on electricity. Considering electricity resilience from a broader perspective in this manner may make it possible to achieve resilience more cost-effectively, since it is implausible to assume that all electricity outages can be prevented.

In Chapter 3, we will take a look at the literature to examine different methods of vulnerability analysis on electric power networks, and how they address protection (through hardening), robustness (through cascading failure), and rapid recovery (through restoration times). We will also analyze how our model may fill a gap in literature.

3. Literature Review

Prior to the terrorist attacks of September 11, 2001, the term “vulnerability assessment” was often used in relation to risk and system safety (Einarsson and Rausand, 1998). Following September 11, however, there was an increased emphasis on vulnerability to security threats from intelligent adversaries.

The U.S. created the Office of Homeland Security and the Homeland Security Council in 2001, which later led to the creation of the U.S. Department of Homeland Security (DHS) in 2002. The first strategic document on homeland security, *The National Strategy for Homeland Security* (Bush, 2002), defined three strategic objectives: to prevent terrorist attacks against homeland targets; to reduce vulnerability to terrorism; and to minimize damage and recover from attacks. These three objectives roughly correspond to threat, vulnerability, and consequence.

In 2004, DHS published *The National Strategy for the Physical Protection of Critical Infrastructures and Key Assets* (DHS, 2003), a guide to the protection of critical infrastructures (e.g., water, energy) and key assets (e.g., national icons, nuclear power plants). This document highlights the importance of protection, response, and recovery, and specifies the roles of the federal government and the states in identifying and securing the critical infrastructures and key assets under their control.

More recently, the *National Infrastructure Protection Plan* was written to provide additional guidance on how to make the nation’s infrastructure safer, more secure, and more resilient (DHS, 2009). In particular, the plan requires implementation of a “long-term risk management program" that includes: hardening, distributing, diversifying, and increasing the resilience of infrastructure against threats and hazards; interdicting potential attacks; and planning for rapid response to disruptions and rapid recovery.

DHS defines vulnerability as a "physical feature or operational attribute that renders an entity, asset, system, network, or geographic area open to exploitation or susceptible to a given hazard" (DHS, 2010). In this definition, vulnerability could include any weakness that a terrorist could exploit, or that makes the system more susceptible to either natural or manmade hazards. Other organizations use a narrower definition of vulnerability. For example, the U.S. Coast Guard (USCG, 2003) defines vulnerability as "the conditional probability of success given a threat scenario occurs."

In this work, we adopt a broader definition of vulnerability, which also includes the short-term response of the system to an attack (in the form of cascading failures), and also the long-term resilience or recovery of the system (i.e. restoration time). Note that this is also consistent with the recent recommendation of NRC that DHS's vulnerability analyses should ideally address issues of system capacity and long-term adaptation (2010). Therefore, after discussing vulnerability assessment methods of electric power networks, we move on to discuss models of cascading failure and restoration times.

3.1. Vulnerability-Analysis Methods

Methods for vulnerability analysis include rating-based methods, risk-based methods, and game-theoretic methods. We discuss all three of these approaches below.

Note also that many risk-based and game-theoretic methods attempt to represent the physical system being analyzed in one of two different ways, using either topological models or flow-based models. Flow-based models aim to represent how a system actually functions. By contrast, topological models consider only the network structure. Thus, topological models can identify redundancies, potential bottlenecks, etc., but cannot take into account factors such as capacity constraints (e.g., whether a particular line has sufficient capacity to serve all needed loads when other parts of the system have been degraded). We discuss rating-based methods, risk-based methods, and game-theoretic methods in turn below.

3.1.1. Rating-Based Methods

Rating-based methods assign scores to various attributes of the system being analyzed. They are not specific to electric power networks, and can generally be applied to a wide variety of systems, facilities, or networks. One such method is the Criticality, Accessibility, Recuperability, Vulnerability, Effect and Recognizability (CARVER) method, originally developed by U.S. Special Operations Forces to help prioritize targets during the Vietnam War. DHS uses CARVER to prioritize critical components and assets as part of the Buffer Zone Protection Plan (Bennett, 2007). In this method, each potential target is scored on the attributes of criticality, accessibility, recuperability, vulnerability, effect, and recognizability; the resulting numbers are simply added to find a final score for each target, effectively assigning the same weight to each attribute. By contrast, Ezell (2007) uses an additive preference model in which different weights are assigned to the various vulnerability attributes to quantify the importance of vulnerabilities.

A similar rating-based method is the Enhanced Critical Infrastructure Protection (ECIP) program, developed by Argonne National Laboratory for DHS (Fisher and Norman, 2010). In this model, facilities score their vulnerabilities in areas such as physical site security, security management, etc. using a questionnaire. The weights corresponding to the various areas were assessed using expert elicitation, and are hardwired into the program. The program is intended for use by facility owners and operators to prioritize protective measures for a facility, and compare its risk-management features (fences, cameras, etc.) with those of other facilities in the same sector.

Rating-based methods are easy to implement, are applicable to virtually any type of infrastructure system, and can provide useful insights to decision-makers. However, they are perhaps best utilized for qualitative tasks, such as identifying threat scenarios or screening critical components, since they generally do not include a realistic representation of the physical system being analyzed, and hence cannot account for factors such as flows within the system, dependencies among components, etc. This weakness is especially limiting when applying rating-based methods important to complex networks such as electric power systems (as opposed to simple facilities, such as factories, that can be modeled as single entities), because the network topologies and dynamic behaviors of complex systems cannot be completely captured by rating the individual components of the system. Moreover, in most rating-based models, threat is not considered explicitly.

3.1.2. Risk-Based Methods

Probabilistic risk assessment has been used to analyze the vulnerabilities of infrastructure systems since the mid-1970s (Rasmussen, 1975). Many risk-based vulnerability-assessment methods were originally developed for assessing system safety and reliability, and are well accepted for that purpose; by contrast, application of risk-based methods to intentional security threats is newer and somewhat more controversial.

Risk-based models attempt to answer three fundamental questions (Kaplan, 1997):

(1) What can go wrong?

(2) How likely is it?

(3) What are the consequences?

The resulting estimate of risk is generally expressed in the form:

$$Risk = f(Threat, Vulnerability, Consequence)$$

Most commonly, risk-based models compute risk as the product of threat, vulnerability, and consequence. Some of the most prominent risk-based vulnerability models are discussed below.

There are numerous government-sponsored vulnerability-assessment methods based on risk. For example, Risk Analysis and Management for Critical Assets Protection (RAMCAP, 2006), developed by the American Society of Mechanical Engineers under sponsorship by DHS, models vulnerabilities using event trees. Moreover, threat is estimated as a function of both target attractiveness, and adversary capability and intent. Risk is then estimated as the product of threat, vulnerability, and consequence.

Similarly, the Transit Risk Assessment Methodology (TRAM), developed by the Port Authority of New York and New Jersey in 2002, also uses event trees to assess vulnerability (TRAM, 2007). While TRAM was originally designed to prioritize surface-transportation assets for protection from possible threat scenarios, similar models can also be applied to other types of infrastructure networks. One such example is the Maritime Security Risk Analysis Model, developed by the U.S. Coast Guard in 2006 to help prioritize the risks of terrorist attacks on ports and waterways (Parfomak and Frittelli, 2007).

Some government models are specific to a particular type of threat, such as DHS's Bioterrorism Risk Assessment model (DHS, 2006). This model uses event trees to prioritize bioterrorism threats based on subjective estimates of their probabilities and consequences. Similarly, DHS also developed the Chemical Terrorism Risk Assessment model and an Integrated Chemical, Biological, Radiological, Nuclear Assessment model (NRC, 2010).

Other government vulnerability-assessment tools and programs are designed for specific sectors. Examples include: the Aviation Domain Risk Assessment; the Dams Sector Analysis Tool; the Emergency Services Self-Assessment Tool; the National Transportation Sector Risk Assessment; the Ports and Waterways Safety Assessment; the Risk Assessment Methodology for Water Utilities; and the Water Infrastructure Simulation Environment. These are all briefly described in the DHS Risk Lexicon (DHS, 2010).

Risk-based models have also been discussed frequently in the academic literature. For example, Ezell et al. (2000) propose a risk-based model to identify the vulnerable components of an infrastructure. The model first identifies vulnerabilities and threats, and ranks both of them. Vulnerability is modeled as a function of access and exposure. Event trees are used to determine sequence of events for various threat scenarios that lead to consequences; for each scenario, risk is then calculated as expected loss (i.e., probability of vulnerability times consequence).

Apostolakis and Lemon (2005) develop a risk-based approach to analyze the vulnerabilities of water, natural-gas, and electric-power distribution networks in the face of relatively minor terrorist attacks. The authors focus on the topological structures of the networks, and also geographic locations of the components. For example, they note that electrical-service ducts are often collocated with (or geographically in proximity to) natural-gas and water networks, creating critical points that are highly vulnerable. The model identifies all possible combinations of failures that may result from a single attack somewhere in the system, but without any assessment of their likelihood. Instead, using expert judgment, the authors estimate the accessibility of each critical point. The screening methodology then identifies and ranks the failure combinations (i.e., minimal cut sets) based on their susceptibility to attack and the value of each target to the decision maker (as calculated using multi-attribute utility theory). For extensions of this work, see Koonce et al. (2008) and Patterson and Apostolakis (2008).

Donde et al. (2005) use a graph-partitioning algorithm to identify critical power lines whose failure may cause severe system disruptions. Similarly, Lesieutre et al. (2006) use graph theory to identify subgraphs that are at risk for unmet demand in the case of extreme events. Bienstock and Mattia (2007) develop a graph-theoretic network model to explore how the robustness of the network can be improved at minimum cost. They consider two types of investments: adding more capacity to an arc; and adding more arcs in parallel. Oliviera et al. (2004) also study transmission-expansion planning, and calculate the improvements needed to avoid possible disruptions. He et al. (2005) analyze voltage stability to identify weak components, while Vulkanovski et al. (2009) generate fault trees for each load in a system, and identify the most important elements in those fault trees using risk-importance measures.

In their flow-based model, Bienstock and Verma (2009) use mixed integer and nonlinear models to identify if there is any small number of arcs whose removal will cause a blackout. Similarly, Pinar et al. (2010) use a bilevel integer program to identify small groups of lines in a network whose removal would be anticipated to cause a severe disruption. By using a special structure in their formulation, they avoid the nonlinearity in the original bilevel-mixed integer nonlinear problem, and approximate the problem as a mixed-integer linear problem. In the outer loop of their optimization, they identify the critical lines, while in the inner loop they measure blackout severity by solving the load-shedding problem that minimally decreases load given an assumed loss of the identified critical lines. Once the critical components that lead to a blackout have been identified, the authors then find the minimum change in generation required to avoid the blackout.

In general, many risk-based models are simple and practical to use; Ezell et al. (2010) argue that the required inputs to risk-based models can be readily obtained through expert elicitation, and note that PRA has been successfully applied to a number of large complex systems. However, Cox (2008) describes some of the limitations of risk-based models when applied to intentional threats. In particular, he notes that threat probabilities may not be well-defined constants, since an adversary might respond to any observed defenses; he also raises similar concerns about the possible ambiguity of vulnerability and consequence. Brown and Cox (2011) similarly warn that risk-based methods may result in misleading recommendations regarding protective actions, since the attack probabilities assumed in such methods may not reflect the attacker's ability to learn from the defender's analysis and/or the observed defenses.

The National Research Council (NRC, 2010), in a recent review of DHS's approach to risk analysis, has stated that the basic idea of representing risk as a function of threat, vulnerability, and consequence is sound, but recommends improvements to the validity and reliability of such models. For that reason, the NRC recommends that DHS incorporate game theory into its vulnerability-analysis methods.

3.1.3. Game-Theoretic and Quasi-Game-Theoretic Methods

Unlike risk-based methods, game-theoretic vulnerability-assessment methods focus on the behavior of a strategic adversary. In particular, many game-theoretic vulnerability-assessment methods for networks are based on interdiction models. These games aim to determine how the attacker can "interdict" various components of the network in order to best achieve an objective. Then, the defender determines how best to operate the remaining network. In optimization terminology, these types of games (which are special cases of sequential or Stackelberg games) are often represented as mixed-integer bilevel programs; see Wood (1993) for an overview of interdiction models. Most interdiction models in the literature are deterministic. However, see

Cormican et al. (1998) for an overview of stochastic interdiction models. In addition, Janjarassuk and Linderoth (2008) reformulate of stochastic network-interdiction problems as deterministic mixed-integer programs, and Morton et al. (2007) apply stochastic interdiction to the problem of nuclear smuggling.

Interdiction models have been extensively applied to transportation, nuclear smuggling, border patrol, etc. Because this field is so broad, we limit the remainder of our discussion in this session to models that are either applied or potentially applicable to electric-power networks.

In addition to truly game-theoretic models, however, we also address here models with conversational games, which "consists of advice, suggestions, and council about how to think strategically" (Smith and von Winderfeldt, 2004), and models with worst-case assumptions regarding threat scenarios. These quasi-game-theoretic models typically do not include any consideration of optimal defenses, and may not even have been intended as models of adversary behavior. However, they go beyond simple risk-based models, since their use of worst-case assumptions in selecting which threat scenarios to consider can still help shed light on possible attacker behavior.

We begin by studying topological models. Albert et al. (2004) develop a topological model to study the structural vulnerability of the North American power grid. They compare the impact of various interdiction strategies, such as removing transmission substations at random, removing those nodes with the highest number of arcs into and out of them (i.e., the nodes of highest "degree"), or removing nodes in decreasing order of estimated load (where load is estimated by the number of paths through each node: i.e., "node betweenness"). They find that even the removal of a single transmission node can cause significant connectivity losses, and that load-based or degree-based removal typically has much greater impact than removal of nodes at random.

In their topological model, Al-Mannai and Lewis (2008) use a game-theoretic approach, in which the defender minimizes the total network risk, calculated as the sum of risk (vulnerability multiplied by consequence) of each component. Vulnerability is then shown as a combined vulnerability function of the attacker and the defender, both of which are dependent on how much resources they allocate. Similar to Albert et al. (2004), Lewis (2009) attempts to correlate the vulnerability of a network with its topology by considering the degree of each node in the network, speculating that networks will generally be more vulnerable to removal of higher degree-nodes. Based on the model of Al-Mannai and Lewis (2008), Lewis recommends that the defender allocate its resources to the most critical components, but notes that the attacker’s optimal strategy may therefore be to attack less critical but undefended components.

Holmgren et al. (2007) use a topological model to analyze effective strategies for defending electric-power networks against intelligent attackers. In their model, the defender can either harden components, or decrease their recovery times. They conclude that the optimal tradeoff between these two measures depends on both the defender’s total level of resources and the nature of the attack scenario. For example, in the case of severe attack scenarios that are likely to cause large consequences, they find that for their assumed parameter values, most of the available defensive resources should be allocated to recovery rather than hardening.

However, some scholars and almost all power systems engineers have pointed out the drawbacks of using topological models to analyze network vulnerabilities. In particular, Hines et al. (2010) find that power grids are generally more vulnerable to flow-based attacks that consider actual flows within the network than to attacks that consider only the topology of the network (such as the degree of each node). Therefore, we now consider flow-based models.

Salmeron et al. (2004) develop a flow-based interdiction model to protect against worst-case attacks on electric-transmission systems. The model is solved as a sequential game, in which the attacker selects an interdiction plan to maximize the cost of operating the network (including the cost of any lost loads), while the defender then operates the remaining parts of the network so as to minimize that cost. The authors solve the resulting optimization problem by a decomposition-based heuristic algorithm. Salmeron et al. (2009) improve on that algorithm, with the result that they can generate faster and better solutions for considerably larger electric-power grids.

Arroyo and Galiani (2005) reformulate the model in Salmeron et al. (2004) as a general nonlinear mixed-integer bilevel-programming problem, making it possible for the attacker and the defender to have different objective functions. Instead of the decomposition-based heuristic used by Salmeron et al. (2004, 2009), Motto et al. (2005) transform Salmeron's mixed-integer bilevel program into a mixed-integer nonlinear program using the duality theorem, and then convert this new problem into a mixed-integer linear program. Using a flow-based model, Yao et al. (2007) extend Salmeron's problem to a tri-level sequential game (i.e., a defender-attacker-defender game) in which the defender is able to anticipate the optimal attack strategy for any given network structure, and design the network accordingly. They also propose a solution procedure for the resulting game.

Bier et al. (2007) use a simple flow-based heuristic interdiction model, in which a greedy attacker interdicts the components with the maximum flow. In this model, there are three nested algorithms. First, the power flow in the network is simulated using a DC load-flow algorithm (Carreras et al., 2002) that minimizes the cost of operating the system. A greedy interdiction algorithm then identifies the most heavily loaded line, and sets its flow to zero (representing a hypothetical attack); the resulting flows in the rest of the network are then computed. Finally, a hardening algorithm identifies a set of potentially-interdicted lines to be protected, as a way of assessing the effectiveness of protection against a greedy attacker. Bier et al. (2007) obtain results similar to those of Salmeron et al. (2004), but note that hardening even a significant fraction of the transmission lines in a network may not be sufficient to dramatically diminish the unmet demand resulting from a greedy attack, concluding that hardening of components is unlikely to be cost effective.

Finally, in their flow-based model, Romero et al. (2012) study the problem of allocating fixed budget to minimize the consequences of an intelligent attack. In order to find an optimal defense strategy, they use Tabu search with an embedded greedy algorithm to simulate the attacker.

Game-theoretic models incorporate the intelligent nature of the terrorist threat, reflecting the fact that attackers can observe and investigate the potential vulnerabilities of a network. Moreover, flow-based game theoretic models enable the attacker to consider the traffic on the network when planning an attack.

However, game-theoretic models also have some drawbacks. For example, it may be unrealistic to assume that attackers are perfectly rational, and have unlimited computational ability. Another concern is the conservatism of game-theoretic models in assuming the attacker will maximize the consequences of an attack (Ezell et al., 2010), which will result in defending against only the most-severe attacks, and may therefore leave the defender vulnerable to less severe attacks.

Moreover, with the exception of Hines et al. (2010), existing game-theoretic and quasi-game-theoretic models also do not address the impact of cascading failures. In fact, the only model we have identified that addresses both cascading failures and restoration times is Anghel et al. (2007), which does not include game-theoretic representation of attacker behavior. In the following sections, we review models of cascading failures and restoration times respectively.

3.2. Modeling Cascading Failures in Electric Power Networks

Even small attacks can have a catastrophic impact on a system if there is a potential for cascading failure. Mili et al. (2004) define the events in a cascading failure as follows:

> The triggering event is a short-circuit that occurs on one of the transmission lines of the system. ...the short-circuit current is sensed by a certain number of relays located within the region of influence of the fault. …each of these relays may unnecessarily open an unfaulted line if it suffers from a hidden failure. …Consequently, the power that used to pass through the tripped lines finds its way through other links in the network, which in turn may overload some of them. …this sequence of line tripping followed by line overloading may propagate throughout the network until either the line overloading vanishes or the stability limits or voltage collapse limits are reached.

Pure topological models inherently cannot deal with cascading failure, because cascading failure depends on links or nodes being overloaded beyond their capacity, not just on the topology of the network. Thus, there are two ways to represent cascading failure in a network. One is to try to infer which components might experience high flows from their topological position in the network, while the alternative approach is to model the flows explicitly. As a result, we classify the cascading models in the literature into two broad categories: topological models (generally with inaccurate hypotheses of how flow works); and more rigorous flow-based models.

Moreover, models of cascading failure can also be categorized as deterministic (where failure of an overloaded component is assumed to occur based on a deterministic condition, such as load exceeding capacity by a given percentage), or probabilistic (where failure of an overloaded component is assumed to occur at random). In either case, failure of overloaded nodes or arcs has the potential to result in cascading failures by causing other system components to become overloaded. We first review deterministic models of cascading failure, and then consider probabilistic models of cascading failure.

3.2.1. Deterministic Models of Cascading Failure

In their topological model, Albert et al. (2004) simulate cascading failures deterministically by removing the ten nodes with the highest loads, recalculating the estimated loads, and repeating the process until the load shed is at least 60%. They note that removal of only 5% of the nodes in a system using this algorithm can result in failure of almost the entire system. Moreover, they find much greater losses using this cascading-failure algorithm than using simpler load-based or degree-based algorithms that remove nodes based on their initial characteristics, without recalculation of loads.

Crucitti et al. (2004) develop a topological model that uses “the total number of most efficient paths” through a node as an indicator for the load served by that node. Cascading failure is assumed to occur when the load served by a given node is more than its predefined capacity, leading to recalculation of the loads at each remaining node, which could cause even more overloaded nodes. The authors find that under some circumstances, failure of even a single node can lead to a total blackout, especially if the original failed node had a high estimated load.

Zhao et al. (2004) develop a topological model to analyze the vulnerability and tolerance of complex networks to cascading failures. In their model, the load carried by each node is approximated by the number of shortest paths passing through that node. The capacity of the node is in turn assumed to be proportional to the original load (for example, 20% more than the original load). The node with the largest number of arcs (i.e. the highest degree node) is assumed to be attacked, leading to a new set of shortest paths. At that point, the node with the highest number of shortest paths is assumed to fail if it exceeds its capacity, with this failing repeated a predetermined number of times. Like Crucitti et al. (2004), Zhao et al. note that disabling one or a few nodes can result in a complete blackout through cascading failure, even if the nodes of the network have relatively high capacities.

Kinney et al. (2005) model the power grid as a weighted graph. In their model, cascading failures are represented dynamically. As in Zhao et al. (2004), the number of paths through each non-disabled node (i.e., the node-betweenness) increases as breakdowns occur (since other nodes are no longer usable), until the capacity of a node is exceeded, resulting in its failure. The authors assume that after a cascading failure, a previously overloaded component has the possibility of working again if the load goes below the capacity of the node. This study highlights the potential severity of small attacks targeted at nodes with either high node-betweenness or high degree, and finds that losing even a single transmission station may reduce the capacity of a network by up to 25%.

Wang and Rong (2009) develop a topological model to analyze the robustness of the U.S. power grid to two different types of attacks: attacks on the nodes with the highest loads; and attacks on the nodes with the lowest loads. They also develop a new method to measure and redistribute the load levels rather than the commonly used node-betweenness measure. According to this method, each node is assigned a predetermined load level, and these load levels are not necessarily the same. If a node is attacked, its load is distributed to its neighboring nodes proportional to their loads. In particular, Wang and Rong try removing loads in both ascending and descending order of load. Surprisingly, they find that when the initial load on the system is small enough, attacks on the least heavily loaded nodes can actually be more harmful.

Dueñas-Osorio and Vemuru (2009) use a node-betweenness measure to estimate the load flows in a network, and analyze the impact of the initial network design on the potential for cascading failure. They conclude that increasing the capacity of the network does not always increase its robustness to cascading failures, and that other types of design changes (such as reducing congestion, making the network more decentralized, or increasing the number of alternative routes between any given origin and destination) can be more useful.

Buldyrev et al. (2009) develop a topological model to analyze the impact of cascading failure in two interdependent networks, such as an electric-power network and a communication network that depends on it. The model randomly removes a fraction of nodes in one network, and assumes cascading failure of the corresponding nodes in the other network (together with the edges that connect the failed nodes in the two networks), which can in turn cause new failures in the original network, and so on. The process continues until there are either no more edges to remove or no more nodes to fail. The authors then examine how much of the supporting network must be protected so that the disabled nodes constitute only a small portion of both networks. The authors find that networks with a more variable degree distribution (i.e., with some high-degree nodes and some low-degree nodes) are generally less robust to random attacks, because failure of high-degree nodes can cause more damage.

As noted above, topological models estimate the flows in a network based on the inherent structure of the network. However, an electric-power system may experience different loads at different times depending on the system characteristics. Therefore, rigorous flow-based models have been developed to simulate how flows within the system change after some components have been disabled.

The Critical Infrastructure Protection Decision Support System (CIP/DSS) includes a sub-model specifically for electricity systems. Jointly developed by Argonne National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratories (Bush, 2005) for use by government and industry, CIP/DSS represents the functional dependencies within and among various infrastructure systems as flows, and then simulates the dynamics resulting from these dependencies. However, it is extremely detailed and computation-intensive, and therefore may not be practical for industrial use.

Similarly, the Critical Infrastructure Modeling System (CIMS) was developed by Idaho National Laboratory in 2005 to identify interdependencies among various infrastructure sectors. CIMS uses discrete-event simulation to help visualize cascading failures, and to explore the possible consequences of infrastructure interdependencies. Its main purpose is to conduct “what-if” analysis to understand the vulnerabilities of infrastructure systems (Dudenhoeffer et al., 2006).

Other models are specifically designed to simulate flows in power networks. Ni et al. (2003) develop an online flow-based model for use by operators to decide when to alleviate stresses on a transmission network. The model deterministically removes circuits if they pass the emergency-overload limit a specified number of times (e.g., once or twice), and then recalculates all flows.

Transmission Reliability Evaluation of Large-Scale Systems is a risk-based model that simulates cascading process as based on some predetermined initial events, in order to identify and rank critical cascading scenarios based on their severity and likelihood. The methodology has been used in transmission-system enhancement projects as a prioritization tool (Hardiman et al., 2003).

In their flow-based model, Zima and Andersson (2004) calculate the impact of line outages on the flows in the remaining lines, and cascade any overloaded lines deterministically after each calculation. The authors also calculate the minimal changes and/or load shedding needed to mitigate cascading failure.

Hines et al. (2010) develop a flow-based model to calculate the impact of cascading failures on blackout size. In this model, after an initial failure, each neighboring component is assumed to be removed from service if its flow exceeds 50% of its capacity for more than five seconds; power flows are then recalculated. To our knowledge, this is the only flow-based model addressing adversarial threats that are explicitly designed to cause cascading failures, although other models could also be used for this purpose.

Despite the development of many deterministic models of cascading failures, cascading failures have historically been considered a major unsolved problem in complex networks such as electricity systems, since it has proven difficult to determine exactly where and when cascading failures will occur. In particular, deterministic flow-based models are incapable of considering the hidden unidentified failures that may lead to cascading failure, since by definition such latent failures are unobservable. Therefore, we now consider the probabilistic models of cascading failure.

3.2.2. Probabilistic Models of Cascading Failure

Since attempts to replicate the physics of what goes on in a network have not been particularly successful, some authors have proposed using probabilistic approaches to account for the difficulty of predicting cascading failures. For example, in their topological model, Liao et al. (2004) compute the probability that a random outage will produce a cascading failure of a certain size, conditional on an assumed set of hidden failures and network stress levels.

Based on historical data, Mili et al. (2004) estimate the likelihood of cascading failure for a given type of relay based on the percentage of relays of that type involved in past cascading failures.

They then use the resulting likelihoods to calculate the probability of system failure using event trees.

In their flow-based model, Chen et al. (2005) assume that an overloaded line is more likely to fail in its first exposure to overload than in subsequent exposures. This is consistent with the idea that cascading failures result from preexisting hidden faults.

Dobson et al. (2001), Carreras et al. (2004), Dobson et al. (2007), and Newman et al. (2011) propose a probabilistic flow-based model in which cascading failure occurs with some probability when one or more lines are at or near their maximum capacities. Their model has two intrinsic dynamics, slow and fast. The slow dynamics represent load growth and response to blackouts on a scale of days, months, or years. On each day of the simulation, the loads of the network are assumed to change by a factor of $\lambda,$ where λ is a uniform random number between λ_{min} and $\lambda_{max},$ with mean value $\bar{\lambda}$ larger than one. Transmission-line capacity is also assumed to increase in response to blackouts.

The fast dynamics represent the possibility of cascading failures on a scale of seconds to minutes. The assumption is that even though disruptions can happen at any time, they are more likely to happen at or near times of peak load, when lines are highly stressed. Each overloaded line is assumed to fail with a specified probability $p,$ after which loads are recomputed, with the process continuing until there are no more overloaded lines.

In order to model how the transmission lines cascade and predict the total number of line failures, Dobson et al. (2005) and Dobson and Carreras (2010) propose using a probabilistic branching process, the parameters of which are generated through observed transmission-line failures in the past. Dobson et al. (2011) test the closeness of their predicted distribution (obtained through their branching-process model) with the empirical distribution of the number of transmission-line failures of their OPA model for 118- and 300-bus systems, and obtain close results in most of the cases. Moreover, Dobson (forthcoming 2012) shows that line outages predicted by his branching process match well with 12 years of transmission-line outage data from a North American utility company.

Inspired by Dobson et al. (2001), Anghel et al. (2007) develop a probabilistic flow-based model that represents both cascading failure and the system operator’s response to disruptions. The authors analyze the optimal tradeoff between the risk of cascading failure and the losses due to intentional load shedding by the operator.

3.3. Modeling Restoration Times

Most of the models discussed above represent the estimated impact of a disruption in a static manner, as a snapshot of the system. However, system owners and operators also care about how long it will take for a system to return to normal operating conditions; likewise, intelligent adversaries may consider the likely durations of the disturbances they cause in deciding which components to target. Thus, models that consider the restoration times of failed components can give a more realistic portrayal of system risks. We again begin with topological models, and then move on to flow-based models.

In their risk-based model, Apostolakis and Lemon (2005) use restoration times as part of their consequence analysis. However, they limit their analysis to minor attacks that would involve only minimal restoration times (e.g., less than a week), so their model may not be relevant to more serious threats. In particular, they note that coordinated attacks on several locations (which would require more time to repair) may also require more complicated minimal cut sets, and therefore may not be computationally feasible to analyze in their model.

Holmgren (2006) analyzes the topological characteristics of two power-grid networks, and proposes three strategies to decrease their vulnerability to both natural hazards and planned attacks: increasing the robustness of the network to cascading failures (by adding two underground power cables); increasing the ability of the network for rapid recovery (by a 15% reduction in restoration times); and increasing both robustness and rapid recovery (with one new underground cable and a 10% reduction in restoration times). He concludes that the combined strategy would yield roughly twice as much reduction in vulnerability as could be expected from either of the individual strategies. However, he also notes that a more realistic consequence analysis would require the use of a flow-based model.

Therefore, we now consider flow-based models involving restoration times. CIP/DSS simulates the impact of disruption over time using system dynamics (Bush, 2005), and can be used to analyze how quickly a system would recover based on various recovery scenarios (e.g., with two repair crews versus three).

Salmeron et al. (2004) weight the importance of each component in their model by the average time required for repair or replacement of that type of component, and use this information in anticipating which components would be most attractive to attackers. However, they do not explicitly simulate changes in system performance over time as a result of restoration efforts after an attack.

Anghel et al. (2007) model the restoration time of a transmission line as being equal to a minimum constant, plus an exponentially distributed additional delay time. The authors simulate the resulting behavior of the system over time, and analyze the optimal level of load shedding during the period before system restoration.

The flow-based model of Romero et al. (2012) simulates restoration process in four stages: (1) all components (and those close to attacked components) have no flow; (2) all lines are repaired except for the ones that are connected to substations (3 days); (3) all substations are repaired except for damaged transformers (15 days); and (4) damaged transformers are replaced by the spare ones (32 days).

3.4. The Summary

Overall, we prefer game-theoretic methods of vulnerability analysis to rating-based or risk-based methods, since game-theoretic methods aim to represent the behavior of a strategic threat. In particular, such models can represent threat as a function of the vulnerability of the network to an attack, and the possible consequences of a successful attack.

In addition, while pure topological models may provide some insights into the structural changes necessary to achieve a less vulnerable system, modeling the flows within the network is critical to understanding what happens in the case of a disruption. Flow-based models are therefore realistic.

Cascading failures can be critical to understanding the response of complex systems that operate at or near their capacity. Hence, representing the possibility of cascading failure in electric power networks can help in identifying the most critical components. Exact modeling of the dynamics of cascading failure is not achievable at present; however, recent probabilistic approaches to modeling of cascading failure may provide a practical solution to this problem. Finally, we believe it is useful to incorporate restoration times into models of vulnerability, to fully represent the overall impact of an attack; and more importantly, we can capture the attacker's strategies in case they consider the restoration times of the components in their objective function.

Table 3.1 summarizes the literature on vulnerability analysis of electric power networks. We classify the models based on the vulnerability method they adopt, the model they use for modeling of cascading failure (if any), and whether they incorporate restoration times.

As can be seen from the Table, most models only satisfy one or two of our criteria. Anghel et al. (2007) develop a risk-based flow model that represents cascading failure probabilistically and includes restoration times, but, they do not consider the behavior of a strategic adversary. Salmeron et al. (2004, 2009) and Holmgren et al. (2007) use game-theoretic models with restoration times, but do not consider cascading failure. Hines et al. (2010) use a quasi-game-theoretic approach and model cascading failure; however they do not consider restoration times. Finally, Dobson et al. (2001), Carreras et al. (2004), Dobson et al. (2007), and Newman et al. (2011) use flow-based models to model cascading failure probabilistically, but do not consider threat scenarios or restoration times.

In the following chapter, we design a game-theoretic, flow-based vulnerability-assessment model that includes both the immediate impact of an attack (in terms of cascading failure), and the long-term impact of the attack (as represented by the restoration times of components). To our knowledge, no model in literature satisfies all of these criteria.

We also hope that our model is practical enough for real-world use. Therefore, we use a simple heuristic approach, instead of a full game-theoretic analysis. We realize that heuristic methods have their pitfalls, and can give misleading results. However, not all facility owners and operators will be able to justify spending significant resources on a vulnerability assessment; operators of relatively small electric-power systems, or systems that serve cities of secondary importance, may feel that the threat of intentional attack is too small to justify a time-consuming analysis. Therefore, we develop a method that is practical enough that any good systems engineer could use it to analyze system risks and evaluate possible defensive investments to protect against intelligent threats. The use of a relatively simple heuristic method also makes it possible to add complexities (such as cascading failure) without compromising the computational tractability of the method.

Table 3.1–Comparison of vulnerability models in the literature

No	**Literature**	**Vulnerability Methods**					**Models of Cascading Failure**				**Models of Restor-ation Times**
		Score-Based Models	Risk-Based Models		**Game-Theoretic Models**		Deterministic Models		**Probabilistic Models**		
			Topo-logical	Flow-Based	Topo-logical	**Flow-Based**	Topo-logical	Flow-Based	Topo-logical	**Flow-Based**	
1	Anghel et al. (2007)			+						+	+
2	Albert et al. (2004)				+		+				
3	Apostolakis and Lemon (2005)		+								
4	Arroyo and Galiani (2005)					+					
5	Bienstock and Mattia (2007)		+								
6	Bier et al. (2007)					+					
7	Buldryrev et al. (2009)						+				
8	CARVER (Bennett, 2007)	+									
9	CIMS (Dudenhoeffer et al., 2009)							+			+
10	CIP/DSS (Bush, 2005)							+			+
11	Crucitti et al. (2004)						+				
12	Donde et al. (2005)		+								
13	Dueñas-Osorio and Vemuru (2009)						+				
14	ECIP (Fisher and Norman, 2010)	+									
15	Ezell et al. (2000)		+								
16	Ezell (2007)	+									
17	He et al. (2005)		+								
18	Hines et al. (2010)				+	+	+	+			
19	Holmgren (2006)				+						+
20	Holmgren et al. (2007)				+						+
21	Kinney et al. (2005)						+				
22	Lewis (2009)				+						
23	Lesieutre et al. (2005)		+						+		
24	Motto et al. (2005)					+					
25	OPA (Dobson et al, 2001, 2007; Carreras et al., 2005; Newman et al., 2011)									+	
26	Pinar et al. (2010)					+					
27	RAMCAP (2006)		+								
28	Romero et al. (2012)					+					+
29	Salmeron et al. (2004, 2009)					+					+
30	TRAM (2002)		+								
31	TRELSS (2003)			+							
32	Vulkanovski et al. (2009)		+						+	+	
33	Yao et al. (2007)					+					
34	**Our model**					+				+	+

4. Methodology

In this chapter, we will describe the methodology used in our model for the resilience of an electric power system against an intelligent adversary. Particularly, we describe how we simulate the network flows, cascading failure, and recovery of the system. We first go over the original model that was developed to identify which transmission lines might be targeted by a greedy attacker. Then, we will explain each extension to the original model; i.e. including components other than transmission lines, modeling both cascading failure and restoration times, and finally using the model to assess the effectiveness of various defensive investments.

4.1. Original Model

Bier et al. (2007) present a heuristic method for modeling intentional attacks on electric transmission systems. The model uses a greedy heuristic to identify transmission lines that may be attacked, and then assesses the effect of hardening those lines on the vulnerability of the system. The model is based on three nested algorithms: (a) a load-flow algorithm to determine how loads are distributed through the network; (b) a greedy algorithm to identify the lines assumed to be interdicted by the attacker at each iteration (the Max Line interdiction algorithm); and (c) a hardening algorithm.

The notation for the model is as follows:

Sets:

B	set of nodes in the network, indexed by i
F	set of transmission lines in the network, indexed by k

Parameters:

G_i	generation at node i
L_i	load supply at node i
$L_{i,\ demand}$	load demand at node i
$L_i\ (t)$	load supply at node i after iteration t of the Max Line algorithm
F_k	negative or positive power flow on line k (to reflect bi-directional flow)
$F_{k,\ max}$	maximum power flow permitted on line k (in absolute value)
F	vector of F_k for all $k\ \epsilon\ \boldsymbol{F}$
P_i	total power at node i (given by G_i - L_i)
P	vector of P_i for all $i\ \epsilon\ \boldsymbol{B}$
$W_{gen,\ i}$	cost of generation at node i
$W_{shed,\ i}$	cost of load shedding at node i
M	DC load flow matrix relating line flows F to power levels P
$k^*(t)$	index of the line with the highest absolute value of power flow at iteration *t* of the Max Line algorithm

Decision variables:

K(t)	set of lines attacked in iteration *t* of the Max Line algorithm
A	ordered set of (sets of) attacked lines ***K(t)***
A(s)	ordered set of (sets of) attacked lines after iteration s of the hardening algorithm
H	set of hardened lines

The load-flow algorithm to calculate optimal DC power-flow dispatch for electrical networks solves the following problem:

$$\sum (G_i W_{gen, i} - L_i\, W_{shed, i}) \qquad (1)$$

subject to the following constraints:

$$0 \le G_i \le G_{i, max} \qquad (2)$$

$$-L_{i, demand} \le -L_i \le 0 \qquad (3)$$

$$-F_{k, max} \le F_k \le F_{k, max} \qquad (4)$$

$$F = MP \qquad (5)$$

In the original model, the cost of generation ($W_{gen, i}$) is set to 1 for all nodes, whereas the cost of load shedding ($W_{shed, i}$) is set to 100, so that the load-flow algorithm prioritizes minimizing load shedding. The greedy interdiction algorithm identifies the transmission line(s) carrying the most load in the optimal solution to the load-flow algorithm. Note that in order to calculate line flows, Bier et al. (2007) use line admittance and voltage angles. As a result, they simulate line interdiction by setting the corresponding line admittance zero (impedance to infinity). We follow the same approach; i.e., setting the admittance zero to simulate line removal. Finally, the hardening algorithm simulates a network upgrade.

At each iteration, the method involves: (a) running the load-flow algorithm; (b) using the greedy algorithm to identify the most loaded transmission line; and (c) hypothetically interdicting or disabling that component. This process is then repeated for a specified number of times to simulate interdiction of a specified number of transmission lines. The hardening algorithm is also applied iteratively, by removing batches of hardened components from the set of candidates for interdiction and then rerunning the interdiction algorithm (either to determine the next set of components to harden, or to evaluate the effectiveness of the hardening). For a general schematic of this model, see Figure 4.1.

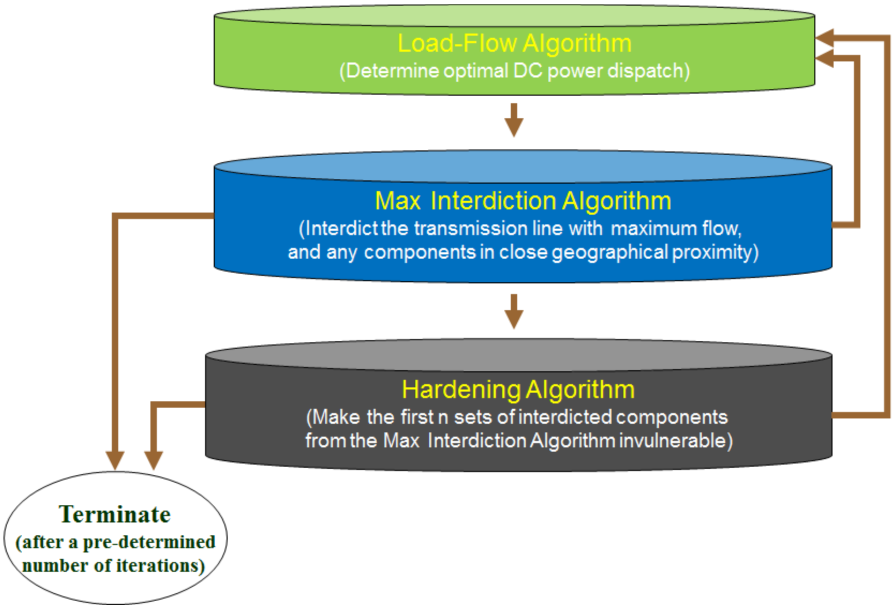
Load-Flow Algorithm
(Determine optimal DC power dispatch)
Max Interdiction Algorithm
(Interdict the transmission line with maximum flow,
and any components in close geographical proximity)
Hardening Algorithm
(Make the first n sets of interdicted components
from the Max Interdiction Algorithm invulnerable)
Terminate
(after a pre-determined
number of iterations)

Figure 4.1-The original model of Bier et al. (2007)

In order to test the efficiency of their algorithms, Bier et al. (2007) use two simple test cases, the IEEE One Area RTS-96 (24-bus system), and the IEEE Two Area RTS-96 (48-bus system, composed of two separate areas connected through three interconnections). 24-bus system is composed of 24 nodes and 38 arcs, while 48-bus system has 48 nodes and 79 arcs. In these test cases, the nodes represent either generators or load points, whereas the arcs represent transmission lines (which may also include substations and transformers).

Similar to the model of Salmeron et al. (2004), the model uses a nested optimization approach, but solves the attacker's outer optimization problem using a greedy heuristic. This approach overcomes some of the computational difficulties stemming from the bi-level optimization approach of Salmeron, and yields surprisingly comparable results. Salmeron et al. develop two interdiction plans for the 24-bus system: one that involves substations (Plan 1); and one involving transmission lines only (Plan 2). Since the model of Bier et al. does not include interdiction of substations, we compare the Max Line interdiction algorithm to Plan 2 of Salmeron et al., as seen in Figure 4.2. Salmeron et al. do not provide results on the amount of load shed for differing numbers of interdicted components, but the result of Salmeron's model for six interdicted components is remarkably close to the result of Bier et al.

Note, however, that there are some small differences in assumptions between the two models, so the comparison is not a perfect one. In particular, Salmeron et al. use different weights for the each component type to provide higher incentives for attacks on components other than transmission lines and generators. Using the same weights for each component, as done by Bier et al., the model of Salmeron et al. might perform better than the Max Line interdiction algorithm. However, Salmeron et al. state that the weights were chosen specifically to increase the impact of their interdiction algorithm, so it seems unlikely that other weights would perform significantly better.

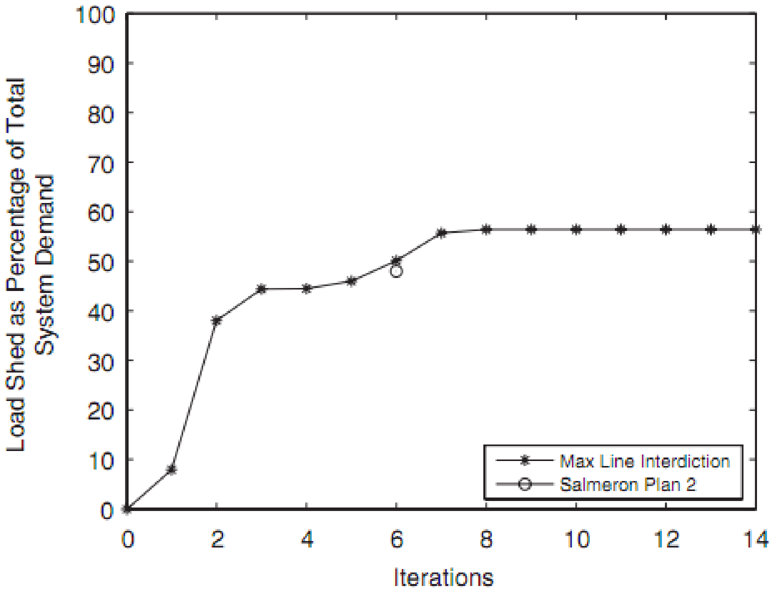

Figure 4.2-Comparison of the Max Line Algorithm to Plan 2 of Salmeron et al. (2004)

Salmeron et al. also develop an interdiction plan for 48-bus system (Plan 3). Again, the results of the Max Line algorithm and Plan 3 of Salmeron et al. are quite close; see Figure 4.3.

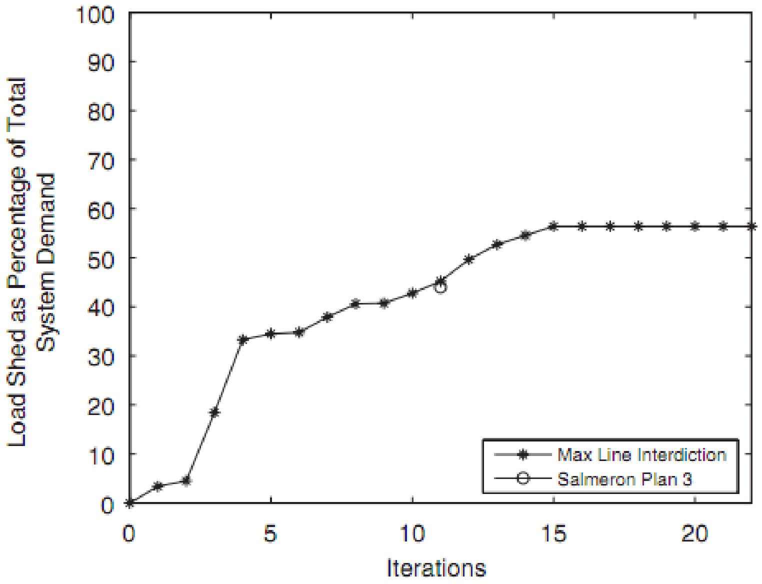

Figure 4.3-Comparison of the Max Line Algorithm to Plan 3 of Salmeron et al. (2004)

In the following sections, we will explain the methodology to expand the original model to include (1) interdicting any components (such as generators, loads, transmission lines, or transformers) rather than just transmission lines (Section 4.2); (2) modeling cascading failures as well as an modeling an attacker with or without cascading knowledge (Section 4.3); (3) modeling restoration times as well as an attacker who considers restoration times (Section 4.4). Then, we will present the complete model in Section 4.5. In addition to the original hardening algorithm, we will also consider investments that make cascading less likely or less consequential, such as adding capacities to transmission lines, or adding new transmission lines to the network, or decreasing the restoration times of critical components such as transformers. Note that, in our model, we will follow the same greedy approach used in the original model of Bier et al. (2007).

We will use the IEEE 24-bus system and 48-bus system, as well as the larger IEEE 118-bus and 300-bus systems. The 118-bus test system represents a portion of the electric power system in the Midwestern U.S. as of December 1962. This system has 118 nodes (including 99 load-serving nodes) and 186 transmission lines. The 300-bus system was designed by the IEEE Test Systems Task Force in 1993 as a standard test case, and has 300 nodes and 411 transmission lines.

4.2. Modeling Attacking Nodes as well as Arcs

Our first extension of the method in Bier et al. (2007) is to allow the attacker to interdict nodes as well as arcs, i.e. transmission lines. In particular, transmission lines are inexpensive to attack (unless underground) and difficult to protect, but also easy to repair. Thus, allowing the attacker to disable other components (generators, loads, or transformers) as well as transmission lines will not only increase the realism of the model, but will also make it possible to model attack scenarios with longer restoration times.

As can be seen from Figure 4-1, the interdiction algorithm in the original model allows the interdiction of transmission lines only. Transformers were effectively ignored (or, more precisely, treated in tandem with the transmission lines with which they were associated). Note that if transmission line is associated with a transformer, the two must have the same flow at every iteration. Moreover, generators and loads were not considered as part of the attacker's interdiction problem, hence not considered for protection. To facilitate analysis of components other than transmission lines, we now define the following notation (with changes from the previous notation shown in italics):

$\boldsymbol{T}$ set of transformers in the network, $\mathbf{T} \subseteq \boldsymbol{F}$

$\mathrm{T}_{\mathrm{k}}(t)$ power flow through transformer k at iteration *t*

$\mathrm{F}_{\mathrm{k}}(t)$ power flow through transmission line k at iteration t

$\boldsymbol{G}$ set of generators in the network, indexed by i, $\boldsymbol{G} \subseteq \boldsymbol{B}$

$\mathrm{G}_{\mathrm{i}}(t)$ power generated by generator i at iteration *t*

$\boldsymbol{L}$ set of loads in the network, indexed by i, $\boldsymbol{L} \subseteq \boldsymbol{B}$

$\boldsymbol{Z}$ set of all possible components to interdict, $\boldsymbol{Z}=\boldsymbol{B}\cup\boldsymbol{F}=\boldsymbol{G}\cup\boldsymbol{L}\cup\boldsymbol{F}$

$\mathrm{Z}_{\mathrm{x}}(t)$ power flow through component x at iteration *t*

$\mathrm{x}^*(t)$ index of the *component* with the highest absolute value of power flow at iteration *t* of *the interdiction algorithm*

$\mathbf{X}(\boldsymbol{t})$ set of *components* attacked in iteration *t* of *the interdiction algorithm*

A ordered set of (sets of) attacked *components* **X(*t*)**, |***A***|/|**Z**| being the percentage of components attacked

A(s) ordered set of (sets of) attacked *components* after iteration *s* of the hardening algorithm

n number of *components* hardened at each iteration of the hardening algorithm

H set of hardened *components*, |***H***|/|***Z***| being the percentage of components hardened

a specified proportion of components to be attacked

h specified proportion of components to be hardened

Note that if transmission line k is associated with a transformer, the two must have the same flow at every iteration:

$$\mathrm{T_k}(t) = \mathrm{F_k}(t)\ \forall \mathrm{k}, t \quad (6)$$

We assume that when the attacker attacks a transformer, the line on which the transformer is located is also disabled until the transformer is restored. Similarly, any attack on a transmission line disables the transformer until the line is restored. However, we allow transmission lines and transformers to have different restoration times.

As before, flows through the network are calculated using the optimal DC power-flow dispatch algorithm. However, since we are now modeling a larger number of component types, we can represent more types of attacker behavior. For example, the attacker could: (1) restrict attention to a particular type of component (transmission lines, generators, loads, or transformers), and attack the component of that type with the most flow; (2) simply choose the component with the most flow (regardless of type); etc.

The resulting interdiction algorithm can be summarized as follows:

0. Set $\boldsymbol{H}$=∅.

1. Set t=1, $\boldsymbol{A}$=∅, and $\boldsymbol{X(t)}$=∅ $\forall t$.

2. The DC load-flow algorithm is run with the constraint that $F_k(t)=0$ ∀k ∈$\boldsymbol{A(t)}$. Optimal power flows on transmission lines, $F_k(t)$, and through transformers, $T_j(t)$, are calculated. The unmet demand on the various buses, $L_{i,demand}$-$L_i(t)$, and the generation at the various generating nodes, $G_i(t)$, are also calculated. $Z_x(t)$ is formed using $G_i(t)$, $L_i(t)$, $F_k(t)$, and $T_j(t)$.

3. The component x*(t) to be interdicted is identified according to the assumed attacker strategy. Two options for attacker strategies are considered here:

   a. Attacking the component x*(t) with the highest absolute flow, regardless of type:

   $$\mathrm{x}^*(t) = \arg\max_x \{\, (|Z_x(t)\,|) \colon \mathrm{x} \in \boldsymbol{Z} - \boldsymbol{H}\}$$

   Note that the set $\boldsymbol{Z} - \boldsymbol{H}$ is defined as including elements of $\boldsymbol{Z}$ that are not in $\boldsymbol{H}$. If the component chosen is a transmission line that also has an associated transformer, we assumed that the attack is against transformer, since transformers have longer restoration times. As before, in the case of equality, the algorithm chooses the component with the lowest index.

   b. Attacking the component x*(t) of a specific type with the highest flow:

   $$\mathrm{x}^*(t) = \arg\max_x \{\, |G_x(t)| \colon \mathrm{x} \in \boldsymbol{G} - \boldsymbol{H})\} \text{ or}$$

$$\mathrm{x}^*(t) = \arg\max_x \{\, |L_x(t)| : \mathrm{x} \in \boldsymbol{L} - \boldsymbol{H})\} \text{ or}$$

$$\mathrm{x}^*(t) = \arg\max_x \{\, |F_x(t)| : \mathrm{x} \in \boldsymbol{F} - \boldsymbol{H})\} \text{ or}$$

$$\mathrm{x}^*(t) = \arg\max_x \{\, |T_x(t)| : \mathrm{x} \in \boldsymbol{T} - \boldsymbol{H})\}.$$

4. The selected x*(*t*) is added to **X(*t*)**. Any components in close geographical proximity to x*(*t*) (e.g., parallel lines in the same geographic location) are also added to **X(*t*)**.

   The selected components in **X(*t*)** are added to the attacked components in the network by making their flows zero for all x∈**X(*t*)**. ***A*** is set to ***A*** ∪**X(*t*)**.

5. If |***A***|/|***Z***| < *a*, then the index *t* is incremented by 1 and the algorithm returns to Step 2. Otherwise, rerun the load-flow algorithm to show the impact of the final attack.

The hardening algorithm can then be run to determine which components to harden and analyze the effectiveness of hardening:

1. The system is initialized at iteration *s*=1 with the set ***H*** empty. The set of attacked components, ***A(s)***, is empty.
2. The set of components attacked at iteration *t* of the interdiction algorithm, $\boldsymbol{X}(\boldsymbol{t}), \forall t$ is empty. The interdiction algorithm is run starting from Step 1 until the number of iterations exceeds *a*|***Z***|, and the interdicted component at each iteration, x*(*t*), are added to ***A(s)***.

3. If $|\boldsymbol{H}|/|\boldsymbol{Z}| < h$, then the first n components of ***A(s)***, x*(1) through x*(n), are added to the set ***H*** of components selected for hardening. The index s is incremented by 1, and the algorithm returns to Step 2 of the hardening algorithm. Otherwise, the algorithm ends.

As in the original model, we assume that the attacker can observe all defensive investments, and, as a result, will avoid attacking hardened components.

We will present the results with respect to first extension in Chapter 5. We will first analyze effectiveness of various attacker preferences for component types including generators, loads, transformers, and all components. This will enable us to compare our all component attacks with the original transmission lines only attacks. We will also explore the change of attacker preferences based on the cost of an attack. Then, we will compare our greedy heuristic with other attack strategies, such as degree-based and random. To measure effectiveness of our greedy hardening strategy, we will first gauge it against a static hardening strategy in which the components to be hardened are determined based on their initial flows. Then, we analyze the effectiveness of our hardening strategy against different types of greedy attackers or degree-based attacker. Finally, we will explore if hardening certain percentage of critical components (such as 2%, 5%, 10% or 30%) against a greedy attacker is effective.

4.3. Cascading Failure

In this section, we incorporate cascading failure into our model. For this purpose, we will use a probabilistic modeling approach to model cascading impact. Then, we describe how we plan to model greedy attacker on the risk of cascading failures. In particular, we consider three assumptions regarding the attacker's level of knowledge about cascading failure. An attacker with no cascading knowledge will preplan their attacks based on the initial flows. On the other hand, an attacker with cascading knowledge will be either a static attacker who preplans the attack with the knowledge of cascading algorithm (but not the real world), and a dynamic attacker who greedily attacks as the results of the previous attack unfolds, similar to the models of Salmeron et al. (2004 and 2008) and Bier et al. (2007).

4.3.1. Motivation

Electric-power networks are highly capacity-constrained systems (Amin, 2002). As a result, even small attacks can have a catastrophic impact if there is a potential for cascading failure, since some components may be overloaded in the aftermath of an attack; failure of even a subset of those overloaded components may lead to more overloaded components, and so on. To illustrate this, see Figure 4.4. Although this figure does not actually show the effects of cascading failure, it does illustrate that after even a small percentage of components has been successfully attacked, more than 10% of the transmission lines would be overloaded and potentially susceptible to cascading failure; moreover, the number of overloaded components increases as the percentage of components attacked increases. This figure shows the results for the IEEE 300-bus system, but qualitatively similar results have been obtained for the 24, 48, and 118-bus systems.

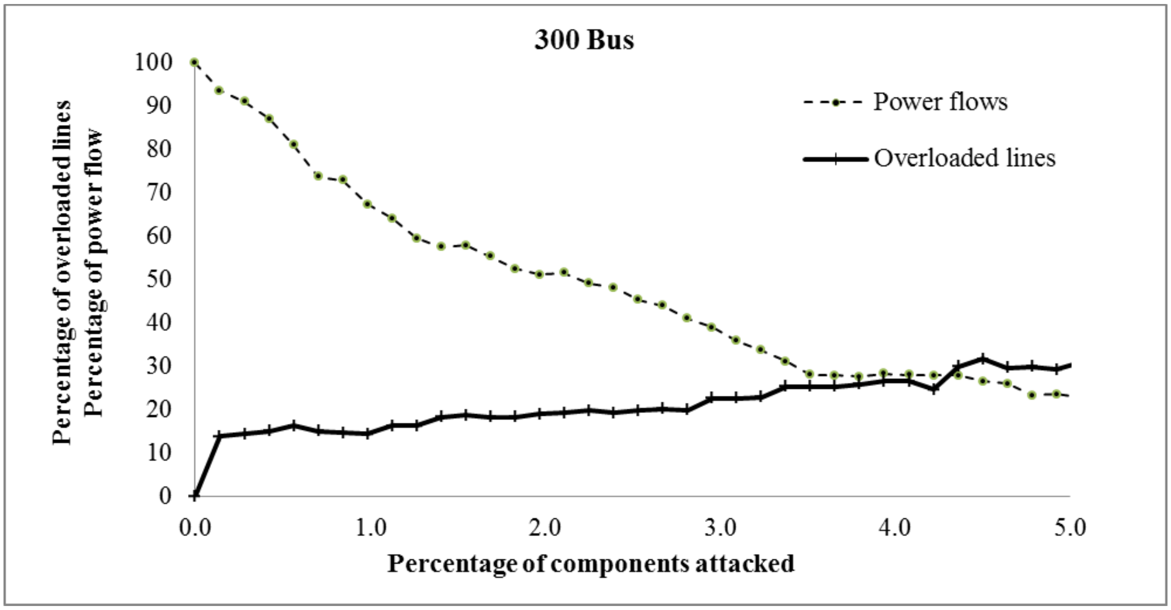

Figure 4.4-Overloaded lines as a function of the percentage of components attacked

Since cascading failure can have a significant impact, any model that aims to represent the failure behavior of electric-power networks should be capable of simulating possible cascading failures after an attack. However, most models of cascading failure are not designed to simulate the behavior of an intelligent adversary. Therefore, we hope to show the impact of an intelligent attack on the potential for cascading failure. This will also enable us to evaluate options for reducing the likelihood of cascading failure (e.g., by increasing the capacity of selected components).

4.3.2. Modeling Cascading Failure

In the literature, there are two main approaches to representing the effects of cascading failure on system vulnerability. In particular, the models of cascading failure can be either deterministic (where failure of an overloaded component is assumed to occur based on a deterministic condition, such as load exceeding capacity by a given percentage), or probabilistic (where failure of an overloaded component is assumed to occur with some probability). We choose to model cascading failure probabilistically, since deterministic models have not yet proven capable of adequately representing the hidden or unidentified failures that may lead to cascading failure. As a result, we believe that the deterministic models in the literature are not likely to perform well in our model. By contrast, the probabilistic model of Dobson et al. (2007) appears to be remarkably realistic, as illustrated by the comparison of model results with NERC blackout data (Dobson et al., 2007) and Western Electricity Coordinating Council data (Carreras et al., forthcoming 2013).

In the OPA model documented in Dobson et al. (2001), Carreras et al. (2004), Dobson et al. (2007), and Newman et al. (2011), cascading failure is assumed to occur with some probability when one or more lines are at or near their maximum capacities. As discussed before, the OPA model has two intrinsic dynamics, slow and fast. We will ignore the slow dynamics of the system, since we do not intend to model long-term changes in supply and demand as well as the changes to the network. Instead, our goal is simply to assess the vulnerability of the system in its current form.

Also, note that the OPA model was designed for random initial outages. Random outages can occur every day, and most of these outages are addressed by operators before they can lead to cascading failure. Hence, the original OPA model requires a minimum number of lines to be overloaded before cascading failure is assumed to initiate. By contrast, our vulnerability model is extended to predict the effect of carefully targeted attacks that are designed to cause maximum damage. As a result, we will assume that cascading failure can occur as the result of a successful attack on even a single carefully chosen component.

In our cascading algorithm, there are three loops. Since our algorithm is probabilistic, the outer loop consists of Monte Carlo simulation to ensure that we have a statistically reliable representation of system behavior. The middle loop determines which component is targeted in each iteration of an attack (up to a predetermined percentage of components). After each iteration of the attack algorithm, the inner loop is then run as often as necessary to identify all cascading failures caused by that attack, and to compute the resulting loads on the remaining (unfailed) components in the system.

We now introduce the following notation:

α the percentage of the unused flow of component at which it would become a potential candidate for cascading failure

β the failure probability of an overloaded component

M minimum number of replications to be performed in the simulation

$\boldsymbol{C_c(t)}$ set of overloaded components at the c^{th} iteration of the cascading algorithm in the t^{th} iteration of the attack algorithm

$\boldsymbol{C(t)}$ set of failed components due to cascading at iteration t of the interdiction algorithm

$\boldsymbol{C}$ set of all failed components due to cascading

Our modified version of the OPA algorithm proceeds as follows:

0. Set the replication index to m=1.

1. Set the attack iteration index to t=1. Also set $\boldsymbol{A}$, $\boldsymbol{C_c(t)}$ and $\boldsymbol{C(t)}$ $\forall t$ to empty.

2. Identify the component to be interdicted, x*(t), based on the attacker's assumed strategy. Add x*(t) to $\boldsymbol{X(t)}$. In the case of equality, choose the component with the lowest index. Any components in close geographical proximity to x*(t) (e.g., parallel lines in the same geographic location) are also added to $\boldsymbol{X(t)}$.

   The selected components in $\boldsymbol{X(t)}$ are removed from the network by setting the absolute values of their flows to zero for all x∈$\boldsymbol{X(t)}$. $\boldsymbol{A}$ is set to $\boldsymbol{A} \cup \boldsymbol{X(t)}$. Also set the cascading index c=1.

3. Determine the set $\boldsymbol{C_c(t)}$ of overloaded transmission line according to

$$\boldsymbol{C_c(t)}=\{\mathrm{k}:|\mathrm{F_k(t)}|\geq(1-\alpha)\mathrm{F_{k,max}}\}$$

4. If the set $\boldsymbol{C_c(t)}$ is not empty, then for $\forall \mathrm{k}\in\boldsymbol{C_c(t)}$, set $\mathrm{F_k}(t)$=0 with probability β (independently). If no flows $\mathrm{F_k(t)}$ are set to zero, go to Step 5. Otherwise, for all k ∈ $\boldsymbol{C_c(t)}$, if $\mathrm{F_k}(t)$=0, then add k to $\boldsymbol{C(t)}$. The set $\boldsymbol{C(t)}$ is also added to the set of all cascaded components, $\boldsymbol{C}$. Set the cascading index c to $c+1$, run the load-flow algorithm with the additional constraint that $\mathrm{F_k}(t)$=0 $\forall$k ∈ $\boldsymbol{C(t)}\cup\boldsymbol{A(t)}$, and return to Step 3.

5. If $|\boldsymbol{A}|/|\boldsymbol{Z}| < a$, then the index t is incremented by 1 and the algorithm returns to Step 2.

6. Increment the replication index *m* by 1. If *m*<*M*, return to Step 1.

In our analysis, we will execute 20 simulation runs, and increase the number of replications, *M*, by 10 until the standard error of the mean for the load loss is less than 3%.

To model the attacker’s knowledge of cascading failure, we will consider three possible assumptions: (1) the attacker has no knowledge of cascading failure (static attacker with no cascading knowledge); (2) the attacker has knowledge of the cascading-failure algorithm, and preplans his attacks accordingly (static attacker with cascading knowledge); or (3) the attacker can observe the results of his attacks, and at each iteration attacks the component that is the most heavily loaded after any cascading failure observed in previous iterations (dynamic attacker).

The first assumption (static attacker with no cascading knowledge) corresponds to a naïve attacker who does not take cascading failure into account, and is therefore non-conservative. In this case, the attacker preplans which component to interdict at each iteration by choosing the component that would be anticipated to have the maximum load:

$$\arg\max_{x \in \boldsymbol{Z}} \{ \mathrm{Z}_{\mathrm{x}}(t) \}$$

The second assumption (static attacker with cascading knowledge) will demonstrate how an adversary might be able to take advantage of detailed knowledge about cascading failure. In this case, the attacker is assumed to preplan the attack by prioritizing the component to be attacked in each iteration to maximize expected flow loss, which is defined to be the sum of the flow on the interdicted component and the expected load loss in the resulting first round of cascade (iteration *c*=1). This is somewhat conservative, but may be plausible, especially if the attacker has inside information. In other words, for each iteration t, the interdicted component is selected based on

$$\arg\max_{\mathrm{x}} \Big\{ \Big(|\mathrm{T_x}(t)| + \beta \sum_{\mathrm{k} \in C_1(t)} |\mathrm{F_k}(t)|\Big), \Big(|\mathrm{F_x}(t)| + \beta \sum_{\mathrm{k} \in C_1(t)} |\mathrm{F_k}(t)|\Big), \Big(|\mathrm{G_x}(t)| + \beta \sum_{\mathrm{k} \in C_1(t)} |\mathrm{F_k}(t)|\Big), \Big(|\mathrm{L_x}(t)| + \beta \sum_{\mathrm{k} \in C_1(t)} |\mathrm{F_k}(t)|\Big) \Big\}$$

Finally, the third assumption (dynamic attacker) simulates a worst-case scenario given that the attacker is greedy. It corresponds to an attacker who plans and implements the various stages of an attack in a sequential manner, after observing the system's reaction to any earlier stages. The attacker begins by attacking the most heavily loaded component, observes how any cascade unfolds, then attacks the next most heavily loaded component, and so on. Thus, instead of choosing all components to be interdicted before beginning an attack, the attacker chooses the next component to interdict after observing the impact of any previous interdiction. As a result, this assumption will enable us to compare our results with the original work of Bier et al. (2007) and our findings with no cascading failure in Chapter 5. While perhaps not realistic, we include this assumption for reasons of conservatism.

We will present our results for modeling cascading failure in Chapter 6. For each of attacker assumptions about cascading failure, we plan to analyze the impact of modeling cascading failure. We will also investigate what happens if the attacker chooses different component types for each of the attacker assumptions about cascading failure. Similar to our analysis in Chapter 5, we will also investigate whether defensive investments decrease the unmet demand. However, this time in addition to just various levels of hardening components (such as 2%, 5%, 10%, or 30% of the components), we will consider investments that are likely to decrease the likelihood or the impact of cascading failure. For this purpose, we will identify the components that are likely to cascade and either double their capacities or add new identical lines.

Doubling the capacities will enable the components to double the maximum load flows allowed. On the other hand, adding new identical lines will change the network structure. We assume that the new added lines are located away from their identical lines. As a result, when one line is attacked, the other one will not be damaged.

As noted above, the cascading algorithm uses two parameters, α and β. In our base case, α is set to 0.01, since the original OPA model considers this value to be a reasonable estimate. Similarly, following the original OPA model, we set β to 0.3. We will evaluate the sensitivity of system performance to the parameters of cascading failure, α and β. For each of these parameters, we will analyze how the extent of unmet demand changes for each of the four systems and three cascading assumptions−but considering only one attack strategy (namely, attacking the most heavily loaded component, regardless of cost), and 0% hardening. These values are chosen based on the discussions with graduate students working with Professors Ian Dobson and Jeff Linderoth.

Finally, we will also compare the results of our model to those of similar vulnerability models that account for cascading failure. To our knowledge, Hines et al. (2010) is the only other model that accounts for both cascading failure and an intelligent adversary. Their model uses various topological and flow-based interdiction algorithms for the 300-bus system (such as maximum flow, minimum flow, highest node-betweenness, and highest degree), but simulates cascading failure deterministically. Therefore, we will compare the results of our algorithm for the 300-bus system with the results of Hines et al.

4.4. Restoration Times

In this section, we will describe how we plan to model component restoration times, so that we can analyze not only the immediate impact of an attack, but the subsequent consequences until the system is fully restored. This will also enable the attacker to consider restoration times as part of their interdiction problem.

4.4.1. Motivation

Of course, the actual damage due to an attack depends on not only its size but also its duration, which can in principle vary from a few minutes to a few months. An intelligent adversary may consider this long-term impact, and target components with long restoration times. So far, only a few models in the literature have considered this as part of the attacker's decision problem; for an example, see Salmeron et al. (2004).

In Figure 4.5, we demonstrate the effect of the attacker's preferences for long restoration times on the defender's optimal strategy for a 300-bus system. (Similar results are found for the other three sample systems.) In this analysis, we allow generators and transformers to have longer restoration times than transmission lines and loads. As the attacker places more emphasis on components with long restoration times, the attack pattern changes to involve more attacks on generators and transformers. From a game-theoretic perspective, the protective strategy of the defender must then adapt accordingly by protecting more generators and transformers.

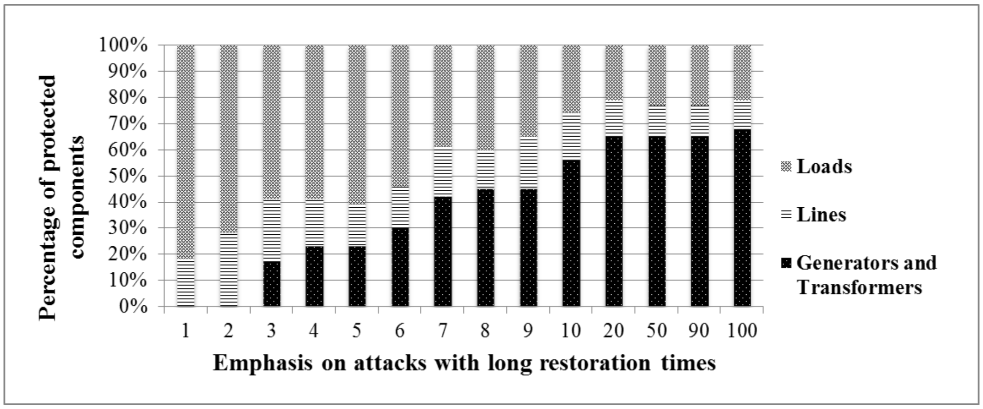

Figure 4.5- Attacker considering restoration times (300-bus system)

4.4.2. Modeling Restoration Times

In our model, we consider restoration times in two places: (1) as part of the attacker's interdiction problem; and (2) to model the system as it returns to normal operating conditions.

As part of the attacker's interdiction problem, we assume that the attacker weights components by the product of their restoration times and their loads. We now define the following new notation:

s_j restoration time for component type j, $j \in \boldsymbol{Z}$

Assuming that the attacker picks the component *x* with the largest product of load and restoration time, the index of the selected component will be given by

$$\arg\max_{x} \{ |T_x(t)|\, s_{\boldsymbol{T}} , |F_x(t)| \mathrm{s}_{\boldsymbol{F}}, |G_x(t)| s_{\boldsymbol{G}}, |L_x(t)| \mathrm{s}_{\boldsymbol{L}} \}$$

Note that if the attacker's strategy is to attack any component and chooses to attack a transmission line that is associated with a transformer, it is assumed that the restoration time is equal to that of transformer.

To model the process of recovery, we introduce the following notation:

$s_{(i)}$ ordered set of component restoration times

t_i duration of system state i, where $t_i = s_{(i)} - s_{(i-1)}$

y_i percentage of unmet demand in state i

Y percentage of energy loss

The system is assumed to remain in steady state until a particular component type is restored. Once a component type is restored, the load-flow algorithm is rerun to calculate the new optimal flow. Then, the system is again assumed to be stable until a new component type is restored; see Figure 4.6.

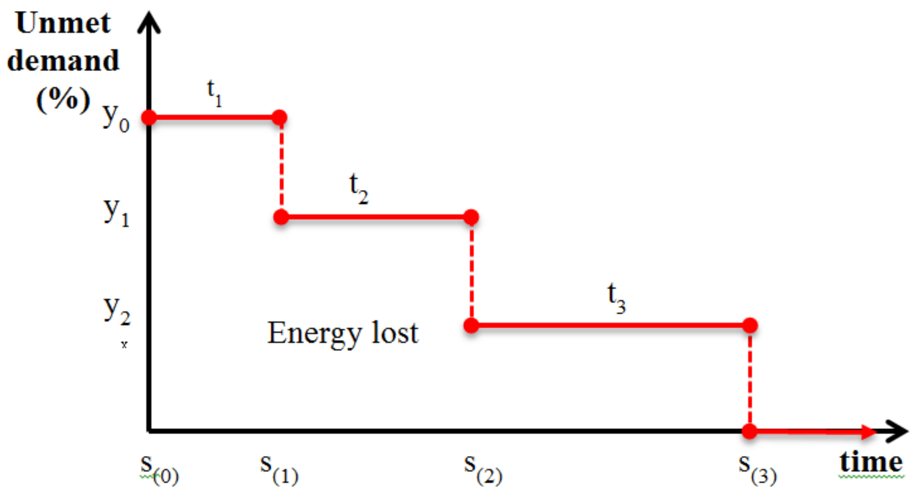

Figure 4.6-Restoration function

Unmet demand occurs when generators cannot meet some of the demand, either because the generators have inadequate capacity or because some system components are out of service as a result of an outage. We can use percentage of unmet demand as a performance metric to measure the impact of a blackout. This metric enables us to measure the immediate impact of a disruption, but not what happens after the initial impact. For this purpose, we will utilize the concept of energy lost (or energy unserved). We assume that the energy lost in a particular system state equals the product of unmet demand in that state and the duration of that state. Then we can calculate the total energy loss of the system as the sum of the energy lost in each state:

$$\sum_{i=1}^{\#\, of\, states} t_i . y_{i-1}$$

Note that this measure of energy lost is in units of time-weighted percentage loss. As a result, the total energy lost is calculated as equivalent days lost.

By default, we assume transmission lines that fail due to cascading failure have the shortest restoration times (one day), loads and transmission lines that fail due to direct attacks on those lines have longer restoration times (three days), generators take even longer to restore (15 days), and finally transformers have the highest restoration times (32 days), although other assumptions could also be used (including, for example, different restoration times for individual components.) Our restoration durations are based on Reliability Test System (RTS)-96 (IEEE, 1999) and Salmeron et al. (2004). For example, in RTS-96, the outage duration of transmission lines is 10-11 hours, whereas Salmeron et al. use three days for attacked transmission lines. As a result, we use three days if the transmission line is attacked, but only one day if the line fails due to cascading. For loads, we assume one day of restoration time, the same as transmission lines. Salmeron et al. (2004) and Romero et al. (2012) use 32 days for the restoration time of transformers based on RTS-96; we also use 32 days for transformers. Finally, RTS-96 provides no data for generators. However, Salmeron et al. (2004) and Romero et al. (2012) use 15 days (almost half the 32 days of outage assumed for transformers).

We will show the results of adding restoration times to the model in Chapter 7. We will first analyze the impact of an attacker who considers restoration times for each of the three cascading assumptions about the attacker. We will also compare the attacker who considers restoration times with various attacker types including attacker who attacks generators or transformers first (but do not consider restoration times), and a degree based attacker. Finally, we will compare the change in total energy loss for various levels and types of investments, including decreasing restoration times of all transformers by 50%.

4.5. Summary of the Comprehensive Model

Figure 4.7 illustrates our new model that combines an analysis cascading failure, restoration times, and different defensive measures. In this model, there are five algorithms: load-flow; greedy interdiction; cascading failure; restoration; and defense (including but not limited to hardening). Before running the model, we must specify the attacker’s assumed interdiction strategy, and the attacker’s level of knowledge about the cascading algorithm. The stages of the model are as follows:

1. Run the optimal load-flow algorithm.

2. Identify the interdicted components based on attacker's attack strategy and level of knowledge about cascading failure.
    a. If the attacker has full knowledge about the cascading algorithm, run one round of cascading algorithm.
    b. Interdict the component based on attacker's assumed strategy (e.g., highest loss of load by (1) not including cascading failure or restoration time; (2) including cascading failure; (3) including restoration time; or (4) including both cascading failure and restoration time).
3. For each proposed set of defensive improvements (hardening, improving robustness, decreasing restoration times, or combinations of these), if any, implement the recommended improvements (based on the attack strategy) until the desired protection level *h* is achieved, and go to Step 1.
4. Simulate cascading failures, and run the load-flow algorithm. Repeat the cascading algorithm until there are no more overloaded components left to cascade. Go to Step 2 if interdicted components are less than $a|\boldsymbol{Z}|$.
5. Restore components in order of their restoration times. After each restoration, rerun the load-flow algorithm. Continue the process until all components have been restored. Calculate the total unmet energy.
6. Repeat the simulation a specified number of times (or until a specified degree of accuracy has been obtained).

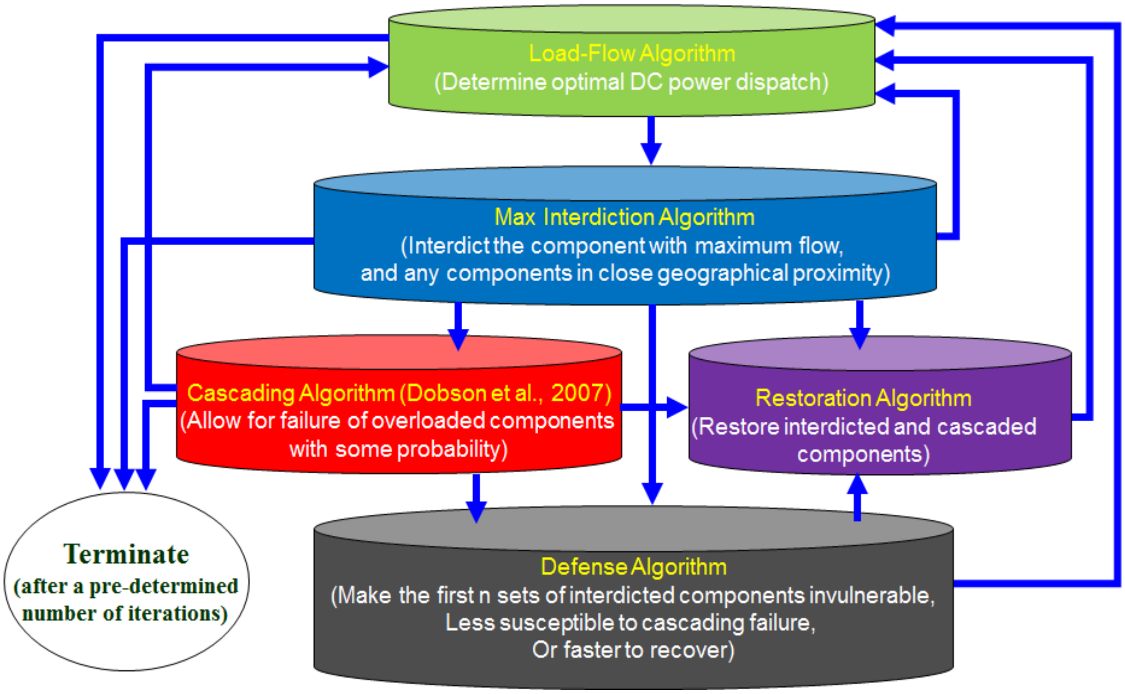
Load-Flow Algorithm
(Determine optimal DC power dispatch)
Max Interdiction Algorithm
(Interdict the component with maximum flow,
and any components in close geographical proximity)
Cascading Algorithm (Dobson et al., 2007)
(Allow for failure of overloaded components
with some probability)
Restoration Algorithm
(Restore interdicted and cascaded
components)
Terminate
(after a pre-determined
number of iterations)
Defense Algorithm
(Make the first n sets of interdicted components invulnerable,
Less susceptible to cascading failure,
Or faster to recover)

Figure 4.7 - Our vulnerability model with cascading failure and restoration times

4.6. Analyzing Effectiveness of Defensive Investments and Impact of Cascading Failure

The metric of energy lost enables us to measure the total impact of an attack including the initial impact of an attack, the impact of cascading failure immediately after the attack, and the long term impact until the system fully recovers, as demonstrated in Figure 4.8. As a result, the impact of any change in attack or defense strategies can be observed using this metric. Moreover, since the original model of Bier et al. (2007) did not include cascading failure, we can also measure the impact of cascading failure, which may be impossible to measure in real life. This idea also makes it possible to compare different investment types, including hardening components (to improve protection) versus decreasing the restoration times of particular components or reducing the impact of cascading failure through adding capacities or new transmission lines.

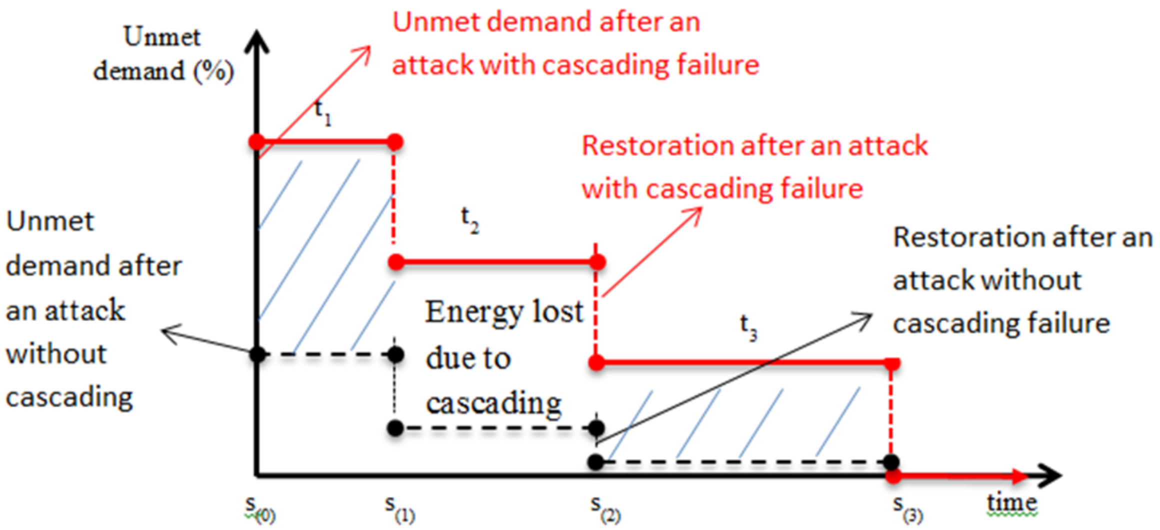

Figure 4.8 - Measuring the total impact of an attack

The base case (with cascading failure) involves no defensive investments, cascading parameters, α = 0.01 and β = 0.3, and restoration times for component types will also be default. In order to identify critical transmission lines that are likely to cascade, we will simulate the system for 100 times for a small (1% of the components) and a large (10% of the components) attack for three of the cascading assumptions for the attacker, and rank the components based on the number of time they cascade.

In order to improve robustness of transmission lines, we will increase the maximum capacity of transmission lines that are most likely to cascade ($F_{k, max}$) by 100%. We will also add new identical lines for the transmission lines that are most likely to cascade, and compare this investment with the investment of adding new capacities.

To improve the recovery time of the system, we plan to model a decrease in the restoration time of transformers, the component type with the longest restoration time. Li et al. (1999) argue that spare transformers should be considered. In 2006, federal energy regulators approved the Spare Transformers Equipment Program (STEP), which requires any participating electric utility to acquire (if necessary) and maintain a certain number of transformers so that transformers can be replaced quickly in case of a terrorist attack (FERC, 2006). Since then, many utility companies have joined the program; however, to our knowledge, there is no study on the expected reduction in restoration times when STEP is implemented. Romero et al. (2012) use 15 days for transformer replacement when a spare is available, and 32 days for repairable transformer failures with no spare. As a result, we assume that it is reasonable to decrease the restoration time of transformers by 50% through STEP, and use a 50% reduction in restoration times as an achievable level of improvement.

4.7. Conclusion

Extending our model to nodes as well as arcs will allow us to consider various types of attacker and defender strategies. For example, the attacker may take into account the cost of an attack, or restrict his attention to a specific component type. Moreover, capacity-constrained networks such as electric power systems are prone to cascading failures, so that even small attacks can cause significant damage; increasing the capacity of some components may reduce the potential for cascading failure. Similarly, different component types have different restoration times, and the attacker may take this into account in choosing which components to target; decreasing recovery times could reduce the impact of an attack. Modeling these factors more realistically will enable us to consider defensive strategies other than hardening.

The methodology presented in this chapter demonstrates the feasibility of addressing these issues. In the following chapters, we will explore how we can utilize our model and analyze some of the results.

5. Results for Attacking Nodes as well as Arcs

In this chapter, we will show how we can use our first extension of the model to analyze different attack strategies. We will also explore effectiveness of our hardening strategy.

In Section 5.1, we explore why we should include nodes as well as arcs. First, we do a comparison of the results of the previous model (which only allows the attacker to attack transmission lines) with our extended model. We then analyze the impact of attacking different component types. Later, we demonstrate how we can utilize cost of an attack. Finally, we compare our greedy dynamic attacker with degree-based and random attackers. In Section 5.2, we explore and compare the change in unmet demand for different types of hardening strategies (static versus dynamic; and degree-based versus dynamic) and hardening levels (2% to 30%).

5.1. Testing the Model for Different Attack Strategies

We test the effectiveness of the extended model for three different attack strategies. We first consider the case in which at each iteration the attacker attacks the components with the highest load, regardless of type. Figure 5.1 compares the amount of unmet demand associated with attacking the most heavily loaded components versus transmission lines only for the 24-, 48-, 118-, and 300-bus systems. Compared to attacking transmission lines only (as in the original model of Bier et al.), allowing the attacker to attack any type of component generally results in more unmet demand (except for some clearly non-optimal results for the 24-bus system, as shown in the upper left-hand portion of Figure 5.1). However, the results suggest that attacking transmission lines only does almost as well as attacking components regardless of type.

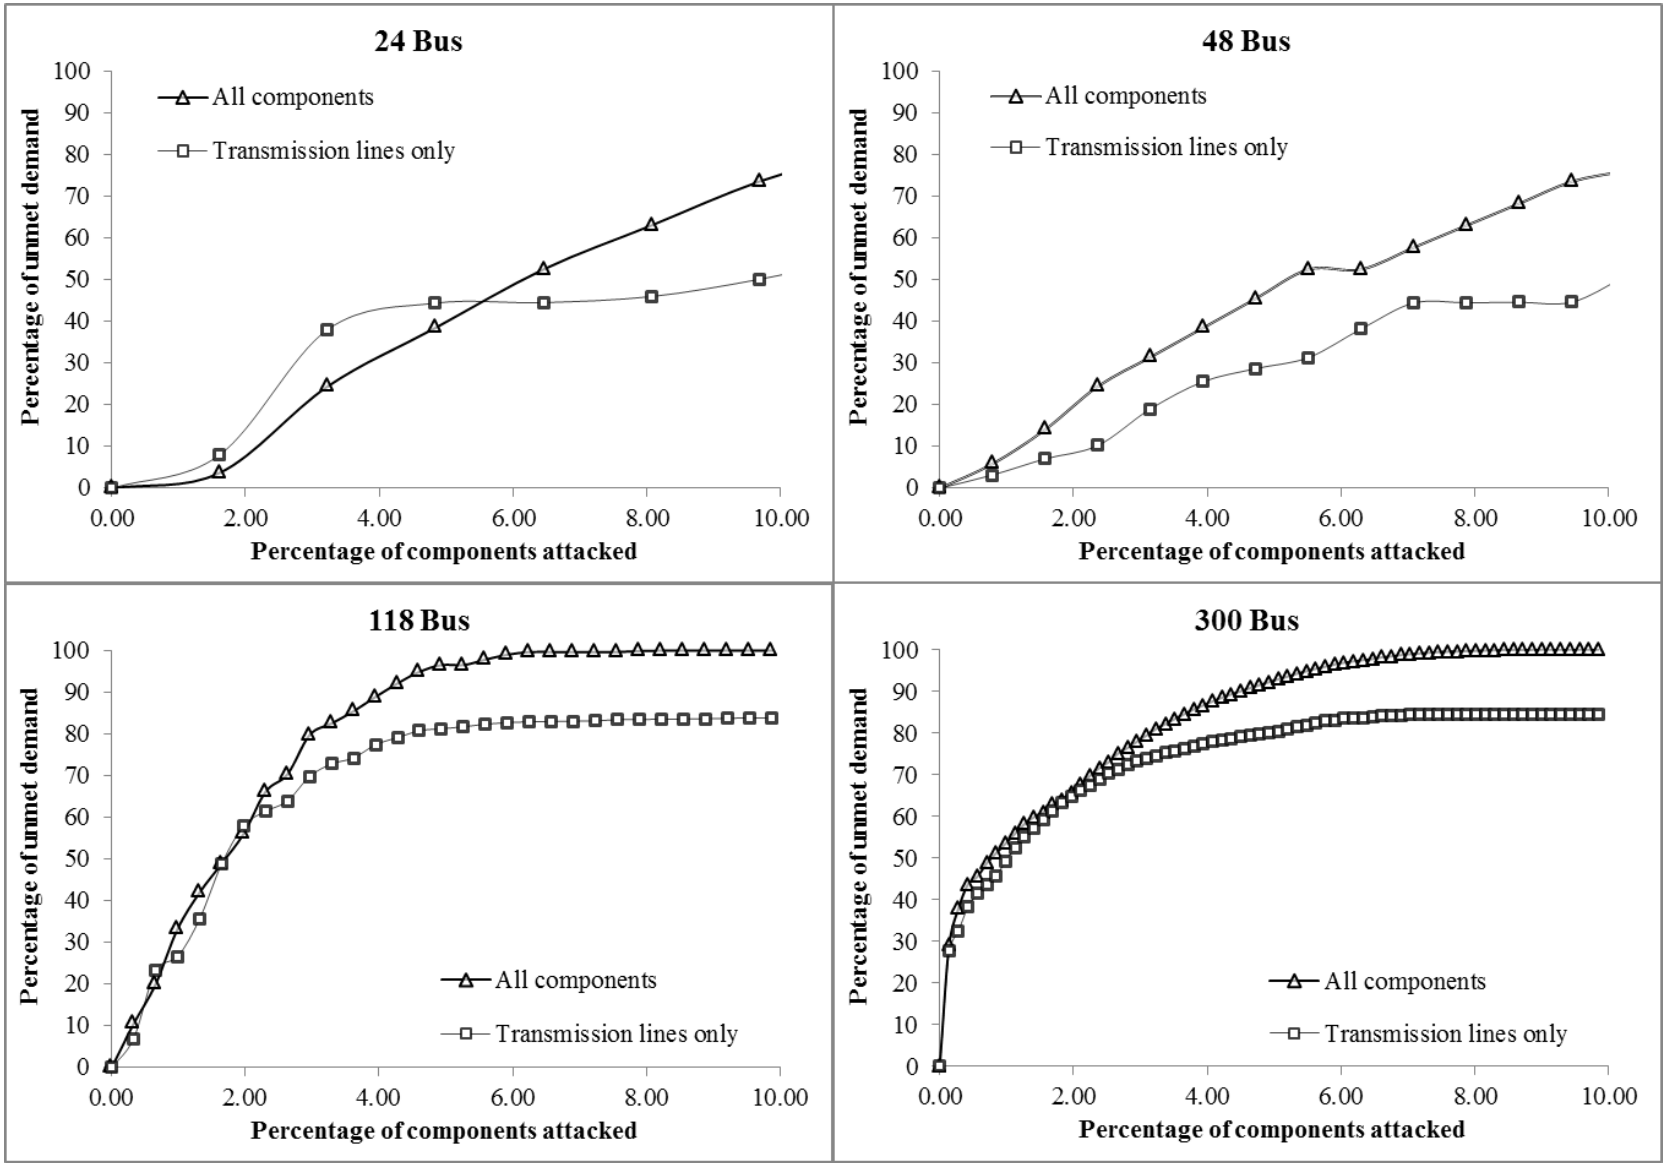

Figure 5.1–Attacking the most heavily loaded components (regardless of type) versus transmission lines only

Of course, the attacker may choose to target some other type of component, instead of transmission lines. Figure 5.2 compares five different attack strategies (any components regardless of type, generators first, transmission lines only, loads only, or transformers first) for the 24-, 48-, 118- and 300- bus systems. Note that in all but the 300-bus system, fewer than 10% of all components is transformers; thus, once all transformers have been attacked, the attacker is allowed to choose any component type. Similarly, the number of generators for 118- and 300-bus systems is less than 10%; hence the attacker chooses other components once all generators are attacked. In Figure 5.2, those attacks are separated from attacks on transformers/generators by a vertical bar, |.

In general, the results suggest that limiting attention to generators first is approximately as good as attacking the most heavily loaded components regardless of type. Restricting attention to transformers first generally does less well than the other strategies; however, attacking large numbers of transformers can still cause substantial unmet demand.

Of course, the attacker can also take the cost of an attack into consideration. For example, the attacker could target the component with the highest ratio of flow to attack cost given by:

$$\mathrm{x}^{*}(t) = \arg\max_{x} \{ |Z_x(t)/\mathrm{c}_x| : \mathrm{x} \in \mathbf{Z} - \mathbf{H} \}$$

where c_x is the cost of attacking component x.

Figure 5.3 and 5.4, respectively, show how the cost of attacking generators affects the attacker's preference for which types of components to attack, and the total unmet demand.

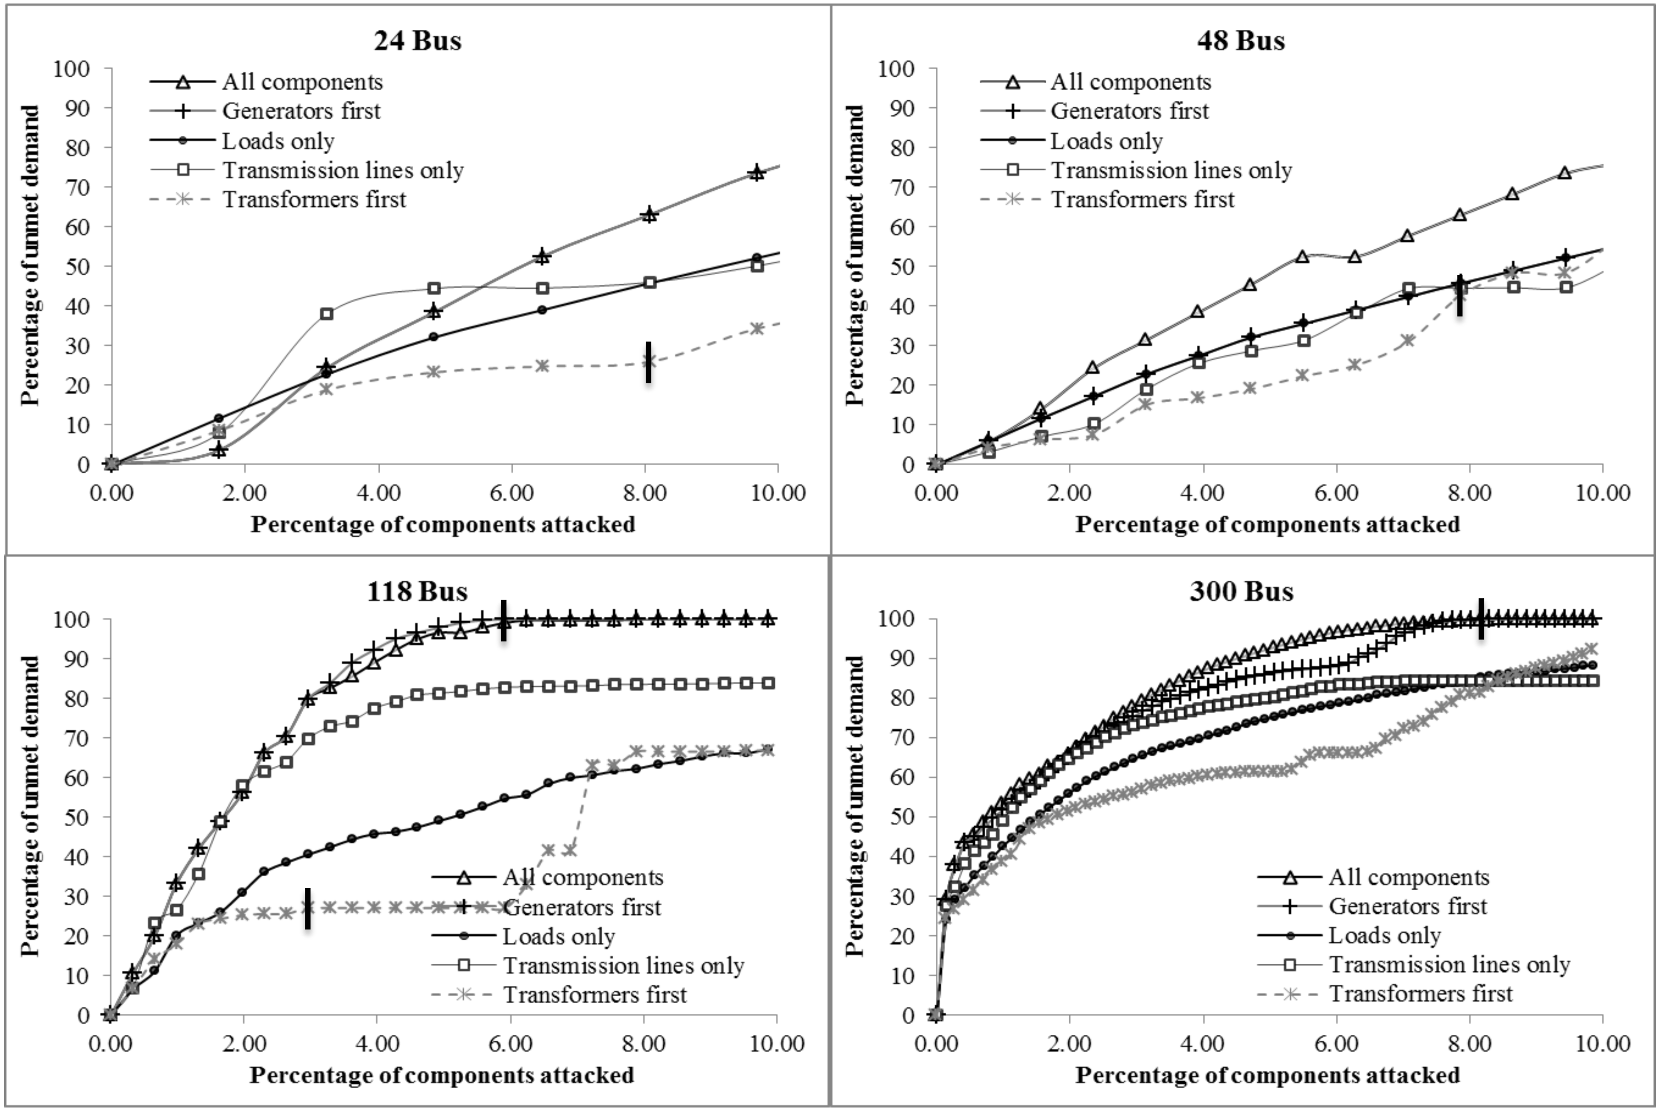

Figure 5.2– Effect of component type attacked

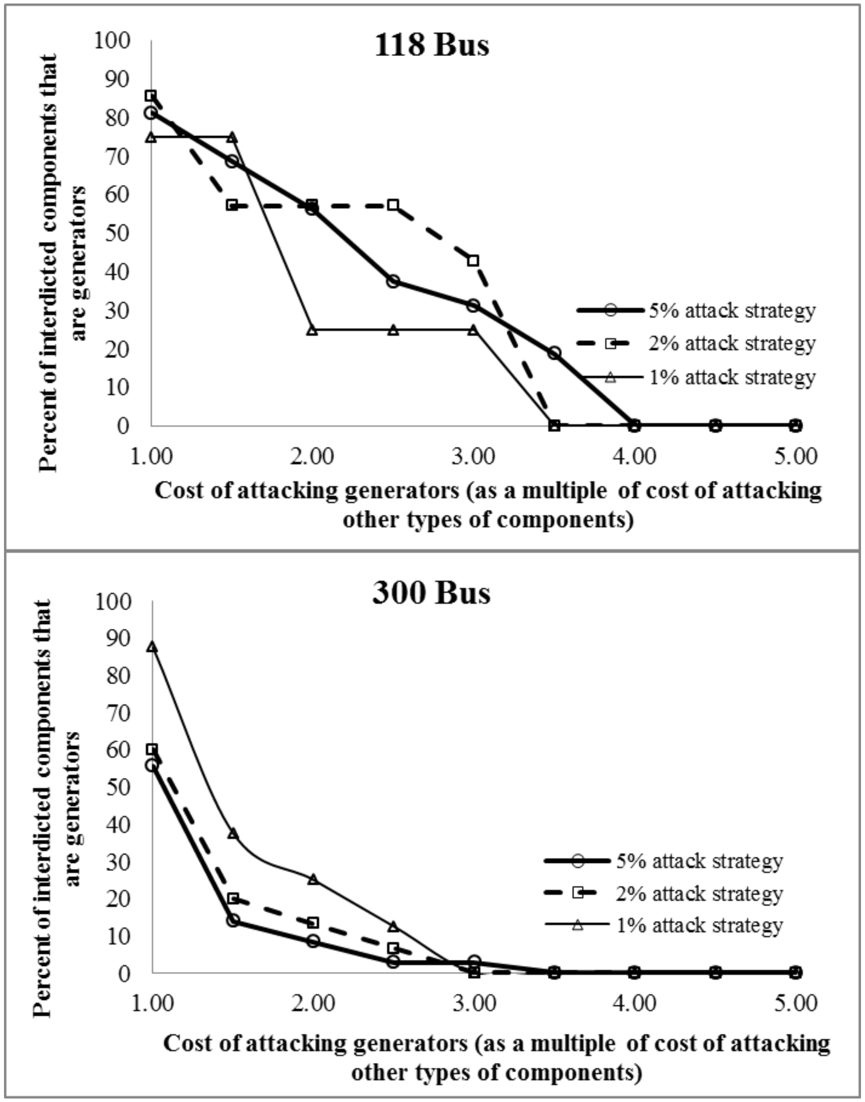

Figure 5.3-Attacker preference as a function of attacking generators

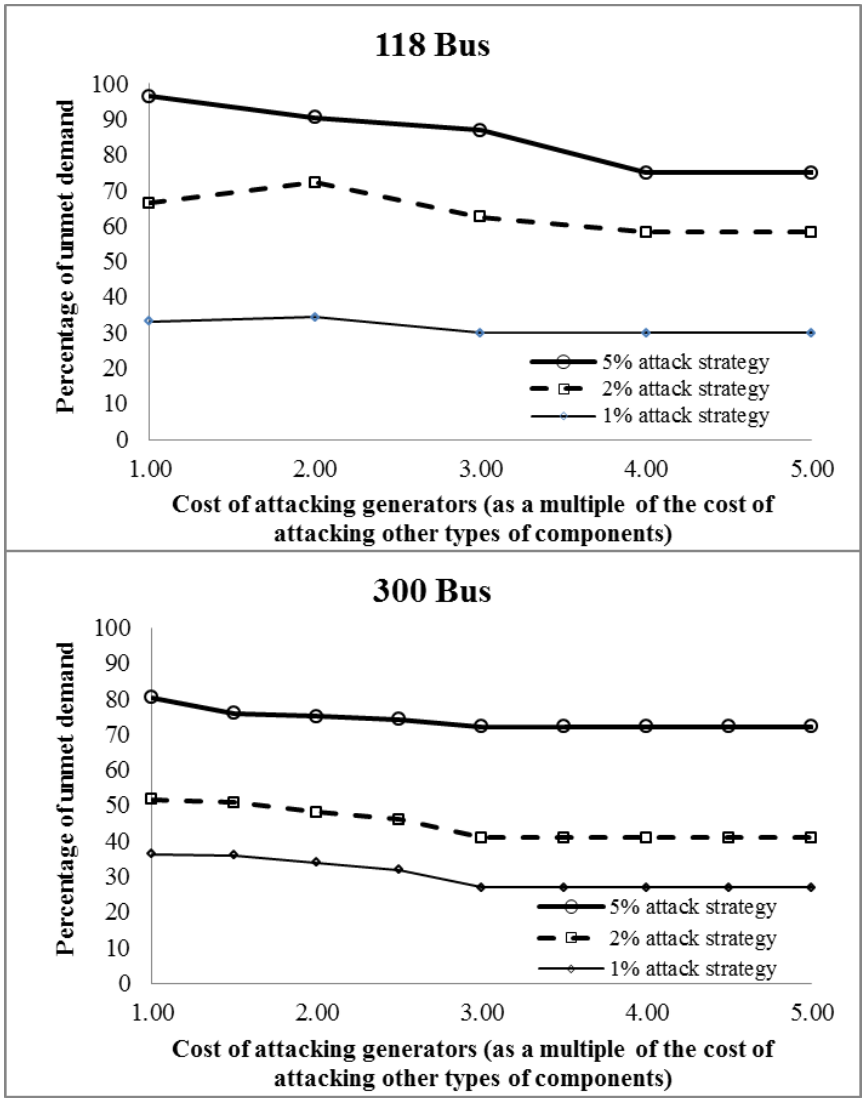

Figure 5.4-Unmet demand as a function of the cost of attacking generators

As can be seen from Figure 5.3, the results predict significant numbers of attacks on generators when those attacks are not too costly, with such attacks being deterred when attacks on generators are several times more costly than attacks on other component types. However, the results for unmet demand in Figure 5.4 suggest that deterring attacks on generators is likely to have only a negligible effect on how damaging an attack of a given size will be. Note that, in this study, we will not further analyze cost of an attacker’s preferences.

To evaluate the effectiveness of our heuristic attack strategy, we now compare that strategy with other possible heuristics. In particular, Lewis (2009) suggests that heavily connected nodes may be especially important to system operability. However, this idea has been criticized by others as a poor heuristic that may be far from optimal. Figure 5.5 compares our greedy heuristic attack algorithm with both a degree-based heuristic algorithm and a random attack strategy for 24-, 48-, 118-, and 300- bus systems. (Note that results for the random attack strategy were computed using sufficient simulation replications to achieve estimates of the average unmet demand that are accurate to within 3%).

Our heuristic algorithm does dramatically better than interdicting the highest-degree nodes. In fact, the degree-based heuristic is sometimes not even significantly better than a random attack, especially for small attacks. Thus, while our greedy heuristic is clearly not optimal (as shown, for example, in Figure 5.5), it is demonstrably much more effective than either a degree-based heuristic or a purely random attack strategy.

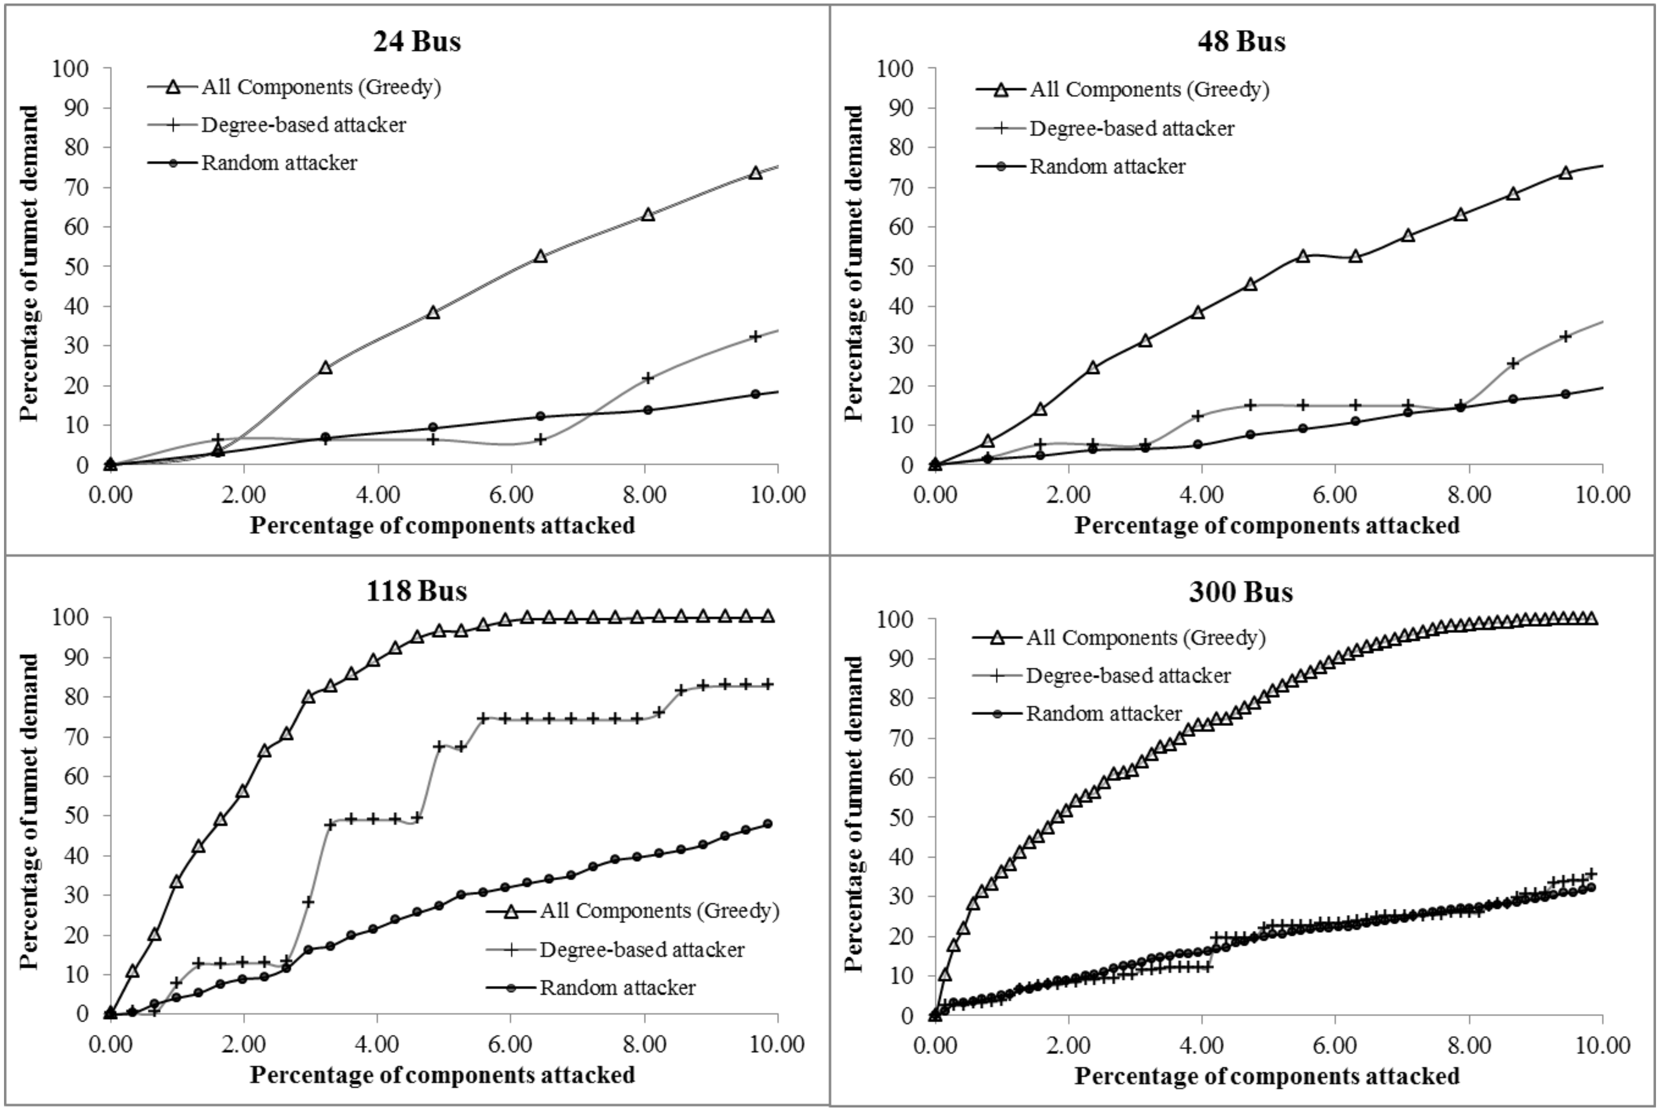

Figure 5.5-Comparison of various attack strategies

5.2. Effectiveness of Hardening

We now explore the effects of hardening components (i.e., making them invulnerable to attack). We first explore the effectiveness of our basic hardening strategy, in which we harden components regardless of type.

Figure 5.6 demonstrates the responses of the 24-, 48-, 118-, and 300-bus systems to various levels of dynamic hardening, ranging from no hardening to hardening 30% of the components. Note that as our base value for dynamic hardening, we set *n*=10 (as in the original model). As can be seen from that figure, hardening is typically of little benefit unless a large fraction of the network is hardened. Moreover, attacks can still cause significant unmet demand even after 30% of the components have been hardened. Note also that hardening a larger fraction of components does not always lead to less unmet demand, which again shows clearly that our algorithm for selecting which component to harden is not optimal. Overall, though, if our results are at least reasonably close to optimal, it seems that hardening even a significant percentage of components is unlikely to dramatically diminish the load shed from an attack, implying that hardening may not be cost effective.

Now, we will compare dynamic hardening strategy with some alternative hardening strategies, including static hardening (in which the defender hardens the components with the highest initial flows in iteration 1 of the interdiction algorithm, rather than finding the component $x^*(t)$ with the highest flow in each of the first *n* iterations), hardening the components of highest degree (as suggested by Lewis), and hardening components of a specific type (such as generators).

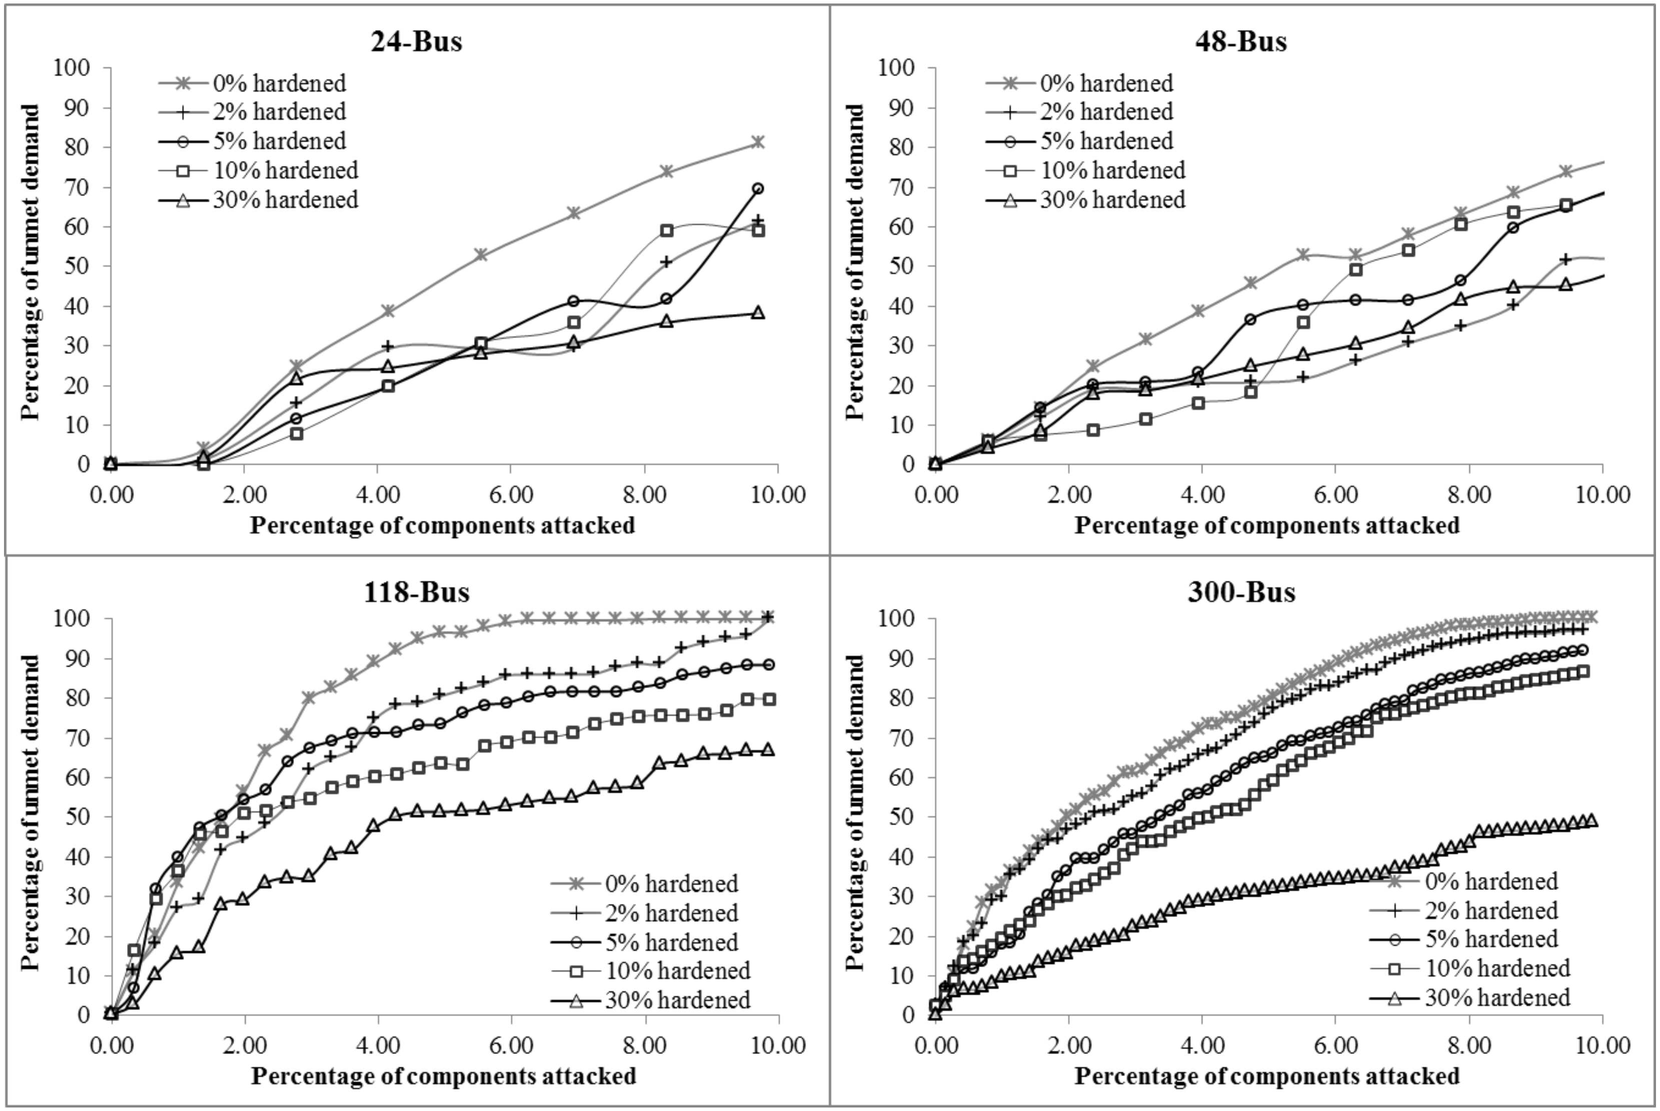

Figure 5.6-Effectiveness of dynamic hardening

First, we compare static hardening with dynamic hardening strategy. Figure 5.7 compares the effects of static and dynamic hardening for 24-, 48-, 118- and 300-bus systems, respectively, for an attack strategy that interdicts 10% of the components for 24- and 48-bus systems, and 4% of components for 118- and 300-bus systems. Note that in order to observe the effect of both hardening types, we use a smaller attack strategy for larger systems. As can be seen from those graphs, our dynamic hardening algorithm is generally much more effective than static hardening.

We now compare our all-component dynamic hardening strategy with other heuristic hardening strategies, including degree-based hardening and hardening of only a specific type of component (such as generators first, or transmission lines only). Figures 5.8 through 5.10 compare our all-component dynamic hardening strategy with other heuristic strategies when 2%, 5%, and 10% of the components are hardened, respectively. Note that similar to attacking a specific type of component scenario (in which there was no hardening), when the attacker has no more generators or transformers to attack, the attacker can choose to attack any other types of components, which is separated by a vertical bar, |.

As can be seen from the figures, virtually no hardening strategy has significant effect against small attacks. For large attacks, dynamic hardening (especially, hardening of all components) is more effective, and sometimes dramatically decreases the unmet demand.

Figures 5.11 through 5.13 demonstrate the difference between our dynamic hardening and degree-based hardening strategies when 2%, 5%, and 10% of the components are hardened, respectively; degree-based hardening often offers less improvement than our dynamic hardening.

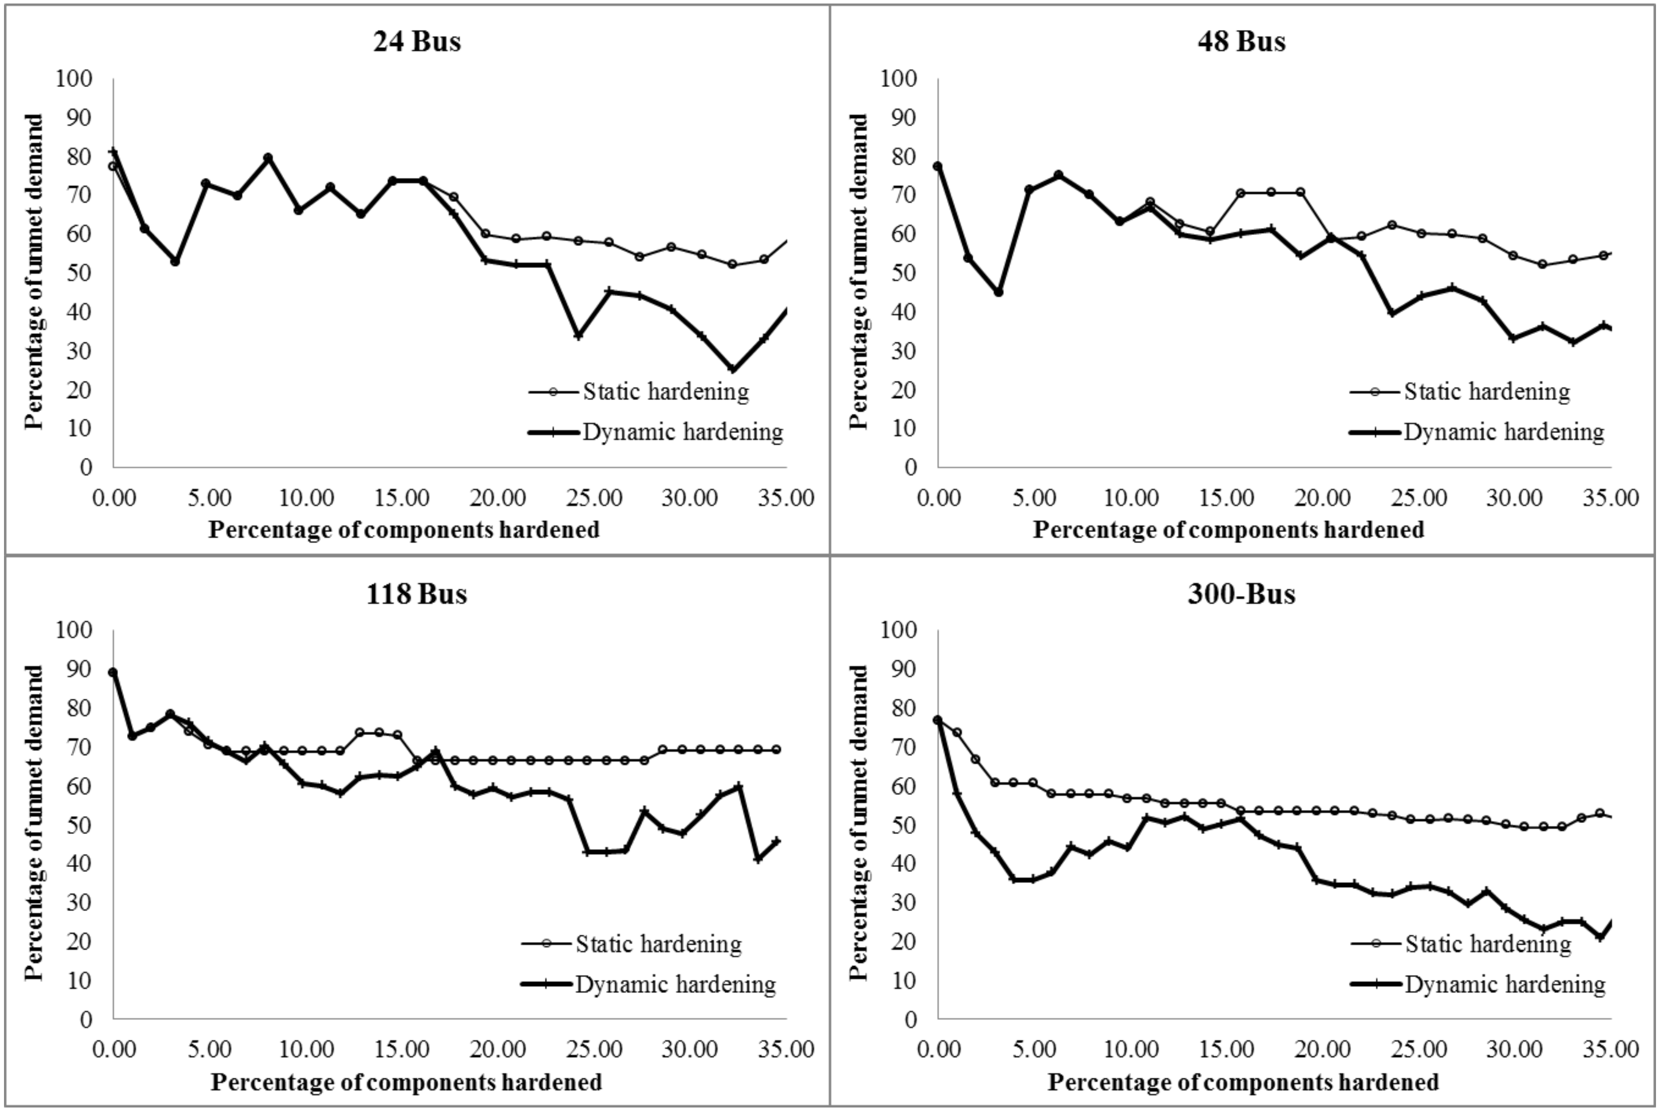

Figure 5.7-Comparison of static and dynamic hardening

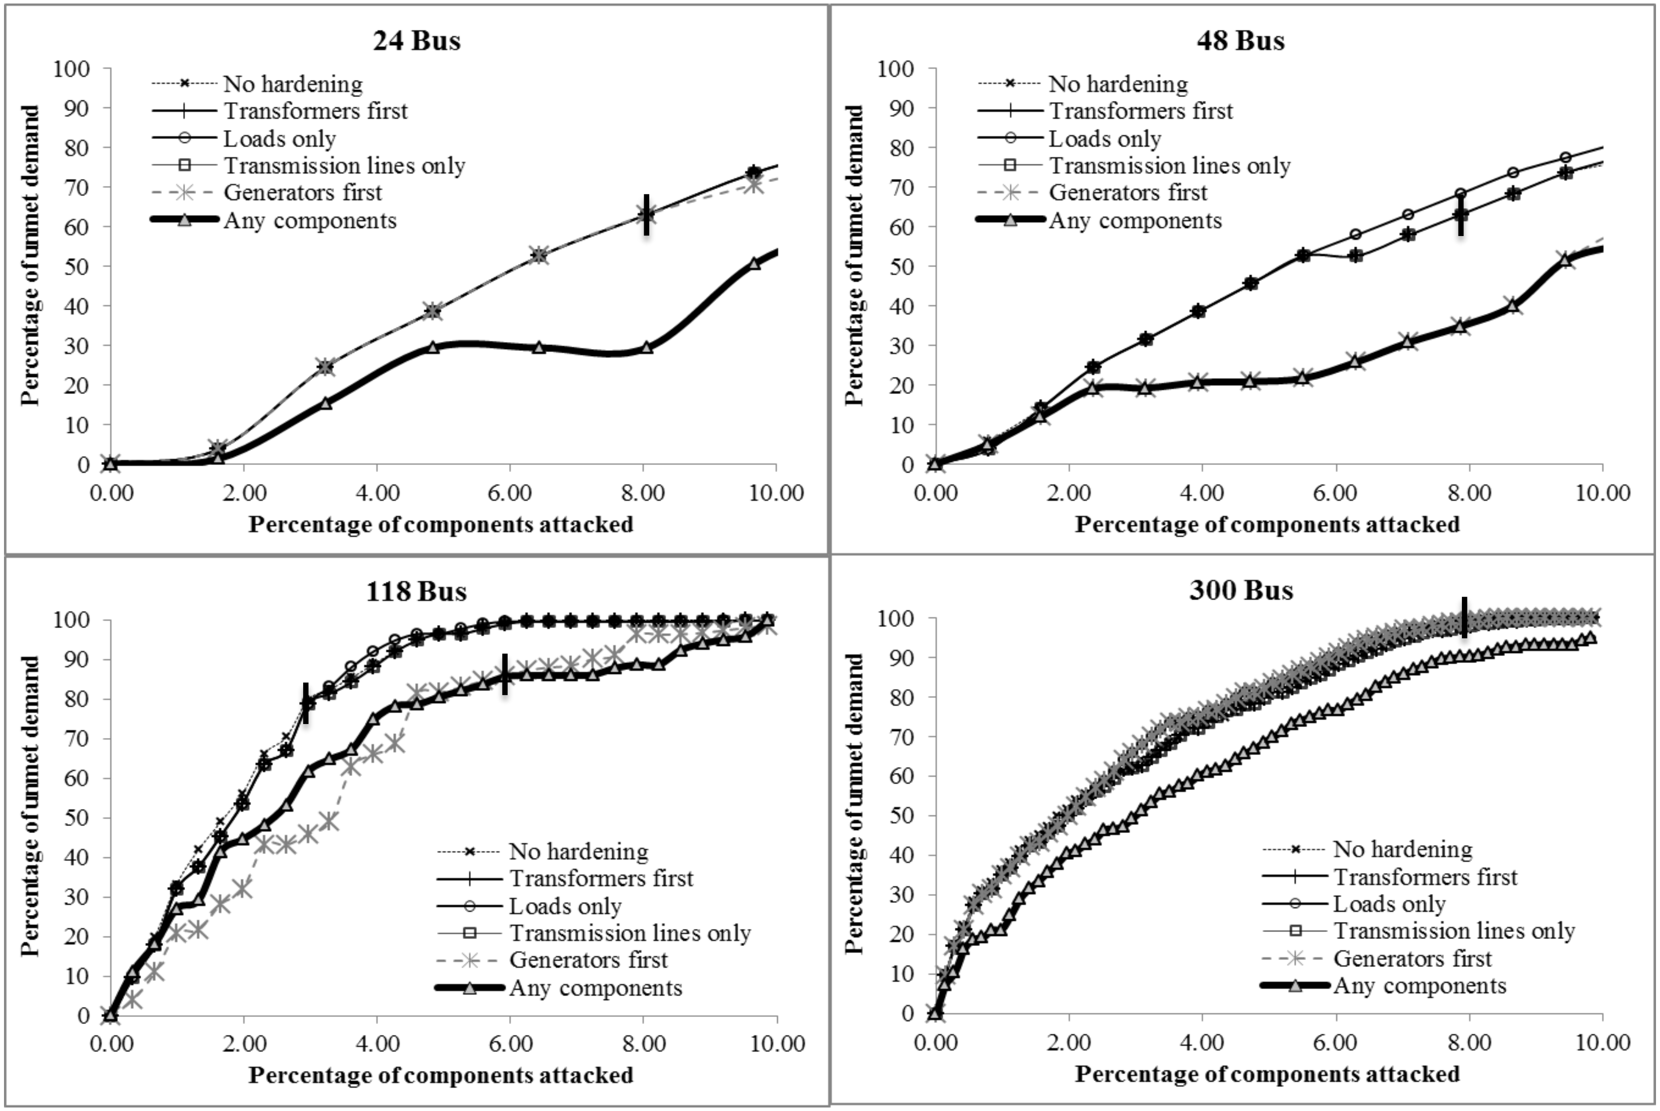

Figure 5.8- Comparison of various dynamic hardening strategies (2% hardened)

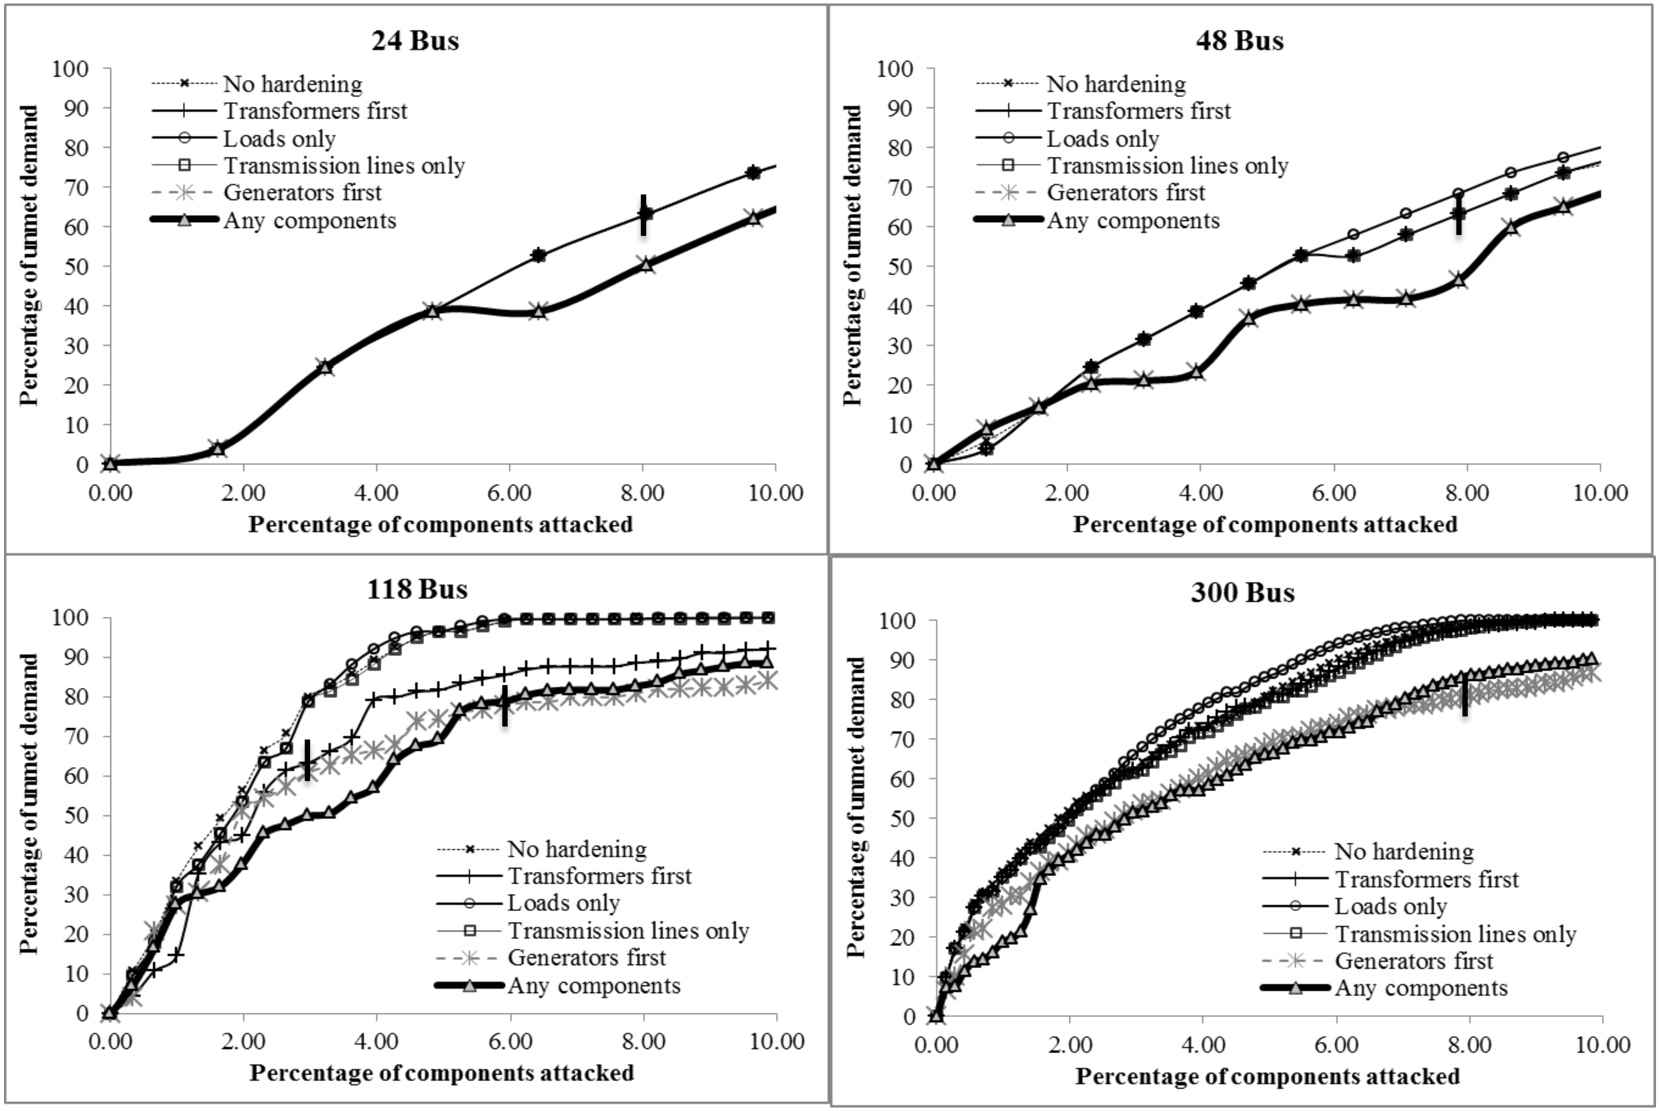

Figure 5.9- Comparison of various dynamic hardening strategies (5% hardened)

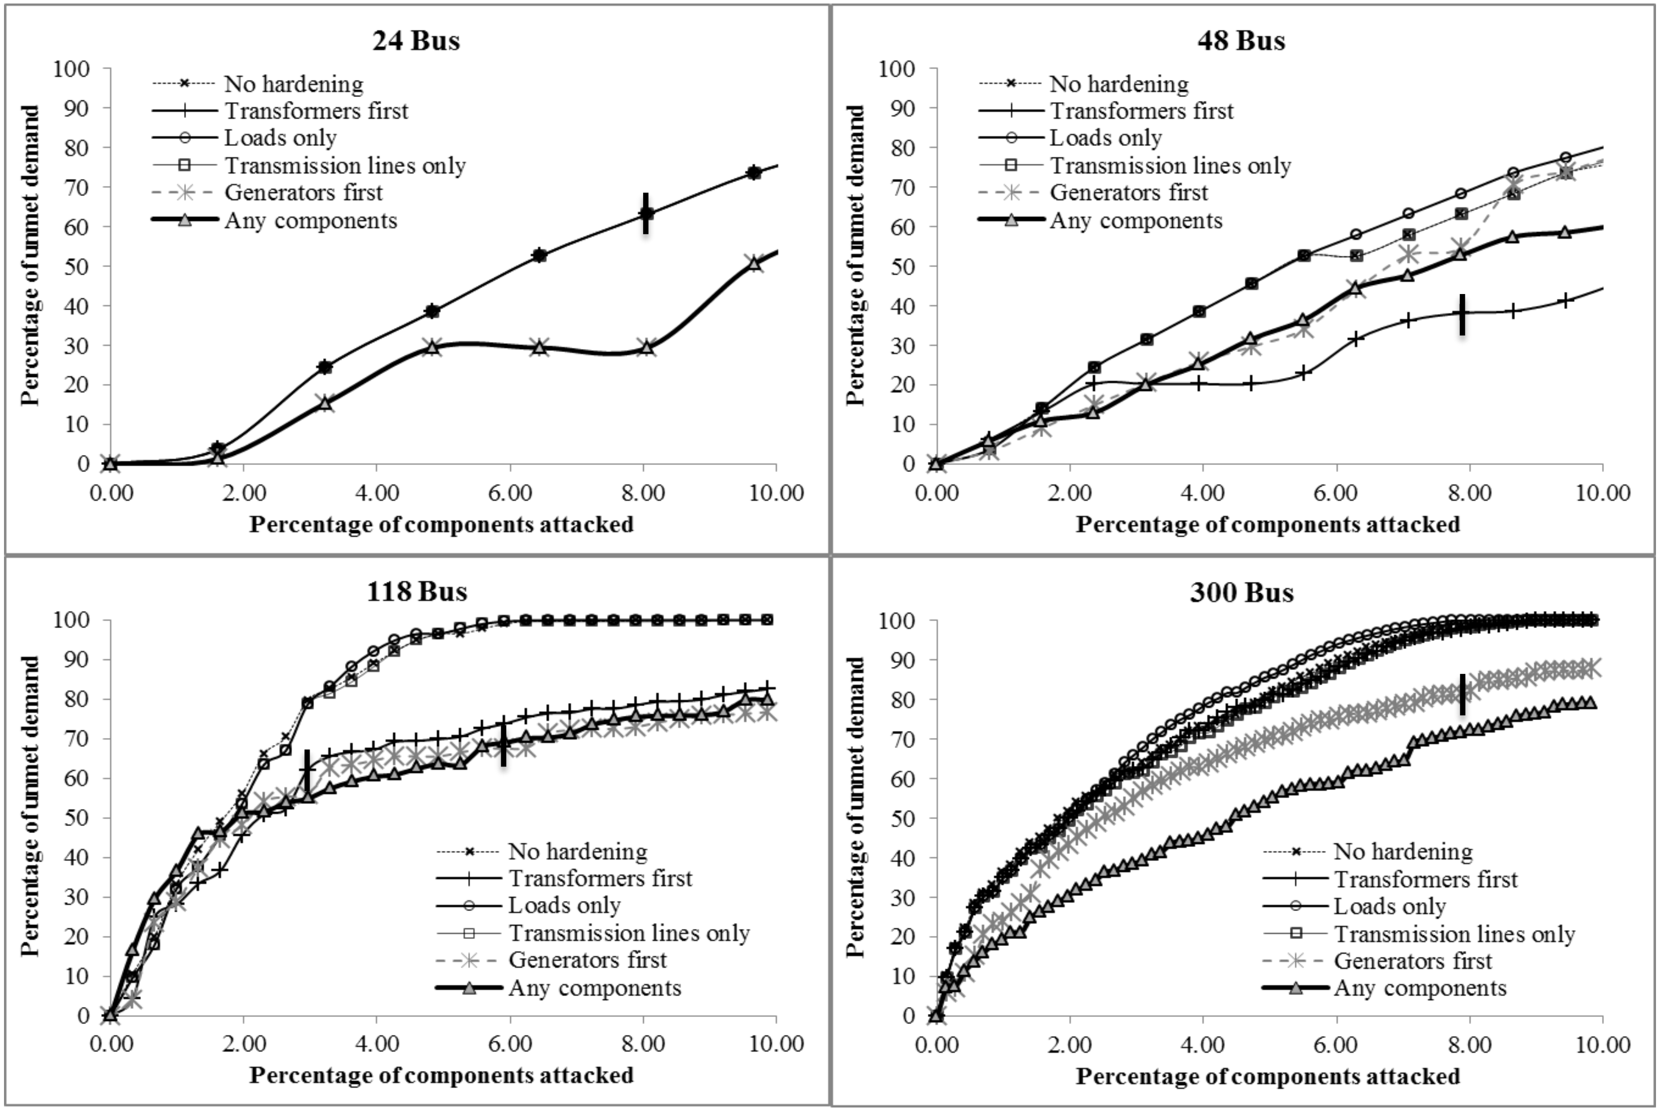

Figure 5.10- Comparison of various dynamic hardening strategies (10% hardened)

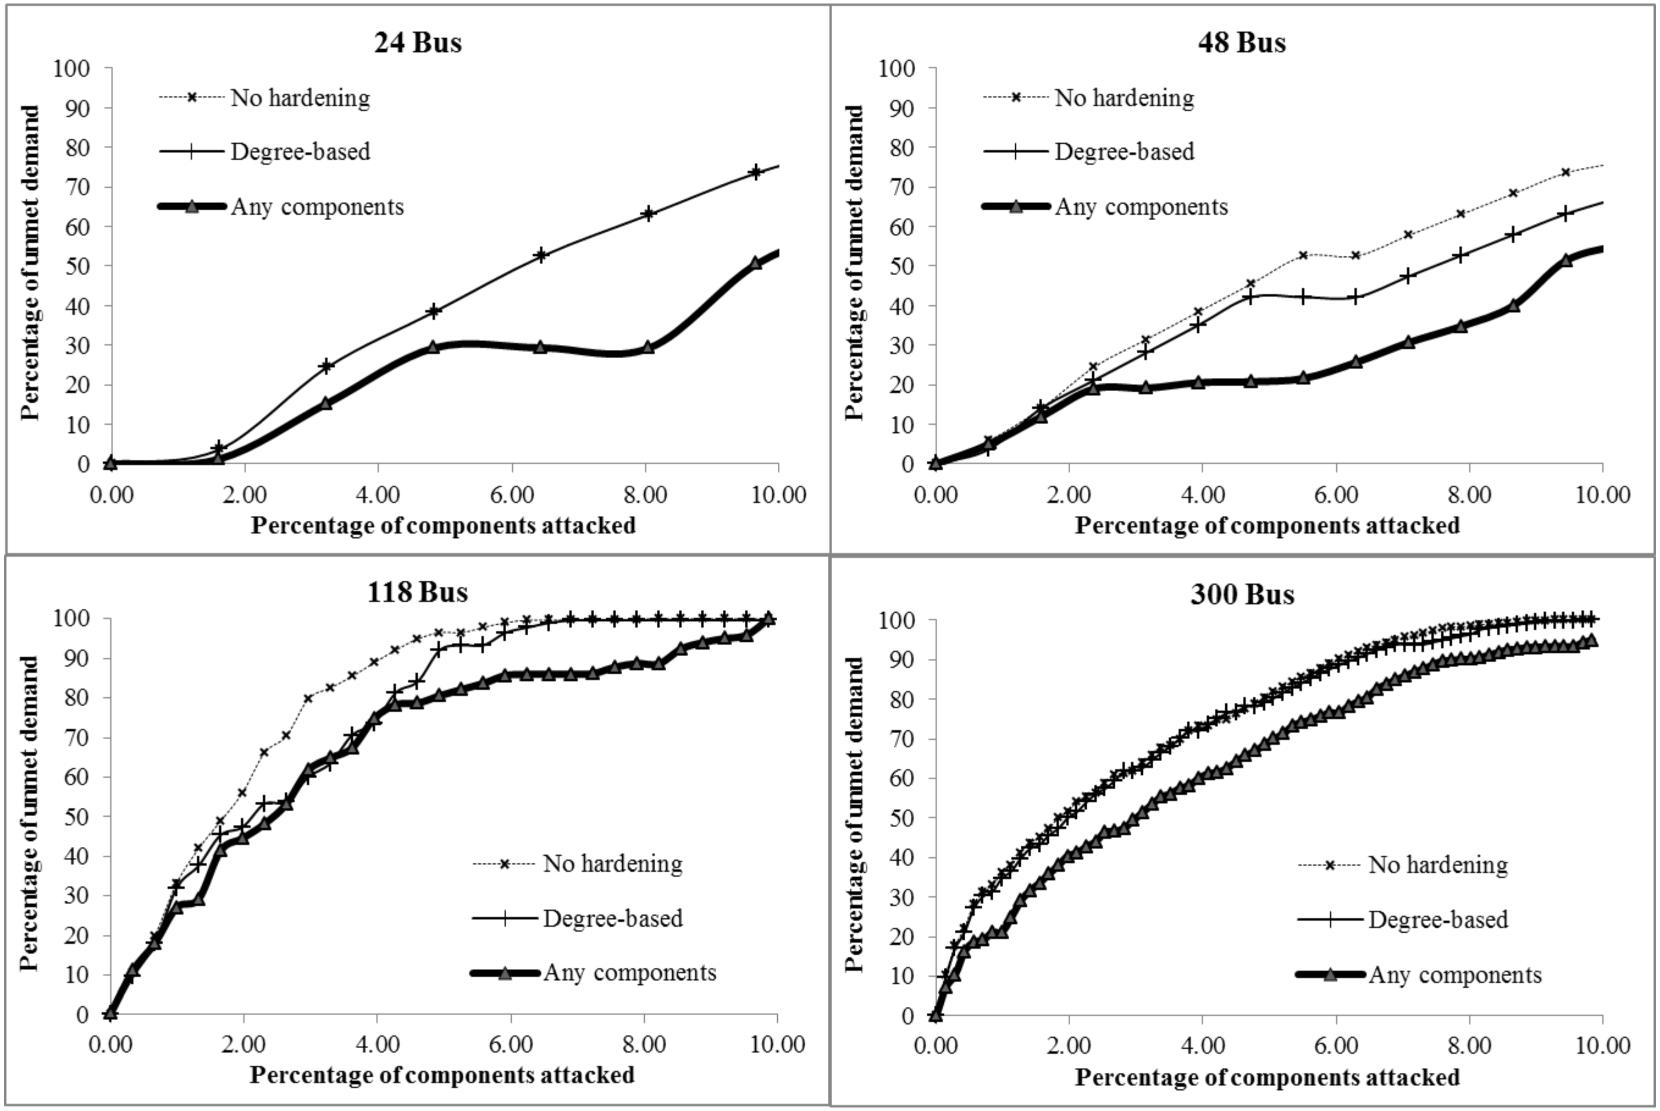

Figure 5.11 –Our dynamic hardening vs. degree-based hardening (2% hardened)

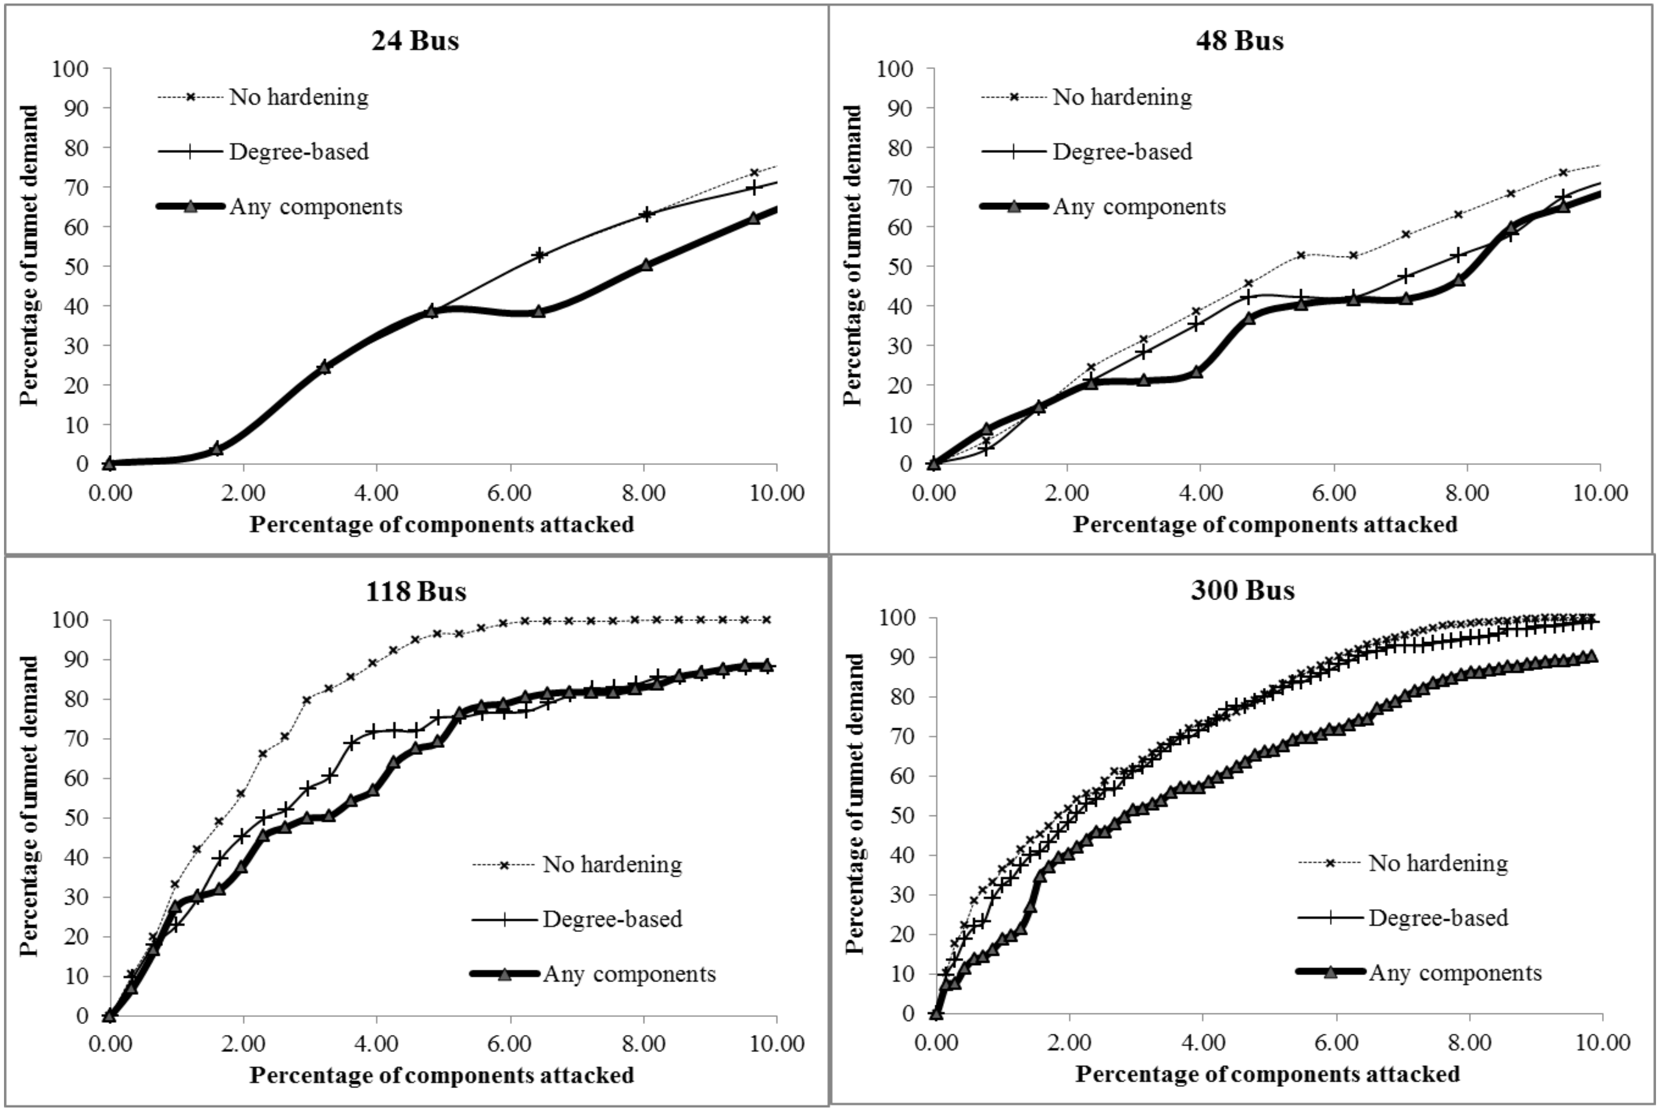

Figure 5.12 –Our dynamic hardening vs. degree-based hardening (5% hardened)

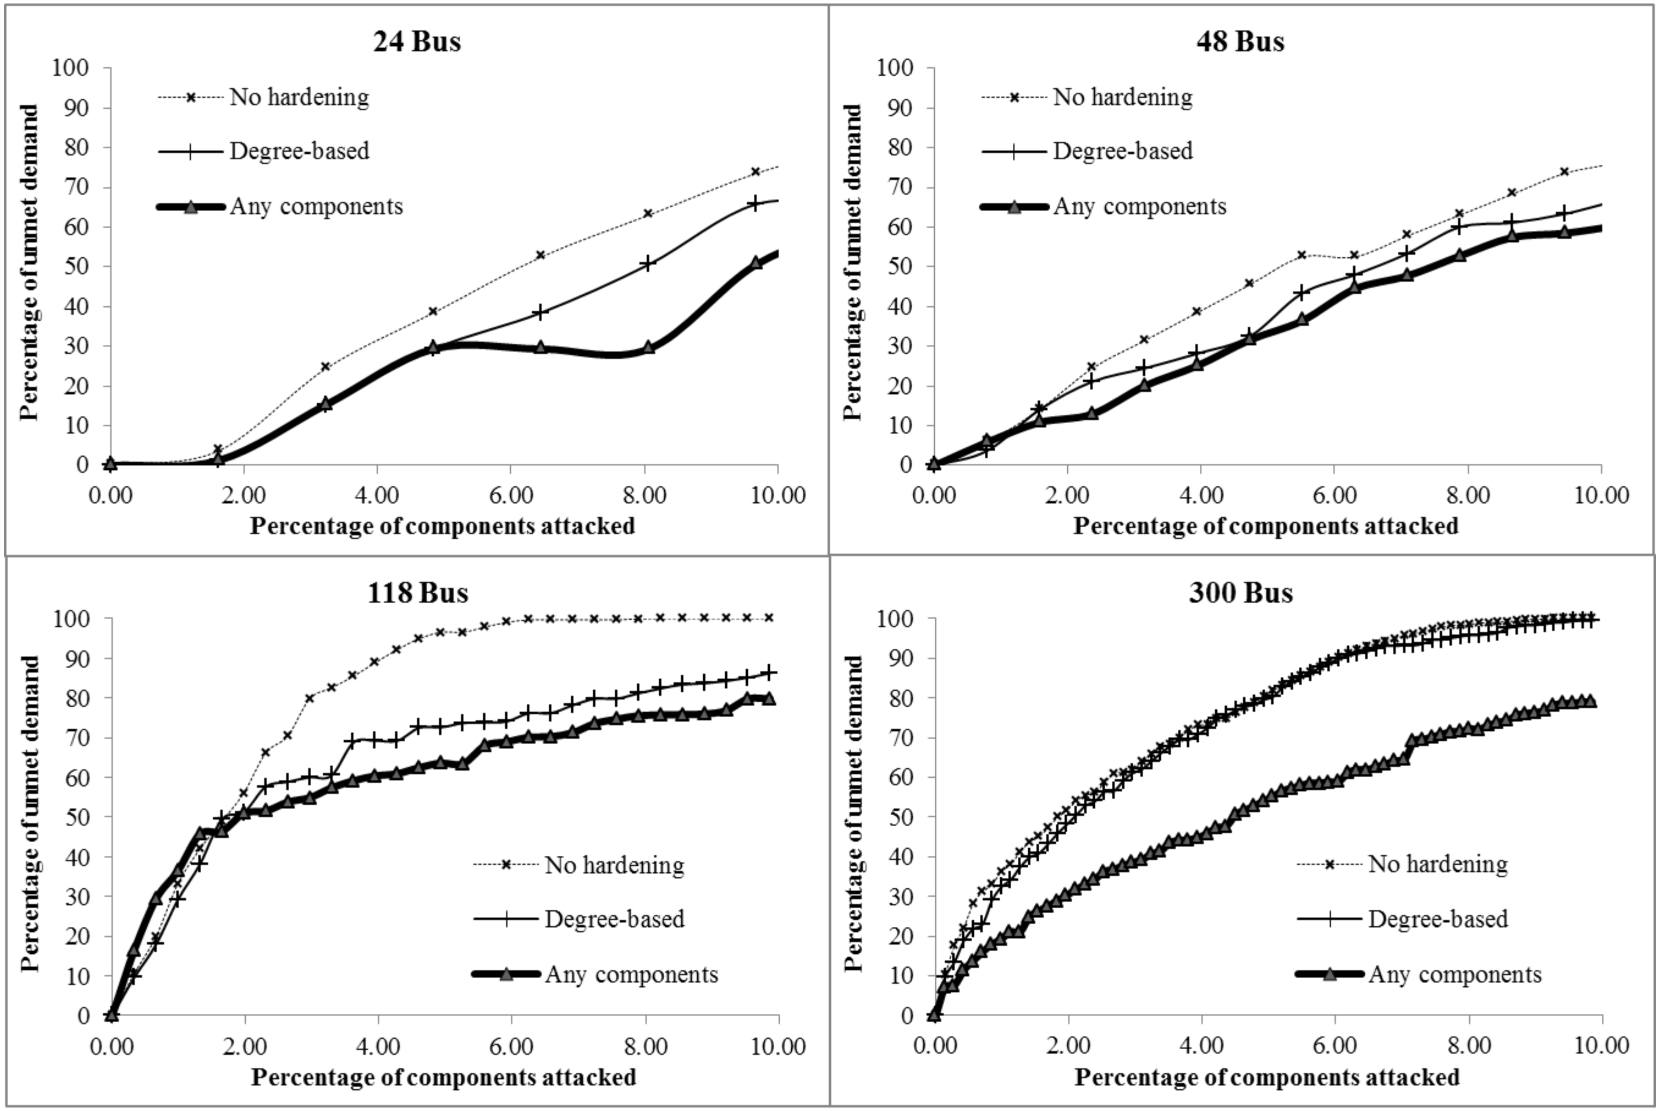

Figure 5.13 –Our dynamic hardening vs. degree-based hardening (10% hardened)

5.3. Conclusion

In this chapter, we analyzed some results given the attacker can attack any components. We are able to analyze various types of attacker and defender strategies. For example, the attacker may take into account the cost of an attack, or restrict the attention of the attacker to a specific component type. The results of the analyses presented in this chapter suggest that defending a complex system (such as an electric-power network) against an attacker with even a modest (greedy or myopic) degree of intelligence may not be highly effective. Therefore, it seems worth exploring other types of defenses beyond hardening.

In Chapter 6, we will explore the impact of cascading failure and an attacker with cascading knowledge. We will also investigate whether adding capacities or new transmission lines improve network's ability against cascading failure.

6. Results for Modeling Cascading Failure

In this chapter, we will explore how the possibility of cascading failure can affect both the behavior and the resulting level of unmet demand. In the first section, we will assume no defensive investments, and demonstrate the impact of cascading failure for different attacker types or strategies. In particular, we will first compare the results with and without cascading failure for three greedy attacker assumptions: a static attacker with no knowledge of cascading failure; a static attacker with knowledge of cascading failure; and a dynamic attacker who is able to observe cascading failures as they occur. Later, we will explore the changes in unmet demand when the attacker aims for a specific component type. Then, we will compare our greedy attacker with two other attacker types (degree-based and random attacker strategies).

In the second section, we analyze the sensitivity of the cascading model its parameters: the percentage of transmission capacity at which cascading failure occurs (α); and the failure probability of an overloaded transmission line (β). In the third section, we compare hardening to measures intended to decrease the impact of cascading failure: adding capacity to transmission lines, and adding new transmission lines. In the final section, we will compare the results of our cascading model with the cascading model of Hines et al. (2010).

Since our cascading model is probabilistic, we use Monte Carlo simulation in our analysis. For each scenario, we run 20 simulations. IF the observed standard error of the mean is higher than 3%, then we increase the number of simulations by 10 until we reach less than 3% standard error of the mean.

6.1. Impact of Cascading Failure

In order to analyze the impact of cascading failure, we first assume that the attacker has no knowledge of cascading failure. In particular, we assume a static attacker with no knowledge of cascading failure (an attacker who attacks the highest flow component based on their initial flows), and that the same components are attacked in both the cascade and no-cascade cases. This enables us to measure the impact of cascading failure on unmet demand. Figure 6.1 illustrates the impact of cascading failure for the 24-, 48-, 118-, and 300-bus systems.

As can be seen from the figures, cascading failure has a significant impact, especially for the 118- and 300-bus systems. For example, in the 118-bus system, the first attack causes an average of 24.8% unmet demand when cascading is included, but only 10.5% without cascading failure. Thus, the larger systems seem more susceptible to cascading failure perhaps due to their tight capacity constraints.

Now, we analyze how the unmet demand can change if the attacker has knowledge of cascading failure. A dynamic attacker with cascading knowledge is assumed to recalculate the system flows after each attack (taking into account any cascading failures that occurred). By contrast, a static attacker with cascading knowledge is assumed to preplan the attack by selecting which component to attack at each iteration to maximize the expected flow loss (including in the first round of cascade). Thus, a static attacker with cascading knowledge is able to exploit cascading failure to some degree, but is not assumed to anticipate all possible cascading failures. By contrast, a static attacker with no cascading knowledge is assumed to preplan its attacks based only on the initial flows in the system. Figures 6.2 and 6.3 illustrate the impact of cascading failure for a static attacker and dynamic attacker, respectively, with knowledge of cascading failure.

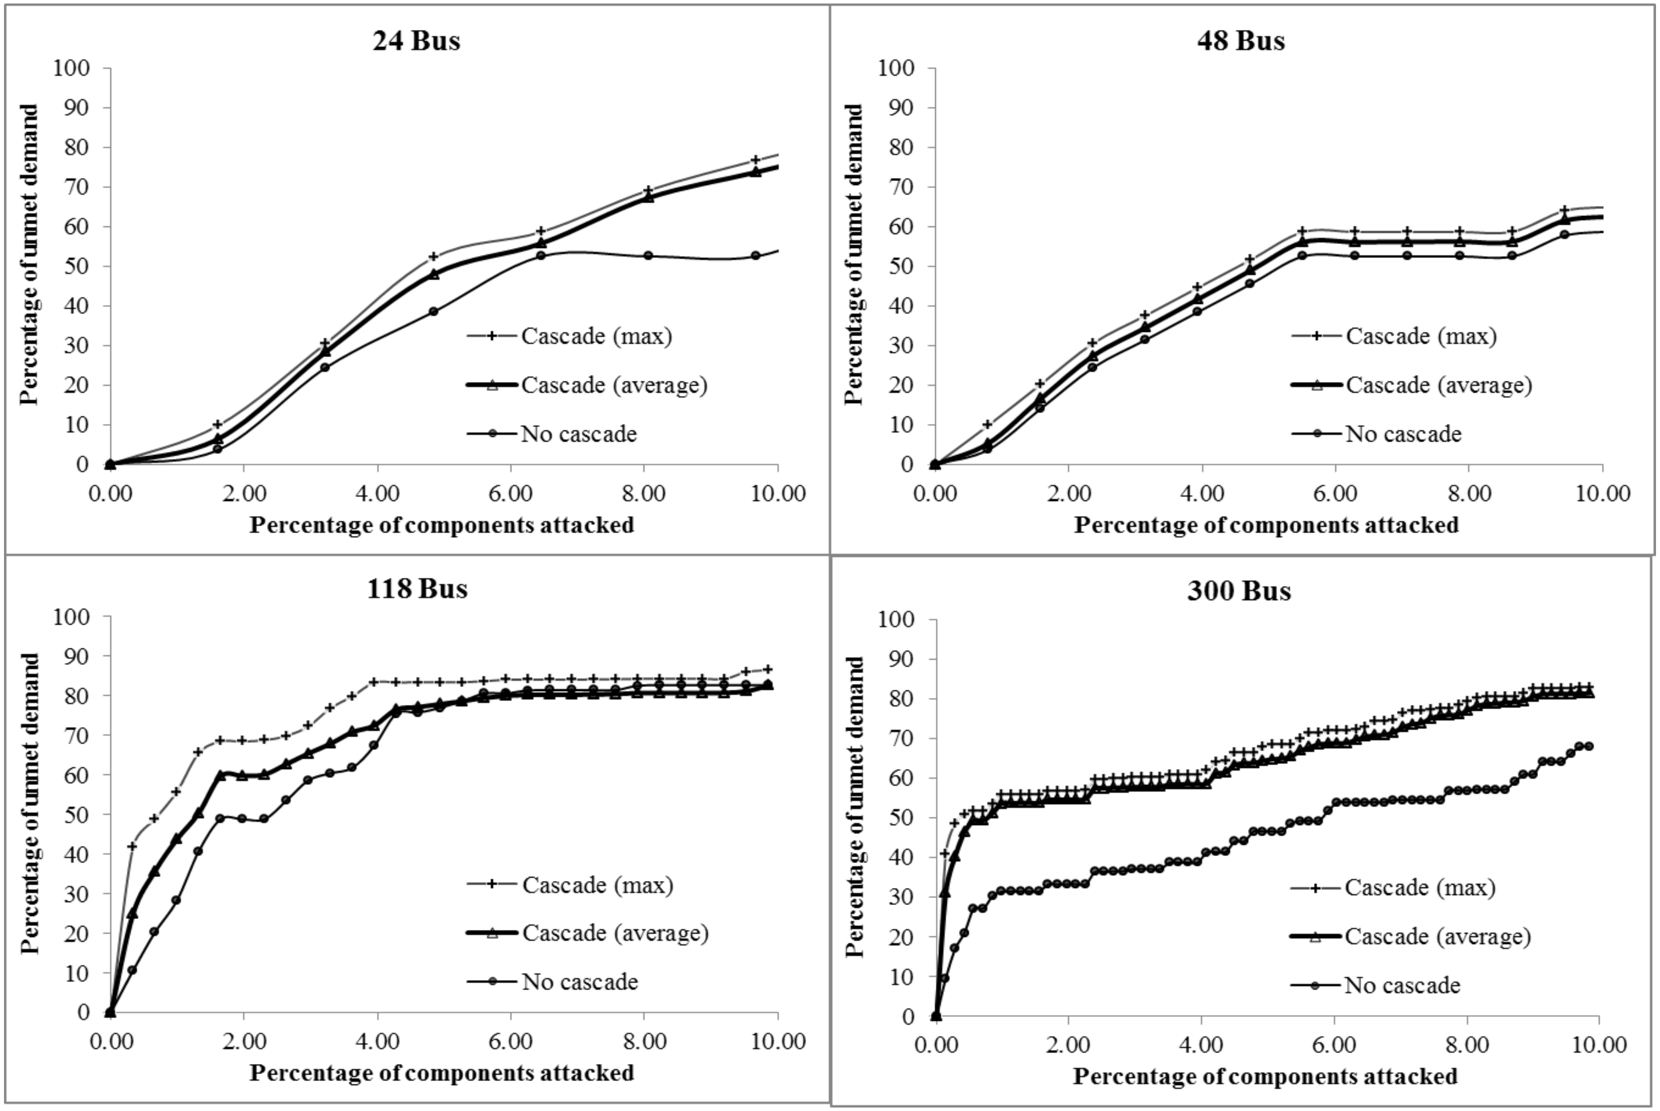

Figure 6.1–Impact of cascading failure for a static attacker with no cascading knowledge

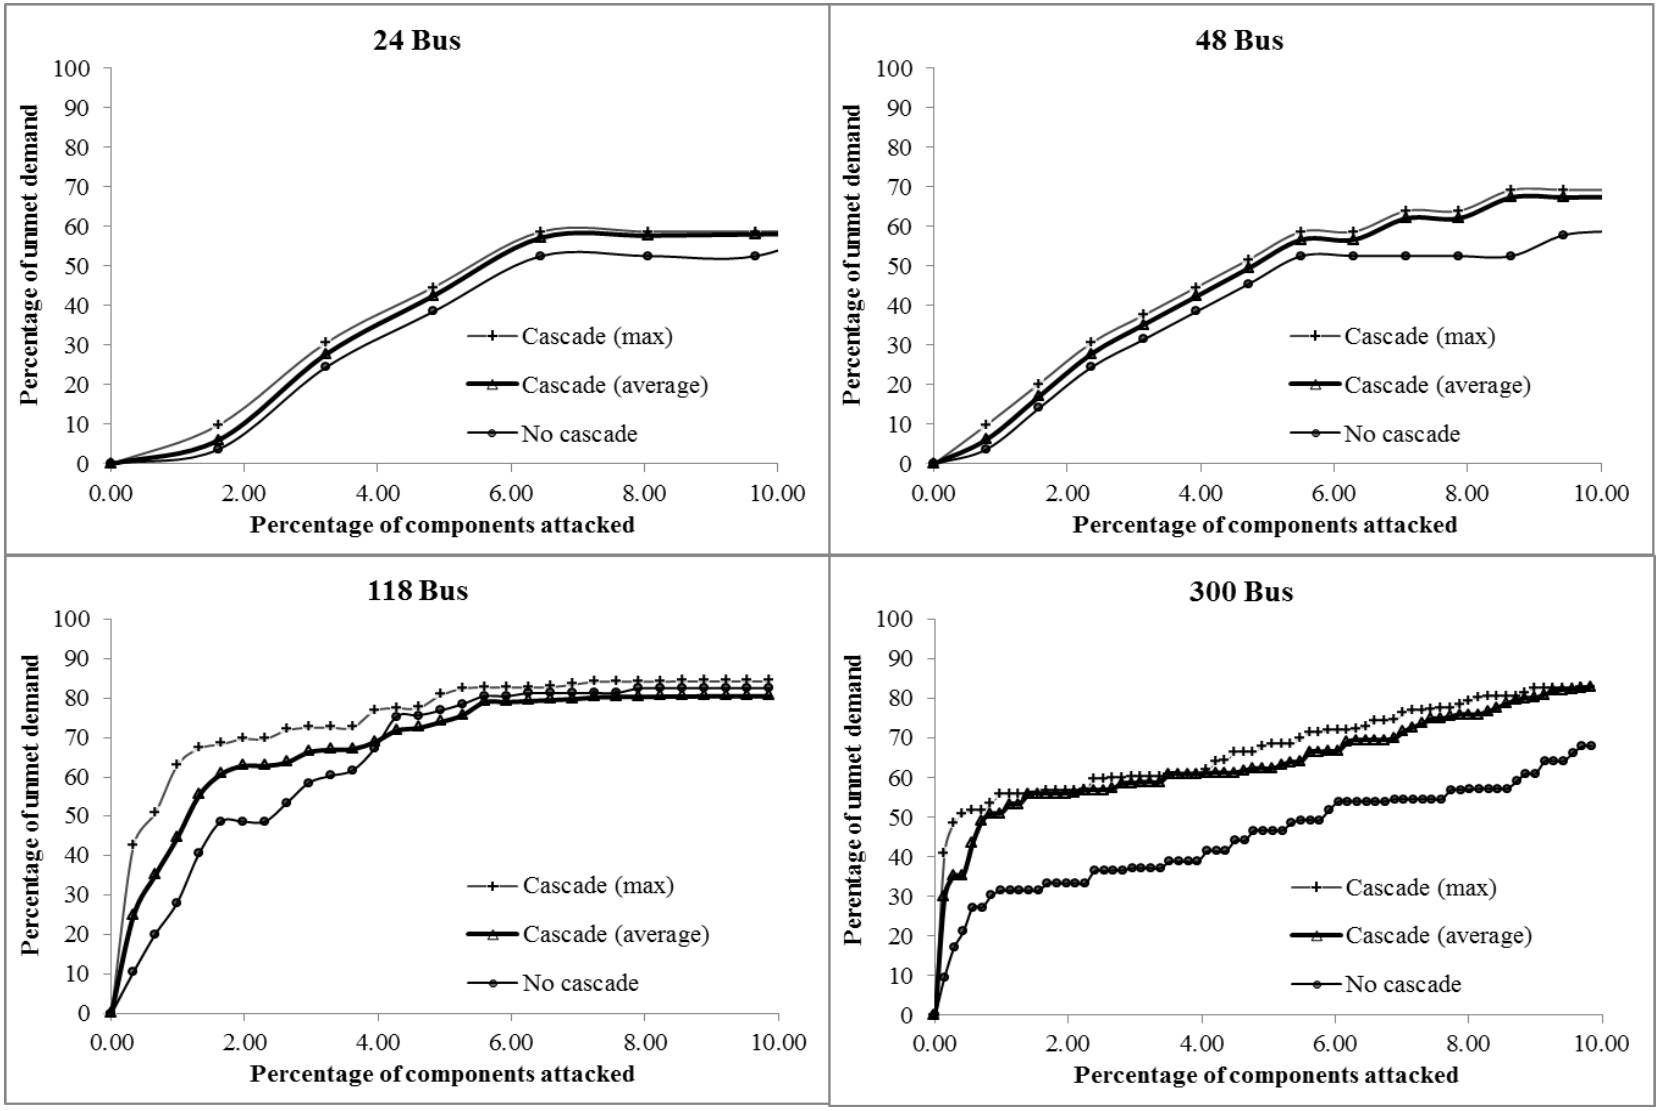

Figure 6.2–Impact of cascading failure for a static attacker with cascading knowledge

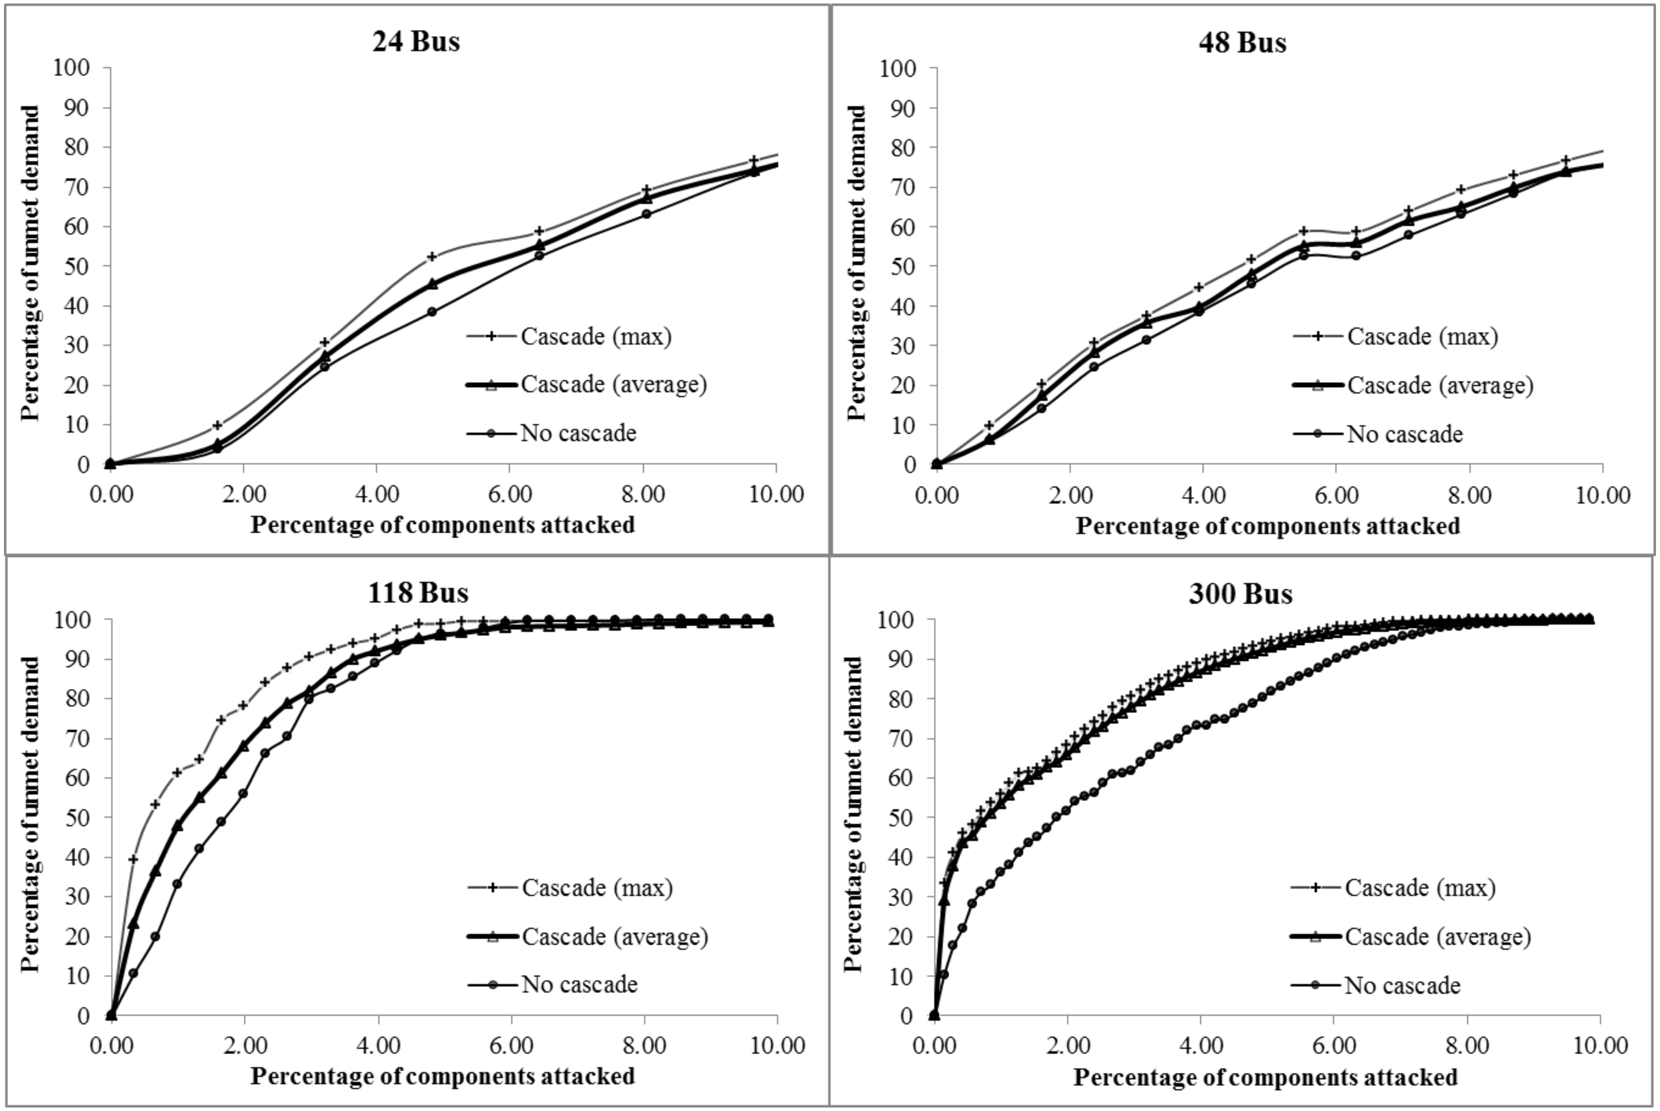

Figure 6.3–Impact of cascading failure for a dynamic attacker with cascading knowledge

Figure 6.4 illustrates the difference between these three assumptions. The dynamic attacker is significantly more effective than either of the static attacker types, at least for significant numbers of attacks. We will use the dynamic attacker with cascading knowledge as our conservative case. By contrast, the two static attacker types achieve similar levels of unmet demand, even though they attack different components. Thus, realistic levels of knowledge about cascading failure apparently cannot be effectively exploited to achieve more damaging attacks, unless the attacker is actually able to observe which components have failed.

Figure 6.5 compares five different attack strategies (any component regardless of type, generators first, transmission lines only, loads only, or transformers first) for a dynamic attacker. Note that in all but the 300-bus system, fewer than 10% of all components is transformers; thus, once all transformers have been attacked, the attacker is allowed to choose any component type. Similarly, the number of generators for 118- and 300- bus systems is less than 10%; hence the attacker chooses other components once all generators are attacked. In Figure 6.5, those attacks are separated from attacks on transformers/generators by a vertical bar, |.

As can be seen from the graphs in Figure 6.5, attacks on all component types and attacks on generators create the most unmet demand. Thus, as in Chapter 5, attacks on generators may be preferred if not too costly. By contrast, attacks on loads, transmission lines, or transformers are significantly less effective. The results for a static attacker, with or without cascading knowledge, are quite similar (i.e., still showing greater impact for attacks on all component types or on generators); see Appendix C.

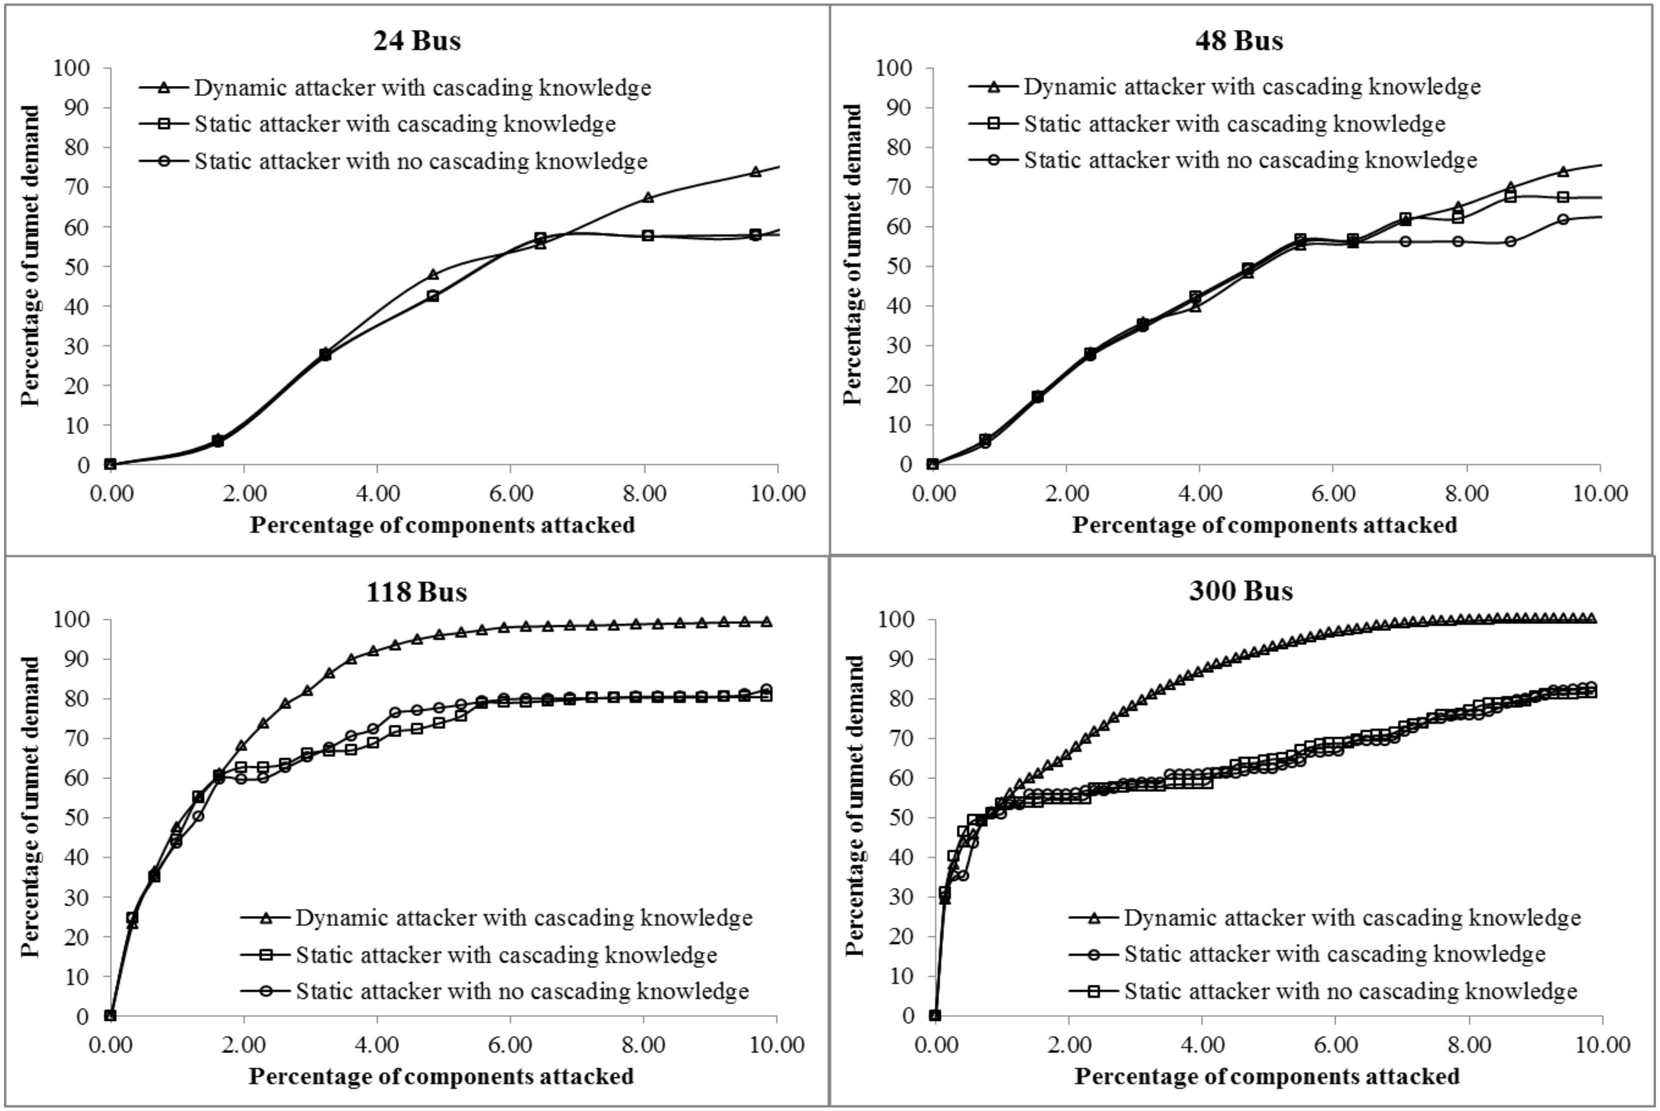

Figure 6.4–Attackers with and without cascading knowledge

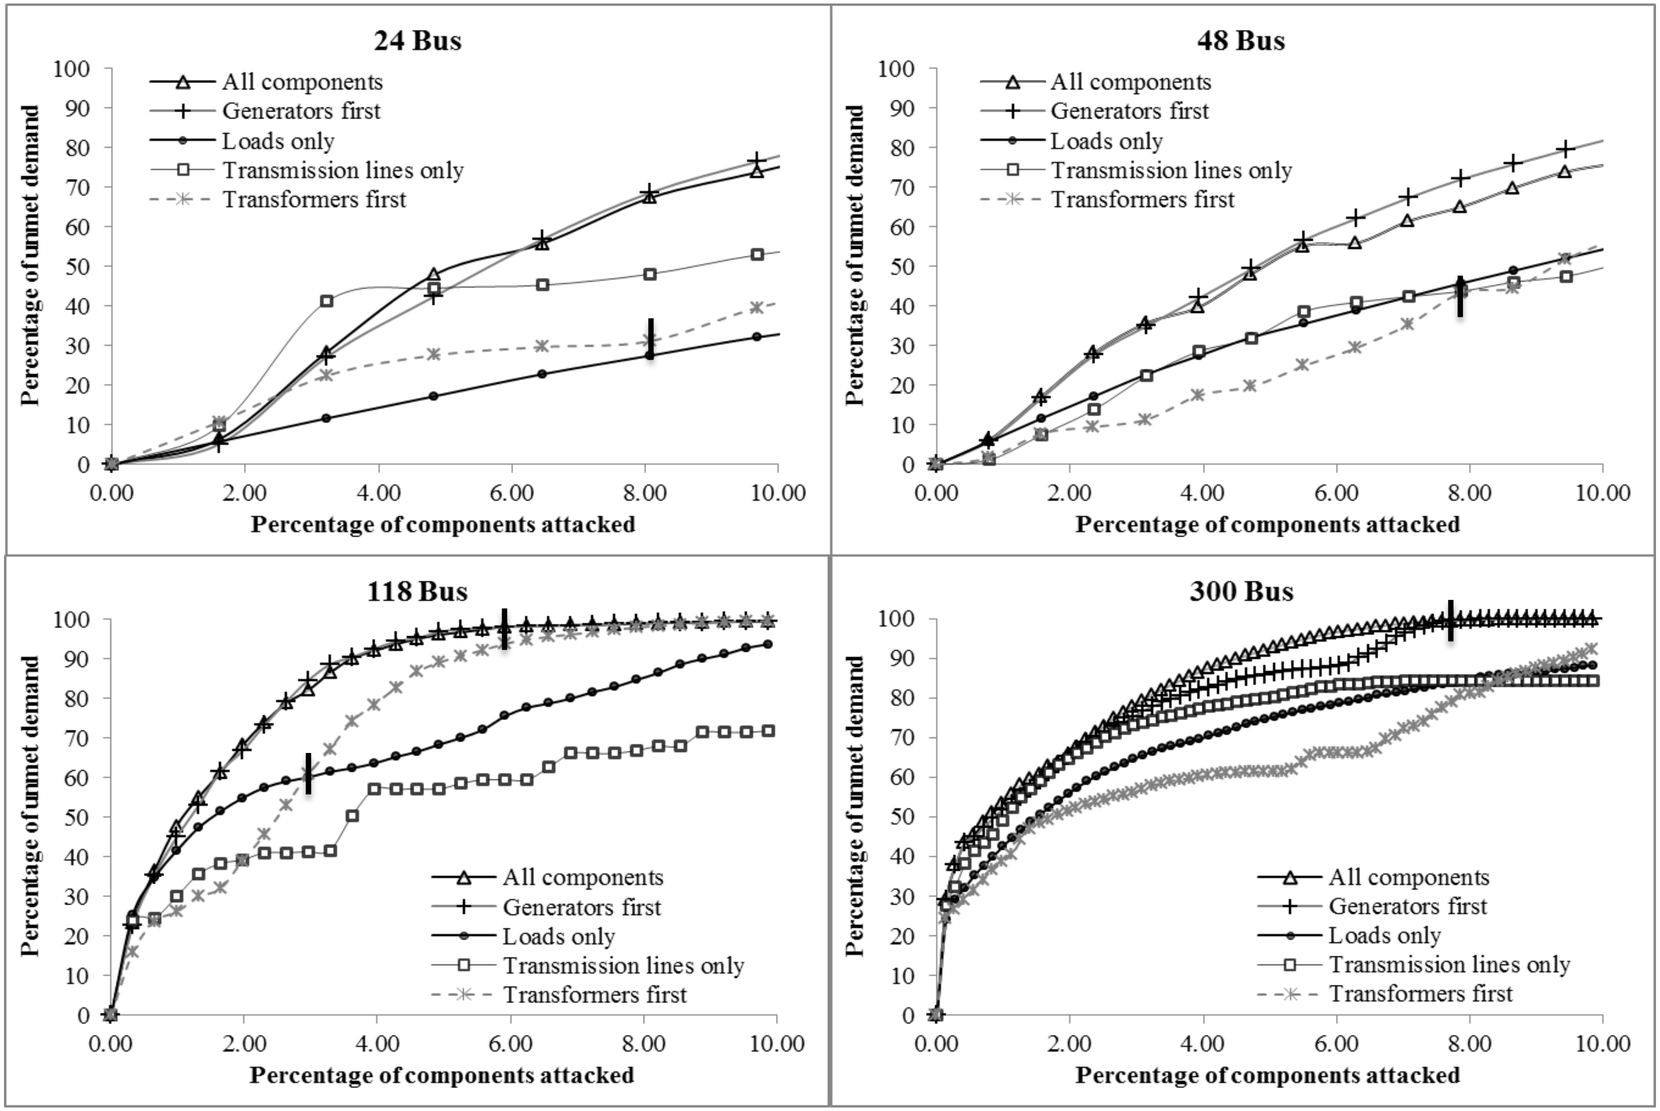

Figure 6.5–A dynamic attacker choosing different component types

In Figure 6.6, we compare our greedy attacker with degree-based and random attacker types for each of our three assumptions about attacker's cascading knowledge. The degree-based attacker generally does less well than even our static attacker types, except in the 118-bus system. As in Chapter 5, the degree-based heuristic is sometimes not even significantly better than a random attacker, especially for small attacks. This confirms the observation from Chapter 5 that a degree-based heuristic is not a conservative representation of the impact of an intelligent adversary.

6.2. Sensitivity Analysis of the Cascading Model

In this section, we will explore the sensitivity of our cascading model to the fraction of transmission capacity at which cascading failure occurs (α), and the failure probability of an overloaded transmission line (β). Since transmission lines are much more vulnerable to cascading failure when operating close to their capacities (Crucitti et al., 2004; Zhao et al., 2004; Kinney et al., 2005), we assume that once a transmission line reaches its critical level of flow α, it becomes a candidate for cascading failure. If that line fails due to cascading (which happens with probability β), the flow on that line is redistributed through the electric-power system according to the DC power dispatch algorithm. This process continues until all components have failed, or no additional components have exceeded a fraction α of their capacities, or those components that have exceeded their critical flow levels have been determined not to fail.

Figure 6.7 analyzes the impact of cascading failure as a function of the fraction α above which components become candidates for cascading failure; assuming a dynamic attacker. The default value is 0.99, but we vary this value from 0.9 to 0.999. As can be seen from Figure 6.7, the percentage of unmet demand is, perhaps surprisingly, not particularly sensitive to the threshold

α. Figure 6.8 shows again that the level of unmet demand is not particularly sensitive to β, the failure probability of overloaded lines. In fact, higher failure probabilities β are not always associated with higher levels of unmet demand. Since sufficient simulations have been done to ensure accurate results, this does not appear to be due to simulation noise, but perhaps instead to the non-optimal nature of our assumed attack strategy.

The results in Figures 6.7 and 6.8 are for the case of a dynamic attacker, who observes any cascading failures resulting from one attack before choosing which component to target next. Appendix D presents results for both static attacker types. Results are similar to those in Figures 6.7 and 6.8, in the sense that they are not highly sensitive to α and β.

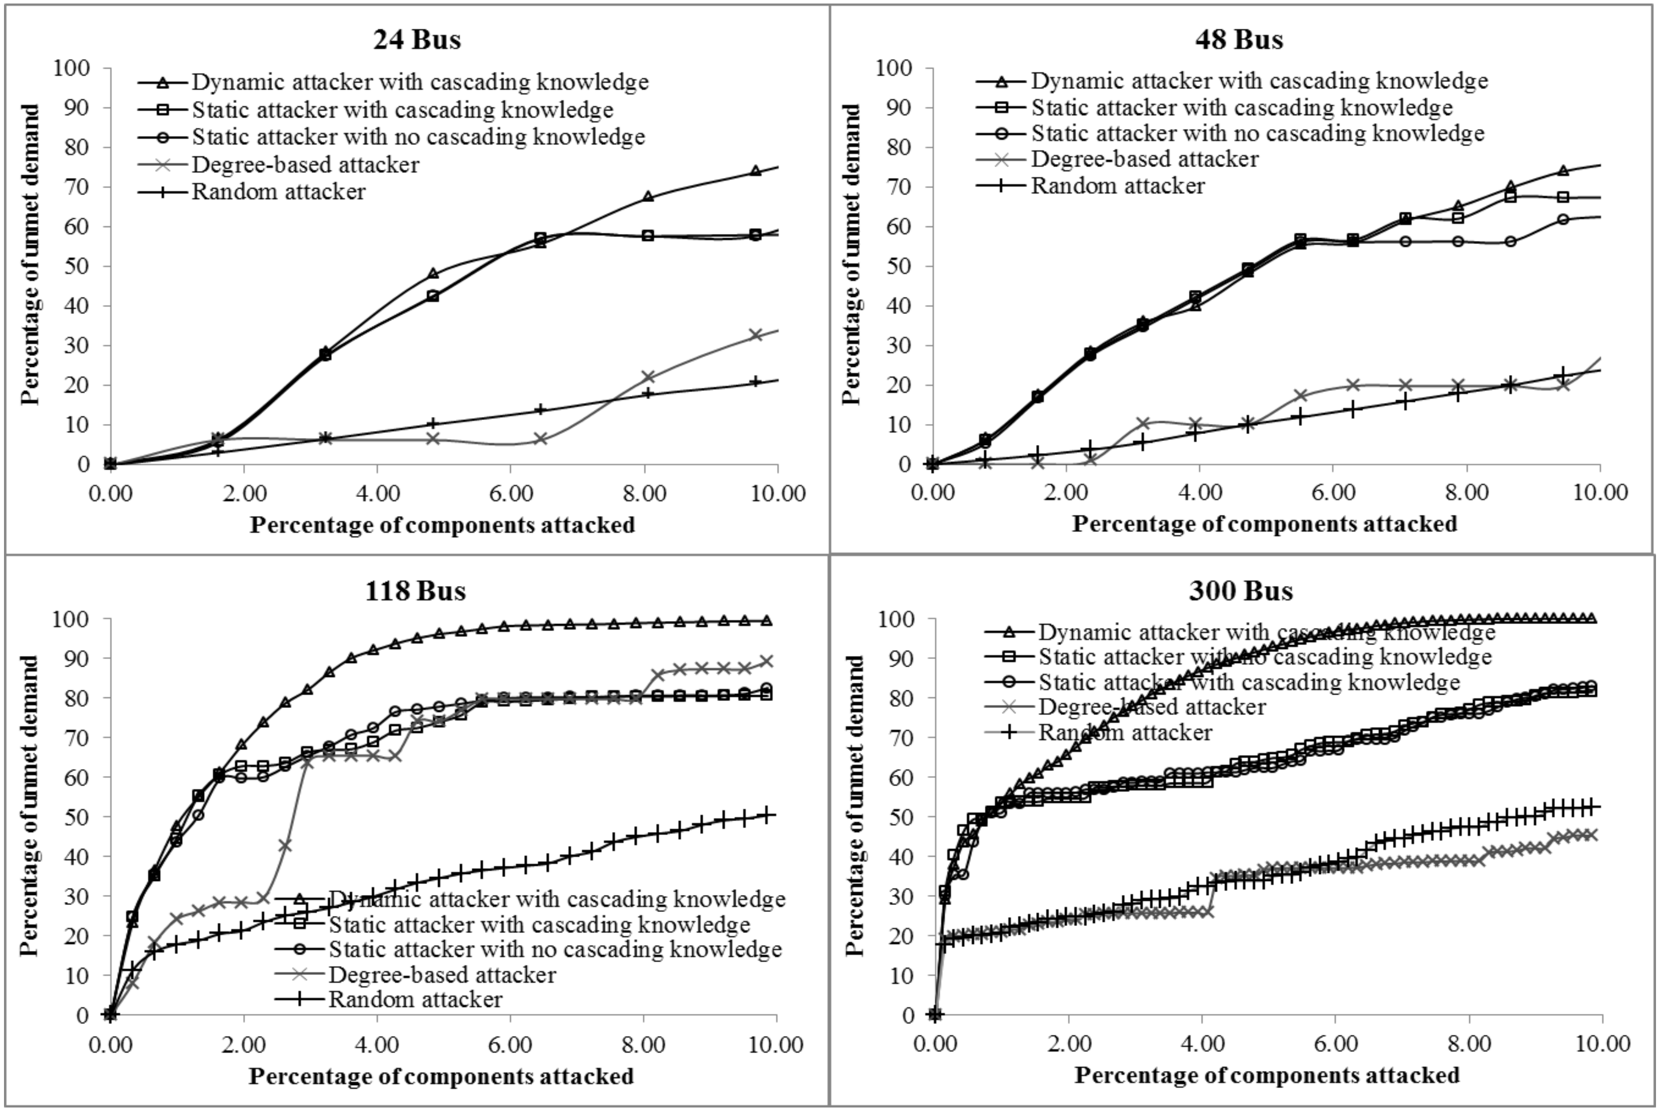

Figure 6.6–Greedy attacker versus degree-based and random attackers

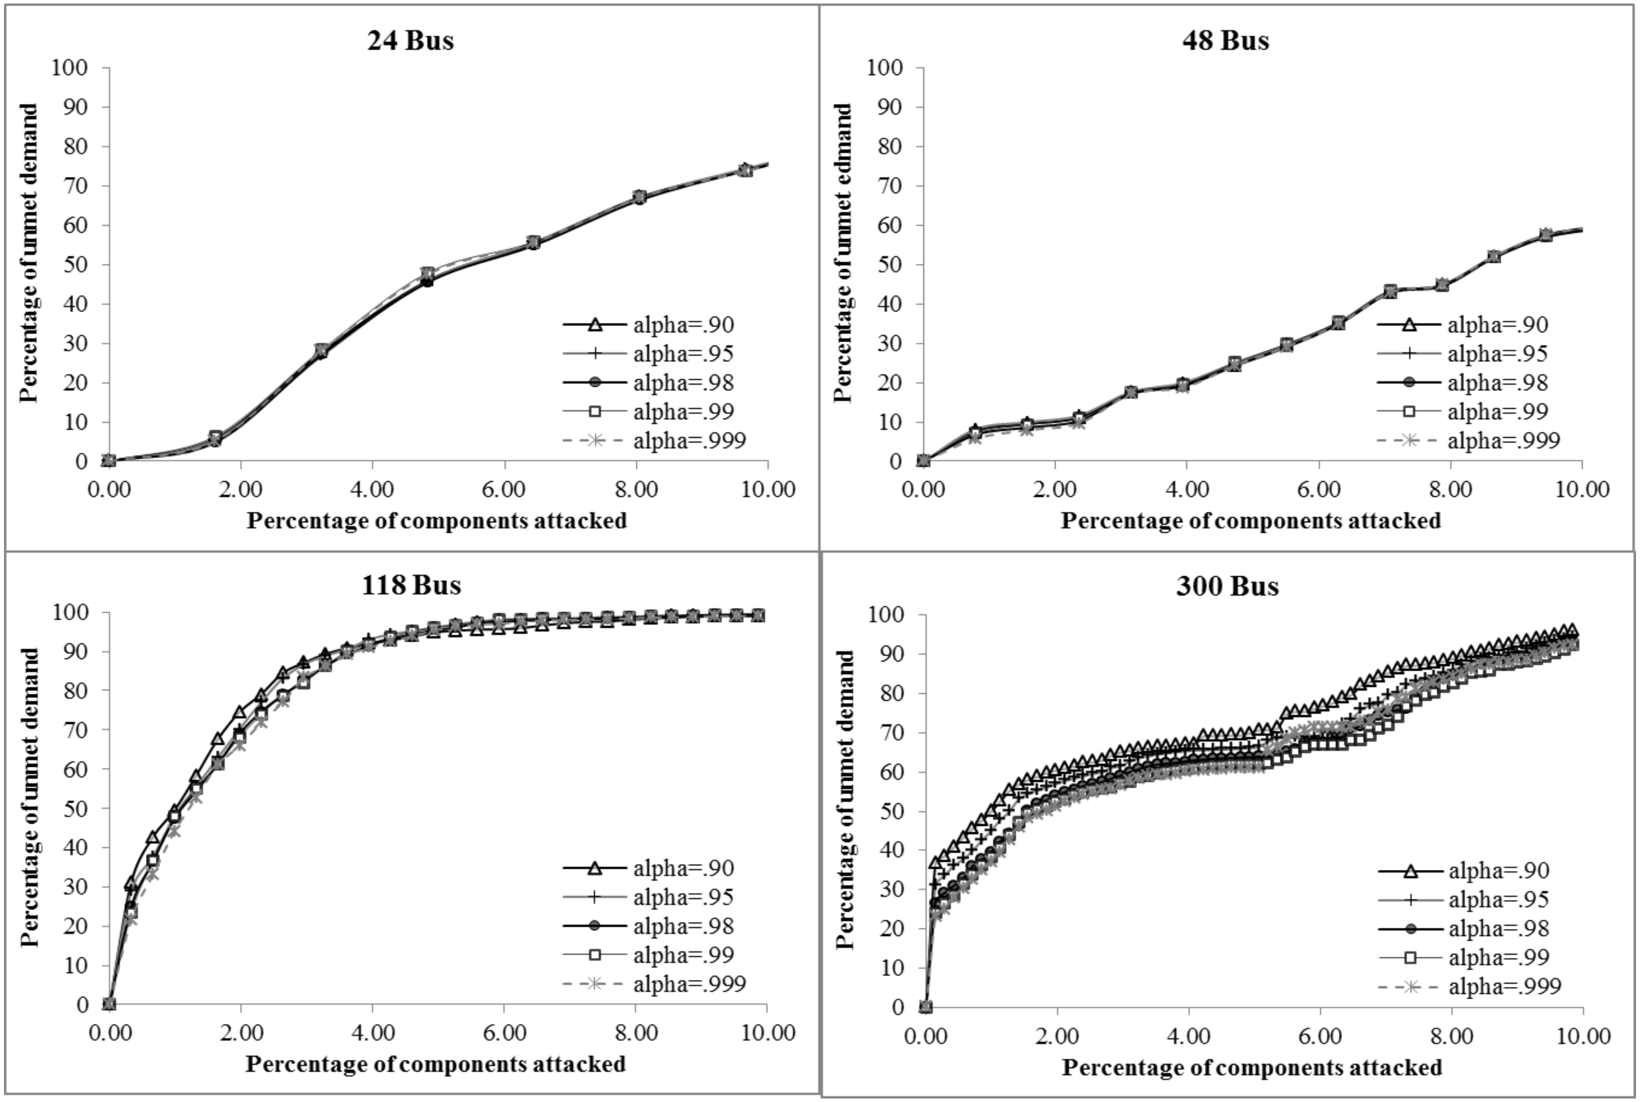

Figure 6.7–Sensitivity of the cascading model to the fraction of transmission capacity above which cascading failure occurs for a dynamic attacker

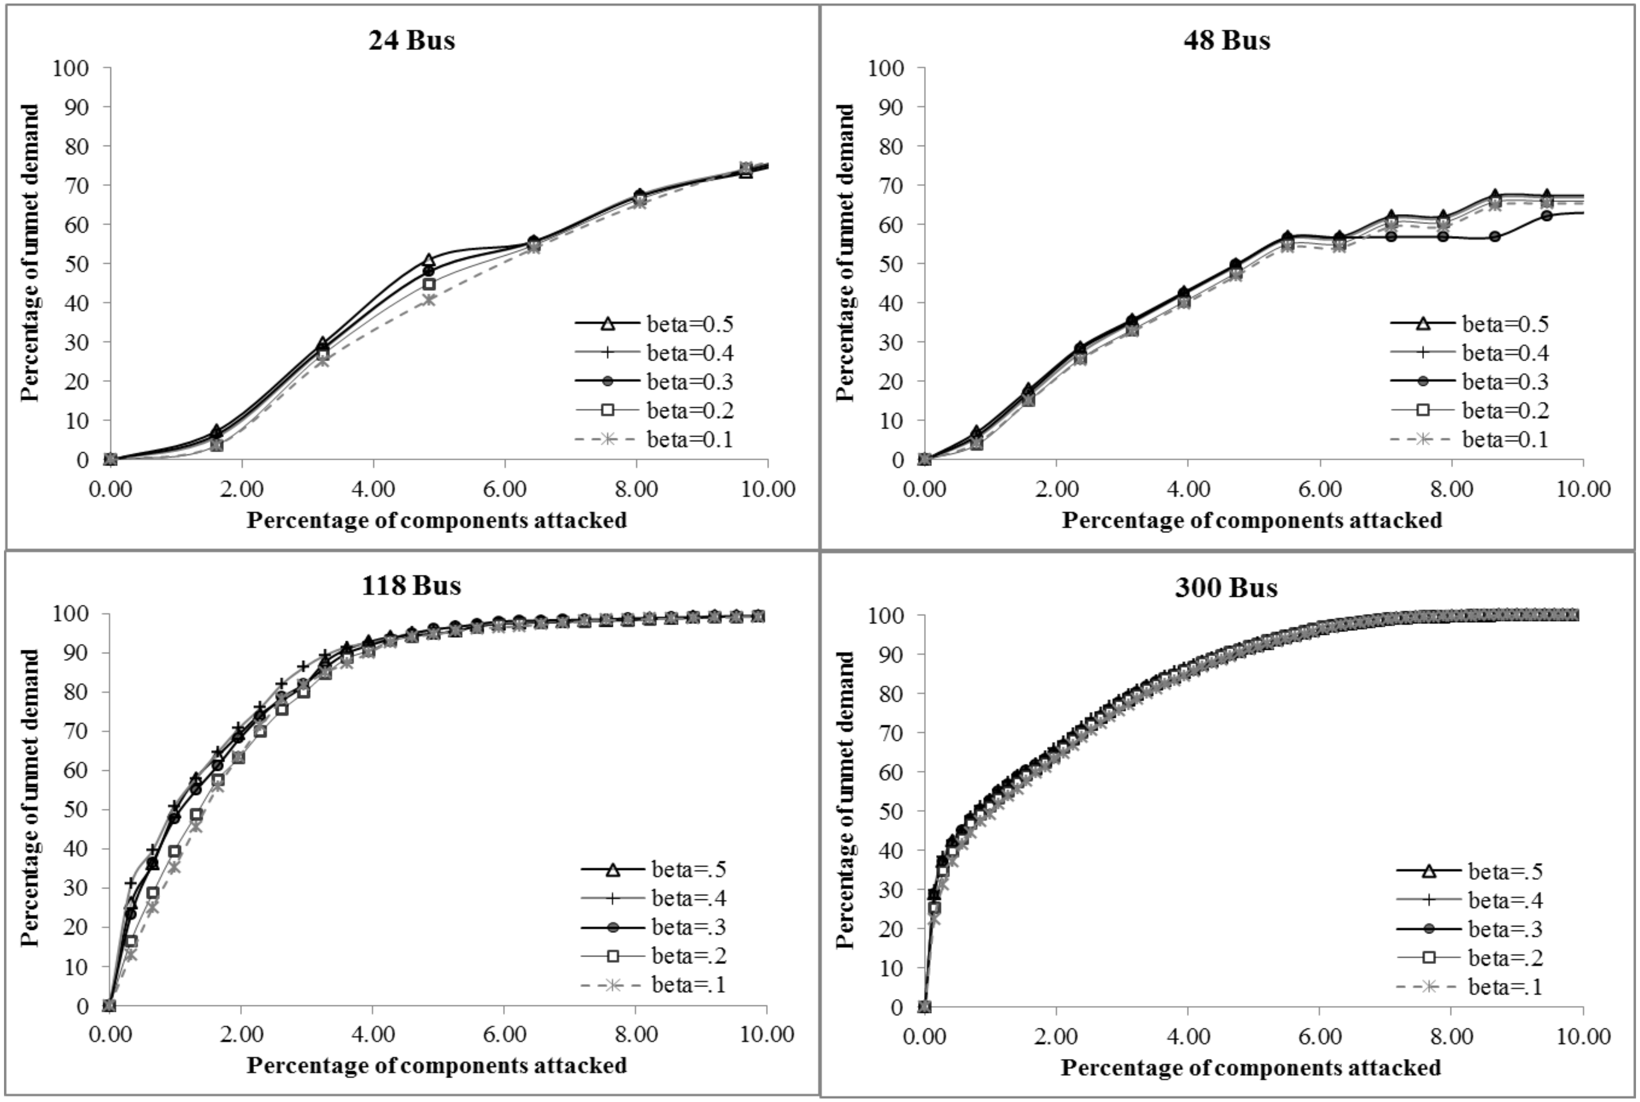

Figure 6.8–Sensitivity of the cascading model to the failure probability of an overloaded component for a dynamic attacker

6.3. Defensive Measures in the Face of Cascading Failure

In this section, we compare the effectiveness of different types of defensive investments. We consider three main investment types; hardening a certain percentage of components (as done in Chapter 5.2 for a system without cascading failure); doubling the capacity of a certain percentage of transmission lines that are the most likely to cascade; and adding new (identical) transmission lines parallel to those transmission lines that are most likely to cascade (but not collocated with them). Note that the last two investment strategies are intended specifically to make the system more robust by decreasing the impact of cascading failure.

6.3.1. Hardening

In Figures 6.9 through 6.11, we analyze the effectiveness of different levels of hardening for a static attacker with no cascading knowledge, a static attacker with cascading knowledge, and finally a dynamic attacker, respectively. For all three attacker types, the level of unmet demand generally decreases as the level of hardening increases, especially for the dynamic attacker. However, as in previous chapters, we still occasionally observe cases where increasing the fraction of hardened components leads to more unmet demand, which again demonstrates that our algorithm for selecting which components to harden is not optimal. Likewise, as in previous chapters, we again observe that significant numbers of components (as many as 30%) must sometimes be hardened in order to achieve much improvement in the level of unmet demand caused by a greedy attacker.

Now, we compare the effectiveness of different hardening strategies (including hardening any component type, generators first, transmission lines only, and transformers first) against a dynamic attacker under three different scenarios: small hardening investment (2%); medium hardening investment (5%); and large hardening investment (10%), shown in Figures 6.12 through 6.14, respectively. As the percentage of components attacked gets larger, the superiority of hardening generators first or all component types (compared to other hardening strategies) becomes significant. For similar analyses of a static attacker with and without cascading knowledge, see Appendix E.

Figures 6.15 through 6.17 compare dynamic hardening with degree-based hardening for small (2%), medium (5%), and (10%) large hardening investments, respectively (note that dynamic hardening implies hardening the components that would be attacked by a dynamic attacker). For a small hardening investment (2%), degree-based hardening actually does better than greedy hardening in the 118-bus system. For the same bus system, degree-based hardening is still comparable to greedy hardening, especially for large attacks. In general, however, greedy hardening performs better than degree-based hardening for most attack sizes and levels of investment. In particular, in the 24- and 48-bus systems, degree-based hardening is comparable to greedy hardening for only large attack sizes; and greedy hardening always performs better than degree-based hardening in 300-bus system. (For the corresponding analyses for a static attacker with and without cascading knowledge, see Appendix F. As for the case of a dynamic attacker, degree-based hardening performs better for some attack sizes and levels of investment.)

Overall, hardening even a significant percentage of components is not predicted to dramatically diminish the load shed from an attack. This implies that hardening may not be cost effective. Therefore, in the following sections, we explore alternative types of defensive investments: doubling the capacity of transmission lines; and adding new transmission lines.

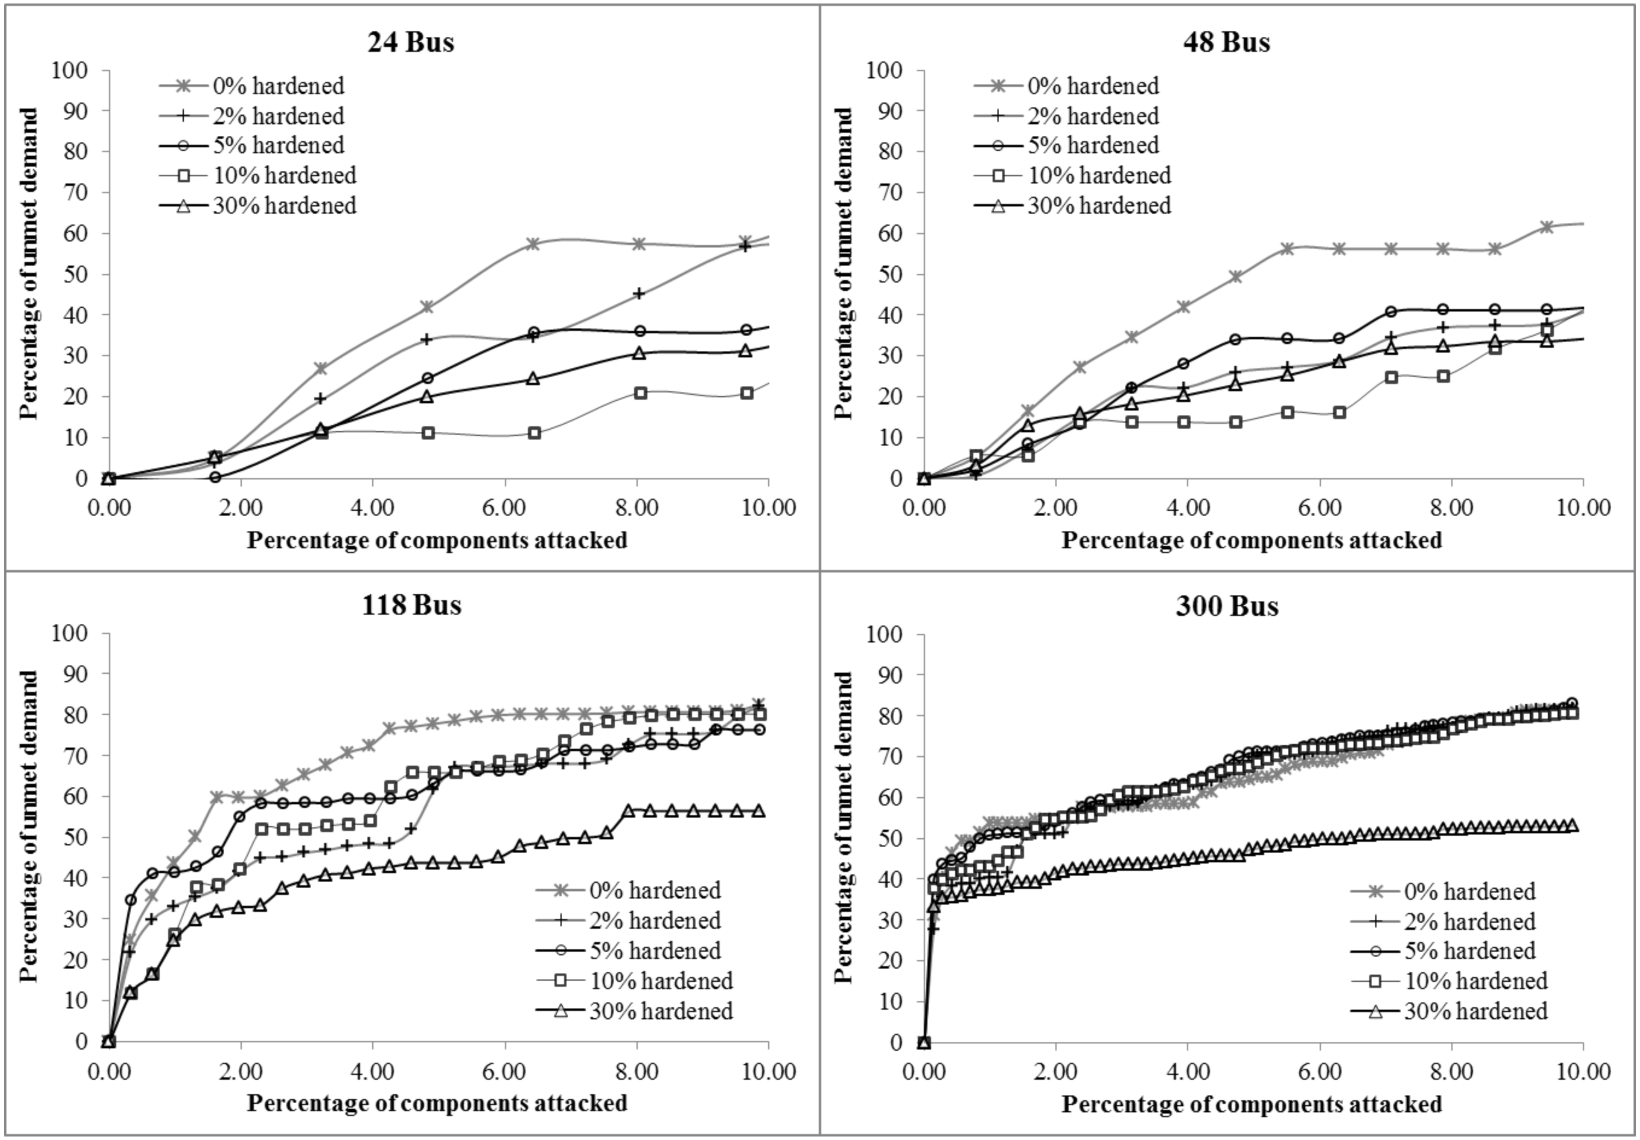

Figure 6.9–The effects of different levels of hardening for static attacker with no cascading knowledge

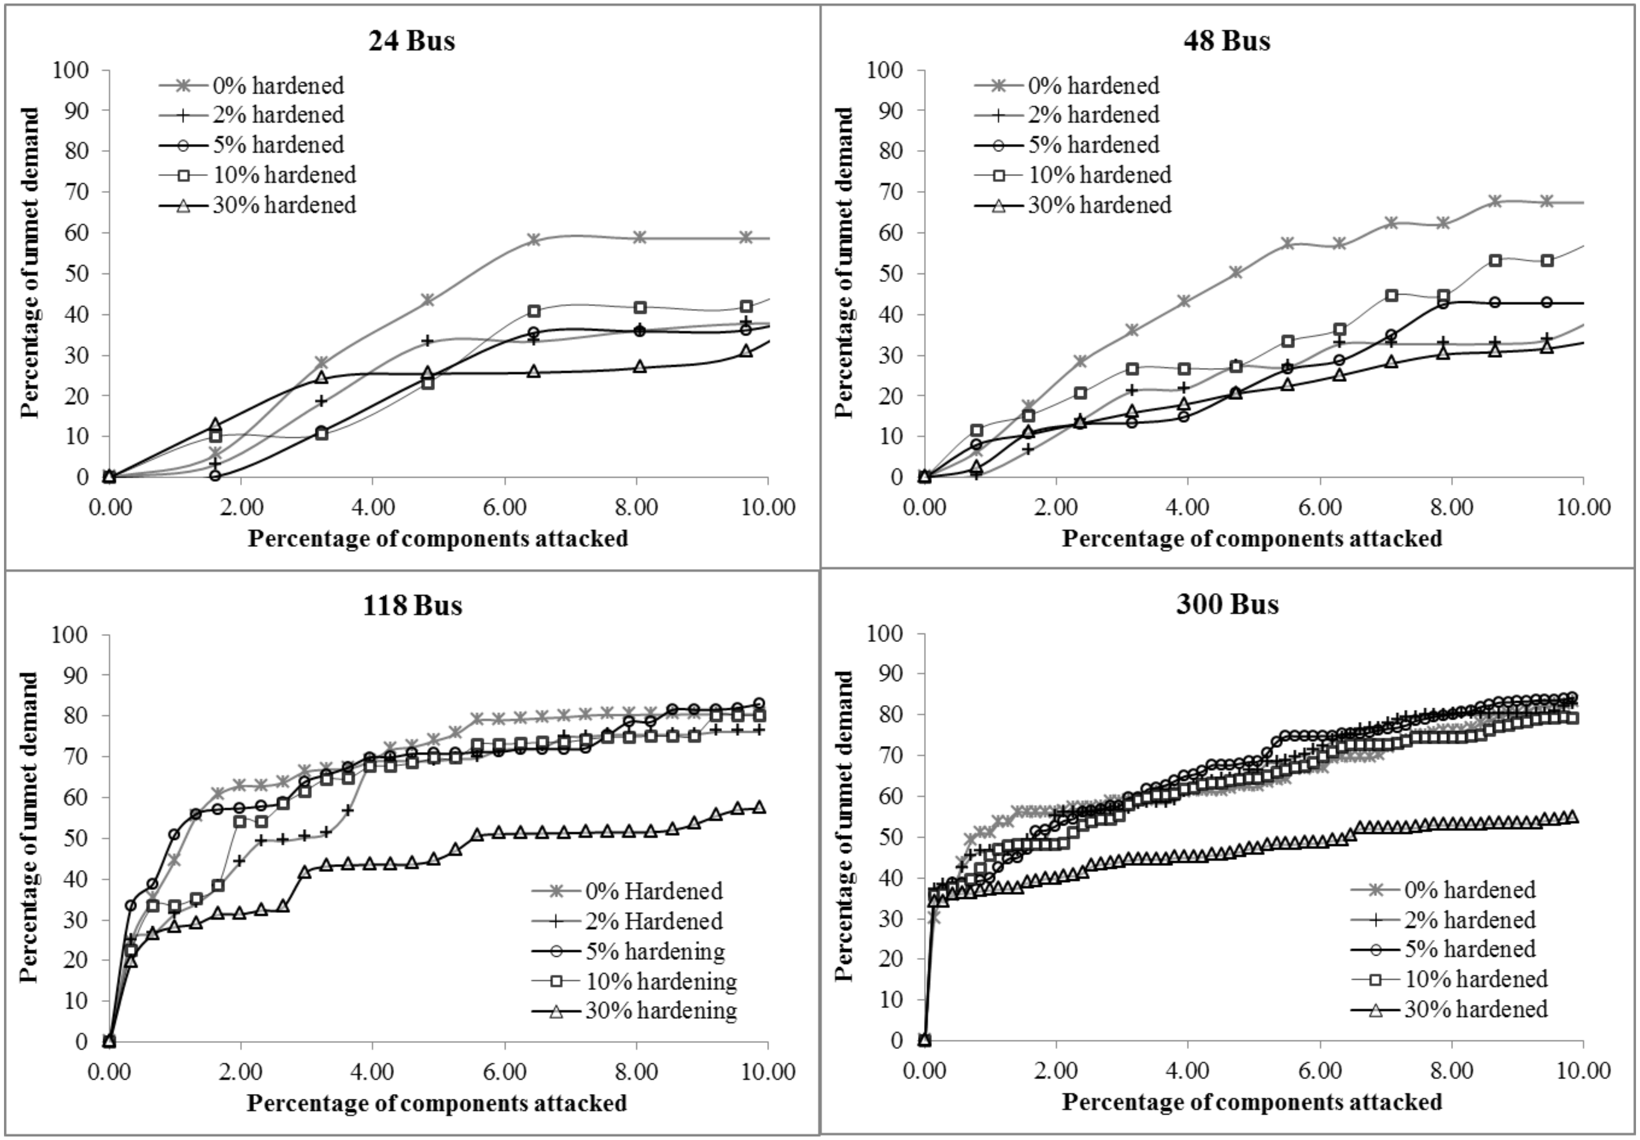

Figure 6.10–The effects of different levels of hardening for static attacker with cascading knowledge

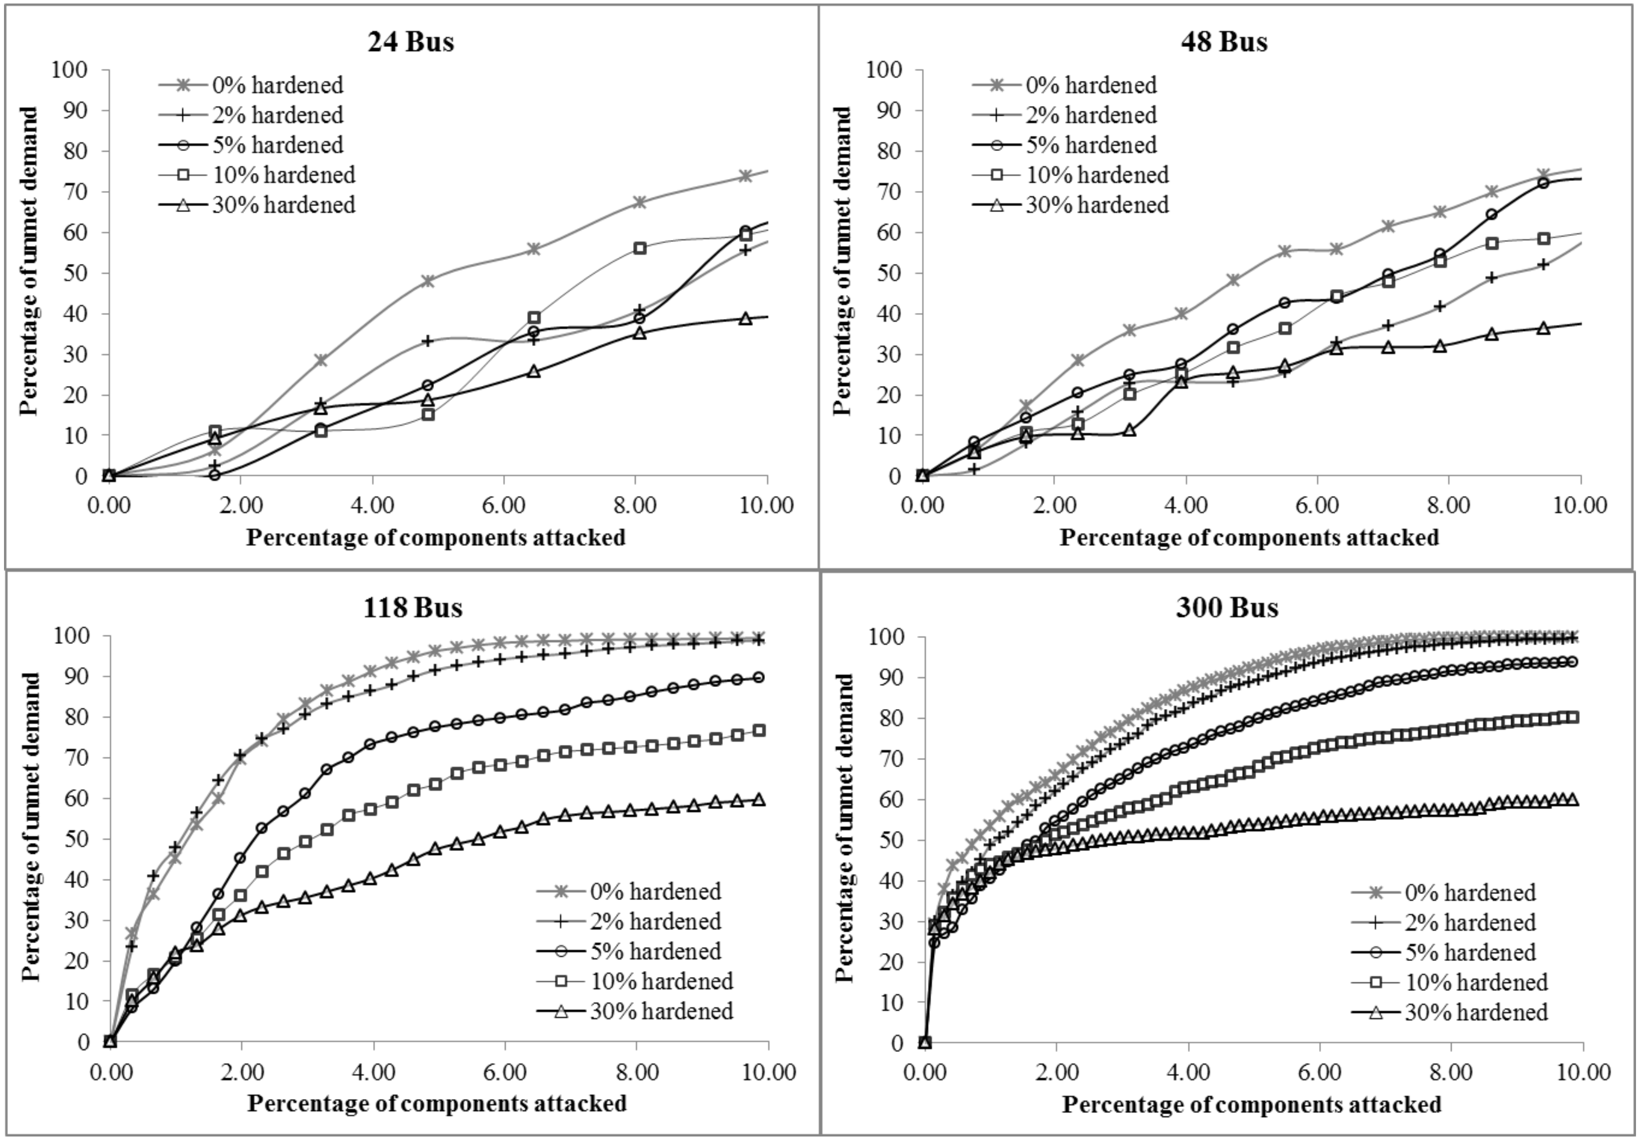

Figure 6.11– The effects of different levels of hardening for a dynamic attacker with cascading knowledge

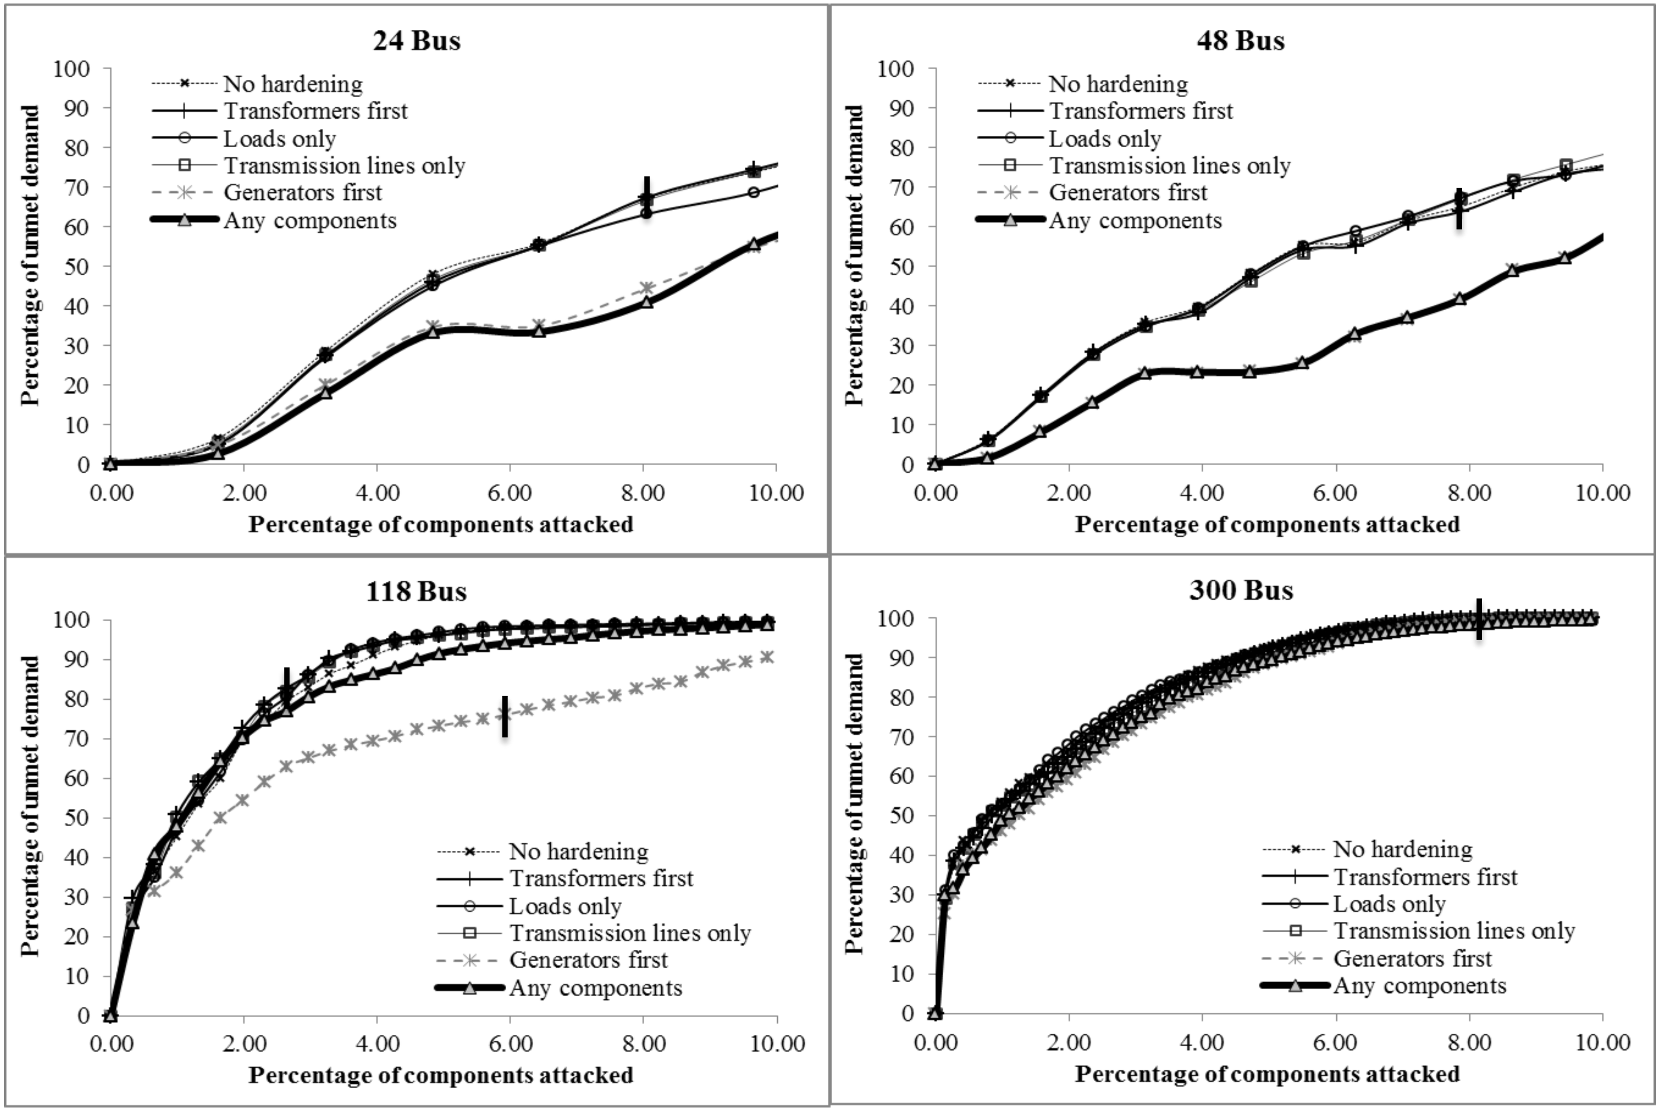

Figure 6.12–Various 2% hardening strategies against a dynamic attacker

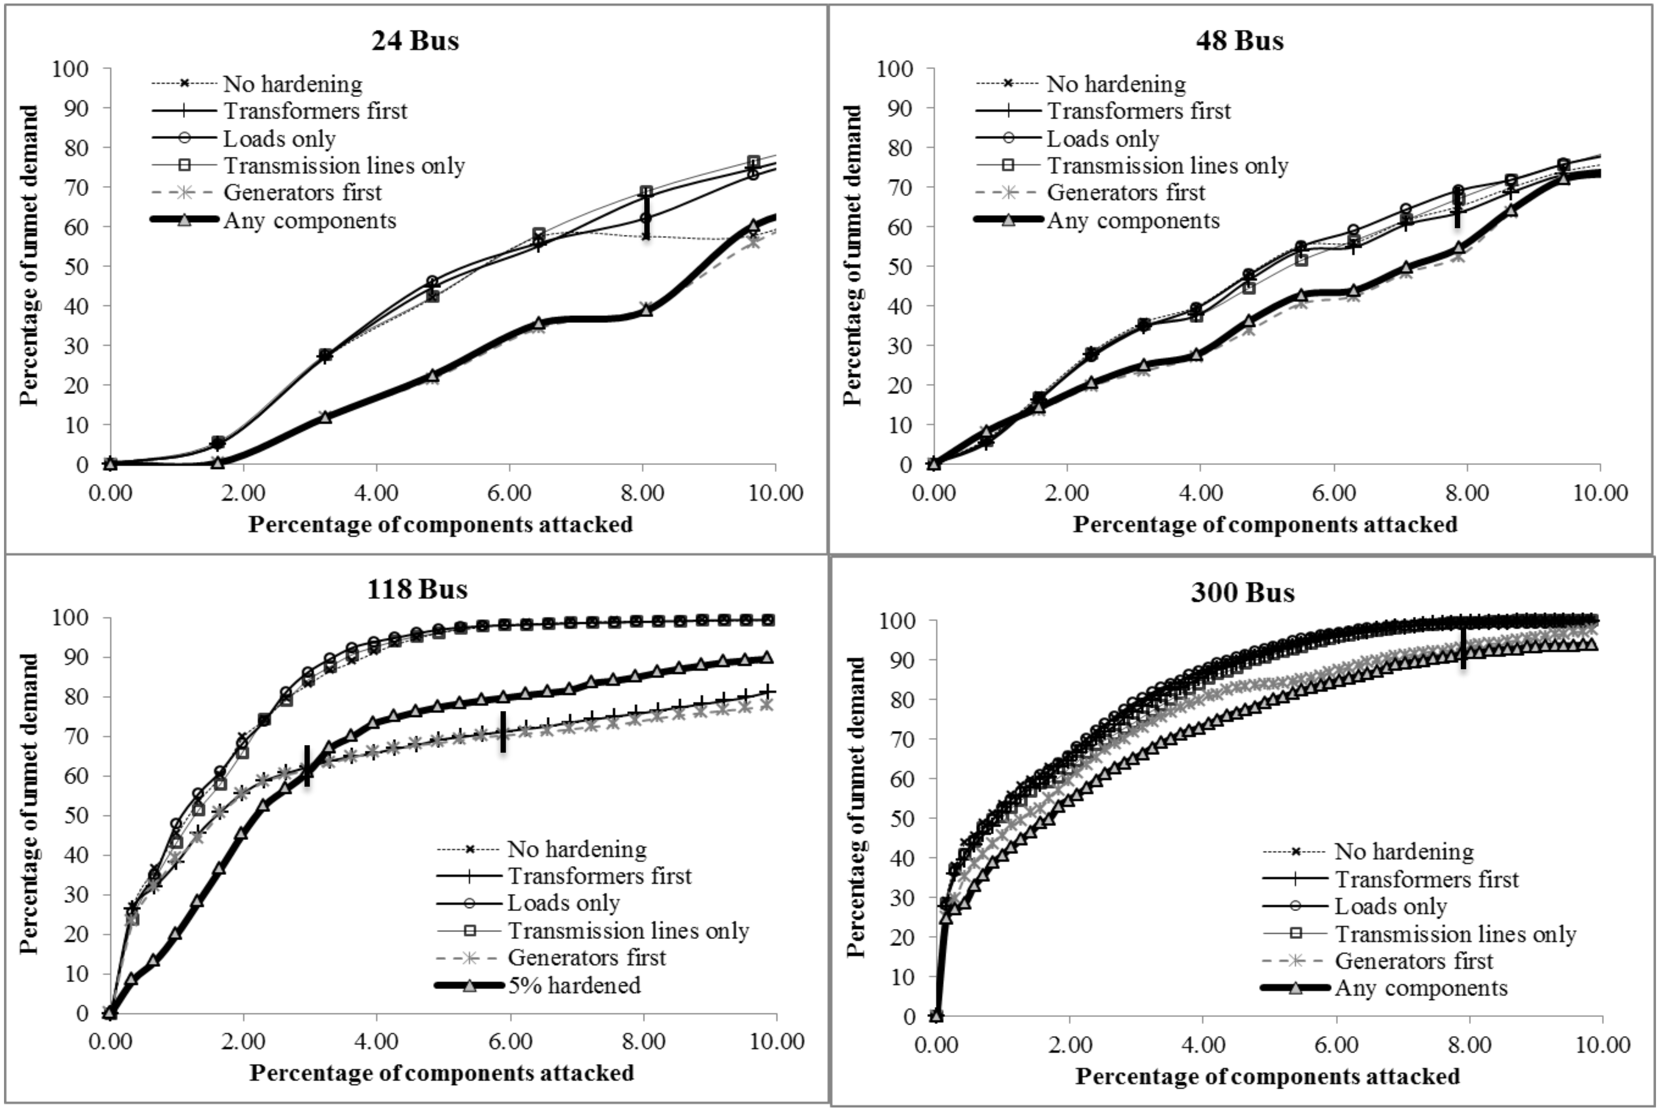

Figure 6.13–Various 5% hardening strategies against a dynamic attacker

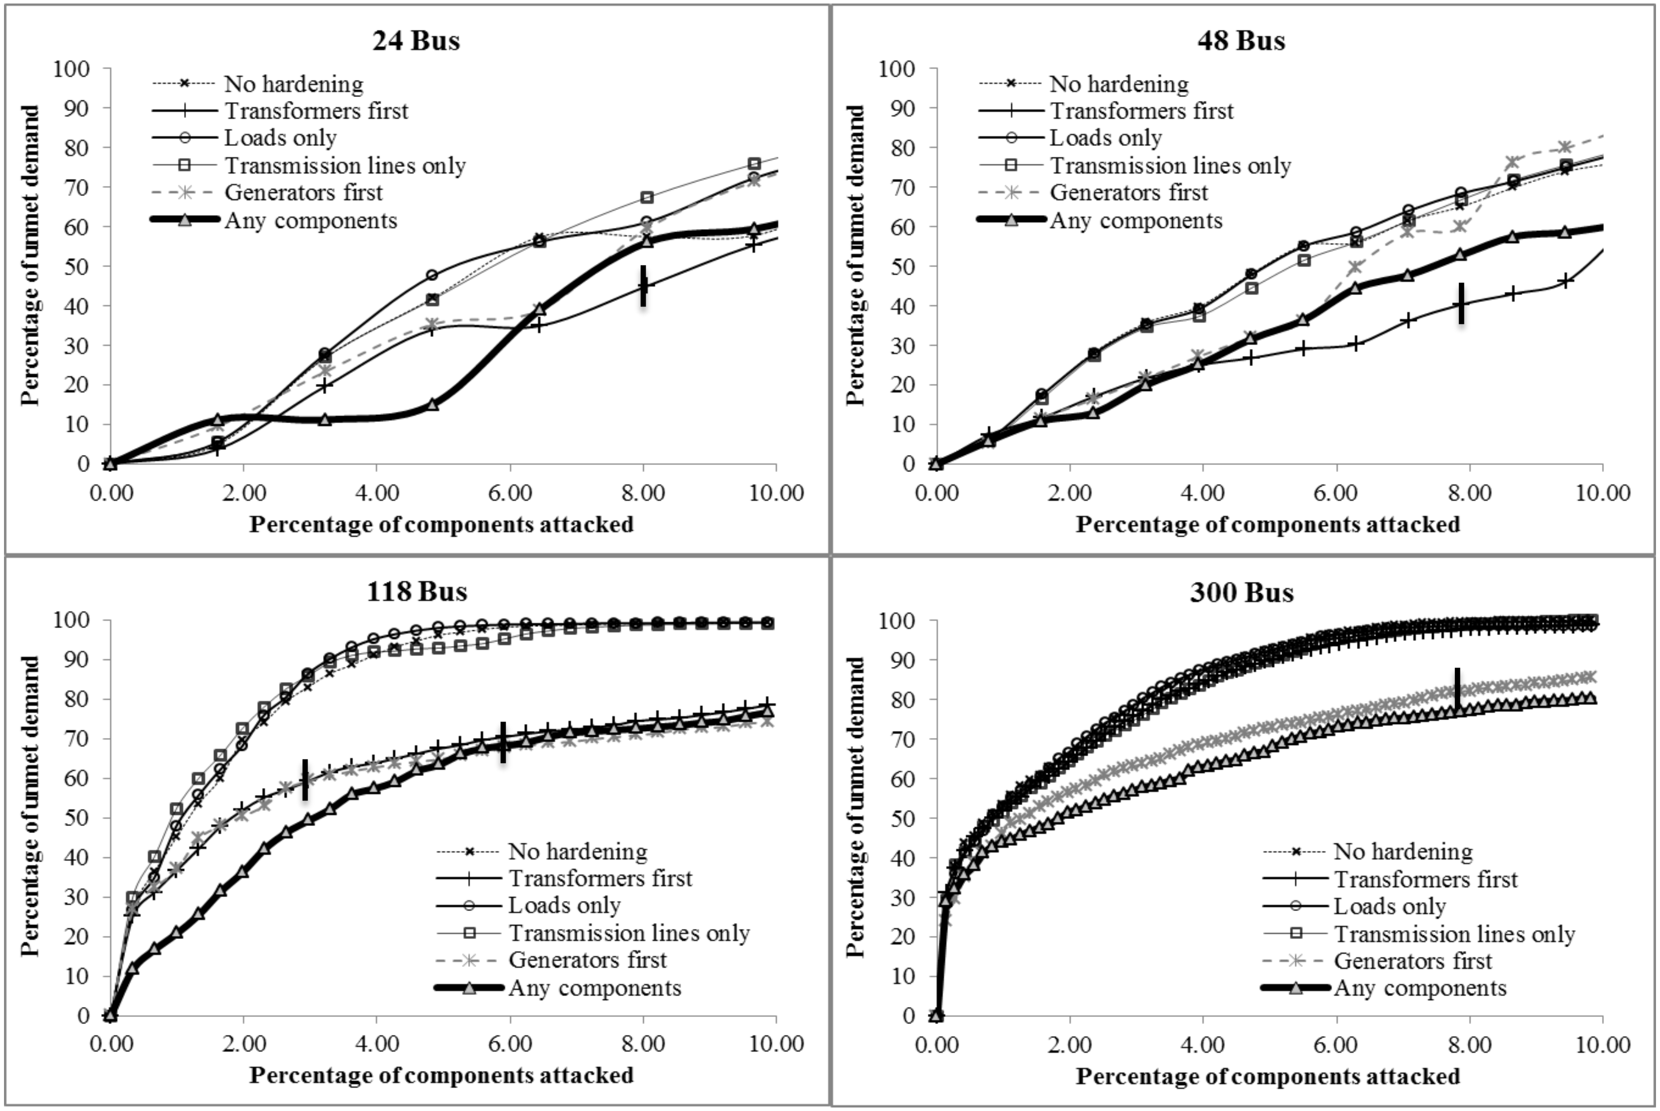

Figure 6.14–Various 10% hardening strategies against a dynamic attacker

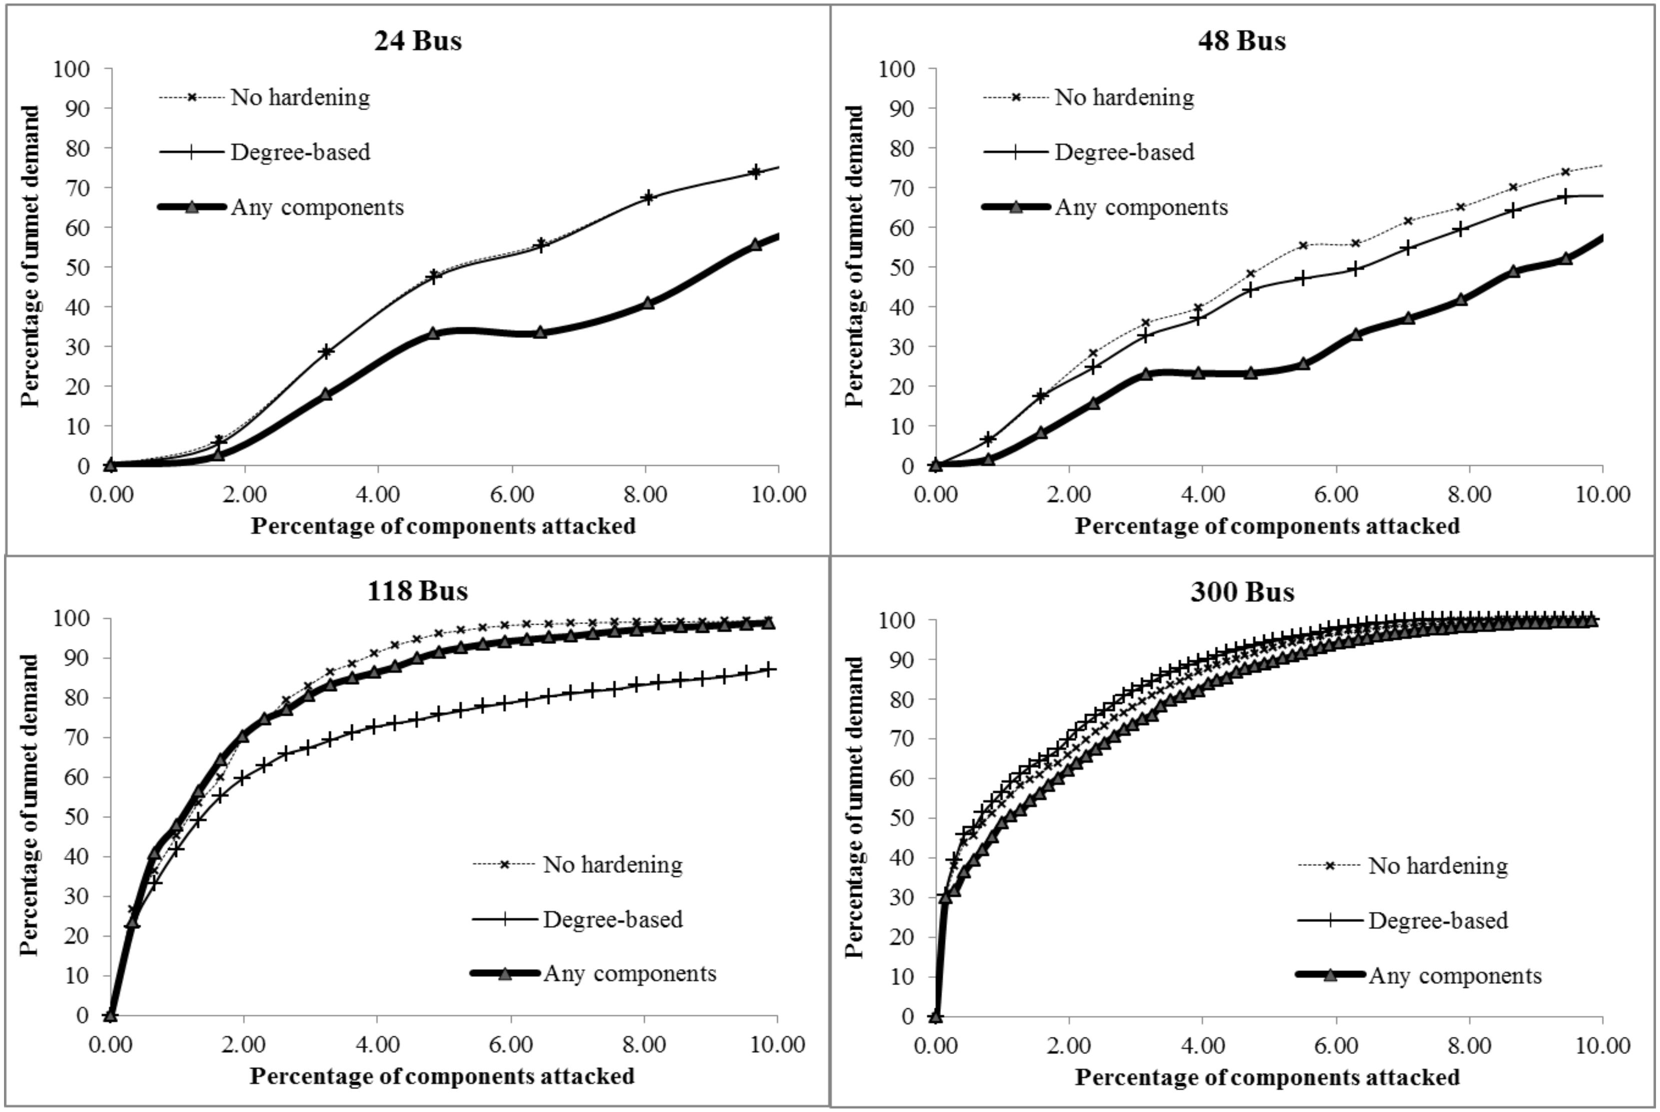

Figure 6.15–Dynamic hardening versus degree-based hardening with cascading failure (2% hardening scenario)

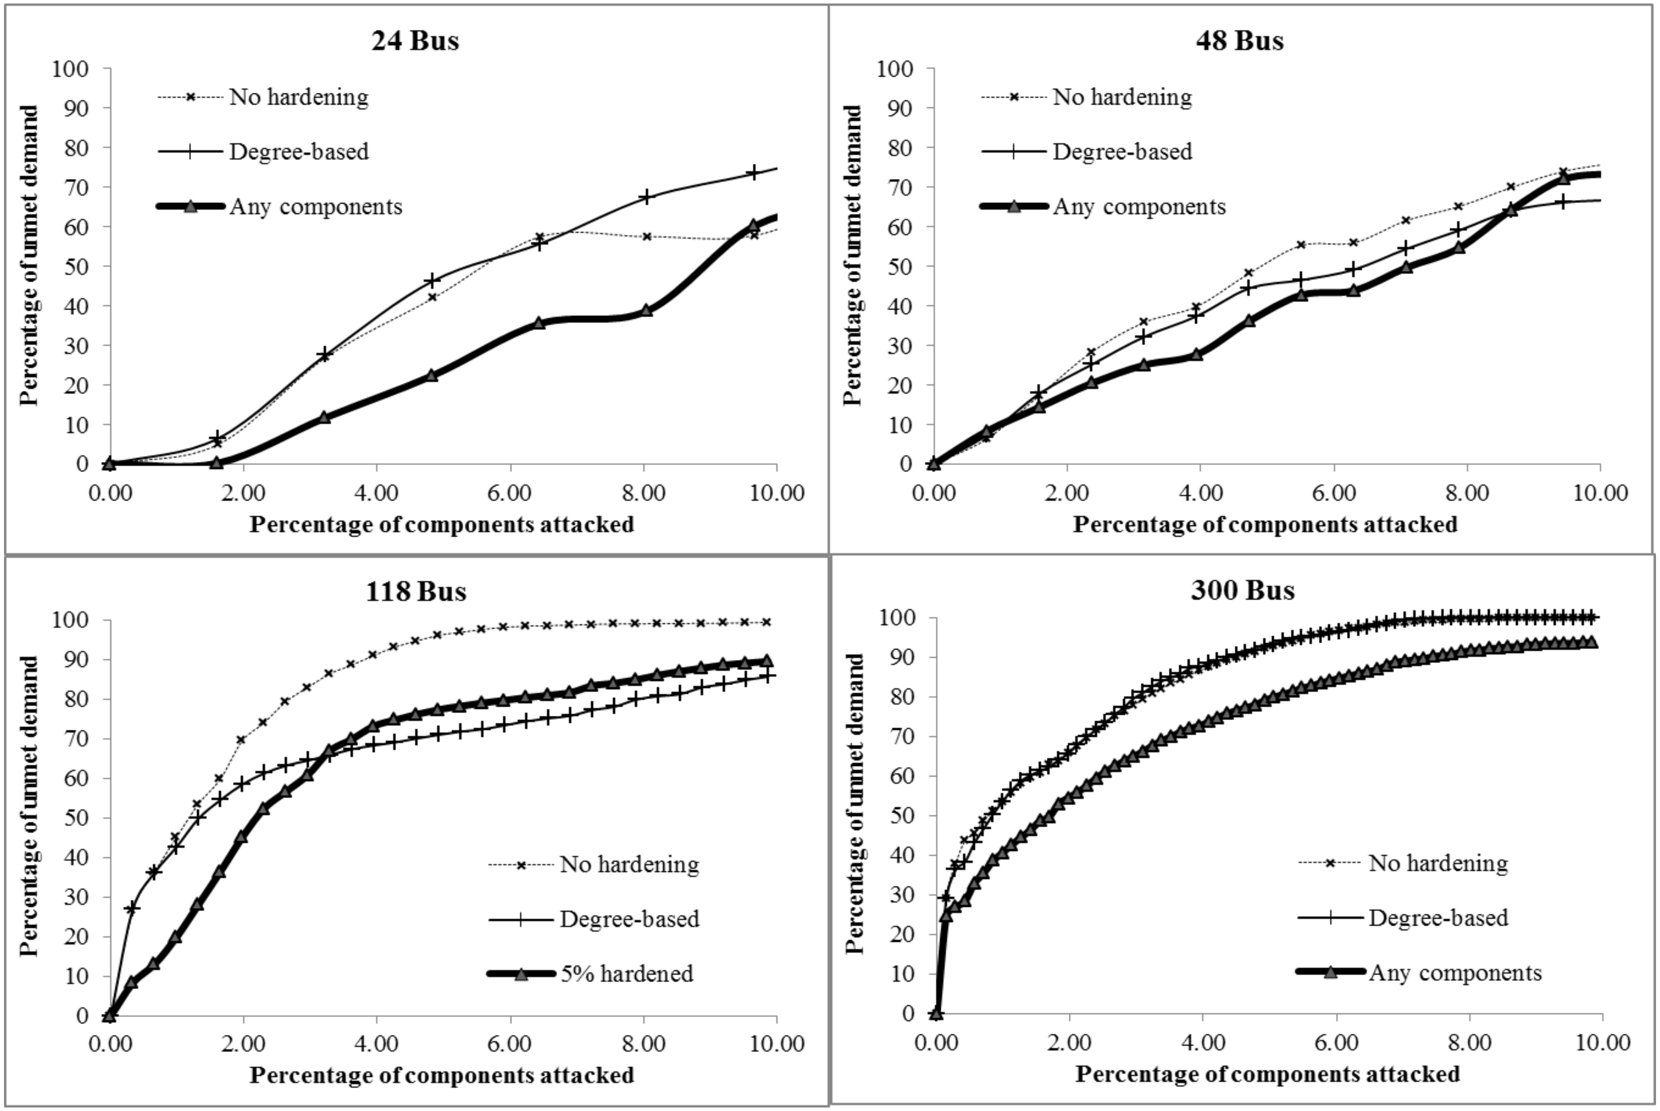

Figure 6.16–Dynamic hardening versus degree-based hardening with cascading failure (5% hardening scenario)

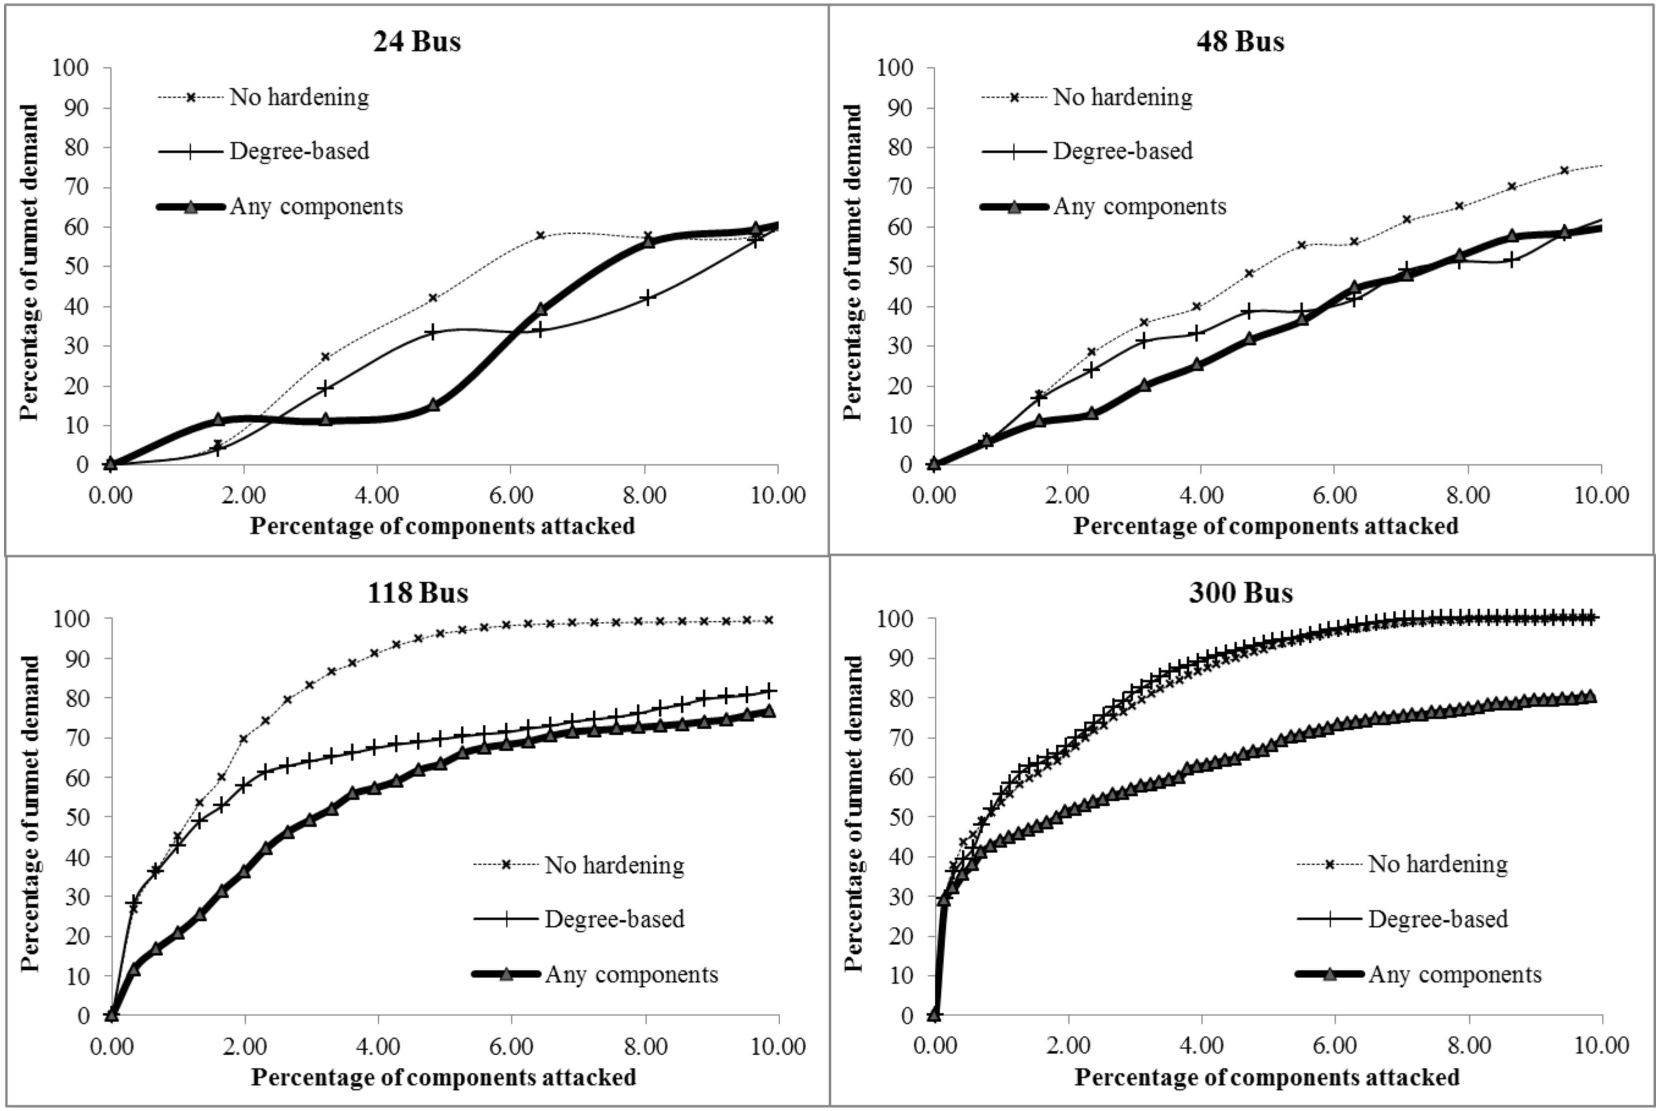

Figure 6.17–Dynamic hardening versus degree-based hardening with cascading failure (10% hardening scenario)

6.3.2. Doubling the Capacity of Transmission Lines

In this section, we double the capacity of a certain percentage of transmission lines that are most likely to cascade. Our goal is to make the system less prone to cascading due to capacity constraints, thus decreasing the amount of unmet demand. In order to identify the critical transmission lines, we assume two scenarios: small attacks (1% of the components); and large attacks (10% of the components). In order to identify the critical transmission lines, we run 100 simulations and rank the transmission lines based on the number of times they cascade. We compare three investment sizes; doubling the capacity of small (2%), medium (5%), and large (10%) numbers of transmission lines.

Figures 6.18 and 6.19 demonstrate the level of unmet demand for various levels of capacity investments in the components that are mostly likely to cascade given small and large dynamic attacks, respectively (Ideally, we would like to confine our investments to the lines that are unlikely to be attacked, but likely to cascade). For comparison purposes, the figures also show the level of unmet demand if the cascading impact has been mitigated completely (no cascade). This enables us to measure the percentage improvement in unmet demand considering cascading failure, from no change in unmet demand (the 0% investment case) to 100% improvement (the no-cascade case).

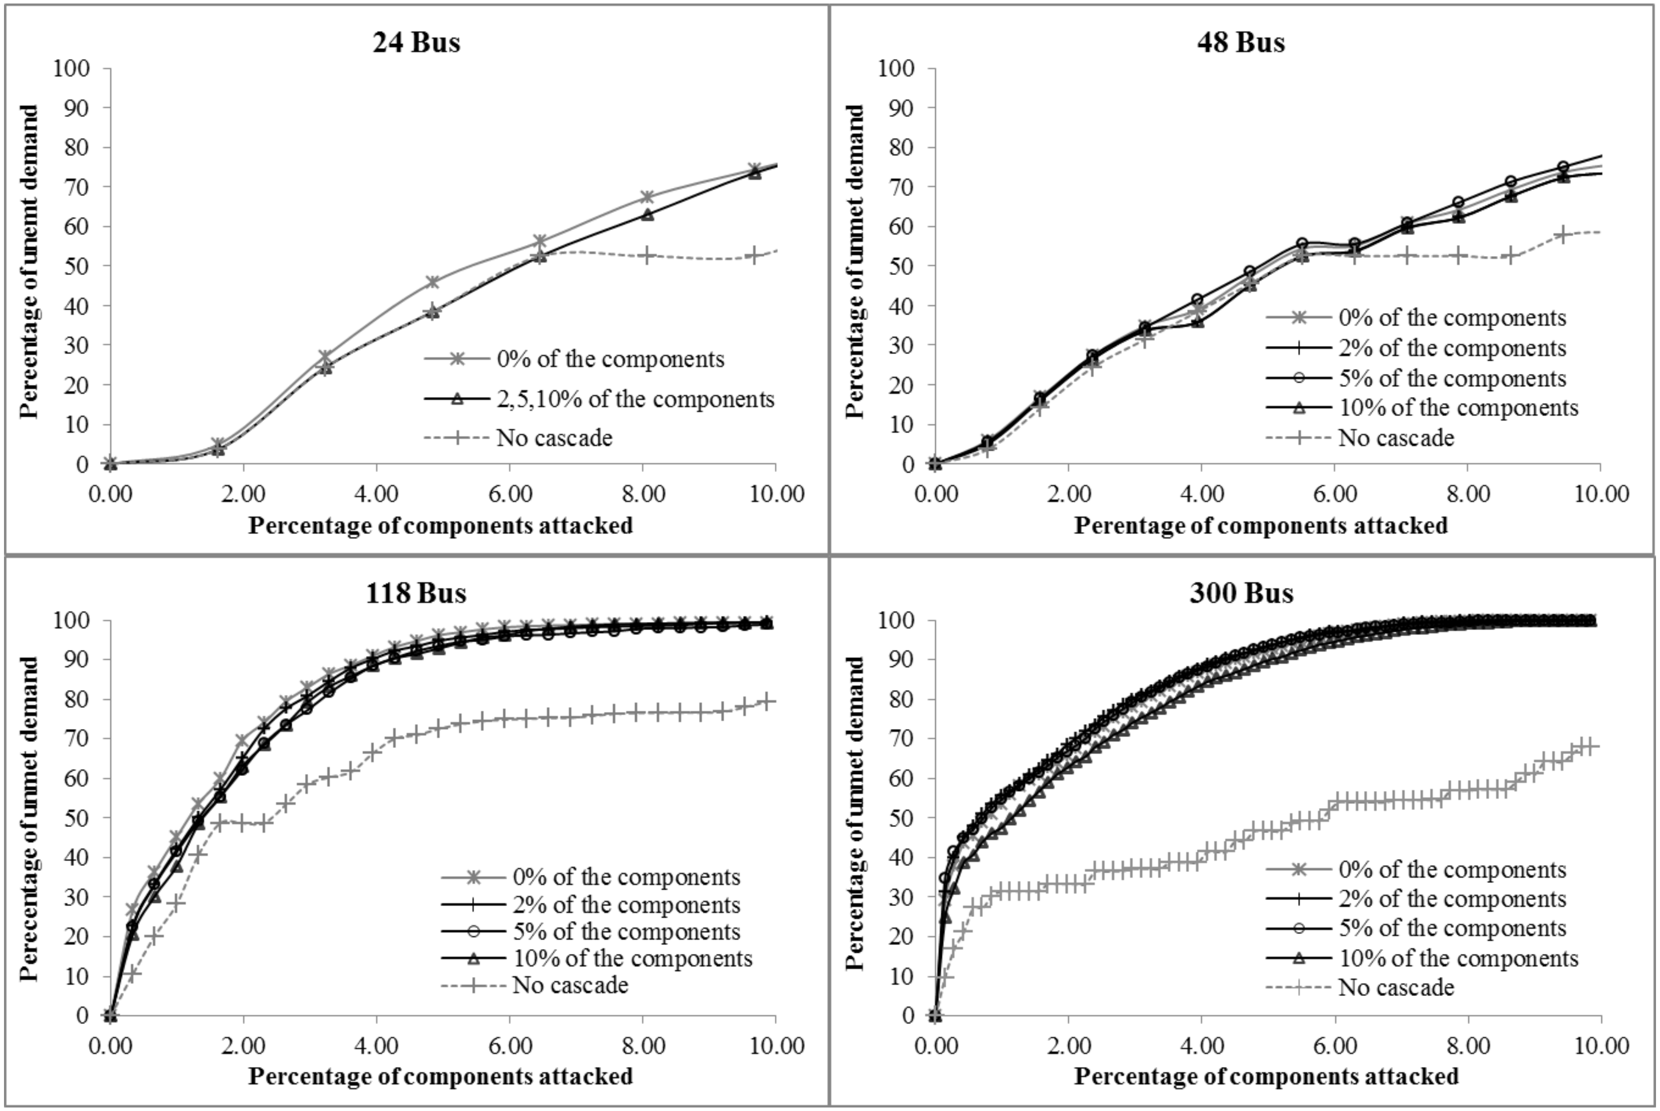

Figure 6.18– Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a small attack (dynamic attacker)

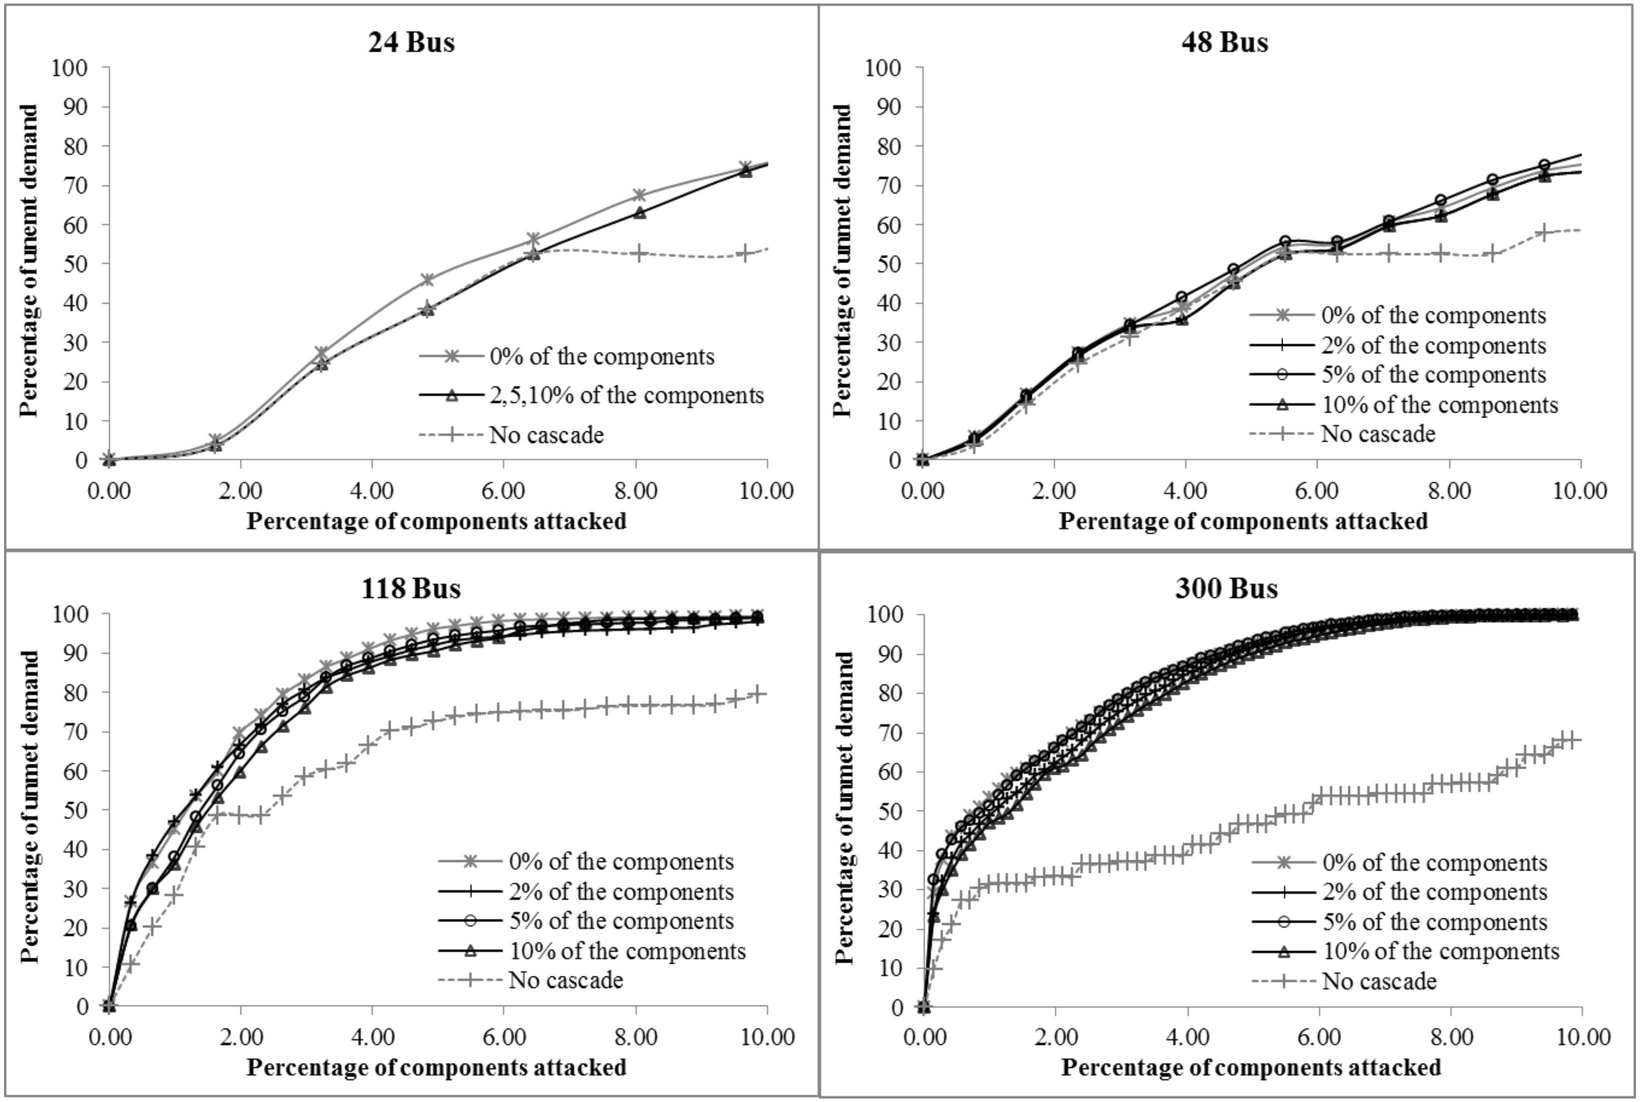

Figure 6.19– Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a large attack (dynamic attacker)

As discussed in Section 6.1, for reasons of conservatism, we present results only for a dynamic attacker (who can observe the cascading failure resulting from earlier attacks before choosing which components to attack next). With a dynamic attacker, the special cases of 0% investment and no cascade are not actually upper and lower bounds on the achievable improvement (unlike in the case of a static attacker); for example, increasing the capacity on some lines could change the flows in the system in such a way that attacking the highest-flow component in the new system could actually do more damage. However, the special cases of 0% investment and no cascade are still presented for comparison purposes.

For the dynamic attacker, it is also possible to have more than 100% (or less than 0%) improvement, since increasing the capacity on some lines changes the flows in the system, which may change the components to be attacked. By contrast, since the static attacker predetermines which components to attack, increasing capacity on some lines cannot make the system better than the no-cascade case.

As can be observed from Figures 6.18 and 6.19, the system improvement as a result of 10% capacity investment against a small attack (1%) is non-negligible. For example, in Figure 6.16, when 1% of the components are attacked, the percentage improvement due to capacity increases ranges from 100% for the 24-bus system to about 27% for the 300-bus system. However, the improvement due to increased capacity is much smaller when a larger percentage of components are attacked. In both Figures 6.18 and 6.19, the improvements are quite small for large attacks. Moreover, assuming small or large attack scenarios may result in improving the capacity of different transmission lines, but yield similar improvements in unmet demand. (Results are similar for the static attacker, either with or without cascading knowledge; see Appendix G.)

As a result, we can conclude that for our model large capacity investments can improve the system's ability to withstand small attacks. However, the unmet demand anticipated to occur from small attacks is modest, so even a significant reduction in unmet demand may not be enough to justify capacity improvements.

6.3.3. Adding New Transmission Lines

As an alternative, we could add new transmission lines parallel to the transmission lines that are likely to cascade. These critical transmission lines are determined as in Section 6.2.2. Note that the added lines are assumed to be not collocated with the original lines; hence, if one of the parallel transmission lines is attacked; the other one is assumed not to be damaged. These additional transmission lines increase the network size, but also provide alternative flow paths.

Figures 6.20 and 6.21 demonstrate the level of unmet demand for various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade given small and large attacks, respectively. As in Section 6.2.3, we compare various levels of investments in new transmission lines with the case when cascading impact is fully eliminated (no cascade).

The results are similar to the results for increasing the capacity of some lines. In particular, we observe little improvement for large attacks, and modest improvement for small attacks. However, as discussed before, the unmet demand anticipated to occur from small attacks is relatively small, so even a significant reduction in unmet demand may not be enough to justify adding new transmission lines. Note also that, in some cases, adding new lines can create a negative impact on the network; for example, see 48-bus system in Figures 6.20 and 6.21.

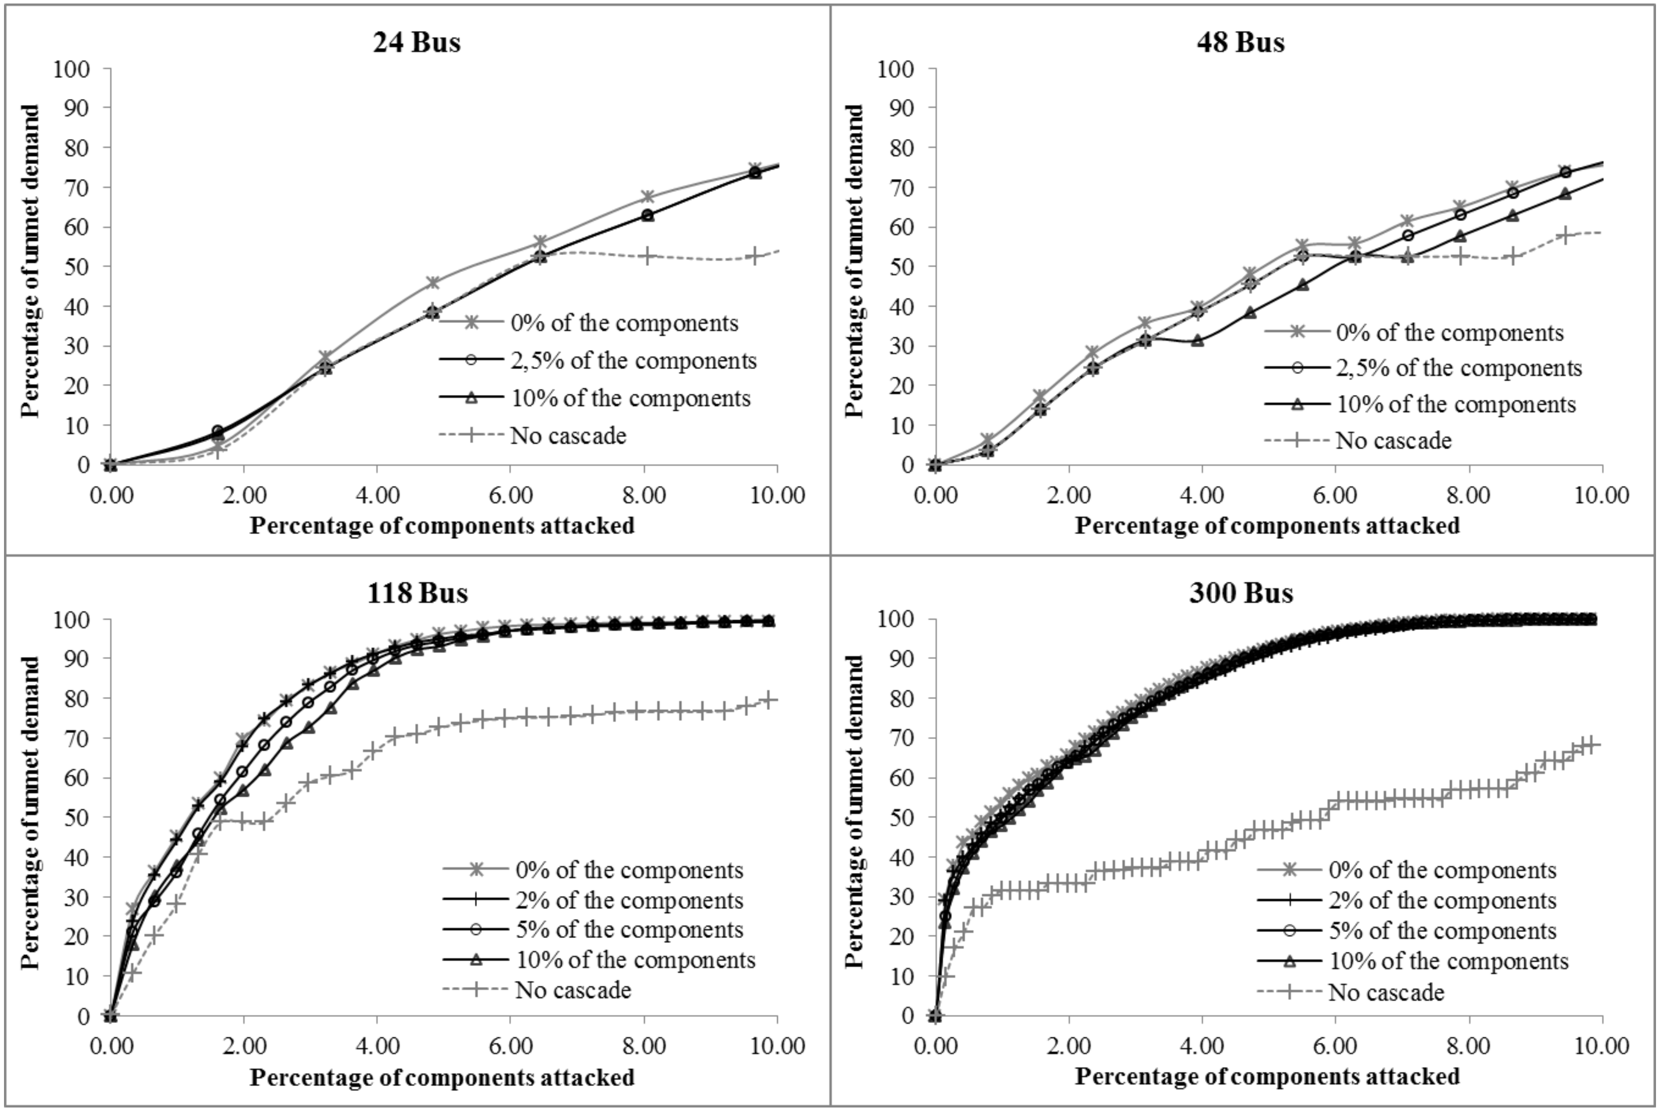

Figure 6.20– Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a small attack (dynamic attacker)

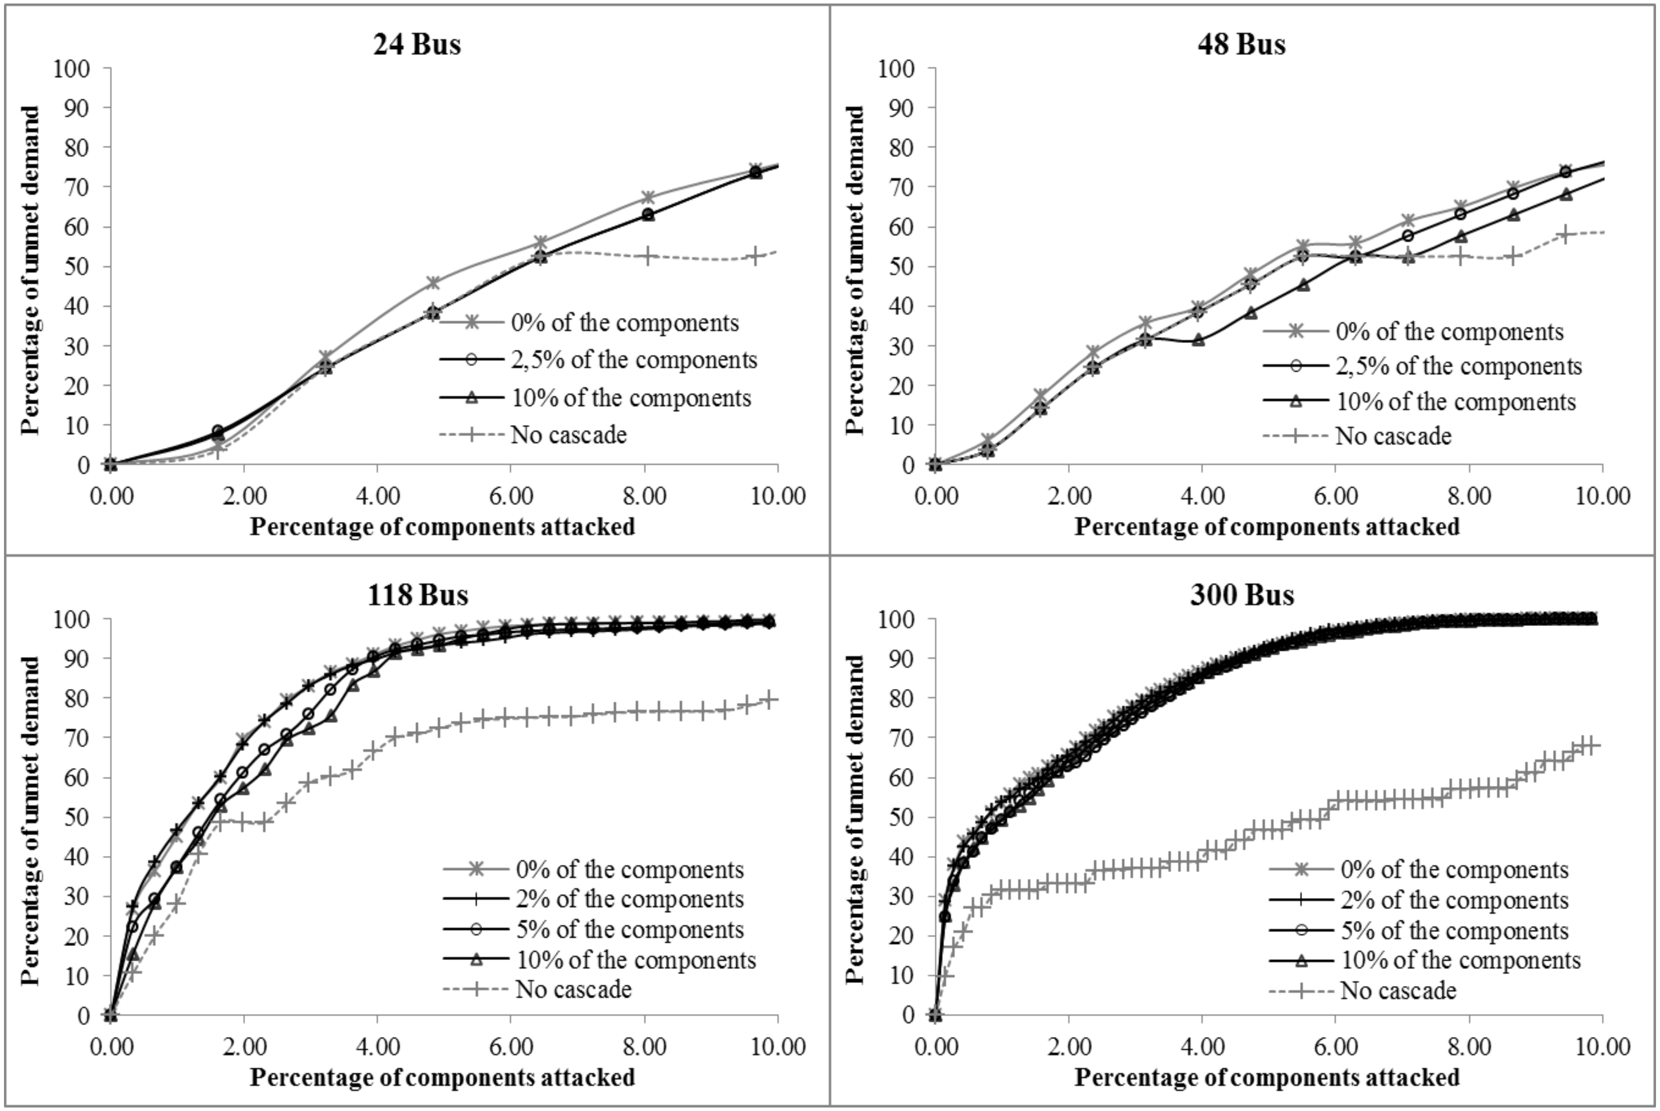

Figure 6.21– Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a large attack (dynamic attacker)

Negative impact can occur for two possible reasons. First, if one of the two parallel lines fails, it is likely that the other line will become heavily loaded, making it a candidate for both attack and cascading failure. Second, since the total capacity of the overall route has now doubled (as there are now two lines) the flow on other lines leading into or out of the region with the new line might increase (i.e., lack of coordination between both ends of line), creating a risk of cascading failure if some of these lines have lower capacities. Overall, however, the effect tends to be small in our examples, and the impact of adding new transmission lines is usually favorable. (Results are similar for the static attacker with or without cascading knowledge, see Appendix H.)

6.3.4. Comparing Three Defensive Investments

In this section, we will compare hardening, increasing capacity, and adding new lines. See Figures 22-24 for small (2%), medium (5%), and large (10%) investments, respectively (Note that the new lines to be added are the ones to protect against a large attack). In most cases, hardening does significantly better than the other investment types. For a small investment, improvements other than hardening have negligible benefit. For medium and large investments, adding new lines does slightly better than increasing capacity, but hardening components still yields significantly greater improvement. (For the case of a static attacker with or without cascading knowledge, see Appendix J.) Note, however, that hardening is possible for all components, whereas capacity improvements and adding new components are considered only for transmission lines. As you might recall from Section 6.3.1, hardening transmission lines only does not create a significant benefit even when 10% of the transmission lines are hardened. As a result, the superiority of hardening over other investment strategies appears to be because transmission lines are not typically the most vulnerable parts of the system studied here.

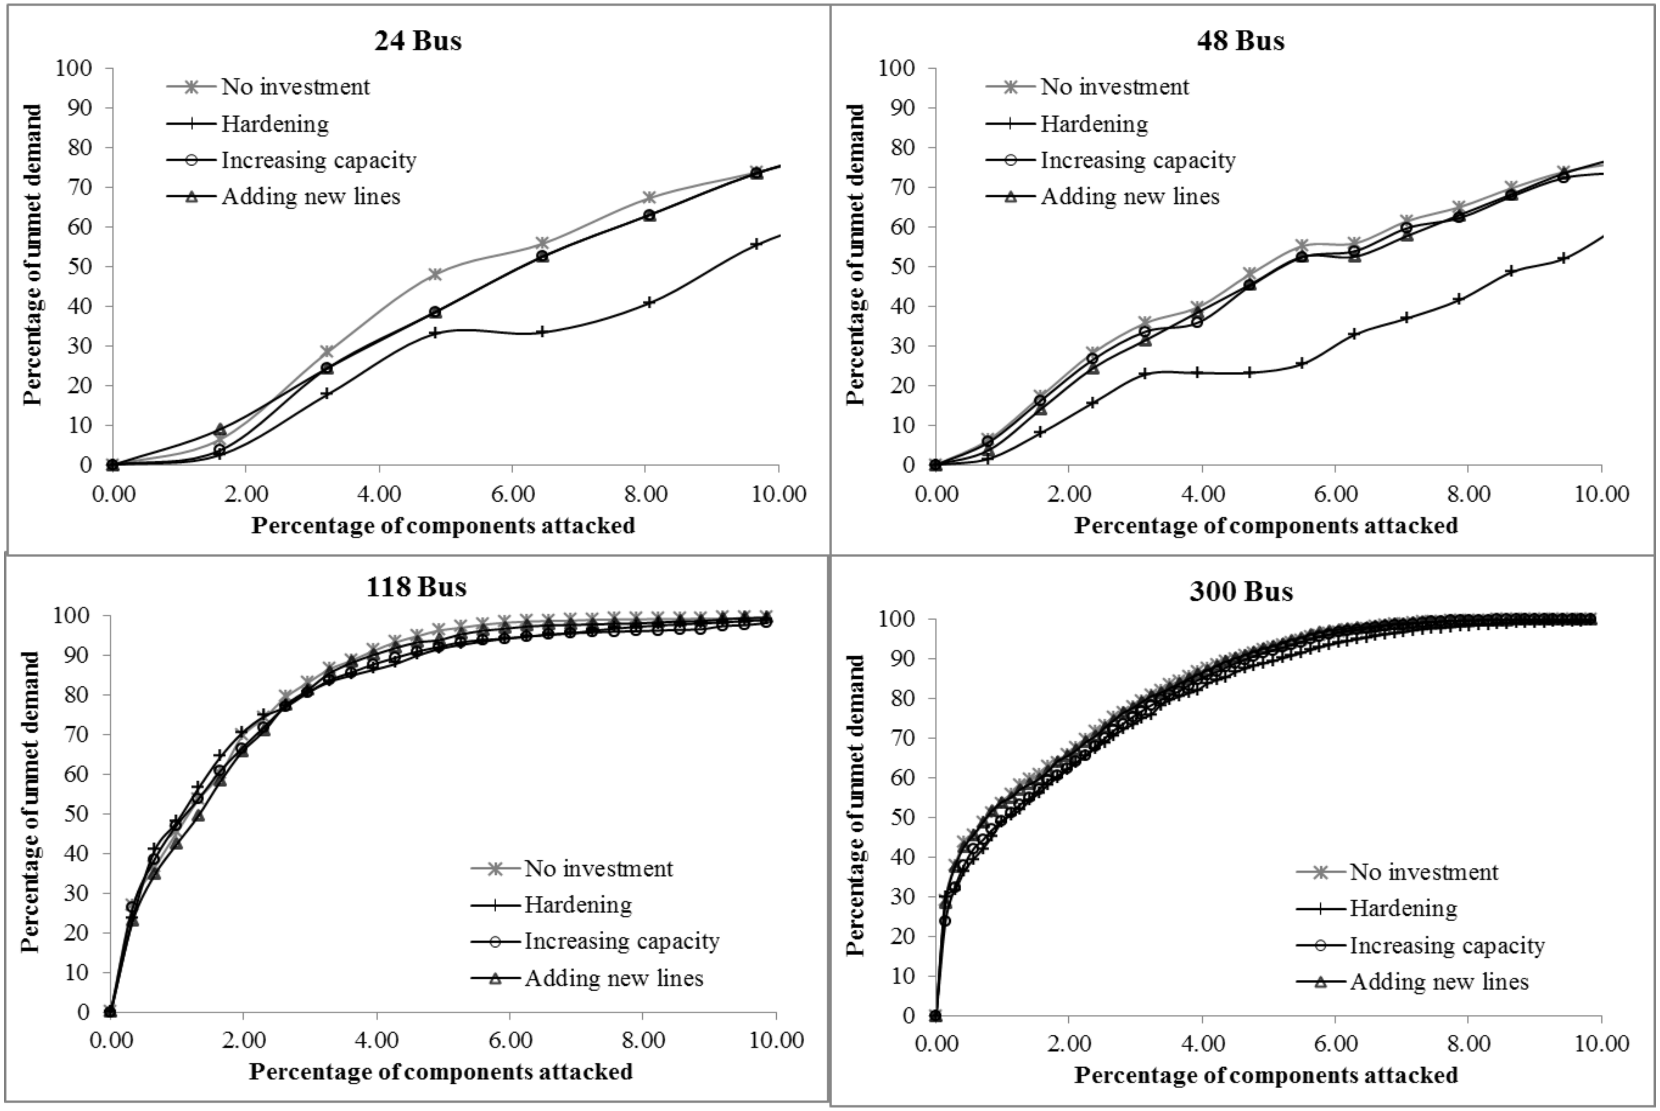

Figure 6.22–Comparing defensive investments against a dynamic attacker (2% investment)

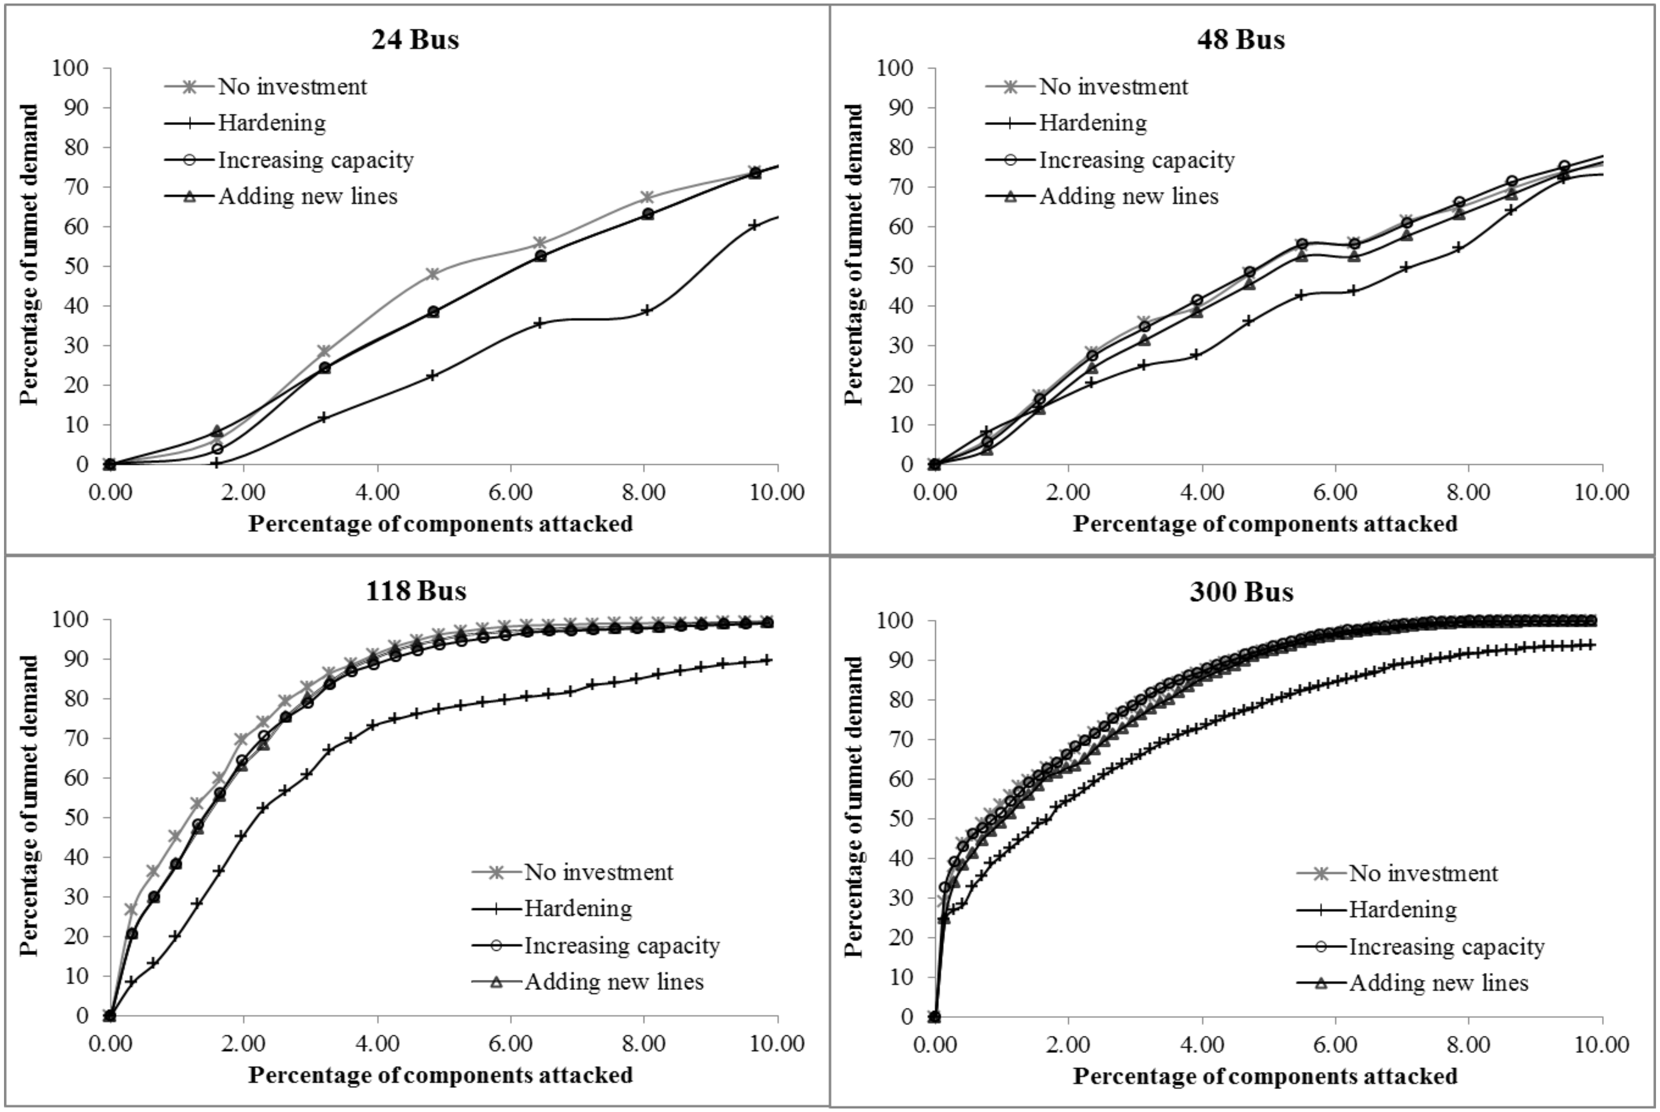

Figure 6.23–Comparing defensive investments against a dynamic attacker (5% investment)

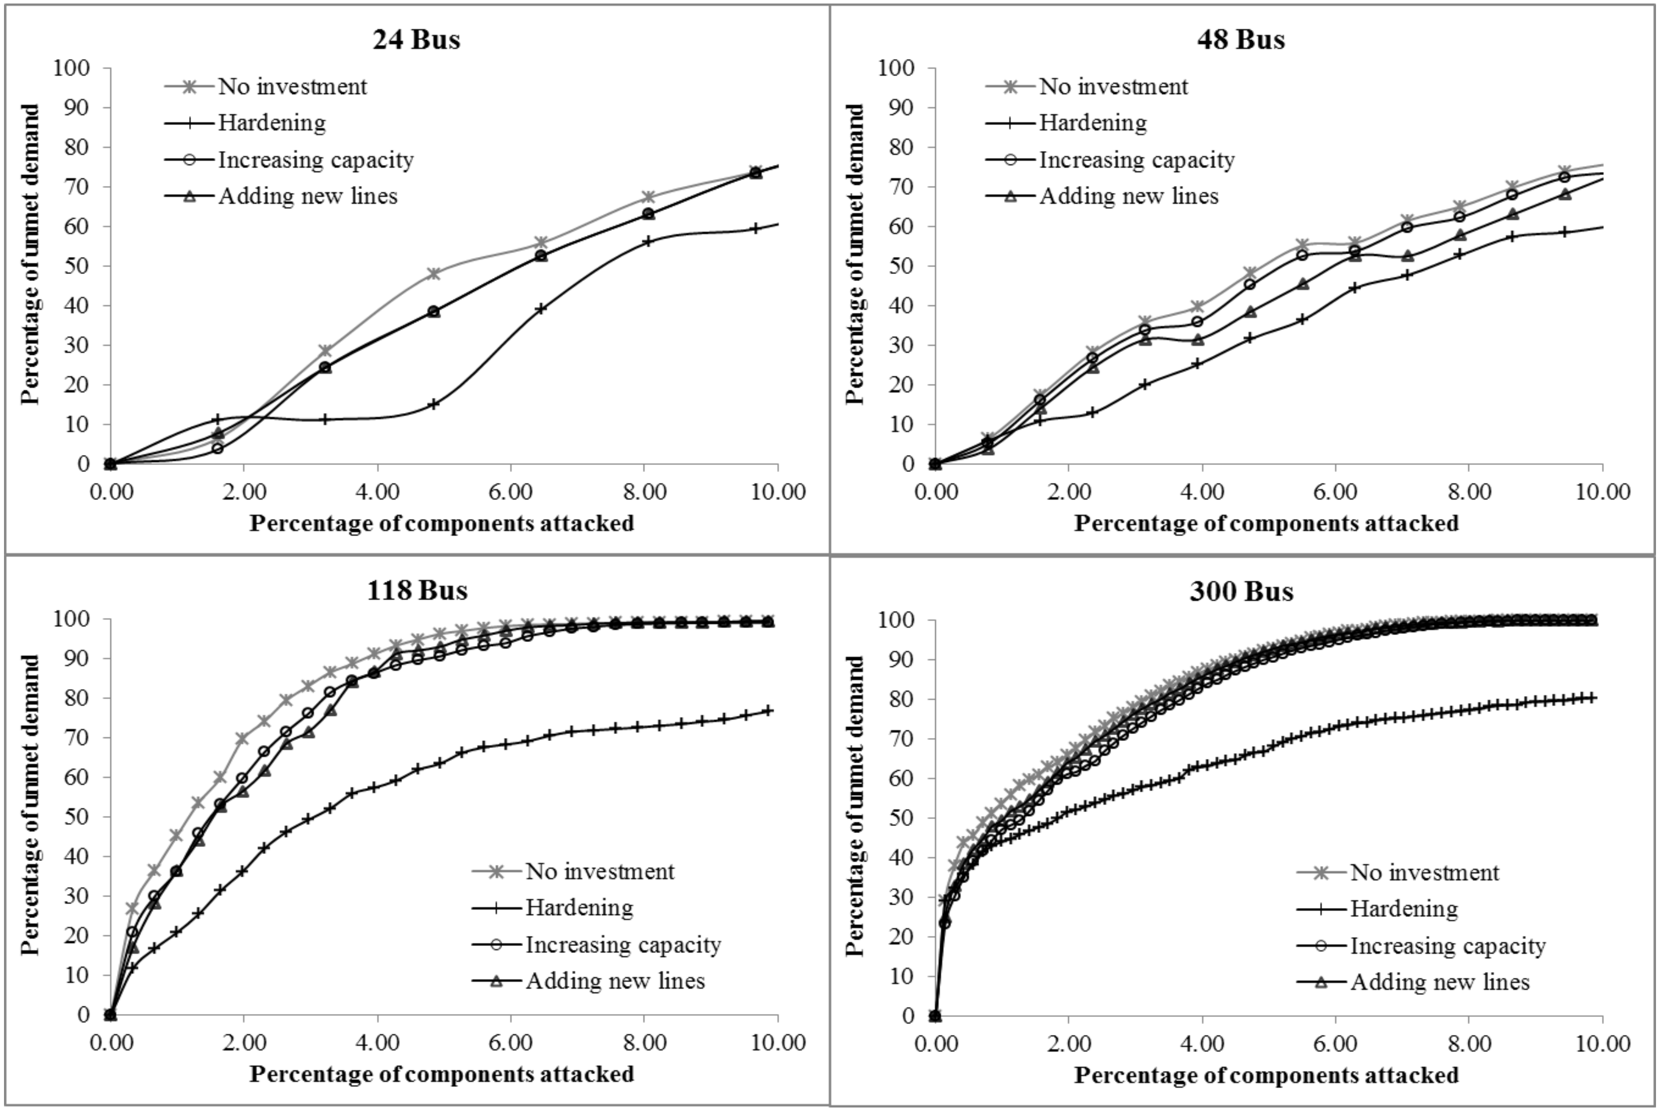

Figure 6.24– Comparing defensive investments against a dynamic attacker (10% investment)

6.4. Comparing Results in the Literature

As discussed in Chapter 3, there are few models in the literature that consider both cascading failure and an intelligent adversary. To our knowledge, Hines et al. (2010) is the only model that provides comparable results to ours.

In their paper, Hines et al. investigate the effectiveness of topological models by comparing their flow-based attack strategy with some topological attack strategies, such as degree-based attacks and attacks against the nodes that are included in the largest number of shortest paths (betweenness-based attacks). Using random failures as their base case, they measure how much the unmet demand as a result of an intentional attack deviates from the average unmet demand due to a comparable number of random failures. As in our model, their flow-based model uses DC-power dispatch, a linearized version of the nonlinear power equations. However, their flow-based attacks are simultaneous rather than sequential, with cascading failures being determined only after all attacks have occurred (unlike in our model, where cascading failure can occur after each individual component is attacked).

The cascading-failure model used by Hines et al. has some other differences from our model. One fundamental difference is that they model cascading failure deterministically rather than probabilistically. In particular, once a component exceeds its rated flow by 50% for at least five seconds, it is assumed to fail irreversibly. In addition, greater the magnitude of an overload, the faster the component is assumed to fail; moreover, any component that exceeds its rated limit for long enough will eventually fail, even if the overload is extremely small (Hines, personal communication, May 2012). Hines also indicated that in his model, components fail one at a time, because of the central role of time in determining which components will experience cascading failure, whereas in our model, multiple components can fail simultaneously in each round of cascading failure if they are all operating close to their capacities. Finally, the model of Hines et al. does not assume any knowledge of cascading failure on the part of the attacker, and in that regard is similar to our static attacker with no cascading knowledge.

Note also that, the results in Hines et al. (2010) are presented as percent deviations from the average unmet demand due to random failures. However, because of the above modeling differences in their treatment of cascading failure, even the effects of random failure may differ significantly between their model and ours. Therefore, in order to do a fair comparison, we obtained the absolute values of unmet demand for their flow-based attack strategies (Hines, personal communication, May 2012). Figure 6.25 compares the absolute values of unmet demand achieved by the flow-based attacker of Hines et al. with the results of our model for both 300-bus systems.

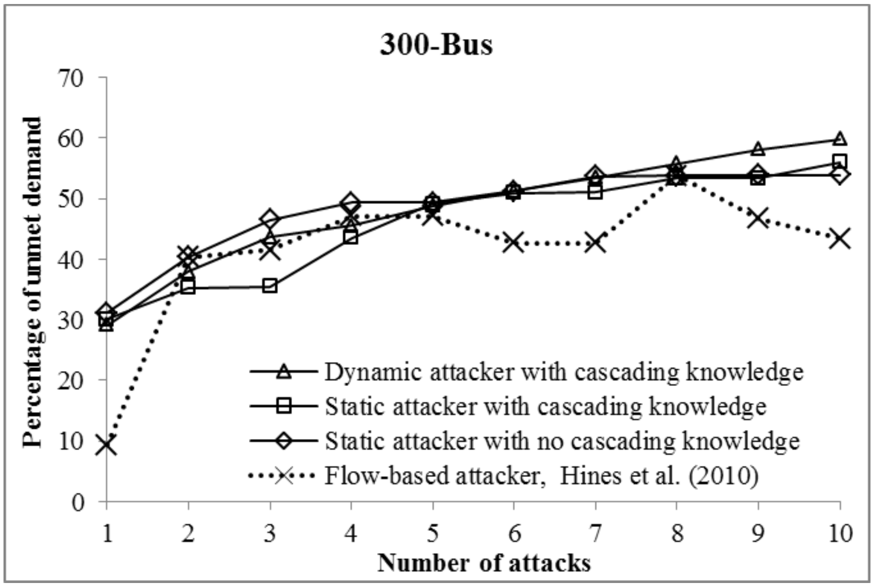

Figure 6.25– Comparison of our results with the flow-based results of Hines et al. (2010)

As can be seen from Figure 6.25, our results are roughly similar to those found by Hines et al. However, there is a significant difference between the unmet demands if the attacker attacks one component only, which mainly could be due to our multiple rounds of cascading assumption. Moreover, the model of Hines et al. generally predicts less unmet demand than our model, providing a rough validation of our model. Overall, the results of the two models are of the same order of magnitude.

6.5. Conclusion

The simplicity of our modeling approach makes it feasible to model cascading failure; something that is not done in many other models. Our results suggest that the impact of cascading failure is non-negligible, especially if the network is highly capacity-constrained (such as in the 118- and 300-bus systems). Moreover, it is fortuitous that our model is not overly sensitive to the values of its parameters, so accurate estimates of α and β (which might be difficult to obtain) are not critical.

Thus, our model provides a tool to analyze how an intelligent adversary might seek to take advantage of cascading failure. We discuss three possible assumptions about attacker knowledge of cascading failure; a static attacker without knowledge of the cascading failure, a static attacker with knowledge of the cascading failure, and a dynamic attacker who is able to observe cascading failures as they occur.

In this chapter, we also compared different investment strategies for decreasing the unmet demand from an attack. As in past work (Bier et al., 2007), we found that electric-power networks generally require substantial levels of investment to obtain any significant improvement in unmet demand. Moreover, hardening of components generally seems to provide greater improvement than investment strategies focused specifically on reducing cascading failure, such as increasing capacity or adding new transmission lines. In fact, due to the complexities of electric-power systems, we found that adding new lines sometimes actually increased the load lost due to cascading failure.

Thus, modeling cascading failure helped us to more realistically analyze the immediate impacts of an attack. However, intelligent attackers could also consider the long-term impact of an attack. Therefore, in the next chapter, we explicitly model restoration times, and compare various defensive investments in terms of total energy lost over time (rather than just the immediate level of unmet demand).

7. Results for Modeling Restoration Times

In Chapter 6, we modeled cascading failure and analyzed the immediate impact of an attack, given differing levels of knowledge about cascading failure on the part of the attacker. However, an attacker may also wish to create an impact that lasts for a long time. In our model, we allow attackers to choose which components to target not only based on their loads, but also based on the product of load and restoration time, as a measure of energy rather than capacity.

In this chapter, we first explore the impact of attack strategies that explicitly take restoration times into account (to cause large energy losses). Then, we compare this greedy attacker to a degree-based attacker. In the final section, we compare our earlier categories of improvement (hardening, increasing the capacity of transmission lines, and adding new transmission lines) with investment to reduce the restoration times of specific types of components (transformers).

As in Chapter 6, we do 20 simulation runs for each scenario, as long as the standard error of the mean (in this case, energy loss) is less than 3%; otherwise, we increase the number of simulations by 10 until we obtain a standard error less than 3%. Note, however, that the simulation is needed only to estimate the impact of cascading failure; our model of restoration times is entirely deterministic. As discussed in Chapter 4, we assume a restoration time of one day for transmission lines that fail due to cascading failure, three days for loads and for transmission lines that fail due to direct attacks on those lines, 15 days for generators, and 32 days for transformers.

7.1. Impact of Considering Restoration Times

In this section, we analyze what happens if the attacker takes the restoration times of components into account in choosing which components to target. To show the impact of restoration times, we first analyze the 118-bus system in detail, and then present the results for other systems. Figure 7.1 demonstrates the impact achieved by an attacker who considers restoration times for the 118-bus system. In particular, we consider a dynamic attacker with cascading knowledge, in order to observe how the attacker behavior changes in our most conservative case.

The graph in the upper right corner demonstrates the restoration of the system after 10% of its components have been attacked. Even though the initial impact of the attack was similar in both cases (as can be seen from the graph on the upper left corner), when the attacker considers restoration times, the long-term impact of an attack is higher. In fact, as can be seen from the graph in the lower right corner, the total energy loss (the area under the curve in the upper right corner) is much higher than when the attacker does not consider restoration times. This example clearly demonstrates why considering restoration times might be important to an attacker, and therefore also to the defender.

The graph in the lower left corner demonstrates the percentage of attacked components that are of each type for a scenario in which 10% of components are attacked. In this example, while the attacker who does not consider restoration times does not attack any transformers to attack, for the attacker who does consider restoration times, more than 20% of the components attacked are transformers.

Figures 7.2 and 7.3 show comparable results for static attackers with and without cascading knowledge, respectively, also for the 118-bus system. Results are roughly similar (more emphasis on attacks against generators and transformers, but less total energy loss). Therefore, in the rest of this chapter, we will present results only for a dynamic attacker with cascading knowledge. [Results for the other attacker types are given in the appendices.] Also note that for the 118-bus system, static attacks chosen with consideration of restoration times actually cause slightly more initial unmet demand than attacks chosen without regard to restoration times, simply illustrating that static attack strategies do not always yield good performance. For completeness, Figures 7.4 through 7.6 show similar results to those in Figure 7.1 for the 24-bus, 48-bus, and 300-bus systems, respectively. As in the 118-bus system, the components attacked tend to shift from lines and loads to generators and transformers. While the initial impact stays roughly similar when the attacker considers restoration times, the total energy loss increases significantly. [See Appendix J for 24-, 48-, and 300-bus systems with a static attacker, and Appendix K for a comparison of dynamic and static attack strategies. The total energy loss for static attack strategies is comparable to the energy loss for the dynamic attack strategies, even though the mix of components attacked is often different.]

Previous figures have shown that when the attacker considers restoration times, generators and transformers are more frequently attacked. Therefore, it seems sensible to compare how the total energy loss from an attacker who considers restoration times differs from the energy loss that can be achieved by targeting just generators or just transformers. In Figures 7.7 through 7.10, we compare these candidate attack strategies for 24-, 48-, 118-, and 300-bus systems, respectively, for the case of a dynamic attacker with cascading knowledge. Similar to the comparisons in the previous chapters, the number of transformers (except for the 300-bus system), and the number of generators for 118- and 300- bus systems is less than 10%; hence the attacker chooses other components once all transformers on generators are attacked. In Figures 7.7 through 7.10, those attacks are separated from attacks on transformers/generators by a vertical bar, |.

As can be seen from Figures 7.7 through 7.10, attack strategies that consider restoration times do about as well as targeting generators first for the 24-bus and 48-bus systems, and about as well as targeting transformers first for the 118-bus and 300-bus systems. [See Appendix L for a similar comparison for static instead of dynamic attackers; results are similar.]

Finally, we compare a dynamic attacker who considers restoration times with a degree-based attacker. In particular, we will compare the degree-based attacker with two types of attack strategies (both considering restoration times); a static attacker with no cascading knowledge (our least conservative attack strategy), and a dynamic attacker with cascading knowledge (our most conservative attack strategy). See Figures 7.11 through 7.14 for the 24-bus, 48-bus, 118-bus, and 300-bus systems, respectively. As can be seen from these figures, in all circumstances, even our static attacker with no cascading knowledge does significantly better than a degree-based attacker in both the initial attack phase and the long run.

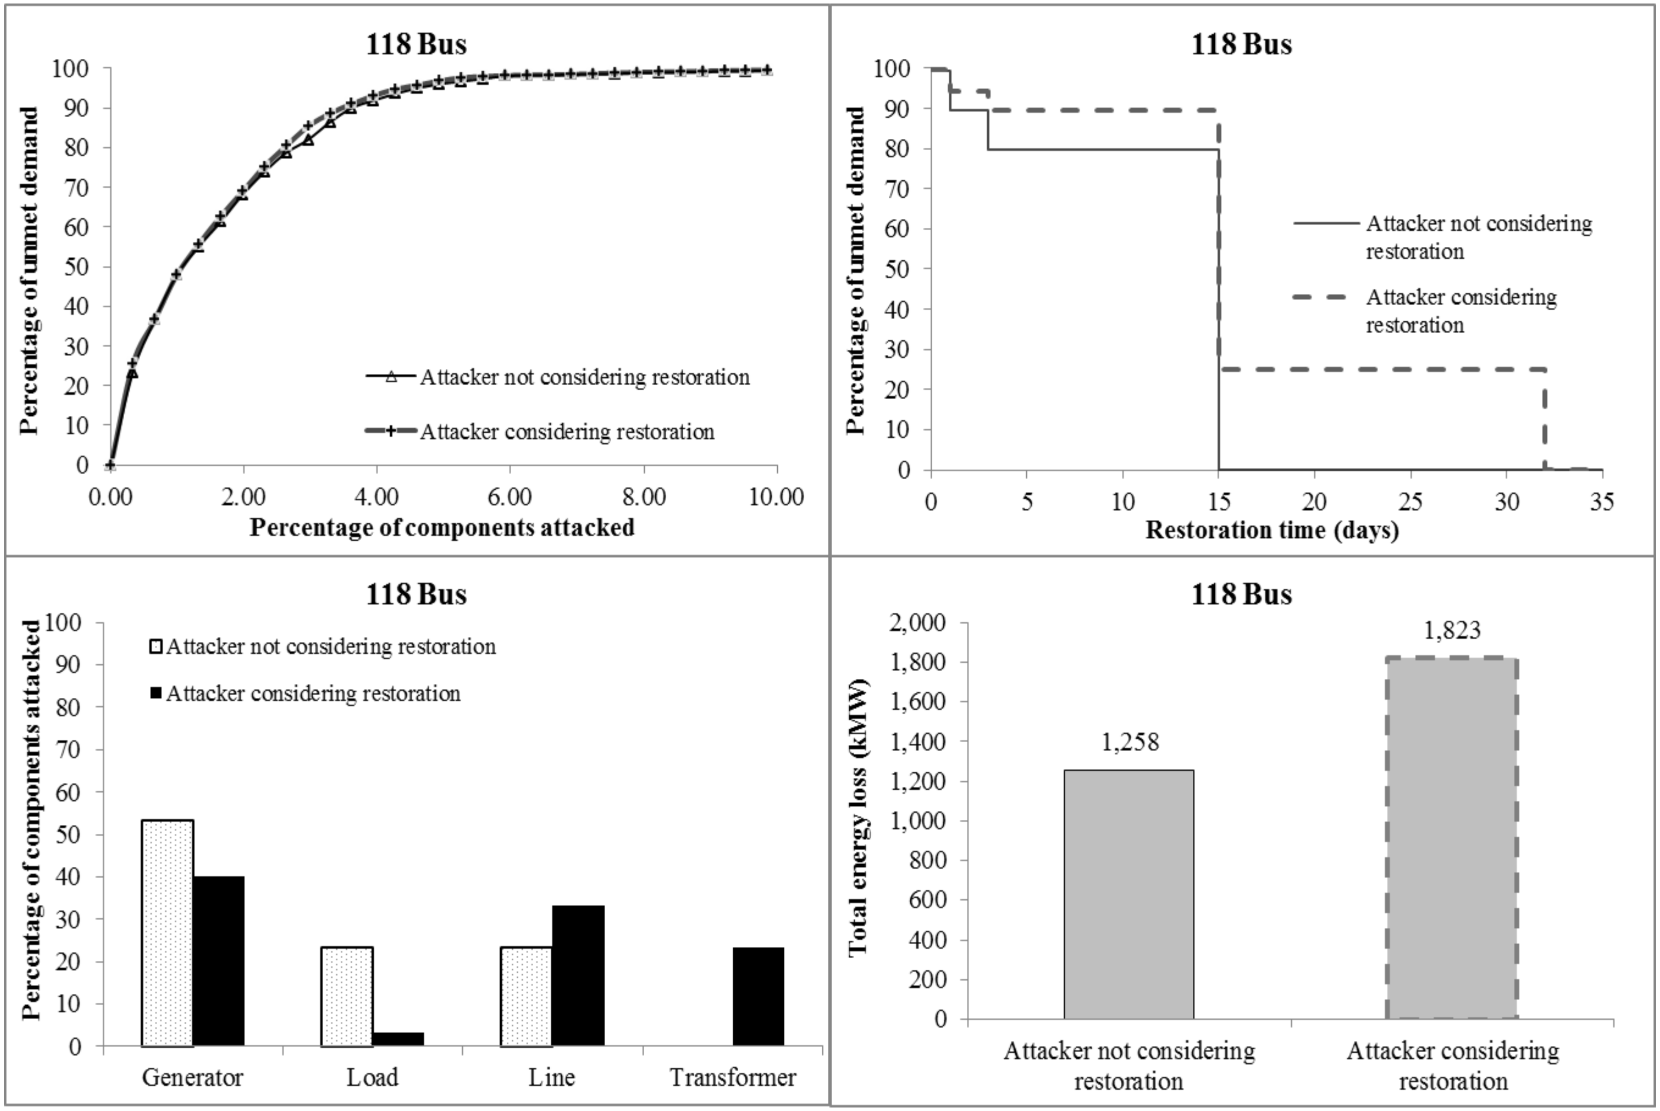

Figure 7.1-Impact of considering restoration times against a dynamic attacker with cascading knowledge (118-bus system)

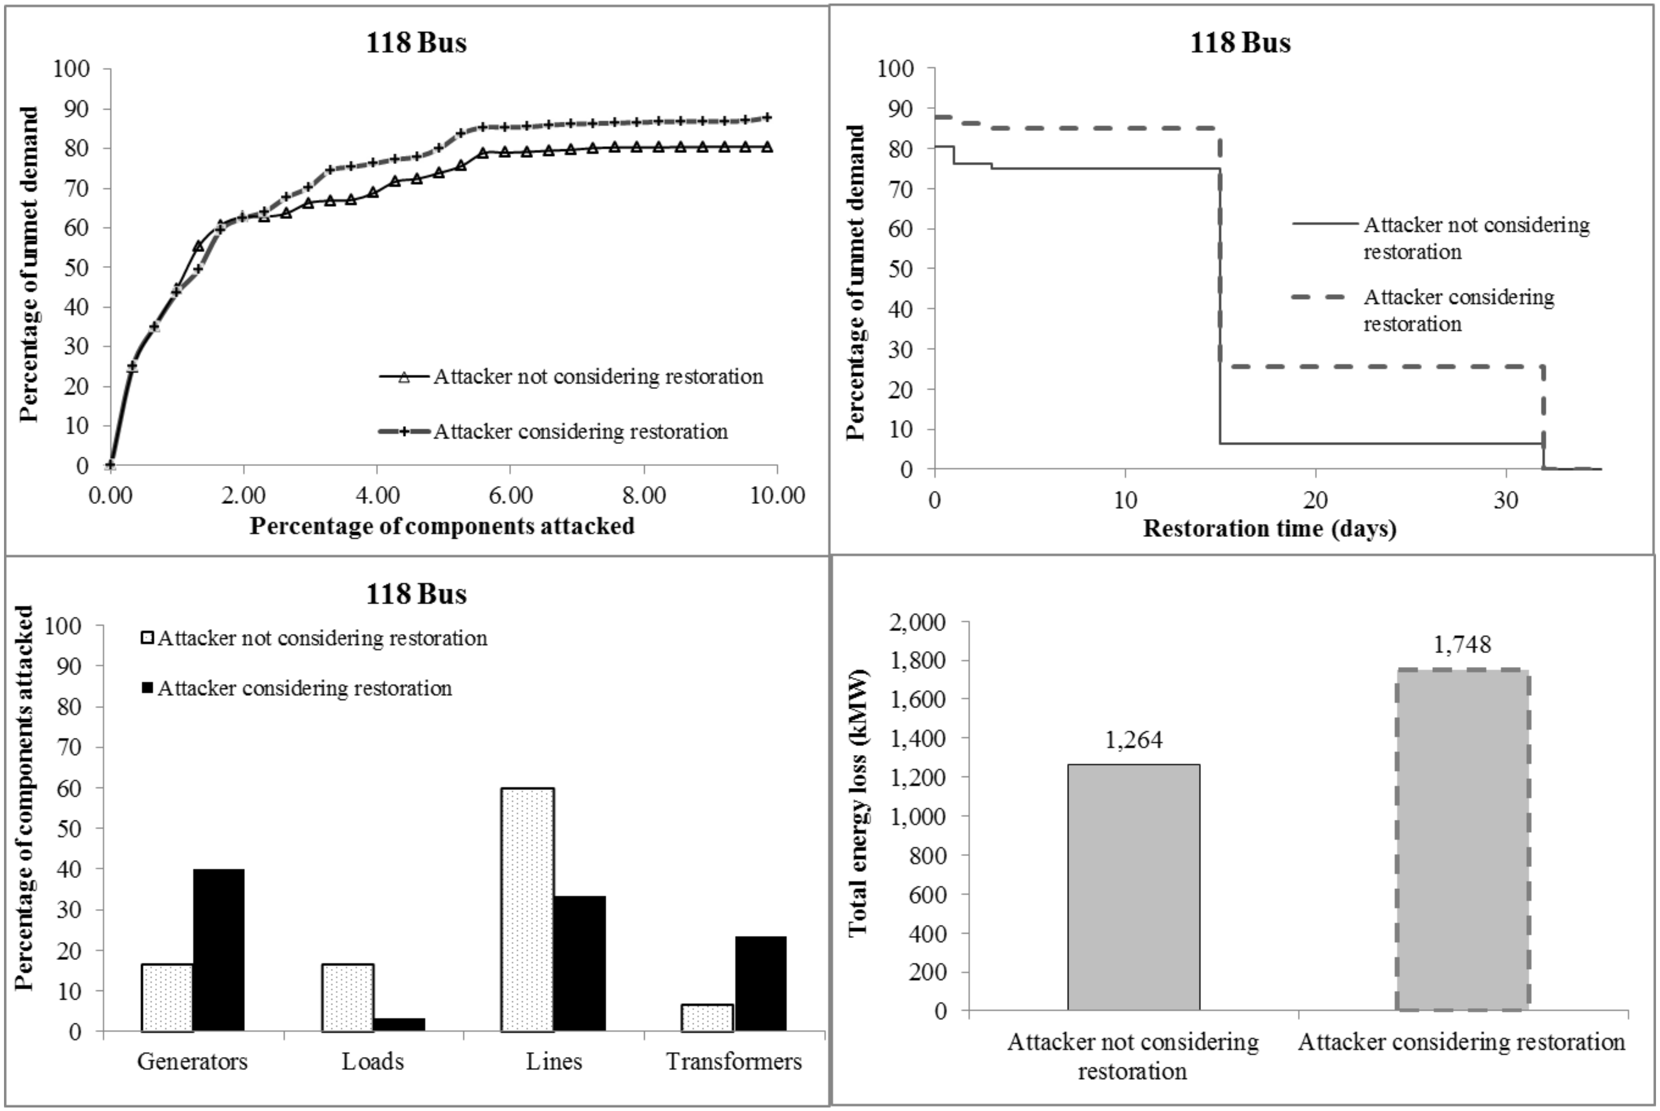

Figure 7.2-Impact of considering restoration times against a static attacker with cascading knowledge (118-bus system)

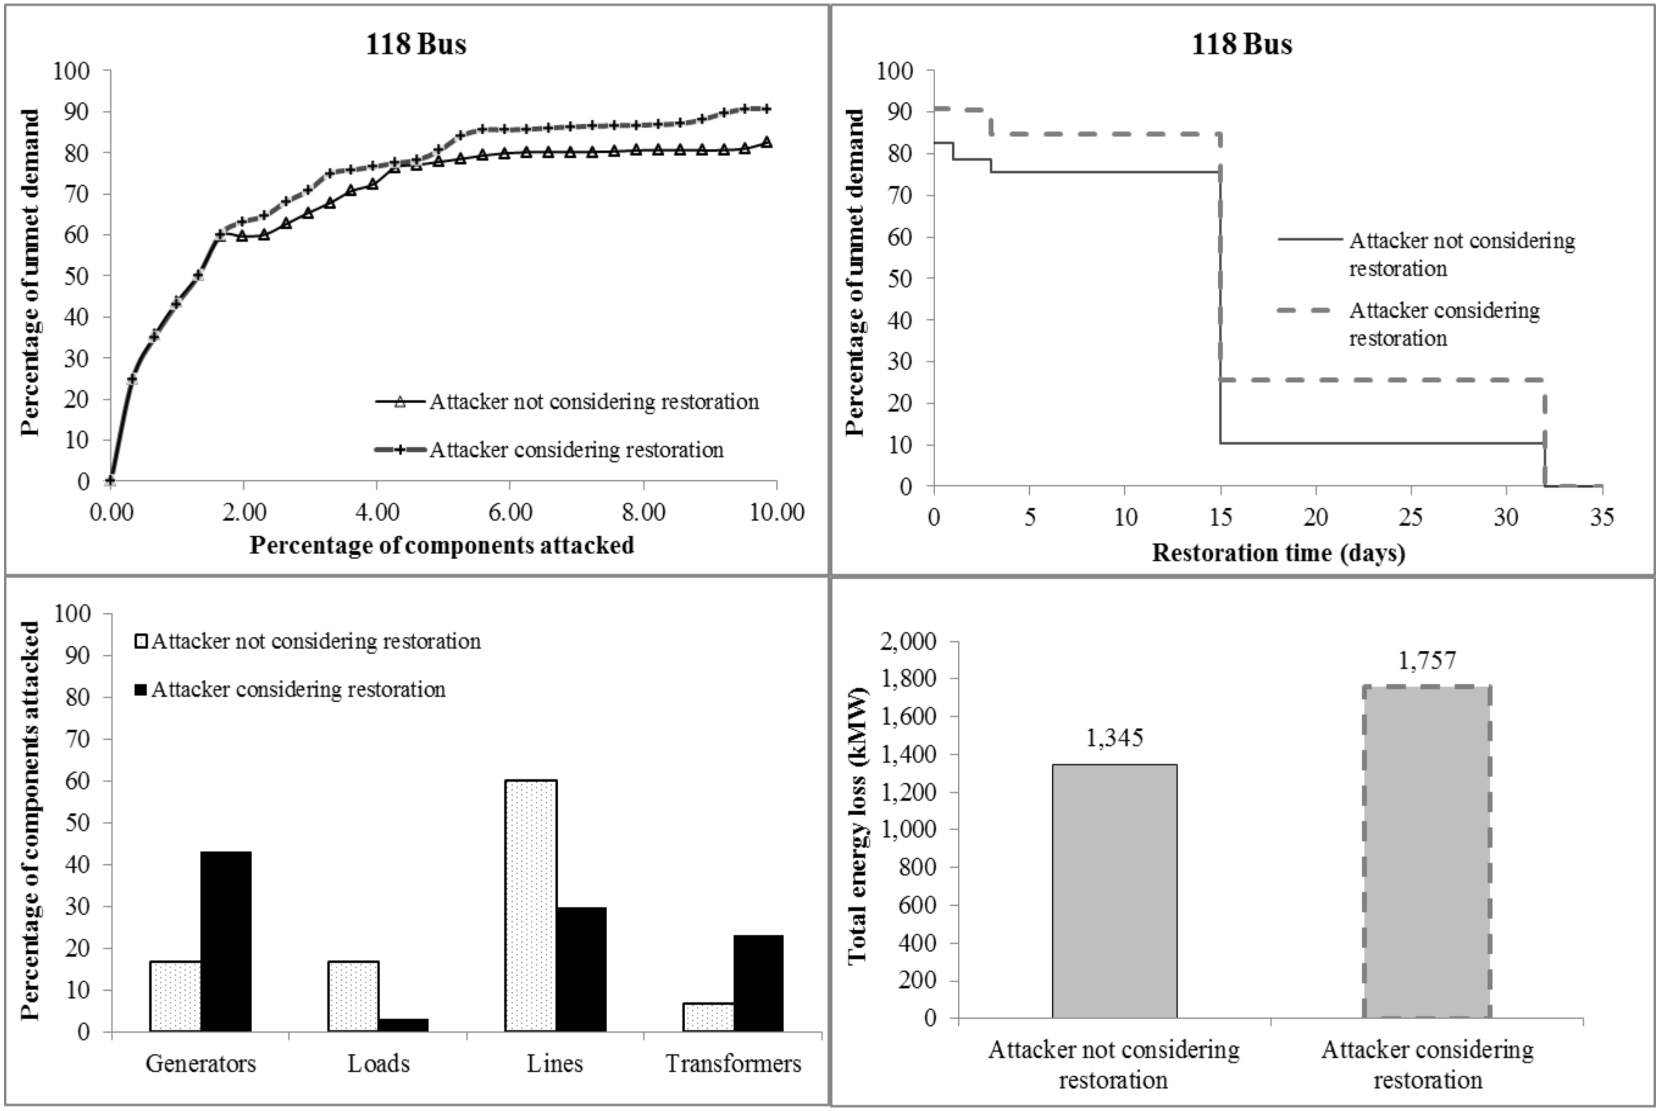

Figure 7.3-Impact of considering restoration times against a static attacker with no cascading knowledge (118-bus system)

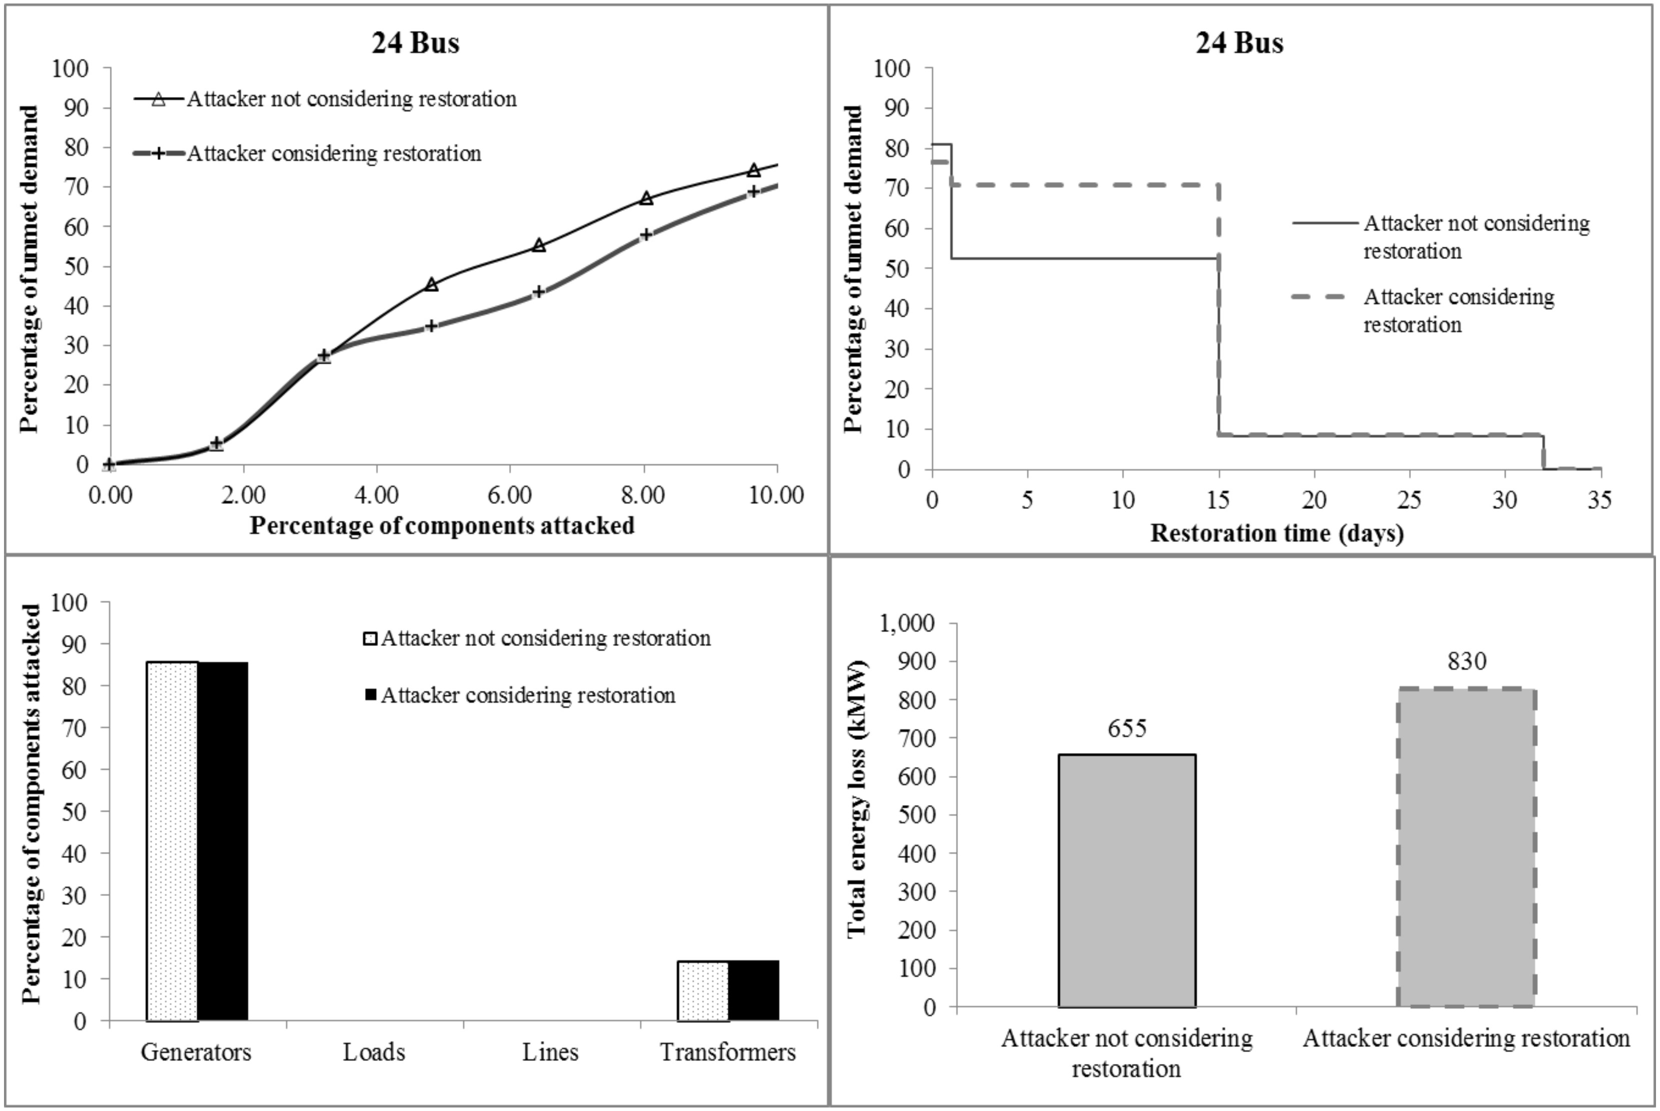

Figure 7.4-Impact of considering restoration times against a dynamic attacker with cascading knowledge (24-bus system)

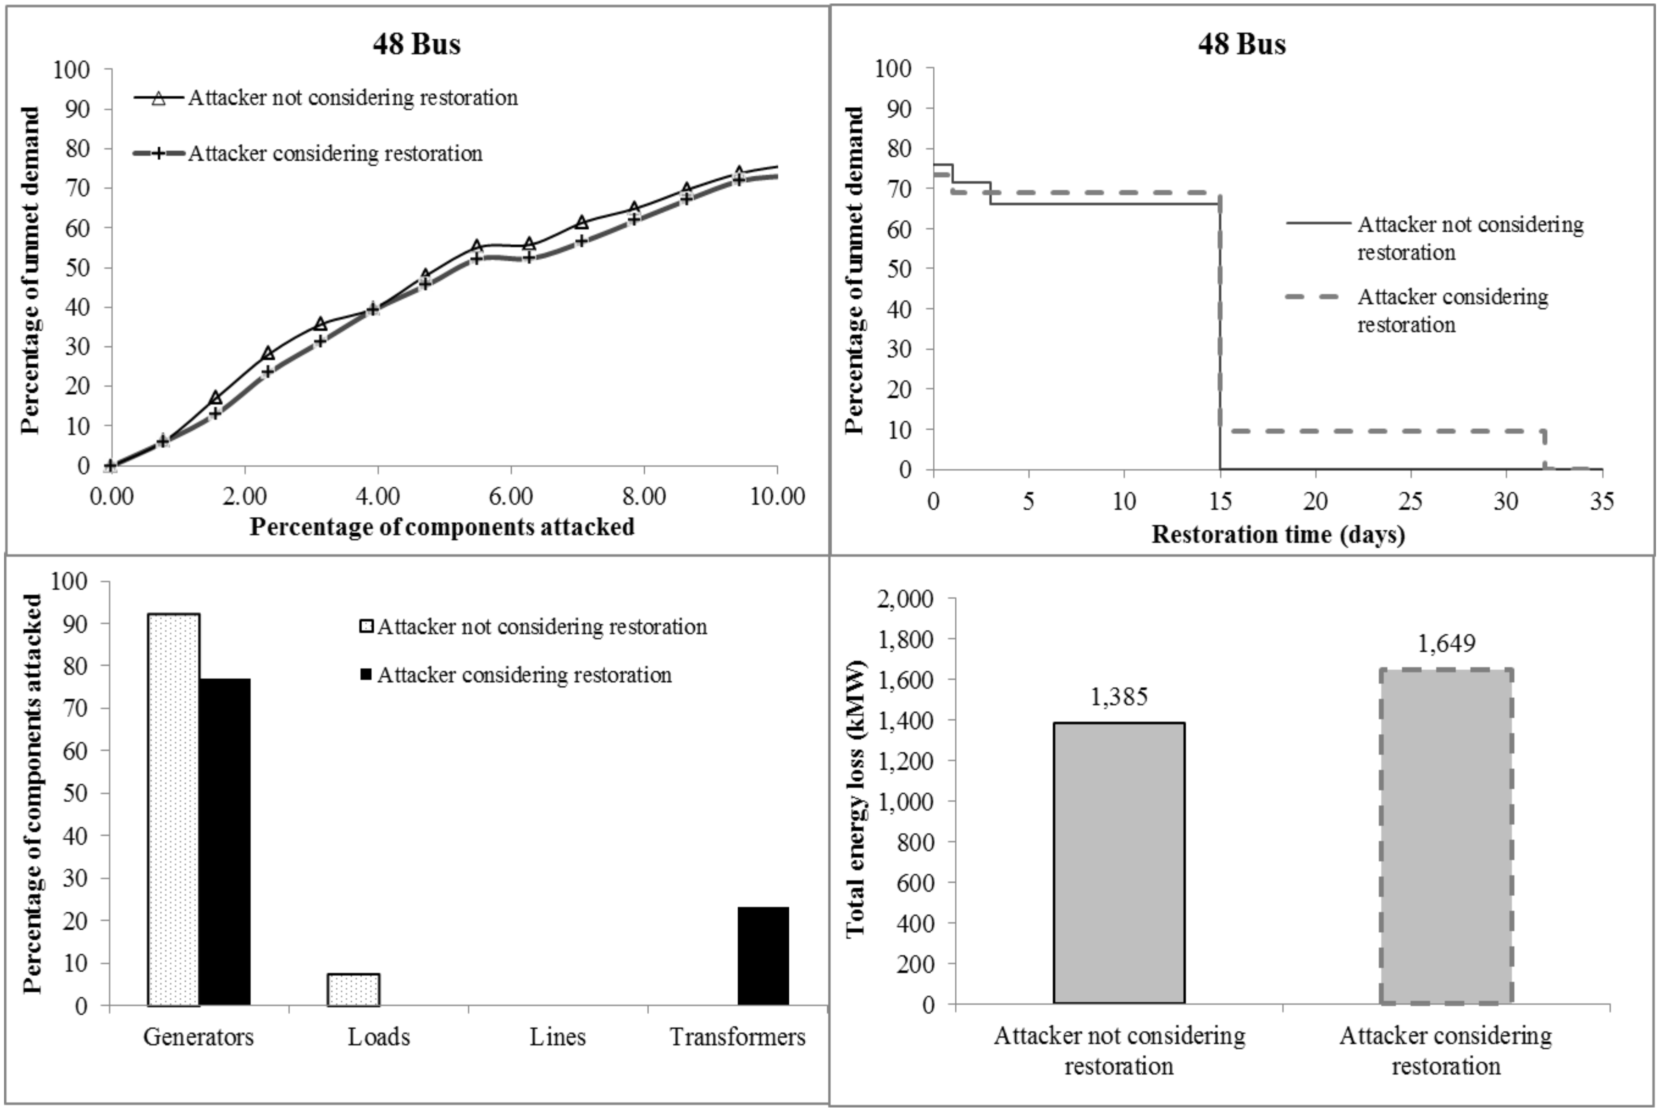

Figure 7.5-Impact of considering restoration times against a dynamic attacker with cascading knowledge (48-bus system)

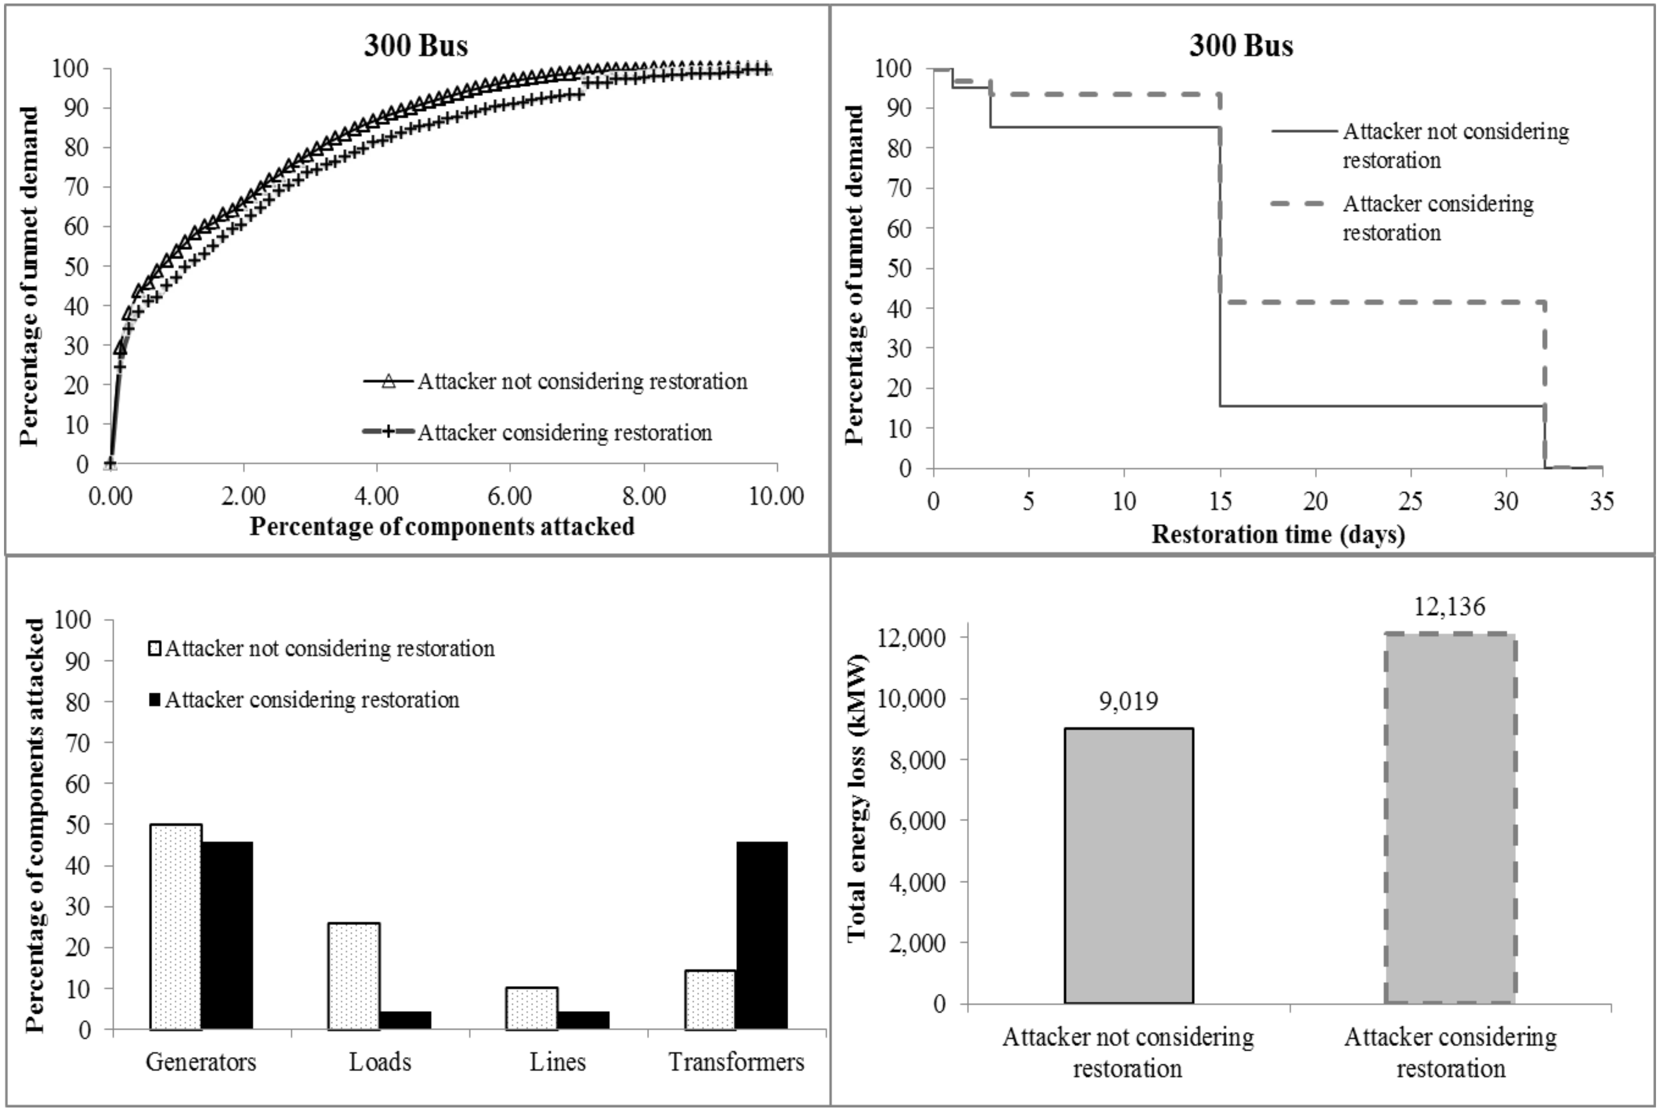

Figure 7.6-Impact of considering restoration times against a dynamic attacker with cascading knowledge (300-bus system)

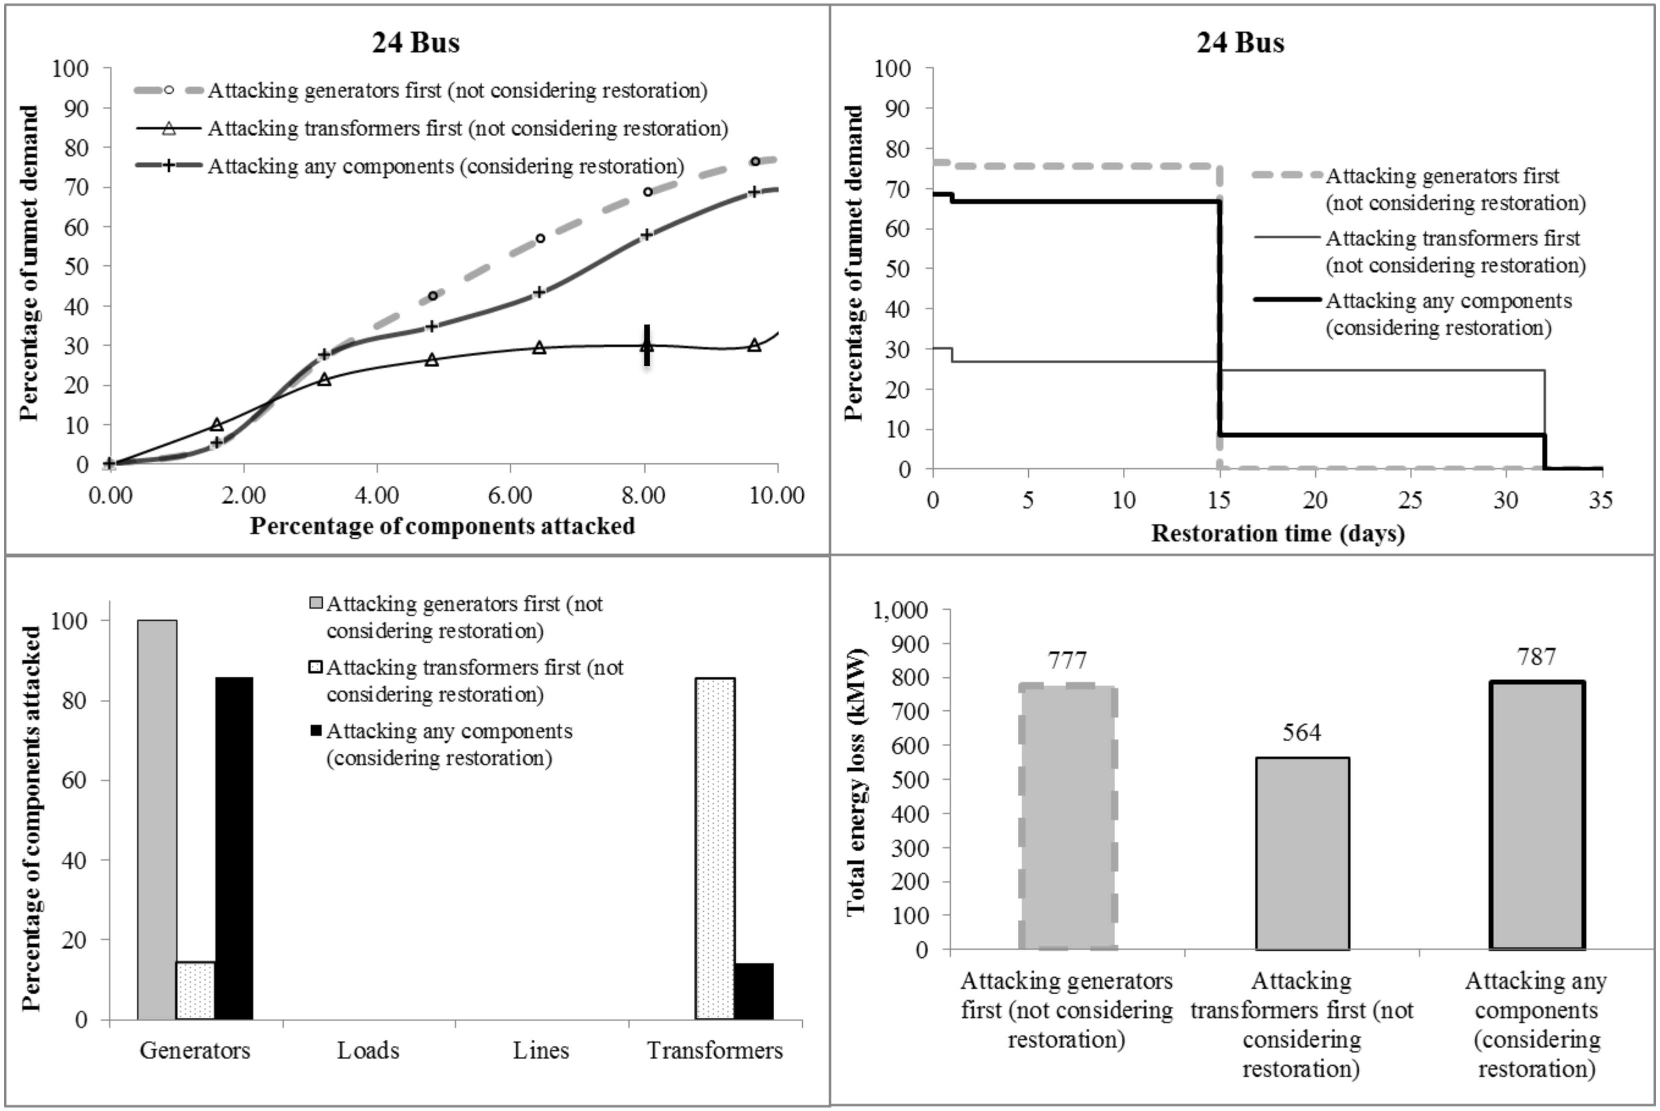

Figure 7.7–Comparison of an attacker attacking generators first or transformers first, and an attacker considering restoration times for a dynamic attacker (24-bus system)

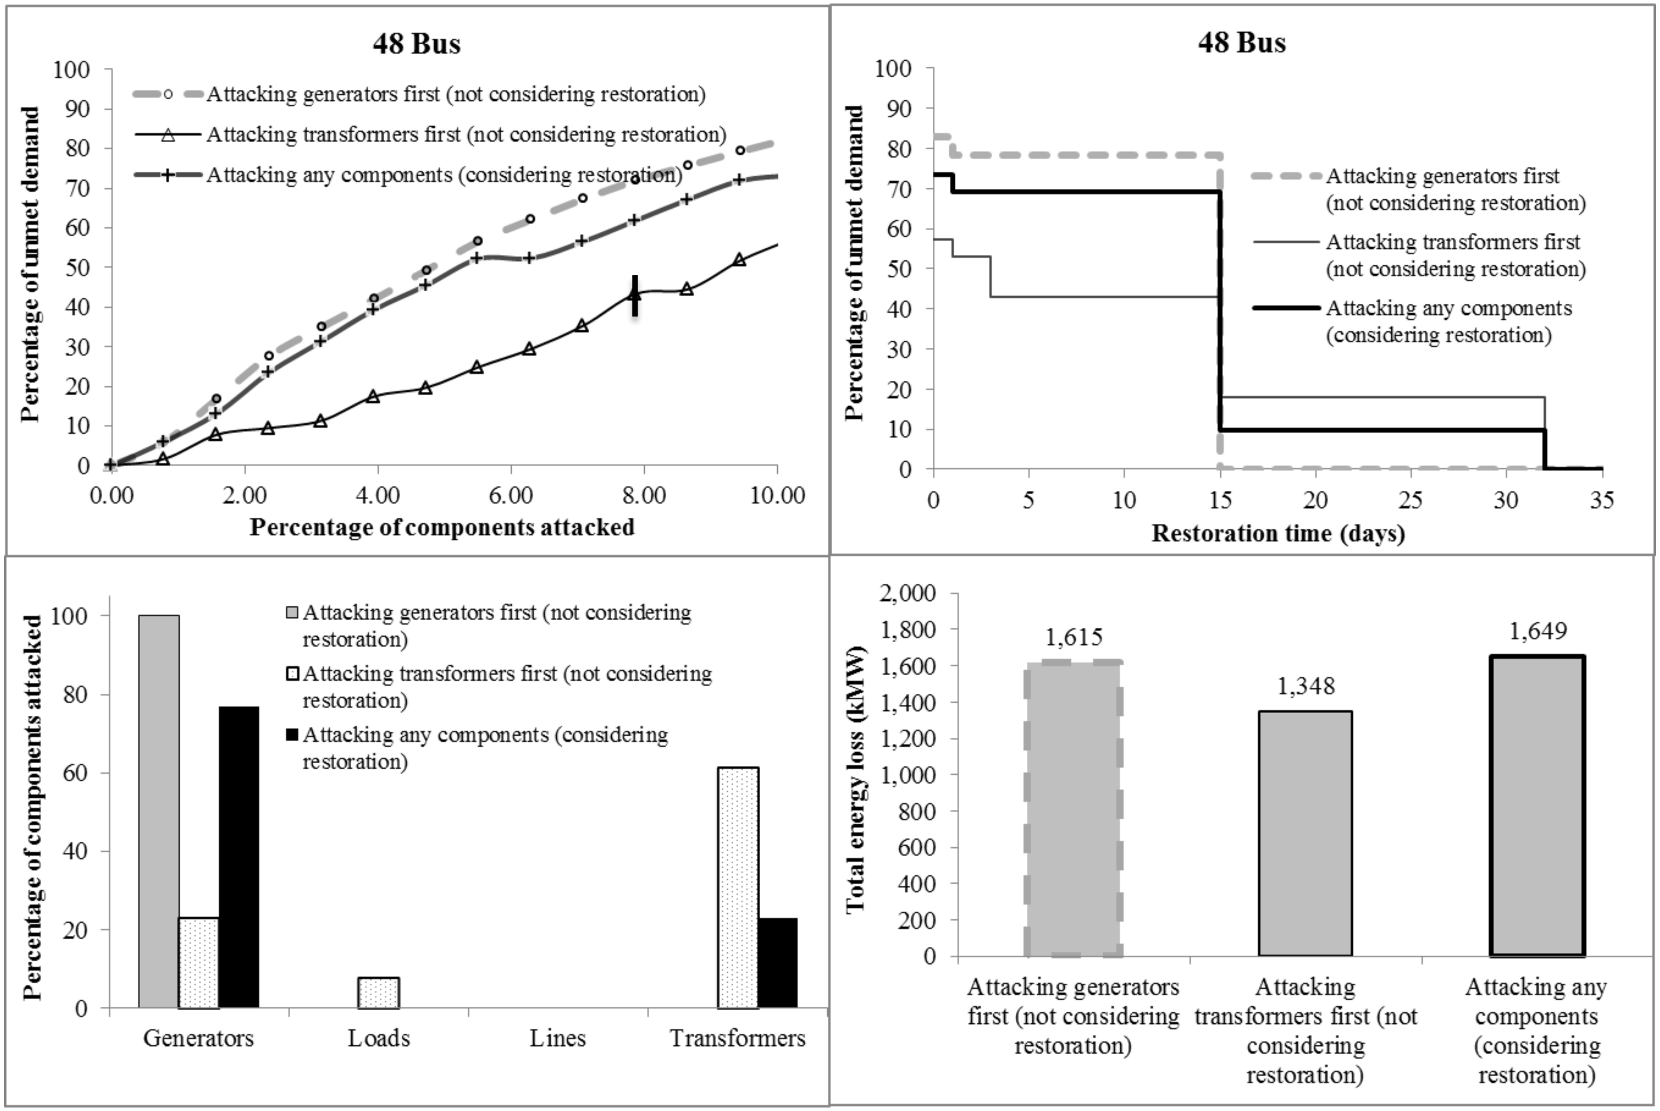

Figure 7.8–Comparison of an attacker attacking generators first or transformers first, and an attacker considering restoration times for a dynamic attacker (48-bus system)

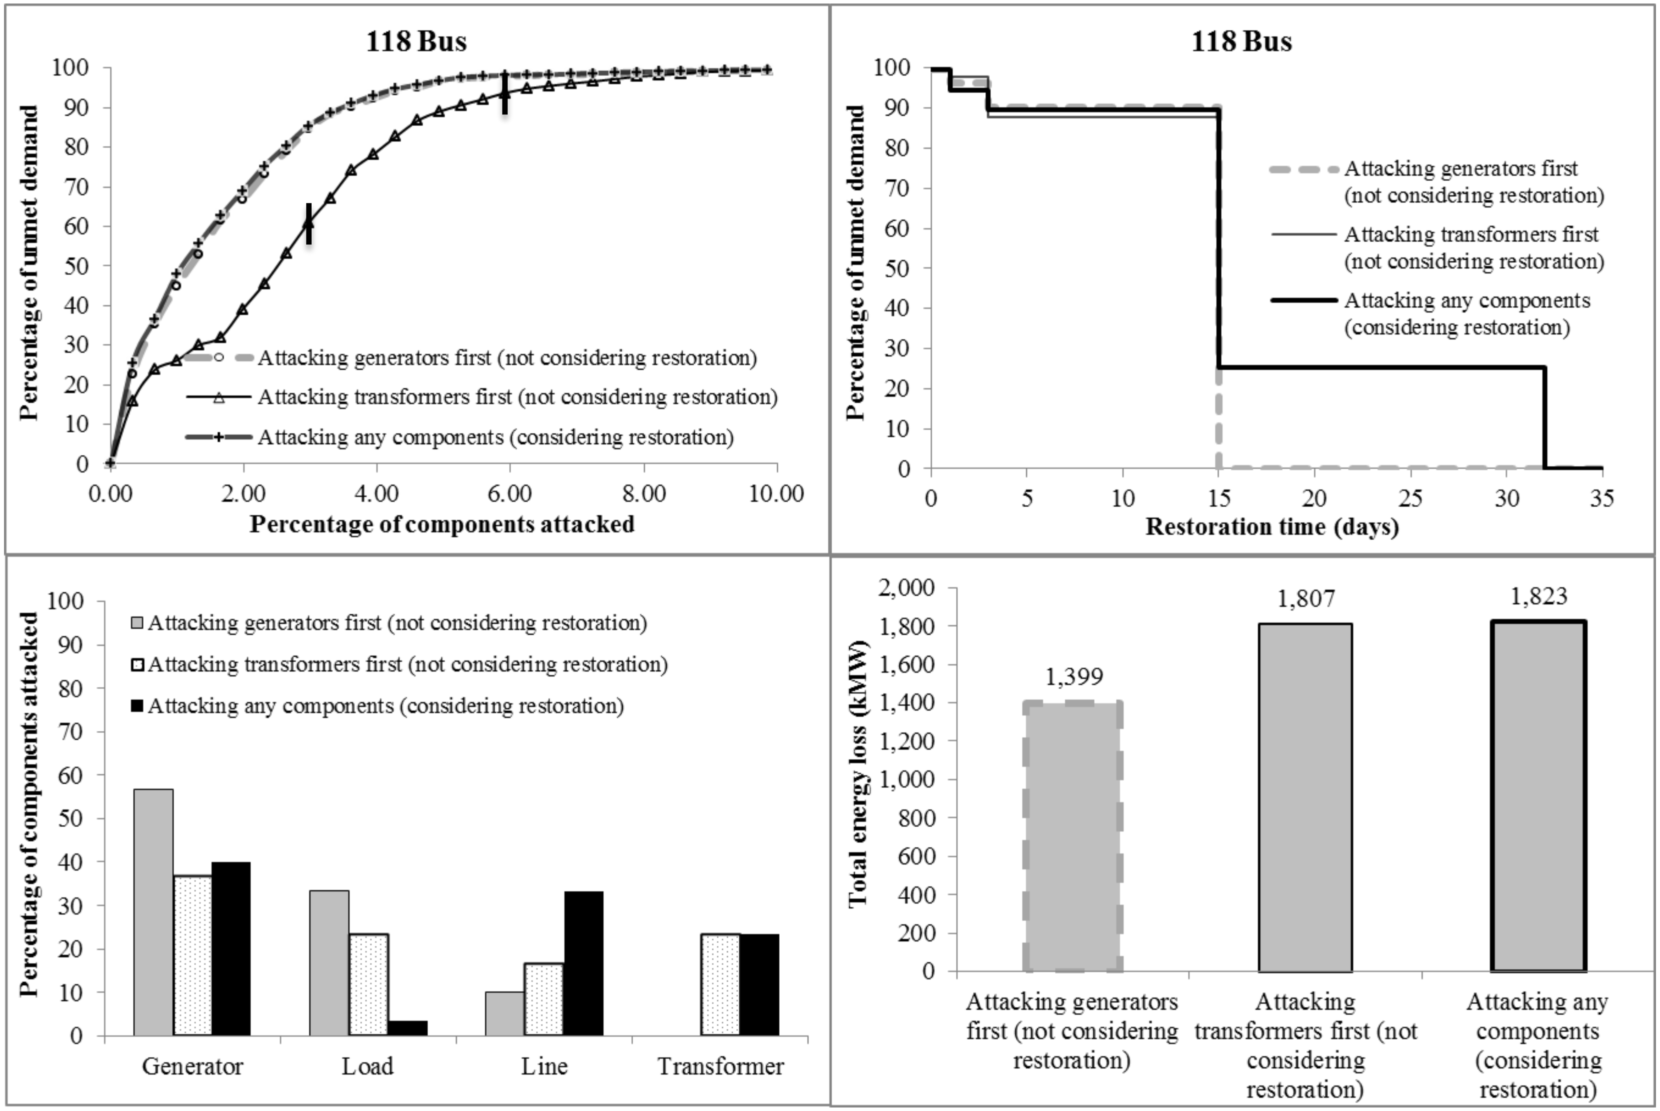

Figure 7.9–Comparison of an attacker attacking generators first or transformers first, and an attacker considering restoration times for a dynamic attacker (118-bus system)

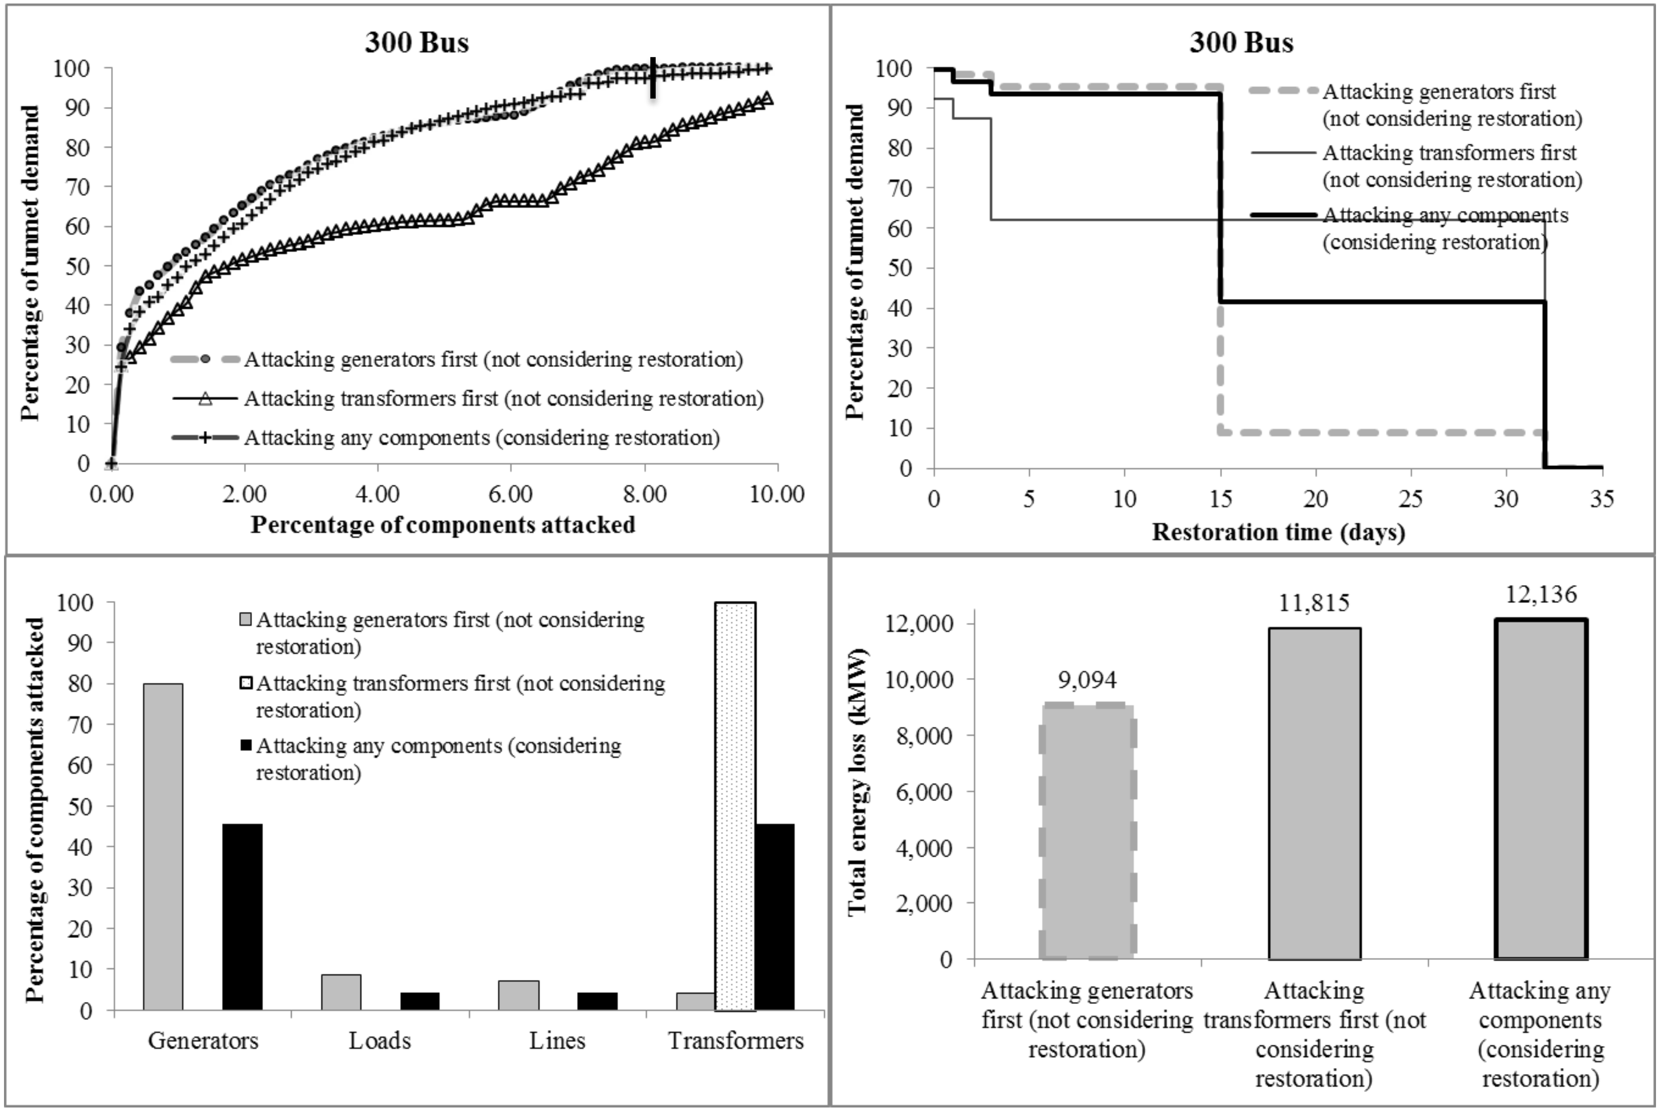

Figure 7.10–Comparison of an attacker attacking generators first or transformers first, and an attacker considering restoration times for a dynamic attacker (300-bus system)

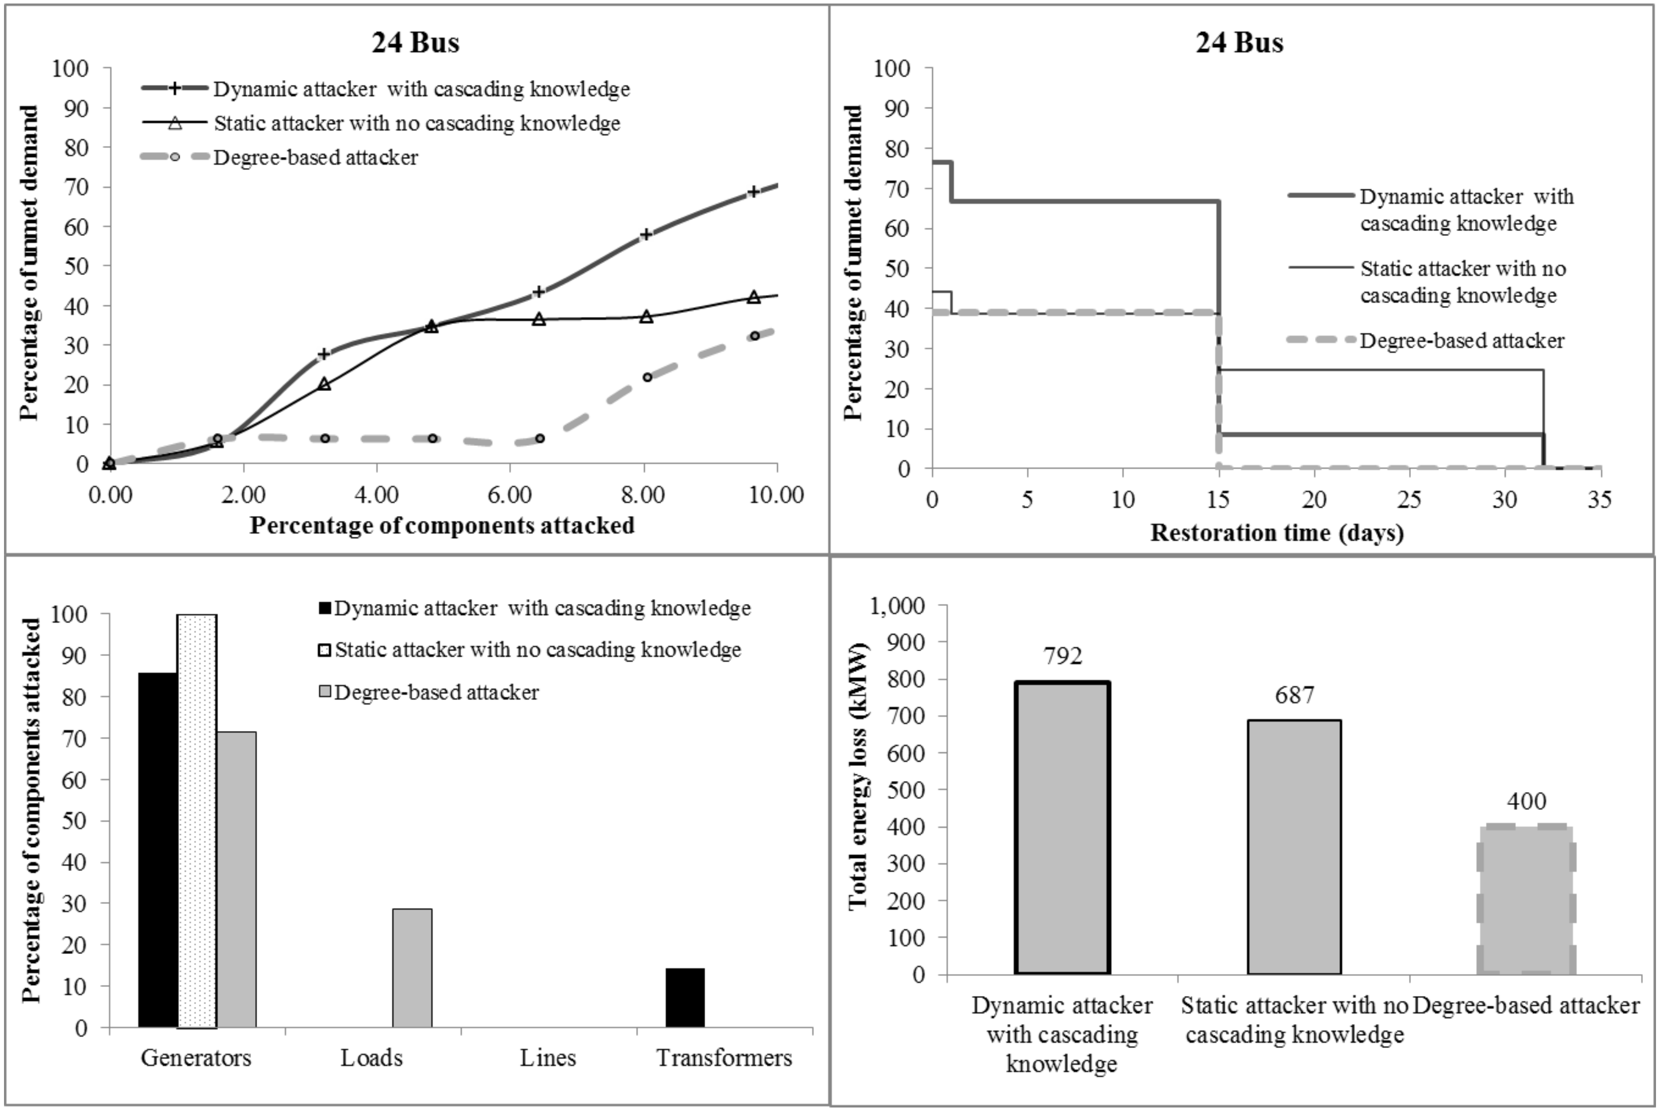

Figure 7.11–Comparison of a degree-based attacker with static attacker with no cascading knowledge and a dynamic attacker with cascading knowledge both considering restoration times (24-bus system)

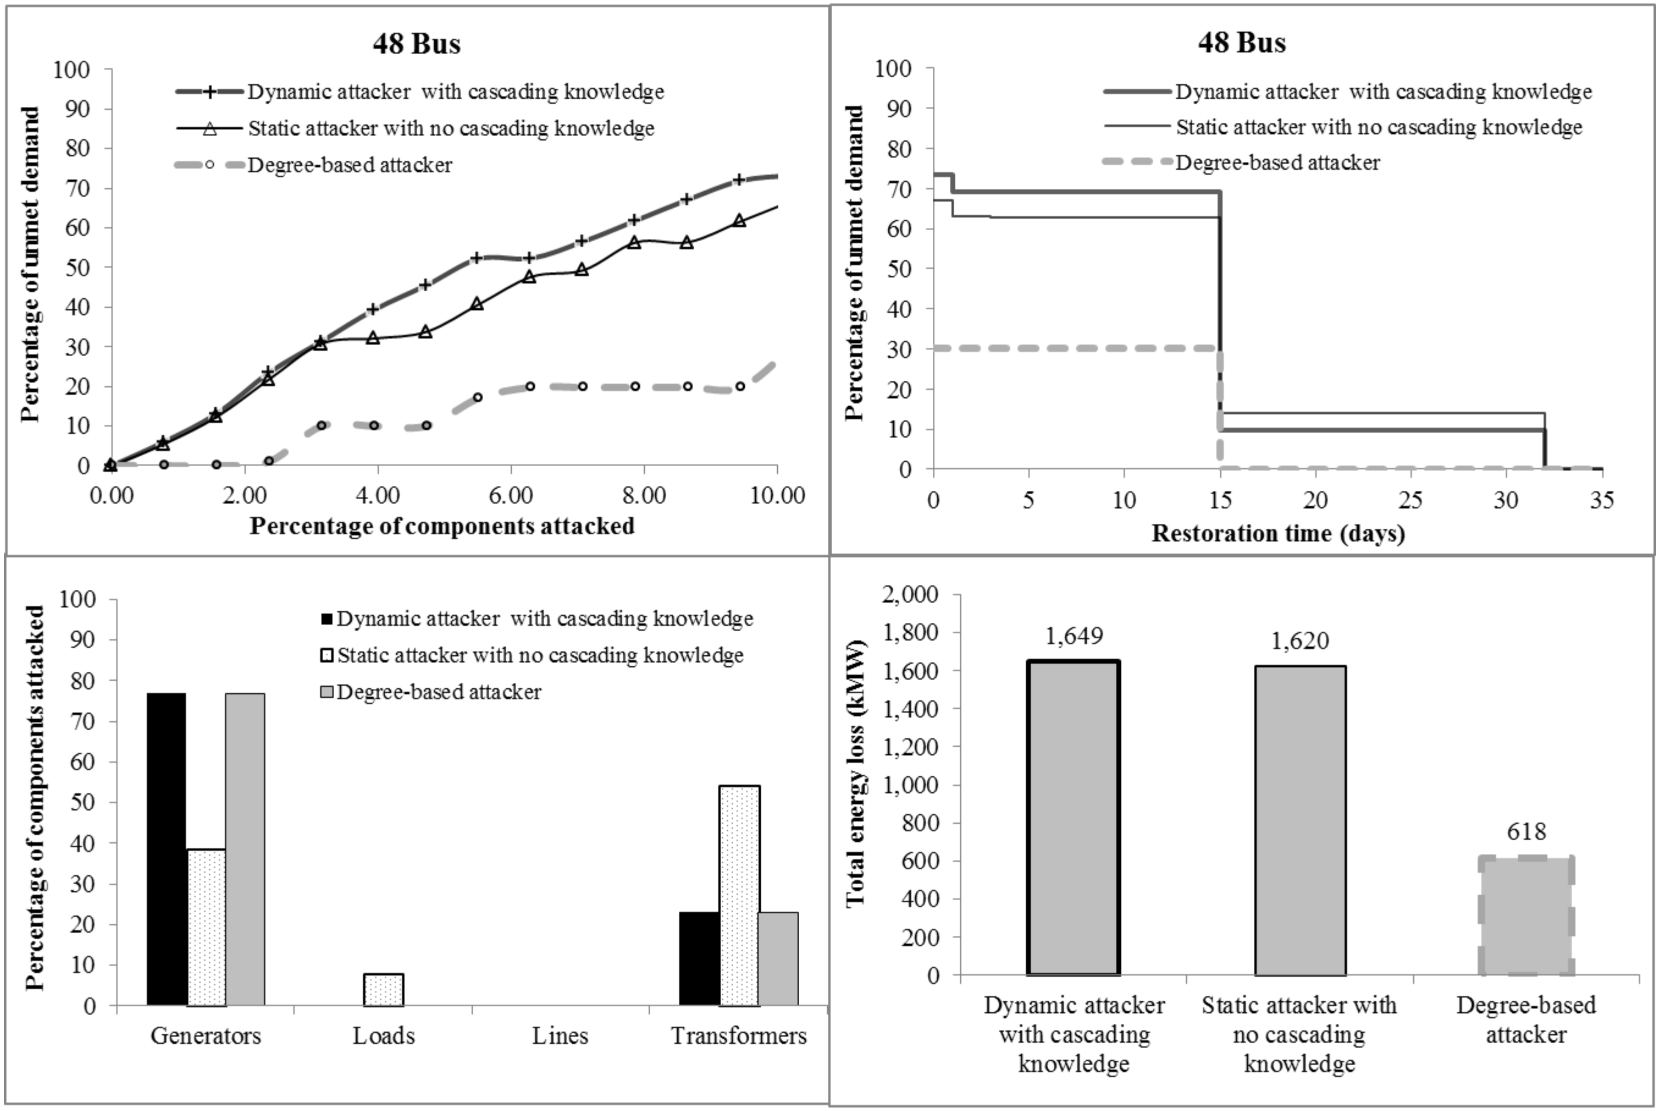

Figure 7.12– Comparison of a degree-based attacker with static attacker with no cascading knowledge and a dynamic attacker with cascading knowledge both considering restoration times (48-bus system)

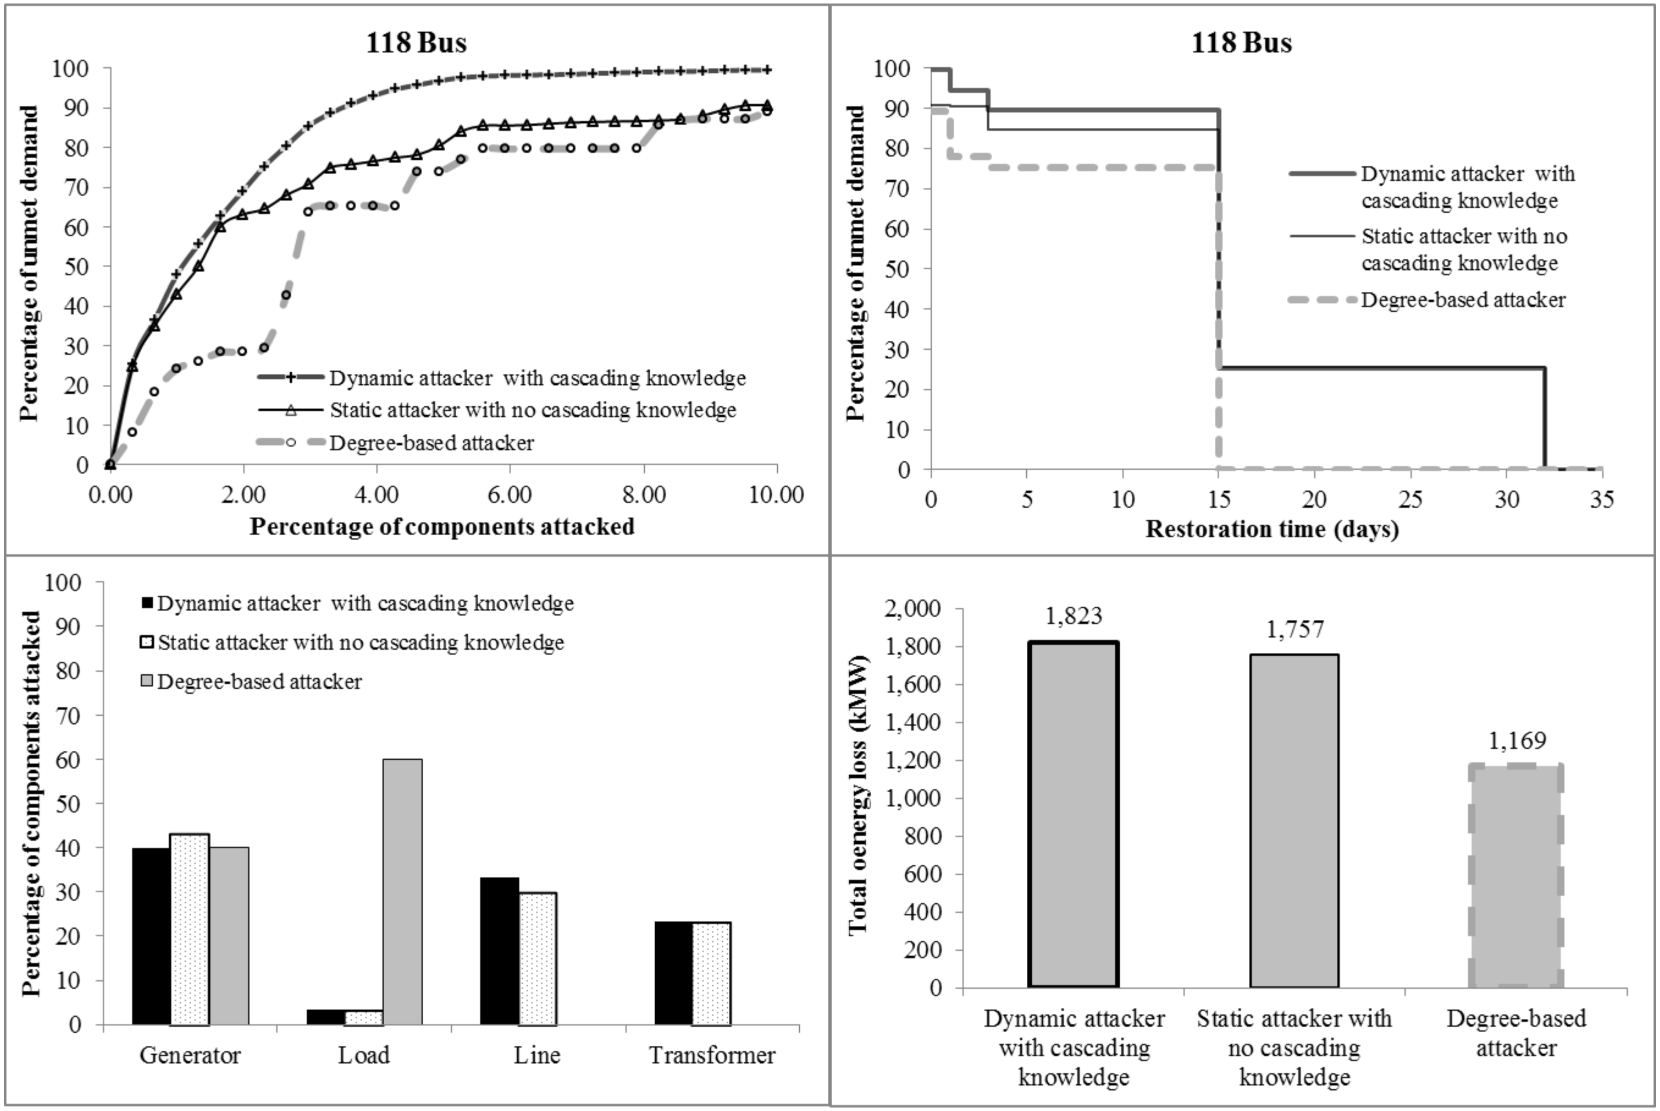

Figure 7.13– Comparison of a degree-based attacker with static attacker with no cascading knowledge and a dynamic attacker with cascading knowledge both considering restoration times (118-bus system)

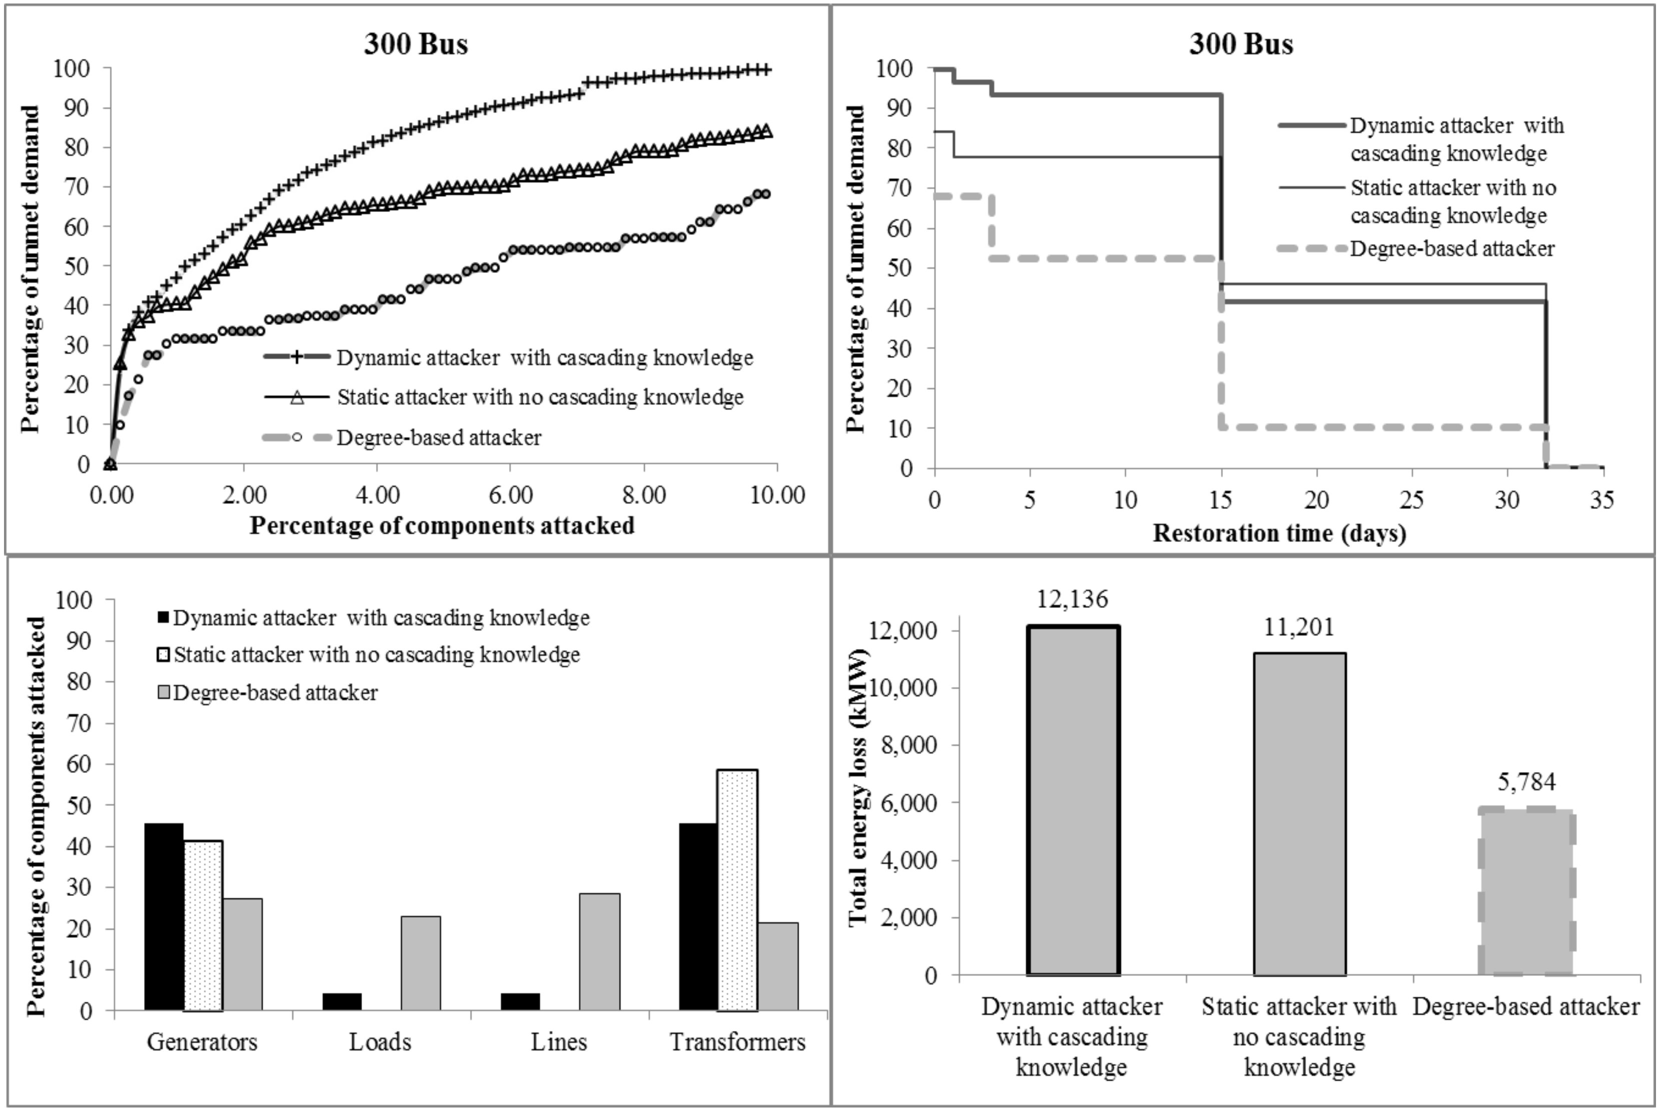

Figure 7.14– Comparison of a degree-based attacker with static attacker with no cascading knowledge and a dynamic attacker with cascading knowledge both considering restoration times (300-bus system)

7.2. Comparison of Defensive Measures in Overall

In this section, we will incorporate the total energy loss in comparing defensive measures. In addition to the previous investments (i.e., hardening, increasing the capacity of transmission lines, and adding new transmission lines), we can also invest in decreasing the restoration times of components. Total energy loss is a metric that captures the tradeoff between initial impact and long term impact, and will thereby enable us to compare improvement in restoration time with other types of investments.

Transformers are critical in terms of their restoration times, so we consider decreasing the restoration times of transformers by 50%. Such investment may simulate a spare transformer program, etc. We then compare this investment type with our base case (no investment) and our earlier candidate improvements.

First, we will look at the case in which a dynamic attacker considers restoration times. In this comparison, we will compare no investment to two other investment types: hardening 2% of components (the best strategy found previously); and decreasing the restoration time of transformers by 50%; see Figure 7.15. For a dynamic attacker, even a small hardening investment (such as 2% of the components) yields more improvement in total energy loss than decreasing the restoration time of all transformers by 50% (except for the 300-bus system).

Now, we will do the previous comparison for a dynamic attacker who does not consider restoration times; see those results in Figure 7.16 and comparison of a dynamic attacker who does or does not consider restoration times in Table 7.1.

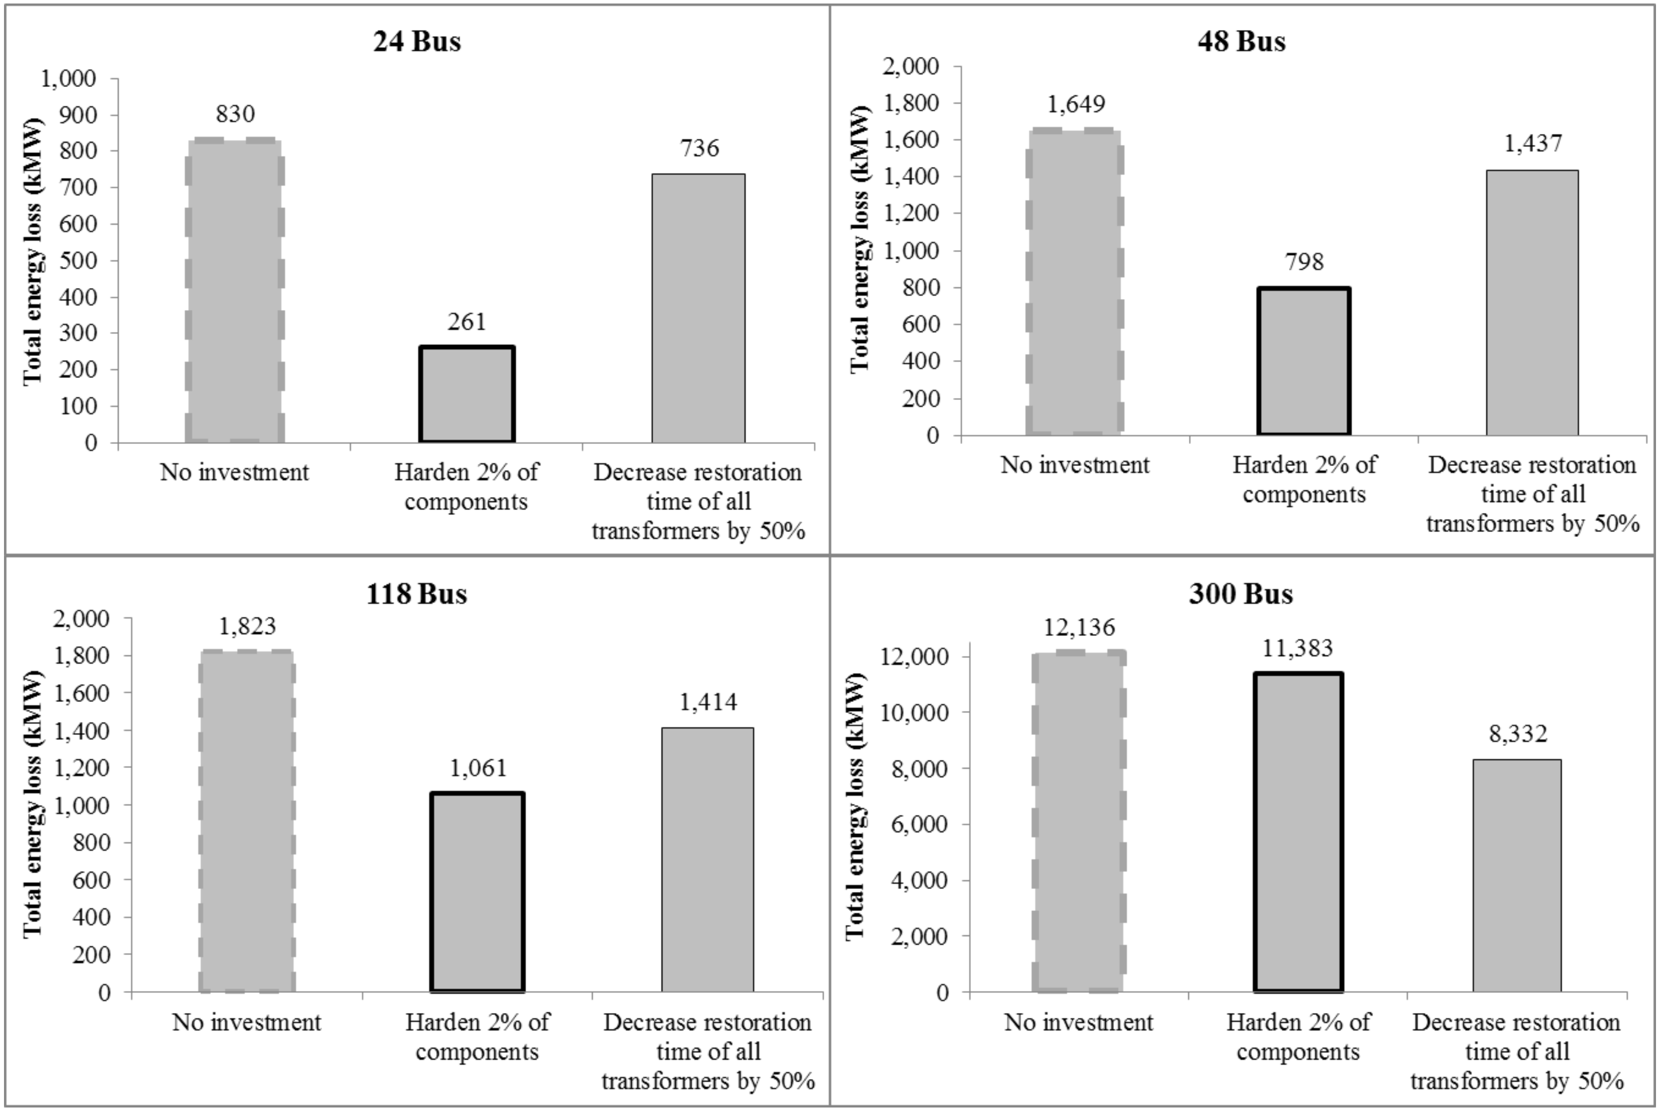

Figure 7.15–Comparison of total energy loss for hardening 2% of the components versus decreasing restoration times of all transformers against a dynamic attacker who considers restoration times

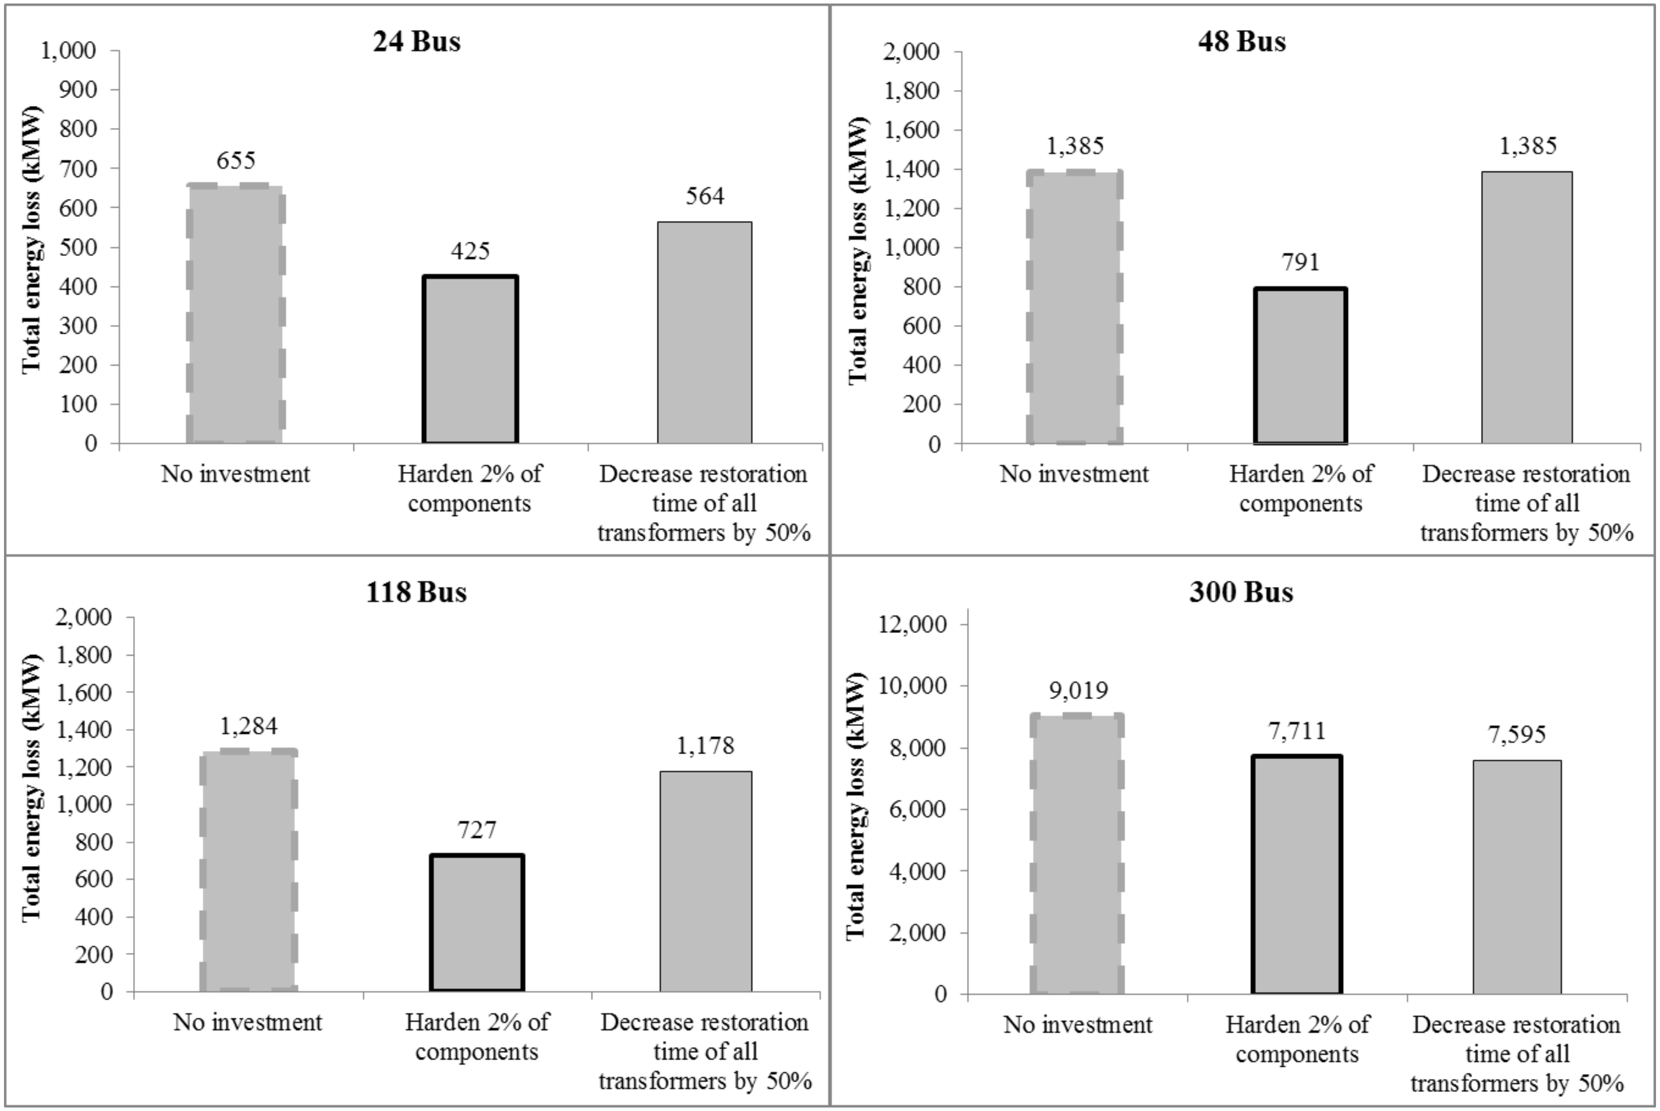

Figure 7.16–Comparison of total energy loss for hardening 2% of the components versus decreasing restoration times of all transformers against a dynamic attacker who does not consider restoration times

Table 7.1-Percentage improvement in total energy loss for an attacker who does or does not consider restoration times

Type of attacker	24-Bus		48-Bus		118-Bus		300-Bus	
	2% Hardening	Decrease restoration time of transformers by 50%	2% Hardening	Decrease restoration time of transformers by 50%	2% Hardening	Decrease restoration time of transformers by 50%	2% Hardening	Decrease restoration time of transformers by 50%
Attacker who considers restoration times	71	19	52	13	42	22	6	31
Attacker who does not consider restoration times	35	14	43	0	43	8	15	16

Of course, reducing restoration times is even less effective when the attacker does not consider restoration times in their choice of component. The effectiveness of investment in transformers may be limited, since the attacker does not attack as many transformers when the attacker does not consider restoration times.

Table 7.1 demonstrates that reducing restoration time of transformers by 50% is more effective against a dynamic attacker who does consider restoration times than against a dynamic attacker who does not consider restoration times. However, even hardening 2% of the components still does significantly better than decreasing the restoration time of all transformers by 50% (except for 300-bus system). Therefore, investment in hardening seems reasonable [For a detailed analysis of impact of considering restoration times for various sizes of investments against an attacker with different cascading knowledge, see Appendix M; hardening does better than improving restoration times in almost all of those scenarios.]

Now, we compare increasing the capacity of transmission lines or adding new transmission lines investments with investment in decreasing restoration times. As you might recall from Chapter 6, adding capacity or new transmission lines did less well than hardening. So, we will compare investment in transformers only against large investments (rather than the 2% and 5% investment considered in Chapter 6) in capacity or adding new transmission lines (i.e., investing in 10% of the transmission lines that are likely to cascade). In Figure 7.17, we compare large capacity and adding new transmission lines investments with decreasing restoration times for a dynamic attacker who considers restoration times for 24-bus, 48-bus, 118-bus, and 300-bus systems, respectively.

As can be seen from the figure, even large investments in increasing capacity or adding new transmission lines did not lead to significant improvements in total energy loss compared to the no-investment case (in part, because the transmission lines are so quick to restore), and decreasing restoration times of transformers sometimes but not always performs significantly better. As expected, the difference between restoration investment and investments in capacity and new transmission lines is usually smaller when the attacker does not consider restoration times, see Figure 7.18. [For a detailed analysis of impact of considering restoration times for various sizes of investments against an attacker with different cascading knowledge, see Appendix N for increasing the capacities of transmission lines, and Appendix O for adding new transmission lines. The improvement in total energy loss achieved by adding capacity or new transmission lines is even smaller against a static attacker than a dynamic attacker.]

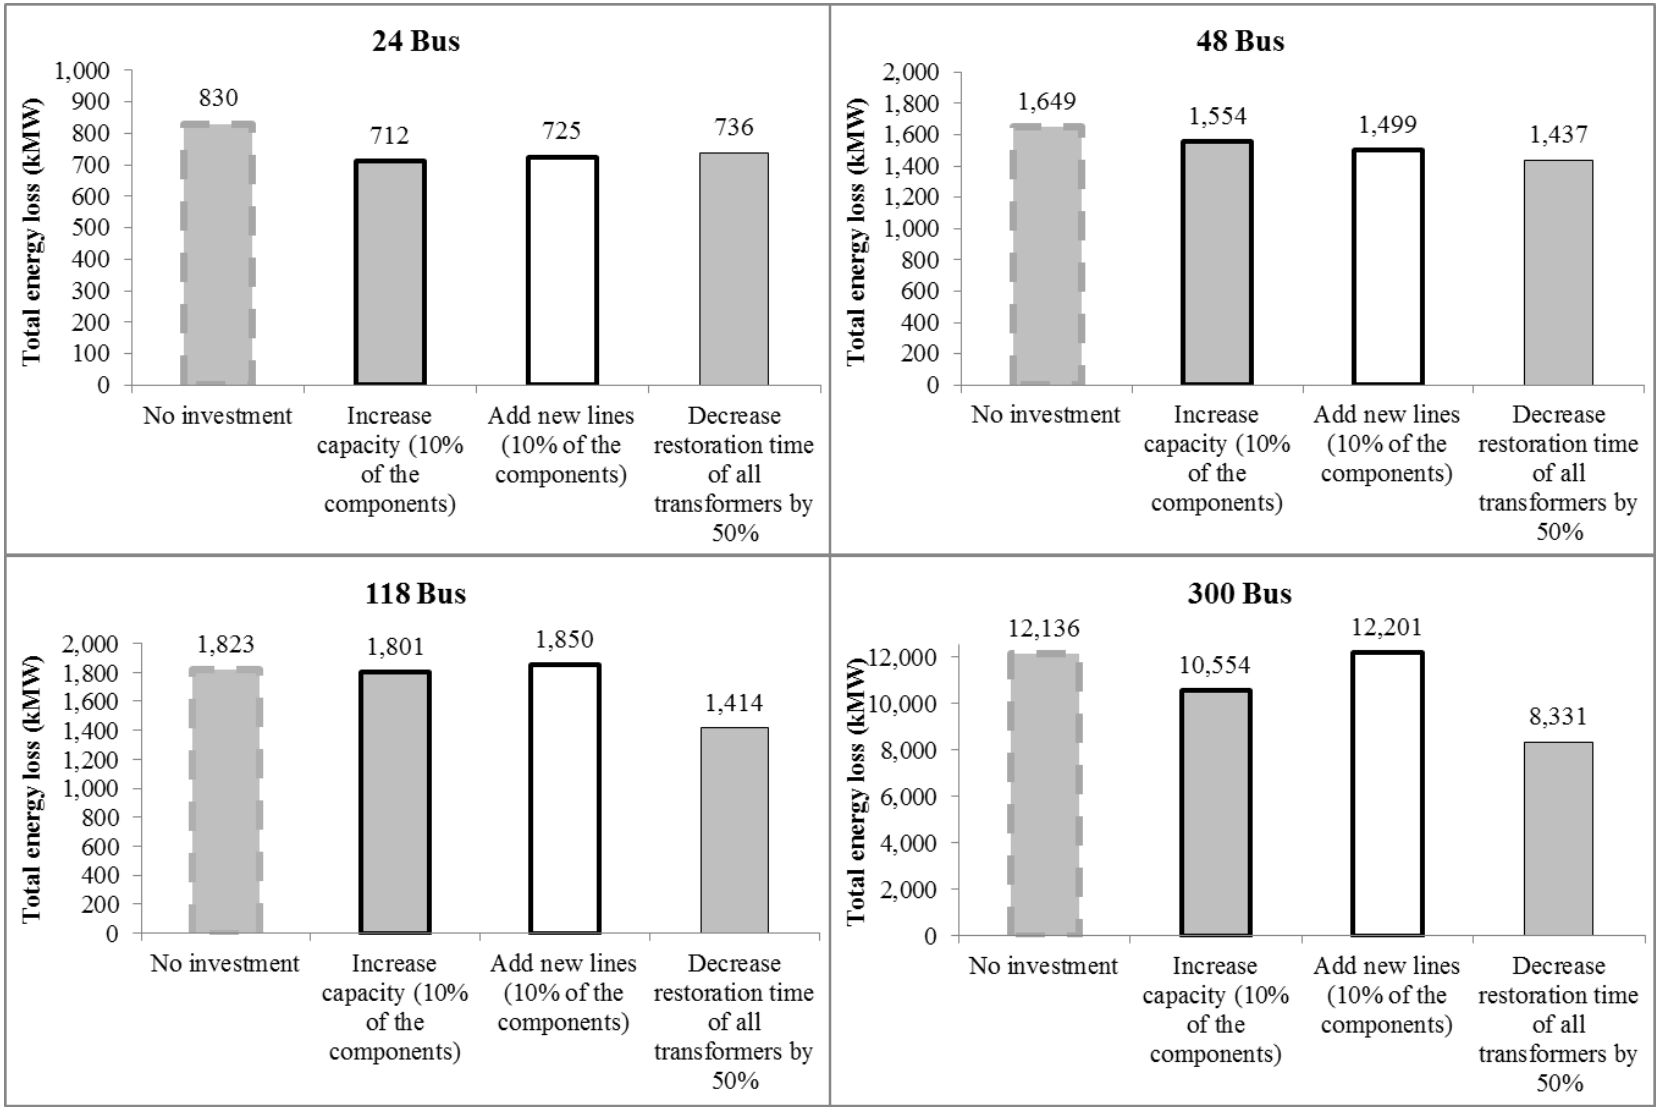

Figure 7.17–Increasing capacity or adding new transmission lines of 10% of the components versus decreasing restoration times of transformers by 50% (an attacker who considers restoration times)

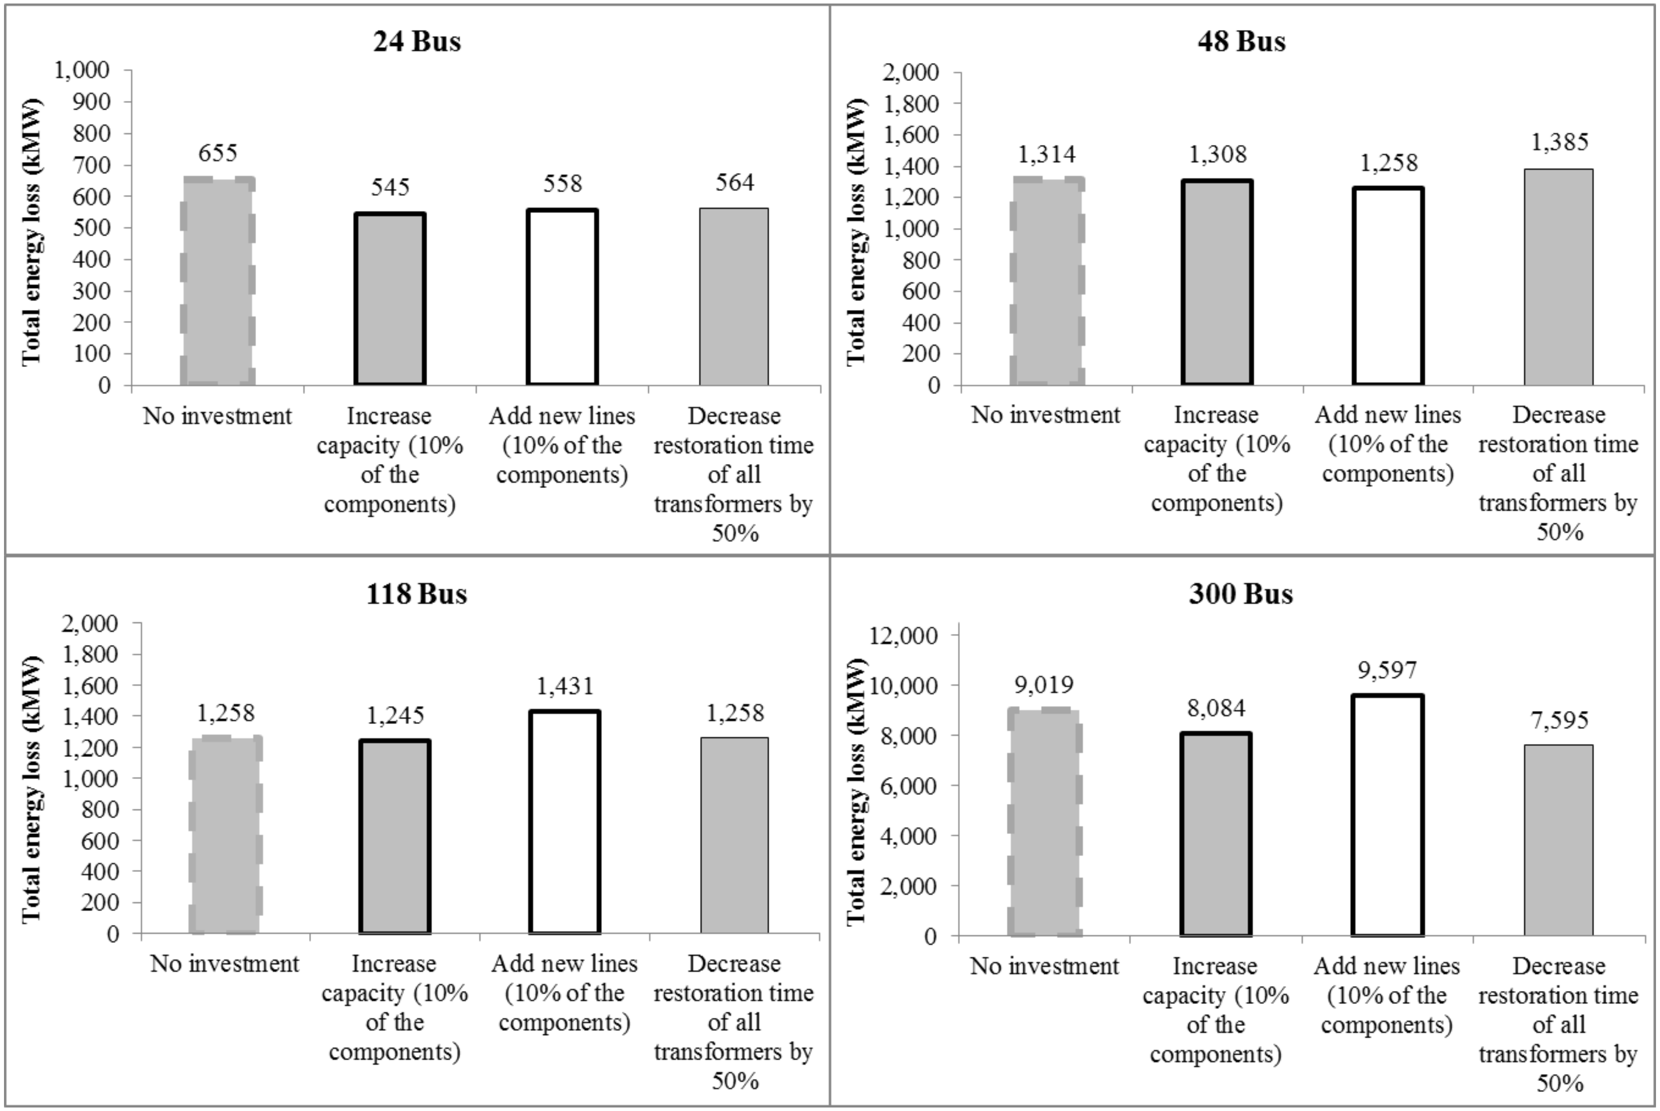

Figure 7.18–Increasing capacity or adding new transmission lines of 10% of the components versus decreasing restoration times of transformers by 50% (an attacker who does not consider restoration times)

7.3. Conclusion

Similar to modeling of cascading failure, the simplicity of our modeling approach makes it feasible to model restoration times, something that is not done in many other models. Modeling restoration times enables us to observe the long term impact of an attack, and how an attacker can take advantage of long restoration times of some components. We use total energy loss as a measure to summarize the overall impact of an attack. This metric integrates both the initial impact of an attack (including any immediate impact due to cascading failure) and the restoration process after the attack, and allows us to analyze the effectiveness of reductions in restoration times.

Our results suggest that when the attacker considers restoration times, the overall impact of an attack is significantly higher than when the attacker does not consider restoration times. This difference emphasizes the point that models of intelligent adversaries may need to consider the recovery process and how it might be exploited. In terms of total energy loss, even a dynamic attacker who does not consider restoration times provides comparable results to preferentially attacking generators or transformers, and does significantly better than a degree-based attacker.

Finally, we compared various defensive investments including decreasing the restoration time of transformers (the component type with the longest recovery period) by 50%. Our results suggest that even though in some cases decreasing restoration times of transformers is comparable to hardening 2% of all components, hardening is generally more effective in decreasing the overall impact of an intelligent attack. Moreover, capacity investment or adding new lines does not provide a comparable improvement in total energy loss.

8. Conclusion and Future Work

In this research, we developed a novel heuristic game-theoretic model that can simulate the entire impact of a disruption to a realistically complex electricity system, from the beginning of an attack, through cascading failure, until the system is fully recovered. It is novel in the sense that it is capable of addressing the issues of an intelligent adversary, cascading failure, and restoration times, all in the same model. Therefore, the model can facilitate the identification of effective defensive investments, by enabling consideration of a broader range of options.

Since our approach uses a greedy heuristic to solve the attacker's optimization problem, rather than solving it as an integer problem in a bilevel optimization program (as is done by Salmeron et al., 2004), the basic model is not optimal, but is simple enough that we can easily add complexities like cascading failure. Moreover, it is simple enough that it can easily be used and understood by practitioners.

In order to determine whether our methodology is accurate enough to be useful, we investigated a few questions: the validity of our model; how sensitive it is to its assumptions; and the consistency of our findings with those of similar models in the literature. We have unfortunately not been successful in obtaining detailed results of more rigorous models, so it is not feasible to assess how far our results are from optimal. However, our greedy heuristic algorithm was clearly more effective than less rigorous attack and hardening strategies, such as degree-based or random strategies. Moreover, we validated that the model behaves in a realistic manner. For example, when we harden more components, reduce restoration times, or decrease susceptibility to cascading failure, the impact of an attack is generally less. Furthermore, we showed that the results of our model are suitably sensitive to the assumptions. For example, when we increase the failure probability of an overloaded component, the system generally becomes more vulnerable. Conversely, the more the attacker knows about cascading failure, the greater the damage is in general.

We have also been able to demonstrate that the level of unmet demand is not overly sensitive to the parameters of our model for cascading failure. This supports the usefulness of our model, because there is no need for precise estimation of parameters that are difficult to estimate.

In order to test the consistency of our results with other models in the literature, we compared our findings with those of Hines et al. (2010), since, to our knowledge, the model of Hines et al. is the only model in the literature that includes both an intelligent adversary and cascading failure. Our results are generally close to the results of Hines et al. (2010), and, in our opinion, more realistic, since the flow-based attack strategy in the model of Hines et al. sometimes does worse when additional components are attacked. Overall, therefore, we conclude that our greedy algorithm may be a practical and feasible alternative for attacker-defender problems in electric power networks.

With regard to defenses, we found that investing in defensive resources generally leads to improved resilience, which shows the realism of our model. Our results indicated that in order to obtain a substantial improvement, the defender may need to harden a large portion of the network, which may not be cost-effective. (Of course, this may be due to the greedy heuristic used to identify components for improvement investments; it may be possible to generate greater improvement with a better coordinated investment strategy – e.g., a group of components that could together provide increased capability in a certain part of the network, rather than choosing components individually.) Overall, however, hardening was generally found to be more effective than investments in increasing capacity, adding new transmission lines, or reducing the restoration times of components (although reducing restoration times sometimes provided comparable results to hardening when we consider the long-term impact of an attack). Moreover, we found that adding new lines can sometimes make the system more vulnerable to attacks.

Close examination of our results may lead to interesting speculations about why different systems or different attack strategies lead to different results. Since our results are based on only four sample systems, it is important not to put too much weight on observations that may turn out to be merely artifacts of a particular system design. However, such speculations might be a source of interesting hypotheses to be tested more rigorously using detailed analysis of realistic networks and/or historical data.

For future work, we plan to implement cost functions for both the attacker and the defender. Our goal is to identify how attack and defense strategies may change in the face of budget constraints. This approach will enable us to identify cost-effective (rather than merely effective) defensive investments, which is often a challenge in security studies.

We also plan to compare our dynamic heuristic attack algorithm with other alternative heuristic methods, including genetic algorithms and capacity-based (rather than load-based) attack strategies. We will continue to develop and explore the use of realistic (rather than optimal) attack strategies that are efficient and simple enough to cause substantial damage to electricity systems over time, and study how to protect against them.

We anticipate that our model could also be applied to other capacity-constrained networks, including transportation systems and structural systems. Moreover, we believe that with some changes, our general approach could be used even for networks that are not highly capacity constrained, such as cyber networks.

REFERENCES

Adger, W. N. (2000). “Social and ecological resilience: Are they related?” *Progress in Human Geography*, 24(3), 347-364.

Albert, R., Albert, I. and Nakarado, G.L. (2004). “Structural vulnerability of the North American power grid.” *Physical Review E*, 69(2), 025103.

Allenby, B. and Fink, J. (2005). “Toward inherently secure and resilient societies.” *Science*, 309(5737), 1034-1036.

Al Mannai, W.I. and Lewis, T.G. (2008). “A general defender-attacker risk model for networks.” *The Journal of Risk Finance*, 9(3), 244-261.

Amin, M. (2002). "Toward secure and resilient interdependent infrastructures." *Journal of Infrastructure. Systems*, 8(3), 67-75.

Amin, M. (2003). "North America’s electricity infrastructure: Are we ready for more perfect storms." *IEEE Security & Privacy*, 1(5), 19-25.

Anderson, R. (2001). “Why information security is hard – an economic perspective.” *Proceedings of Seventeenth Computer Security Applications Conference*, 358-365.

Anghel, M., Werley, K.A. and Motter, A.E. (2007). “Stochastic model for power grid dynamics.” *40th Hawaii International Conference on System Sciences*, Big Island, Hawaii, January, 2007.

Andersson, G., Donalek, P., Farmer, R., Hatziargyriou, N., Kamwa, I., Kundur, P., Martins, N., Paserba, J., Pourbeik, P. and Sanchez-Gasca, J. (2005). “Causes of the 2003 major grid blackouts in North America and Europe, and recommended means to improve system dynamic performance.” *IEEE Transactions on Power Systems*, 20(4), 1922-1928.

Apostolakis, G. E. and D. M. Lemon (2005). "A screening methodology for the identification and ranking of infrastructure vulnerabilities due to terrorism." *Risk Analysis* 25(2), 361-376.

Apt, J. and Lave, L. B. (2003). “Electric gridlock: A national solution.” *Public Utilities Fortnightly*, 141(18), 14-16.

Apt, J., Lave, L.B., Talukdar, S., Morgan, M.G. and Ilic, M. (2004). “Electrical blackouts: A systemic problem.” *Issues in Science and Technology*, 20(4), 55-61.

Arroyo, J. M. and Galiana, F. D. (2005). “On the solution of the bilevel programming formulation of the terrorist threat problem.” *Power Systems, IEEE Transactions on*, 20(2), 789-797.

Ash, J. and Newth, D. (2007). “Optimizing complex networks for resilience against cascading failure.” *Physica A: Statistical Mechanics and its Applications*, 380, 673-683.

Bennett, B.T. (2007). *Understanding, Assessing, and Responding to Terrorism: Protecting Critical Infrastructure and Personnel*, Wiley-Interscience, Hoboken, NJ.

Bienstock, D. and Mattia, S. (2007). “Using mixed-integer programming to solve power grid blackout problems.” *Discrete Optimization*, 4(1), 115-141.

Bienstock, D. and Verma, A. (2009). “The N-k problem in power grids: new models, formulations and numerical experiment.” *SIAM Journal on Optimization*, 20(5), 2352-2380.

Bier, V.M., Gratz, E.R., Haphuriwat, N.J., Magua, W, Wierzbicki, K.R. (2007). “Methodology for identifying near-optimal interdiction strategies for a power network transmission system.” *Reliability Engineering and System Safety*, 92(9), 315–323.

Blume, S. W. (2007). *Electric power system basics: for the nonelectrical professional*, Wiley-IEEE Press, Hoboken, NJ.

Brooks, D. L., Dugan, R. C., Grebe, T. E. and Sundaram, A. (2002). “Disturbances recorded by power quality monitors during the ‘West Coast Outage’” *Proceedings of 1996 IEEE Power Engineering Society Transmission and Distribution Conference*, 315–320.

Brown, G.G. and Cox Jr, L.A.T. (2011). “How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts.” *Risk Analysis*, 31(2), 196-204.

Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O’Rourke, T. D., Reinhorn, A. M., Shinozuka, M., Tierney, K., Wallace, W. A. and von Winterfeldt, D. (2003). "A framework to quantitatively assess and enhance the seismic resilience of communities." *Earthquake Spectra*, 19(4), 733-752.

Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E. and Havlin, S. (2009). “Catastrophic cascade of failures in interdependent networks.” *Nature*, 464(7291), 1025-1028.

Burnham, J. T. (1995). “Bird streamer flashovers on FPL transmission lines.” IEEE Transactions on Power Delivery, 10(2), 970-977.

Bush, B., Dauelsberg, L. R. , LeClaire, R. J., Deland D. R. and Samsa, S. M. (2005). “Critical infrastructure protection decision support system (CIP/DSS) project overview.” Report no. LA-UR-05-1870, Los Alamos National Laboratory, Los Alamos, NM.

Bush, G.W. (2002). *The National Security Strategy of the United States of America*, Executive Office of the President, Washington, D.C.

Carreras, B. A., Lynch, V. E., Dobson, I. and Newman, D. E. (2002). “Critical points and transitions in an electric power transmission model for cascading failure blackouts.” *Chaos*, 12, 985-994.

Carreras, B.A., Lynch, V.E., Dobson, I. and Newman, D.E.(2004). “Complex dynamics of blackouts in power transmission systems.” *Chaos*, 14, 643.

Carreras, B.A., Newman, D. E., Dobson, I., and Degala, N. S., (forthcoming 2013), “Validating OPA with WECC data.” *46th Hawaii International Conference on System Sciences*, Hawaii, January, 2013.

Chang, S. E., McDaniels, T. L., Mikawoz, J. and Peterson, K. (2007). “Infrastructure failure interdependencies in extreme events: Power outage consequences in the 1998 ice storm.” *Natural Hazards*, 41(2), 337-358.

Chen, J., Thorp, J.S. and Dobson, I. (2005). “Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model.” *International Journal of Electrical Power & Energy Systems*, 27(4), 318-326.

Clarke, L. (2002). “Panic: Myth or reality?” *Contexts*, 1(3), 21-26.

Capodieci, P., Diblasi, S., Ciancamerla, E., Minichino, M., Foglietta, C., Lefevre, D., Oliva, G.,Panzieri, S., Setola, R. and Porcellinis, S. De (2010) “Improving resilience of interdependent critical infrastructures via an on-line alerting system.” *IEEE-Complexity in Engineering*, 88-90.

Comfort, L. K. (1999). *Shared Risk: Complex Systems in Seismic Response*, Pergamon, New York, NY.

Cormican, K.J., Morton, D.P. and Wood, R.K. (1998). “Stochastic network interdiction.” *Operations Research*, 46(2), 184-197.

Cox, L. A. Jr. (2008). “Some Limitations of ‘Risk= Threat× Vulnerability× Consequence’ for Risk Analysis of Terrorist Attacks.” *Risk Analysis*, 28(6), 1749-1761.

Crane, A. T. (1990). “Physical vulnerability of electric systems to natural disaster and sabotage.” *Studies in Conflict & Terrorism*, 13 (3), 189-190.

Crucitti, P., V. Latora, et al. (2004). "Model for cascading failures in complex networks." *Physical Review E*, 69(4), 045104.

Department of Homeland Security (DHS) (2003). The National Strategy for the *Physical Protection of Critical Infrastructures and Key Assets,* U.S. Department of Homeland Security, Washington, D.C.

Department of Homeland Security (DHS) (2006). *Bioterrorism Risk Assessment*. U.S. Department of Homeland Security, Washington, D.C.

Department of Homeland Security (DHS) (2008). *DHS Risk Lexicon*, Risk Steering Committee, U.S. Department of Homeland Security, Washington, D.C.

Department of Homeland Security (DHS) (2009). *National Infrastructure Protection Plan*, U.S. Department of Homeland Security, Washington, D.C.

Department of Homeland Security (DHS) (2010). *DHS Risk Lexicon*, Risk Steering Committee, U.S. Department of Homeland Security, Washington, D.C.

Department of Transportation (DOT). (2004). *Effects of Catastrophic Events on Transportation Systems Management and Operations- August 2003 Northeast Blackout Great Lakes Region*, Volpe National Transportation Systems Center, Cambridge, MA.

Dobson, I., Carreras, B., Lynch, V. and Newman, D. (2001). "An initial model for complex dynamics in electric power system blackouts." p. 2017 (Published by the IEEE Computer Society, 2001)

Dobson, I., Carreras, B. A., Newman, D. E. (2004), "A branching process approximation to cascading load-dependent system failure." *37th Hawaii International Conference on System Sciences*, Hawaii, January, 2004.

Dobson, I., Carreras, B.A., Lynch, V.E. and Newman, D.E. (2007). "Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization." *Chaos*, 17, 026103.

Dobson, I., Kim, J. and Wierzbicki, K.R. (2010). "Testing branching process estimators of cascading failure with data from a simulation of transmission line outages." *Risk Analysis*, 40(3), 650-662.

Dobson, I (forthcoming 2012). "Estimating the propagation and extend of cascading line outages from utility data with a branching process." *IEEE Transactions on Power Systems*.

Dobson, I. and Carreras, B. A. (2010), 'Number and propagation of line outages in cascading events in electric power transmission systems." *48th Annual Allerton Conference on Communication, Control, and Computing*, University of Illinois at Urbana-Champaign, IL, September, 2010.

Donde, V., Lopez, V., Lesieutre, B., Pinar, A., Yang, C. and Meza, J. (2005). "Identification of severe multiple contingencies in electric power networks." *IEEE*, 59-66.

Dudenhoeffer, D.D., Permann, M.R. and Manic, M. (2006). "CIMS: a framework for infrastructure interdependency modeling and analysis." *Proceedings of 2006 Winter Simulation Conference*, pp. 478-485.

Dueñas-Osorio, L. and Vemuru, S. M. (2009). "Cascading failures in complex infrastructure systems." *Structural Safety* 31(2), 157-167.

Einarsson, S. and Rausand, M.(1998). "An approach to vulnerability analysis of complex industrial systems." *Risk Analysis*, 18(5), 535-546.

Electricity Consumer Resource Council (2004). *The Economic Impacts of the August 2003 Blackout*, ELCON, Washington, D.C.

Ezell, B.C., Farr, J.V. and Wiese, I. (2000). “Infrastructure risk analysis model.” *Journal of Infrastructure Systems*, 6, 114.

Ezell, B.C. (2007). “Infrastructure Vulnerability Assessment Model (I-VAM).” *Risk Analysis*, 27(3), 571-583.

Ezell, B.C, Bennett, S.P., von Winderfeldt, D., Sokolowski, J. and Collins, A.J. (2010). “Probabilistic risk analysis and terrorism risk.” *Risk Analysis*, 30(4), 575-589.

Farrell, A. E., Lave, L. B. and Morgan, G. (2002). “Bolstering the security of the electric power system.” *Issues in Science and Technology*, Spring, 49-56.

Farrell, A. E., Zerriffi, H. and Dowlatabadi, H. (2004). “Energy infrastructure and security.” *Annual Review of Environment and Resources*, 29, 421-469.

Federal Energy Regulatory Commission (FERC). (1978). *The Con Edison Power Failure of July 13 and 14, 1977*, U.S. Government Printing Office, Washington, D.C.

Federal Energy Regulatory Commission (FERC). (2006). *Docket Nos. EC06-140-000 and EL06-86-000*, U.S. Government Printing Office, Washington, D.C.

Federal Power Commission (FPC). (1965). *Northeast Power Failure: November 9 and 10, 1965*, U.S. Government Printing Office, Washington, D.C.

Fiksel, J. (2003). “Designing resilient, sustainable systems.” *Environmental Science and Technology*, 37(23), 5330-5339.

Fisher, R. E., & Norman, M. (2010). “Developing measurement indices to enhance protection and resilience of critical infrastructure and key resources.” *Journal of Business Continuity & Emergency Planning*, 4(3), 191-206.

Gunderson, L. H., Holling, C. S., Pritchard Jr., L. and Peterson, G. D. (2002). “Resilience of large-scale resource systems” *Resilience and the Behavior of Large-Scale Systems* eds. Gunderson, L. H. and Pritchard Jr., L., Island Press, Washington, D.C.

Haimes, Y. Y., Matalas, N. C., Lambert, J. H., Jackson, B. A. and Fellows, J. F. R. (1998). “Reducing vulnerability of water supply systems to attack.” *Journal of Infrastructure Systems*, 4(4), 164-177.

Haimes, Y. Y. (2009). "On the definition of resilience in systems." *Risk Analysis*, 29(4), 498-501.

Hansson, S. O. and Helgesson, G. (2003). “What is stability?” *Synthese*, 136(2), 219-235.

He, T., Kolluri, S., Mandal, S., Galvan, F. and Rastgoufard, P. (2005). “Identification of weak locations in bulk transmission systems using voltage stability margin index.” *Applied Mathematics for Restructured Electric Power Systems*, 25-37.

Hines, P., Cotilla-Sanchez, E. and Blumsack, S. (2010). “Do topological models provide good information about electricity infrastructure vulnerability?” *Chaos*, 20, 033122.

Holling, C. S. (1973). "Resilience and stability of ecological systems." *Annual Review of Ecology, Evolution and Systematics*, 4(1), 1-23.

Hollnagel, E. (2006). "Resilience - the challenge of the unstable." (Chapter 4) *Resilience Engineering: Concepts and Precepts* eds. Hollnagel, E., Woods, D. D. and Leveson, N., Ashgate Publishing, Hampshire, England, 9-17.

Holmgren, Å.J. (2006). “Using graph models to analyze the vulnerability of electric power networks.” *Risk Analysis*, 26(4), 955-969.

Holmgren, Å. (2007). “A framework for vulnerability assessment of electric power systems.” *Reliability and Vulnerability in Critical Infrastructure: A Quantitative Geographic Perspective* eds. Murray, A., Grubesic, T., Springer, New York, NY, 31-55.

Holmgren, A.J., Jenelius, E. and Westin, J. (2007). “Evaluating strategies for defending electric power networks against antagonistic attacks.” *Power Systems, IEEE Transactions on*, 22(1), 76-84.

Horne III, J. F. and Orr, J. E. (1998). “Assessing behaviors that create resilient organizations.” *Employment Relations Today*, 24(4), 29-39.

Ibáñez, E., Gkritza, K., McCalley, J., Aliprantis, D., Brown, R., Somani A. and Wang, L. (2010). “Interdependencies between energy and transportation systems for national long term planning.” *Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modeling, and Intelligent Engineering*, eds. Gopalakrishnan, K. and Peeta, S., Springer-Verlag, Berlin, Germany, 53-76.

IEEE (1999). “The IEEE reliability test system—1996,” *IEEE Transmission Power Systems*, 14, 1010-1020.

Janjarassuk, U. and Linderoth, J. (2008). “Reformulation and sampling to solve a stochastic network interdiction problem.” *Networks*, 120-132.

Kaplan, S. (1997). “The words of risk analysis.” *Risk Analysis*, 17(4), 407-417.

Kappernman, J. G. and Albertson, V. D. (1990). “Bracing for the geomagnetic storms.” *IEEE Spectrum*, 27(3), 27-33.

Kerry, M., Kelk, G., Etkin, D., Burton I. and Kalhok, S. (1999). “Glazed over: Canada copes with the ice storm of 1998.” *Environment: Science and Policy for Sustainable Development*, 41(1), 6-11.

Kinney, R., Crucitti, P., Albert, R. and Latora, V. (2005). “Modeling cascading failures in the North American power grid.”, *The European Physical Journal B-Condensed Matter and Complex Systems*, 46(1), 101-107.

Koonce, A.M., Apostolakis, G.E. and Cook, B.K. (2008). “Bulk power risk analysis: Ranking infrastructure elements according to their risk significance.” *International Journal of Electrical Power & Energy Systems*, 30(3), 169-183.

Lave, L. B., Apt, J., Farrell A. and Morgan, M. G. (2007). “Increasing the security and reliability of the US electricity system.” (Chapter 4) *The Economic Impacts of Terrorist Attacks* eds. Richardson, H. W. and Gordon, P., Edward Elgar Publishing, Cheltenham, UK, 57-70.

Lecomte, E. L., Pang, A. W. and Russell, J. W. (1998). *Ice storm '98*, Institute for Catastrophic Loss Reduction, Toronto, Canada and Institute for Business & Home Safety, Boston, MA.

Lesieutre, B.C., Roy, S., Donde, V. and Pinar, A. (2006). “Power system extreme event screening using graph partitioning.”,*IEEE*,. 503-510.

Liao, H., Apt, J. and Talukdar, S. (2004). “Phase transitions in the probability of cascading failures.” *Conference at Carnegie Mellon University*, Pittsburgh PA USA.

Lewis, T.G. (2009). *Network science: theory and applications*, Wiley Publishing.

Lovins, A. B. and Lovins, L. H. (1982). *Brittle power: Energy strategy for national security*, Brick House Publishing, Andover, MA.

Makansi, J. (2007). *Lights out: The electricity crisis, the global economy, and what it means to you*, Wiley, Hoboken, NJ.

Mileti, D. S. (1999). *Disasters by Design: A Reassessment of Natural Hazards in the United States*, Joseph Henry Press, Washington, D.C.

Mili, L., Qiu, Q. and Phadke, A.G. (2004). “Risk assessment of catastrophic failures in electric power systems.” *International journal of critical infrastructures*, 1(1), 38-63.

Morton, D.P., Pan, F. and Saeger, K.J. (2007). “Models for nuclear smuggling interdiction.” *IIE Transactions*, 39(1), 3-14.

Motto, A. L., Arroyo, J. M. and Galiana, F. D. (2005). “A mixed-integer LP procedure for the analysis of electric grid security under disruptive threat.” *IEEE Transactions on Power Systems*, 20(3), 1357-1365.

National Infrastructure Advisory Council (NIAC) (2009), *Critical Infrastructure Resilience Final Report and Recommendations*, U.S. Department of Homeland Security, Washington, D.C.

National Research Council (NRC) (2010). *Committee to Review the Department of Homeland Security Approach to Risk Analysis*, The National Academies Press, Washington, D.C.

Newman, D.E., Carreras, B.A., Lynch, V.E. and Dobson, I. (2011). “Exploring complex systems aspects of blackout risk and mitigation”. *IEEE Transactions on Reliability*, 60(1), 134-143.

Ni, M., McCalley, J.D., Vittal, V. and Tayyib, T. (2003). “Online risk-based security assessment.” *Power Systems, IEEE Transactions on*, 18(1), 258-265.

North American Electricity Reliability Corporation (NERC) (2002). *Security Guidelines for the Electricity Sector*, Princeton, NJ.

North American Electricity Reliability Corporation (NERC) (2010). *2010 Annual Report on Bulk Power System Reliability Metrics*, Princeton, NJ.

Parfomak, P.W. and Frittelli, J. (2007) *Maritime Security: Potential Terrorist Attacks and Protection Priorities*, Congressional Research Service (CRS) Report for Congress, Washington, D.C.

Patterson, W. C. (2007). *Keeping the lights on: Towards sustainable electricity*, Earthscan/James & James, London, UK.

Patterson, S.A. and Apostolakis, G.E.(2008). “Identification of critical locations across multiple infrastructures for terrorist actions.” *Reliability Engineering & System Safety*, 92(9), 1183-1203.

Pinar, A., Meza, J., Donde, V. and Lesieutre, B. (2010) “Optimization strategies for the vulnerability analysis of the electric power grid.” *SIAM Journal on Optimization*, 20(4), 1786-1810.

Rabkin, M., Brodesky, R., Ford, F., Haines, M., Karp, J., Lovejoy, K., Regan, T., Sharpe, L. and Zirker M. (2004). *Transit Security Design Considerations*, Volpe National Transportation Systems Center, U.S. Department of Transportation, Cambridge, MA.

RAMCAP (Risk Analysis and Management for Critical Asset Protection) Framework (2006). Available at http://www.personal.psu.edu/jsd222/SRA311/RAMCAPframework_Risk_Analysis_and_Manage.pdf.

Rasmussen, N.C. (1975). “Reactor safety study: An assessment of accident risks in US commercial nuclear power plants.”, WASH-1400, Nuclear Regulatory CommissionWashington, D.C.

Risley, A. and Roberts, J. (2003). “Electronic security risks associated with use of wireless, point-to-point communications in the electric power industry.” *Proceedings of DistribuTECH Conference*, 1-16.

Romero, N., Xu, N., Nozick, L. K. and Dobson, I. (2012). “Investment planning for electric power systems under terrorist threat.” *IEEE Transactions on Power Systems*, 27(10), 108-116.

Rose, A. and Liao, S. Y. (2005). "Modeling regional economic resilience to disasters: A computable general equilibrium analysis of water service disruptions." *Journal of Regional Science*, 45(1), 75-112.

Salmeron, J., Wood, K. and Baldick, R. (2004). “Analysis of electric grid security under terrorist threat.” *IEEE Transactions on Power Systems*, 19(2), 905-912.

Salmeron, J., Wood, K., Baldick, R., (2009). “Worst-case interdiction analysis of large-scale electric power grids.” *IEEE Transactions on Power Systems*, 24(1), 96-104.

Scanlon, J. (1999). “Emergent Groups in Established Frameworks: Ottawa Carleton's Response to the 1998 Ice Disaster.” *Journal of Contingencies and Crisis Management*, 7(1), 30-37.

Smith, J.E. and Von Winterfeldt, D. (2004). “Decision Analysis in Management Science." *Management Science*, 50 (5), 561-574.

Talukdar, S. N., Apt, J., Ilic, M., Lave, L. B. and Morgan, M. G. (2003) “Cascading failures: Survival versus prevention.” *The Electricity Journal*, 16 (9), 25-31.

Tierney, K. and Bruneau M. (2007). "Conceptualizing and Measuring Resilience." *TR News*, 250,14-17.

TRAM *(*Transit Risk Assessment Methodology), Security Analysis and Risk Management Association (2007). Available at http://www.sarma-wiki.org/index.php?title=Transit_Risk_Assessment_Methodology_(TRAM)

U.S. Coast Guard (USCG) (2003), “Implementation of National Maritime Security Initiatives”, *Federal Register*, 68(126), 39240-39250.

Wang, J.W. and Rong, L.L. (2009). “Cascade-based attack vulnerability on the US power grid.” *Safety Science*, 47(10), 1332-1336.

Wildavsky, A. B. (1988). *Searching for safety*, Transaction publishers, Piscataway, NJ.

Wood, R. K. (1993). "Deterministic network interdiction." *Mathematical and Computer Modelling*, 17(2), 1-18.

Woods, D. D. (2006). “Essential characteristics of resilience.” (Chapter 2) eds. Hollnagel, E., Woods D. D. and Leveson, N., *Resilience engineering: Concepts and precepts*, Ashgate Publishing, Aldershot, UK.

Vugrin, E. D., Warren, D. E., Ehlen M. A. and Camphouse, R. C. (2010). "A Framework for assessing the resilience of infrastructure and economic systems." eds. Gopalakrishnan, K. and Peeta, S., *Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modeling, and Intelligent Engineering*, Springer, Berlin, Germany, 77-116.

Volkanovski, A., Cepin, M. and Mavko, B. (2009). “Application of the fault tree analysis for assessment of power system reliability.” *Reliability Engineering & System Safety*, 94(6), 1116-1127.

Yao, Y., Edmunds, T., Papageorgiou, Alvarez, R. (2007). “Trilevel optimization in power network defense” *IEEE Transactions on Systems, Man and Cybernetics-Part C:Applications and Reviews*, 37(4), 712-718.

Zhao, L., Park, K. and Lai, Y.C. (2004). “Attack vulnerability of scale-free networks due to cascading breakdown.” *Physical Review E*, 70(3), 035101.

Zima, M. and Andersson, G. (2004). “Wide area monitoring and control as a tool for mitigation of cascading failures.” *Eight International Conference on Probability Methods Applied to Power Systems*, Ames, Iowa.

Zimmerman, R., Restrepo, C. E., Simonoff, J. S. and Lave, L. (2007). "Risk and economic cost of a terrorist attack on the electric system" (Chapter 14) eds Richardson, H. W., Gordon, P. and Moore, J. E. II., *The Economic Costs and Consequences of Terrorism*, Edward Elgar Publishers, Cheltenham, UK, 273-290.

Appendix A. Definitions of resilience in the literature

Resilience is “the capacity to cope with unanticipated dangers after they have become manifest, learning to bounce back.’’ (Wildavsky 1991, p. 77)

Resilience is “the ability of a system to withstand stresses of ‘environmental loading’... [it is] a fundamental quality found in individuals, groups, organizations, and systems as a whole.” (Horne and Orr 1998, p. 31)

Resilience is “the capacity to adapt existing resources and skills to new situations and operating conditions.’’ (Comfort 1999, p. 21)

Resilience is “both the inherent strength and ability to be flexible and adaptable after environmental shocks and disruptive events.” (Tierney and Bruneau 2007, p. 17)

Resiliency is “the capability of an asset, system, or network to maintain its function during or to recover from a terrorist attack or other incident.” (DHS, 2006)

Resilience is the “ability to resist, absorb, recover from or successfully adapt to adversity or a change in conditions.” (DHS, 2008)

Resilience is the “ability of systems, infrastructures, government, business, and citizenry to resist, absorb, recover from, or adapt to an adverse occurrence that may cause harm, destruction, or loss of national significance.” (DHS, 2008)

Resilience is the “capacity of an organization to recognize threats and hazards and make adjustments that will improve future protection efforts and risk reduction measures.” (DHS, 2008)

Resiliency is “defined as the capability of a system to maintain its functions and structure in the face of internal and external change and to degrade gracefully when it must.” (Allenby and Fink, 2005)

“Regional economic resilience is the inherent ability and adaptive response that enables firms and regions to avoid maximum potential losses.” (Rose and Liao, 2005)

“Engineering resilience […] is the speed of return to the steady state following a perturbation […] ecological resilience […] is measured by the magnitude of disturbance that can be absorbed before the system is restructured….” (Gunderson et al., 2002)

Social resilience as the ability of groups or communities to cope with external stresses and disturbances as a result of social, political, and environmental change (Adger, 2000).

Resilience is “the essence of sustainability […] the ability to resist disorder.” (Fiksel, 2003)

“Local resiliency with regard to disasters means that a locale is able to withstand an extreme natural event without suffering devastating losses, damage, diminished productivity, or quality of life and without a large amount of assistance from outside the community.” (Mileti, 1999)

Appendix B. Some major blackouts in North America

Date of initiation	Region	Duration	People affected (more than million)	Reason	Source
11.09.1965	Northeast (80 thousand square miles)	2.5 hours	30	Human error, and cascading relay operations and line outages	FPC (1965); Risley and Roberts (2003)
06.05.1967	PJM (15 thousand square miles)	Several hours	4	Scheduled switching not performed, line overload	Apt et al. (2004)
05.17.1977	Miami (15 thousand square miles)	Several hours	1	Unknown cause	Burnham (1995); Amin (2003); Apt et al. (2004)
07.13.1977	NYC	25 hours	8	Lightning on a substation	FERC (1978); Crane (1990)
01.01.1981	Idaho, Utah, Wyoming	7 hours	1.5	Human error, and series of utility breakdowns	Amin (2003); Apt et al. (2004)
03.27.1982	Oregon	1.5 hours	0.9	High-voltage line failure	Amin (2003)
03.13.1989	Quebec	9 hours	6	Geomagnetic storm	Kappernman and Albertson (1990)
12.14.1994	Arizona to Washington state	Several hours	2	Relays/controllers coordination	Brooks et al.; (2002); Amin (2003)
07.02.1996	14 Western states	1-2 hours	2	A high-voltage line touched a tree branch	Amin (2003)
08.10.1996	11 Western states and 2 Canada provinces	6 hours	7	Two high-voltage lines fell in Oregon and caused cascading outages	Amin (2003)
01.05.1998	New York, New England, Canada	Several days	2.2	Ice storms	Lecomte et al. (1998); Amin (2003)
12.08.1998	San Francisco. California Bay area	8 hours	0.5	Human error	Amin (2003)
07.06.1999	NYC	19 hours	0.3	Heat related failure of feeder cables	Amin (2003); Holmgren (2007)
08.14.2003	Great Lakes-NYC	2 days	50	Cascading outage as a result of a combination of electrical, computer and human failures	Andersson et al. (2005); Holmgren (2007)

Appendix C. A static attacker choosing different component types
(a static attacker with no cascading knowledge)

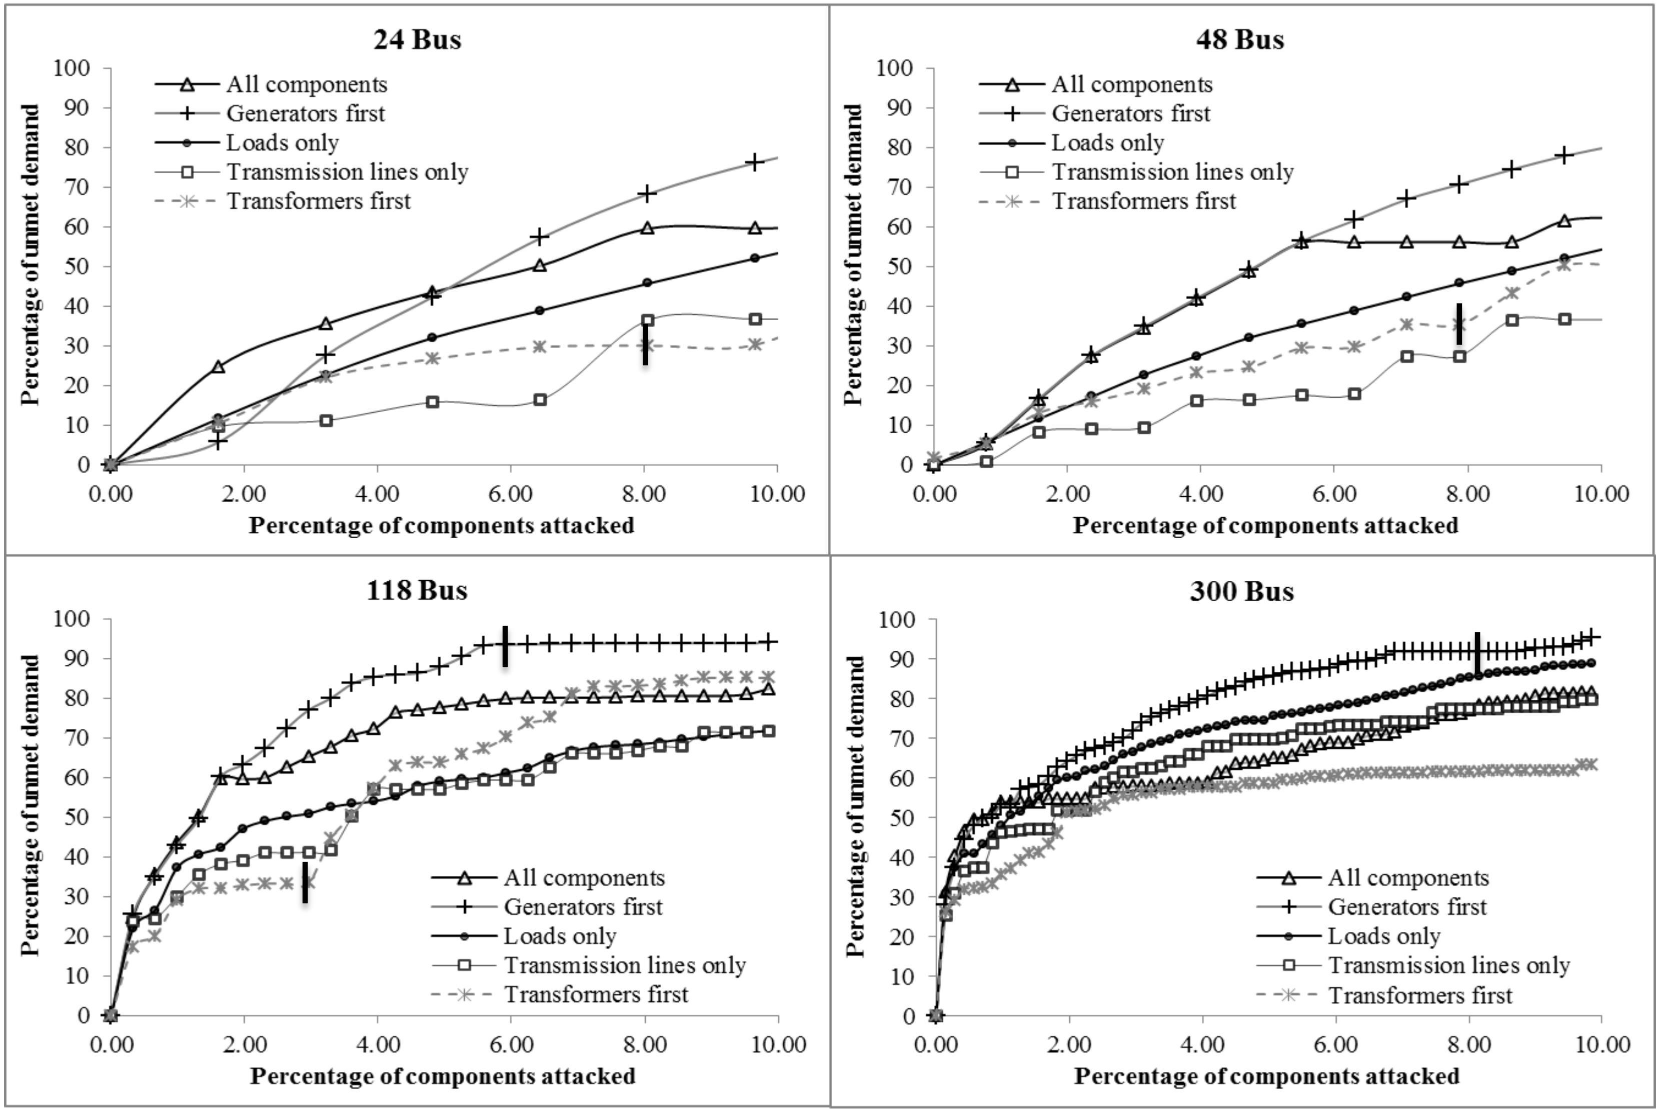

Appendix C. A static attacker choosing different component types (a static attacker with cascading knowledge)

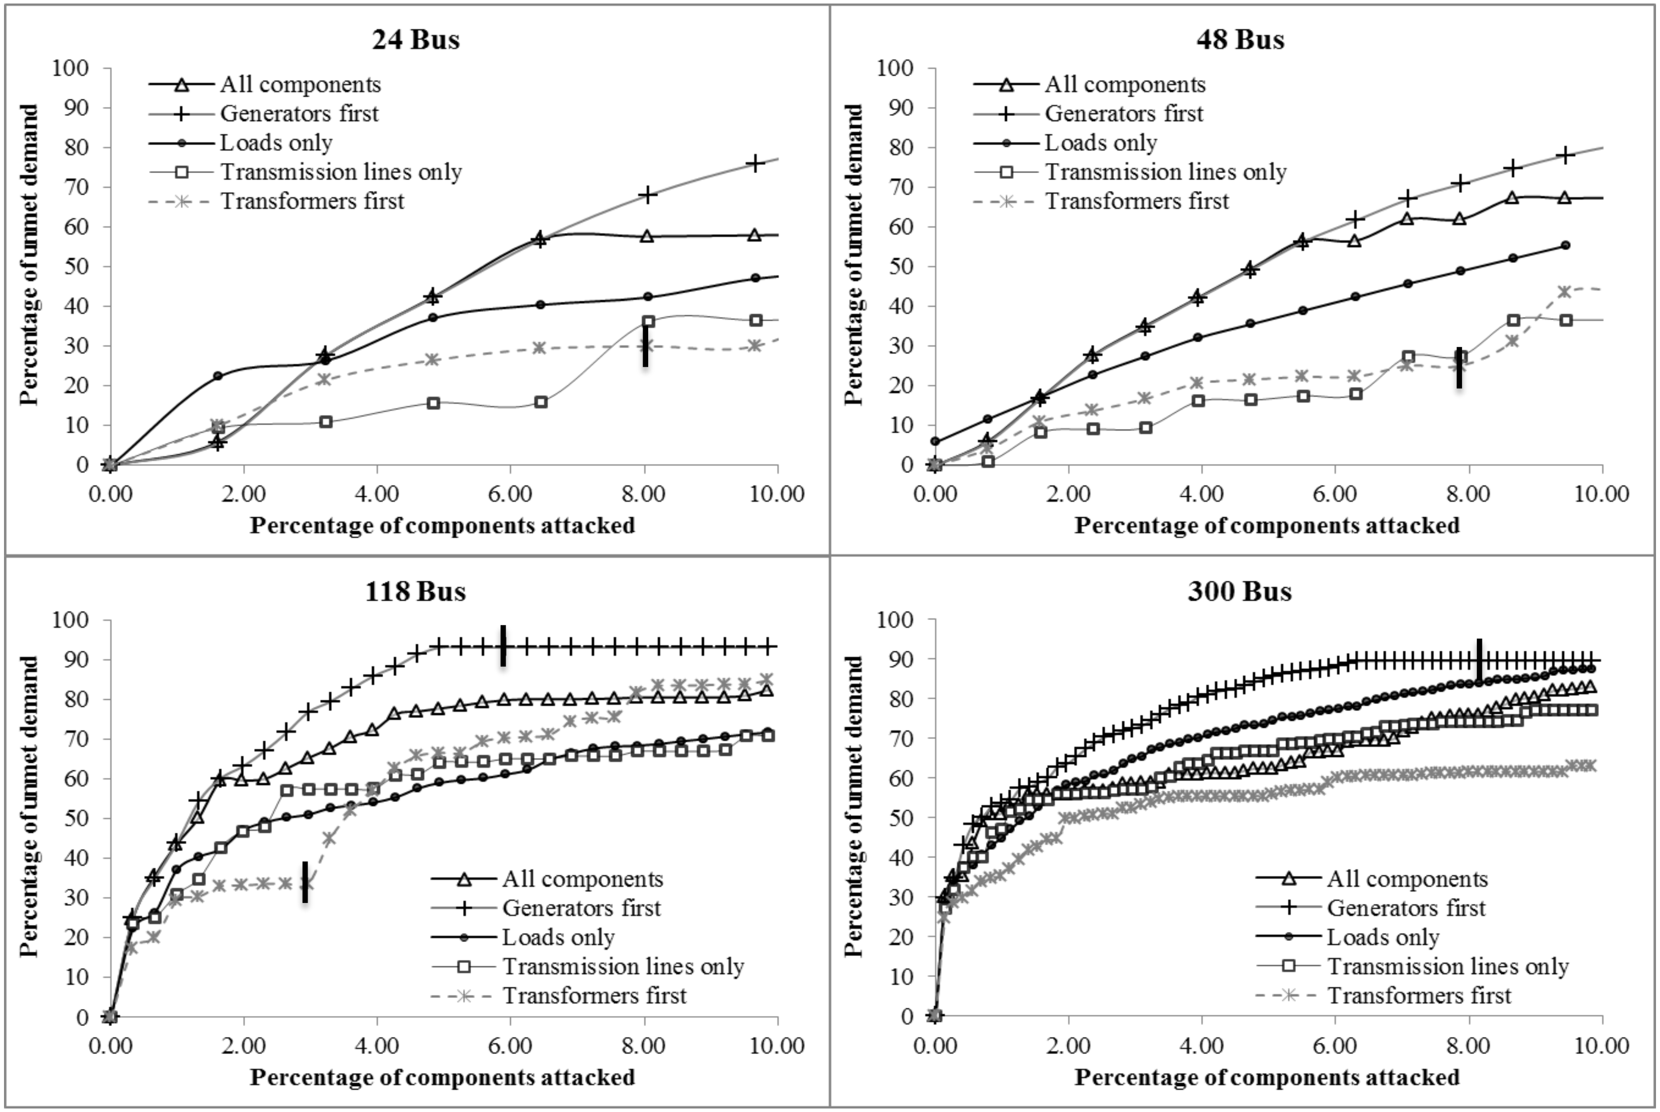

Appendix D. Sensitivity of the cascading model
(to the fraction of transmission capacity above which cascading failure occurs for a static attacker with cascading knowledge)

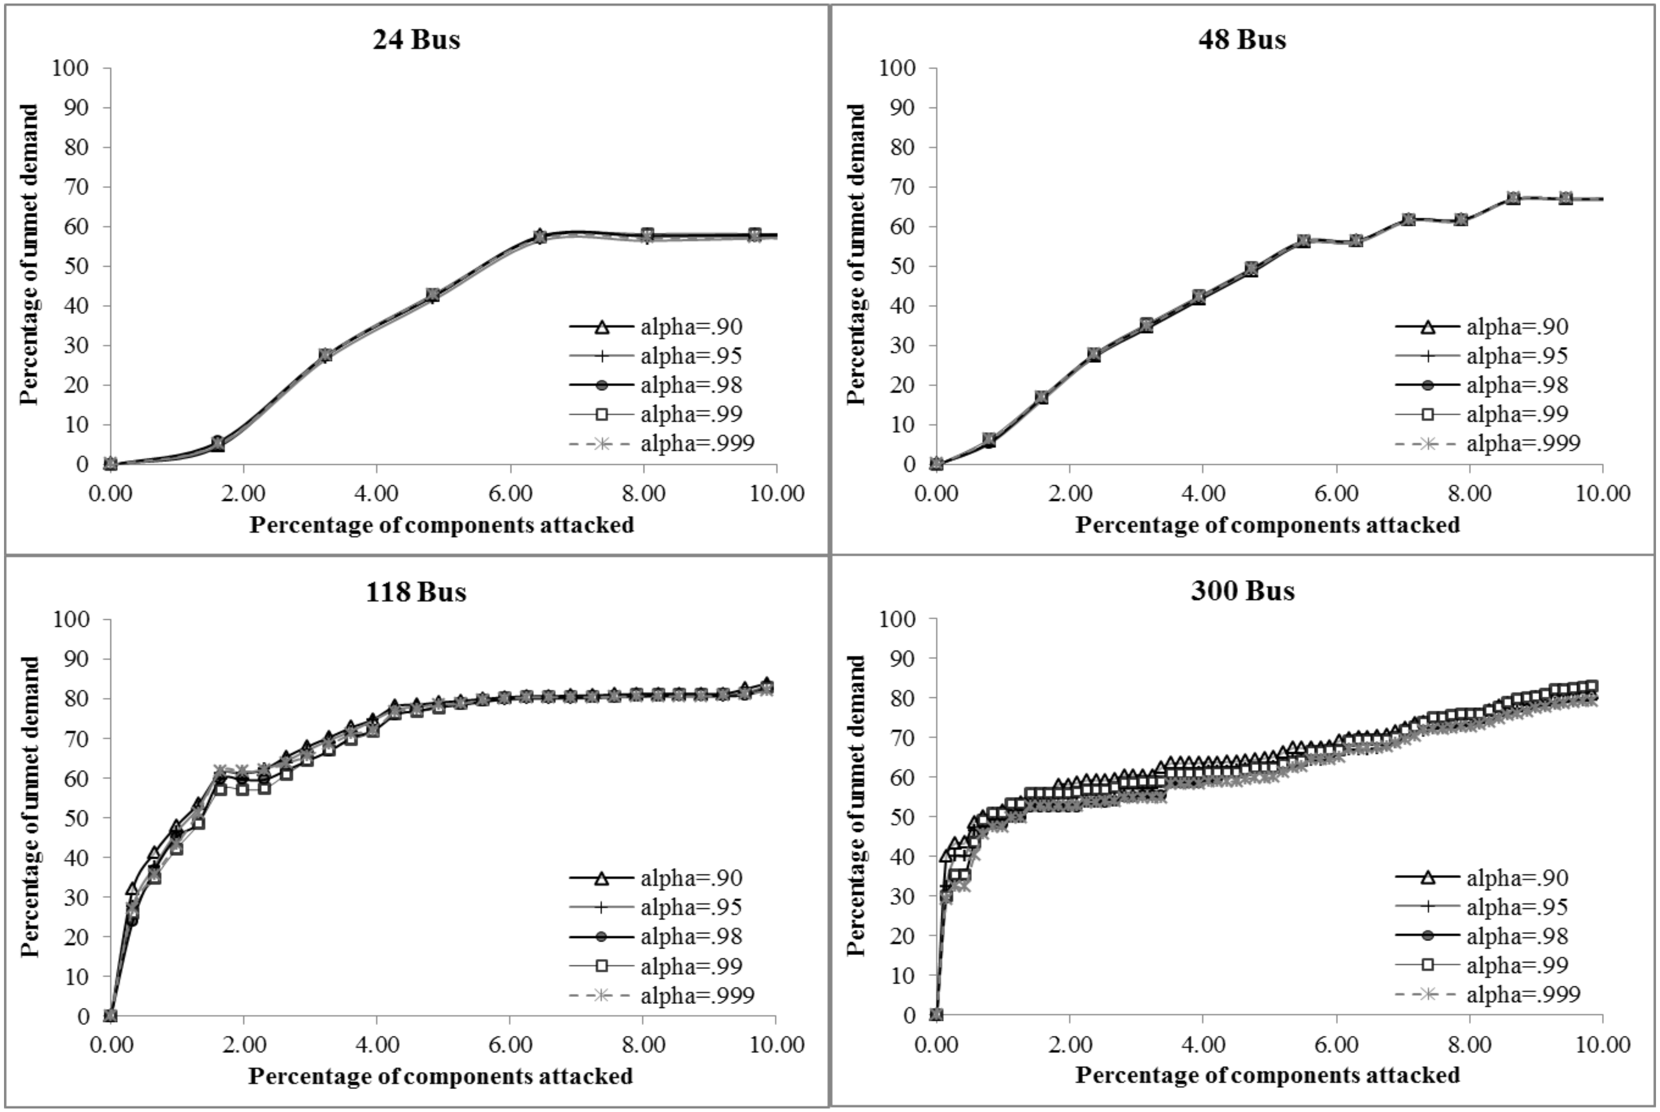

Appendix D. Sensitivity of the cascading model
(to the failure probability of an overloaded component for a static attacker with cascading knowledge)

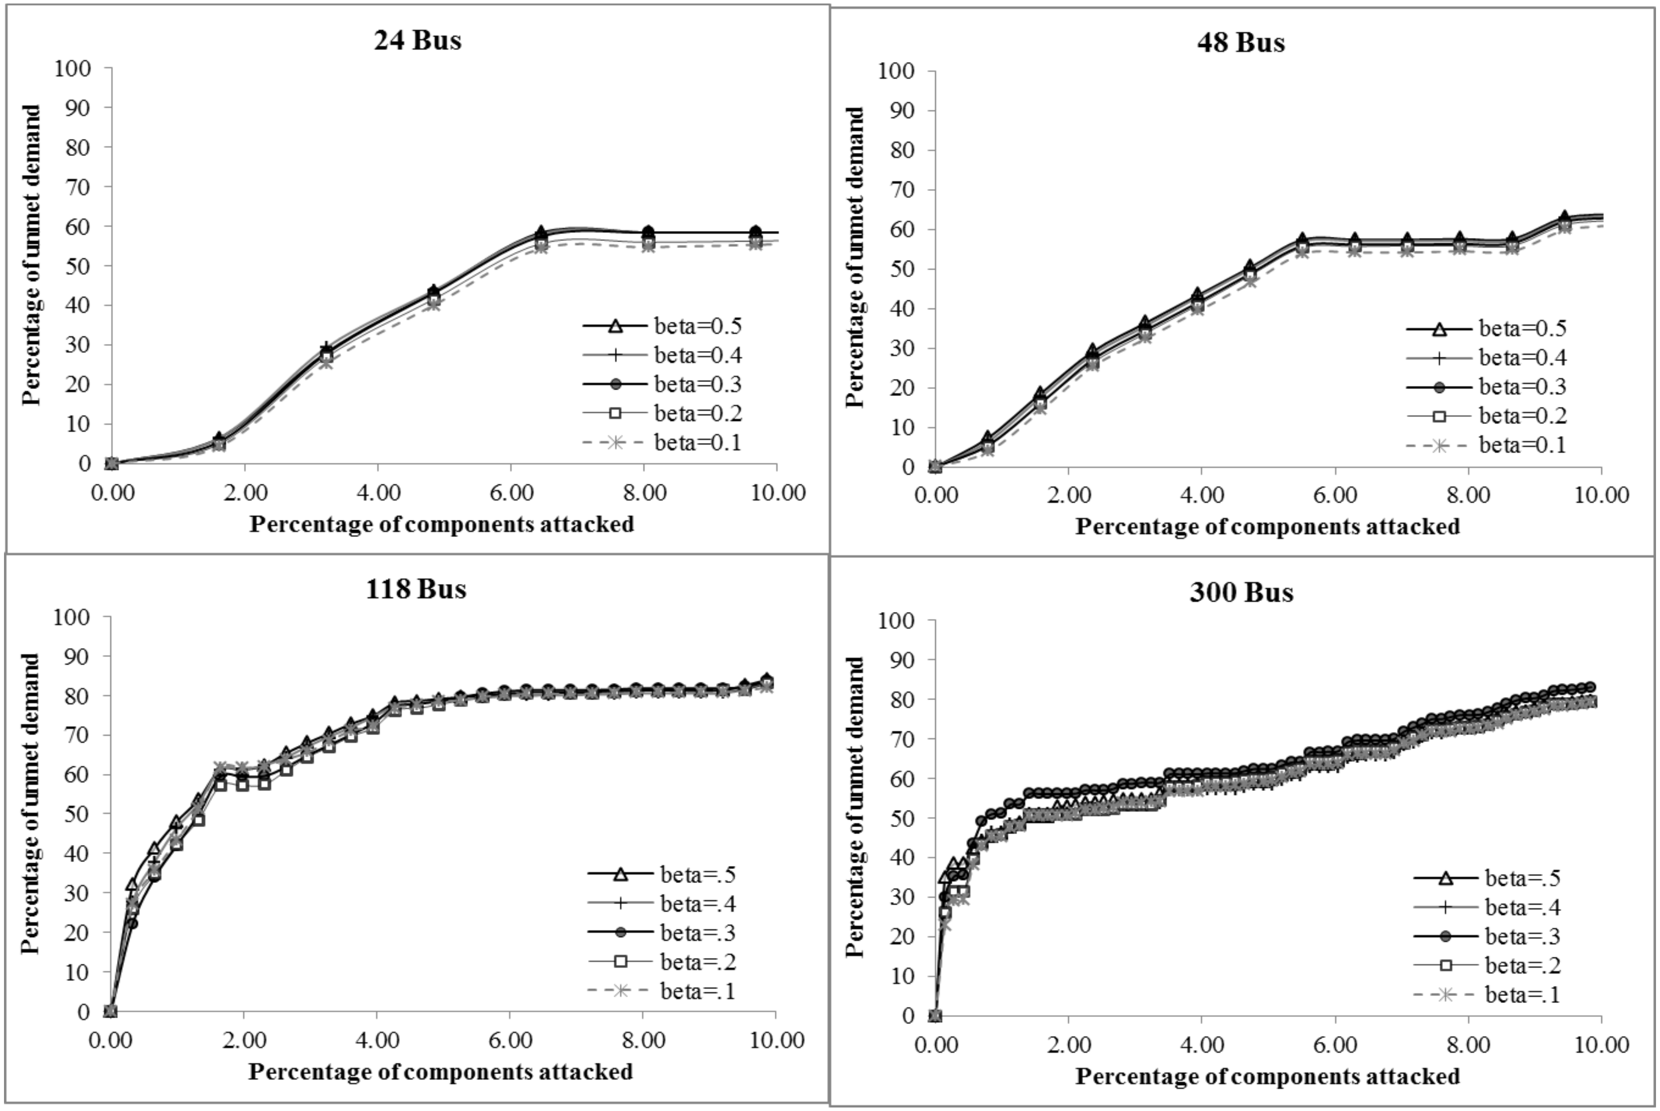

Appendix D. Sensitivity of the cascading model
(to the fraction of transmission capacity above which cascading failure occurs for a static attacker with no cascading knowledge)

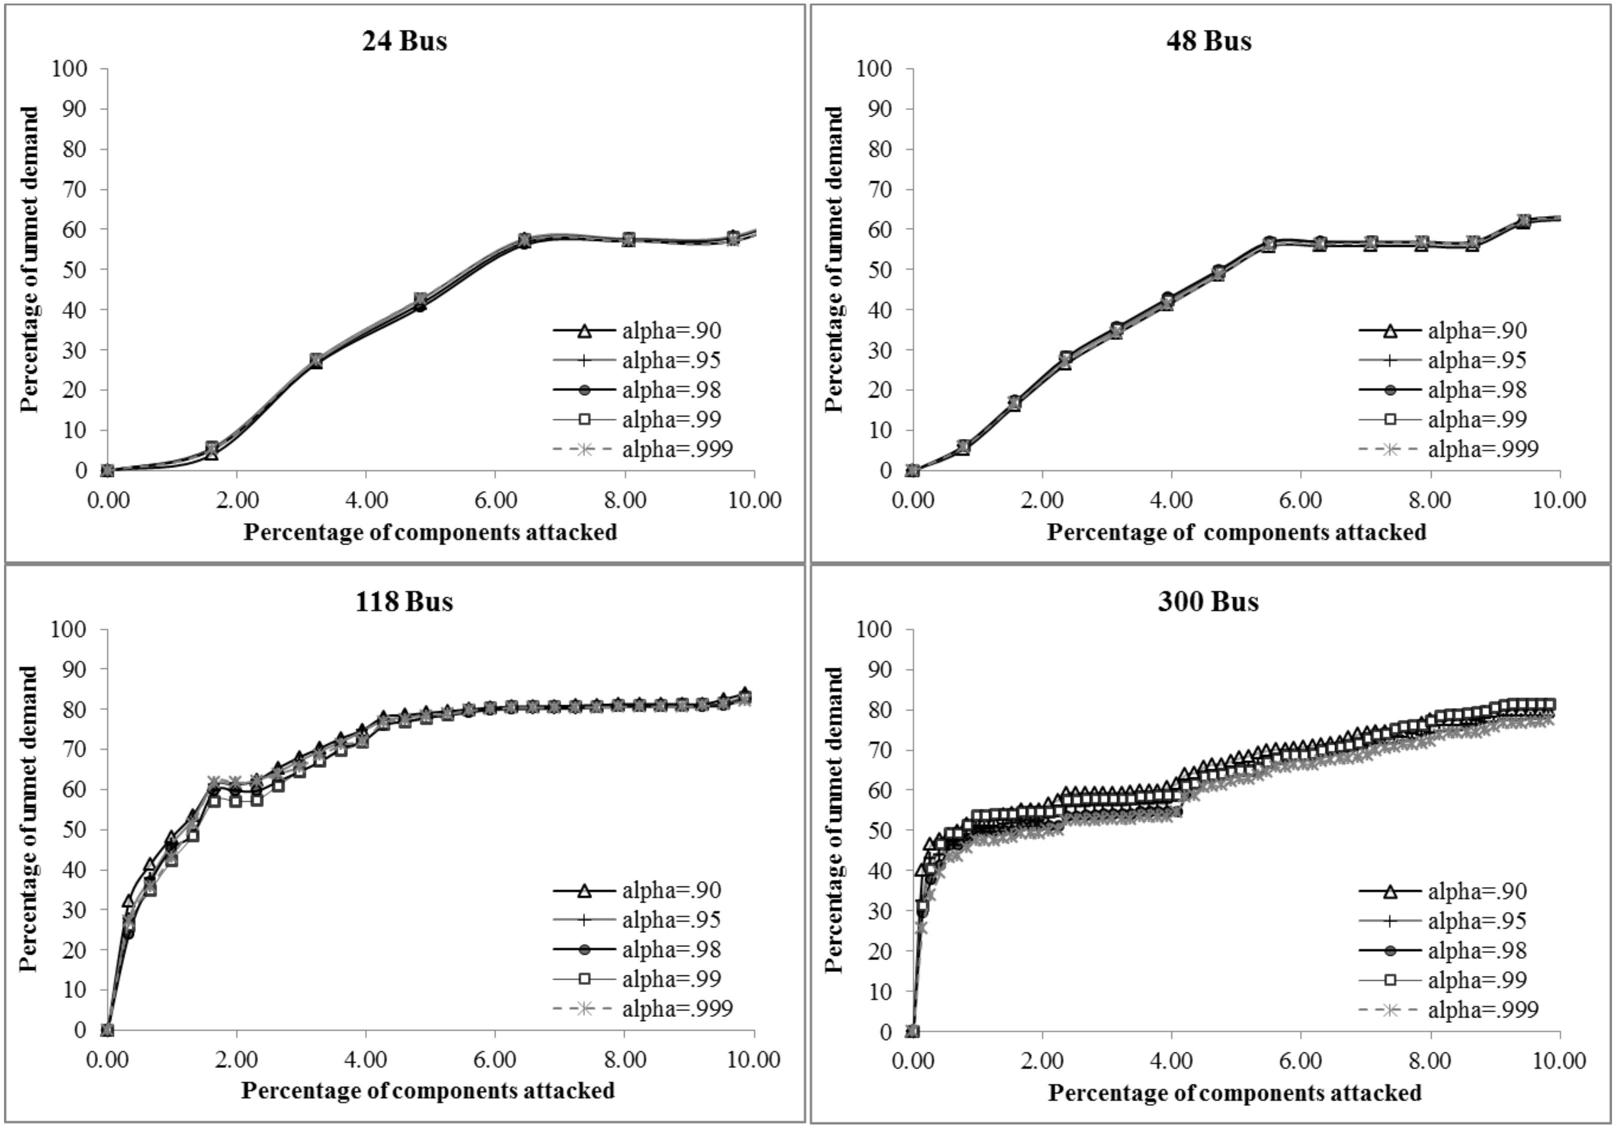

Appendix D. Sensitivity of the cascading model
(to the failure probability of an overloaded component for a static attacker with no cascading knowledge)

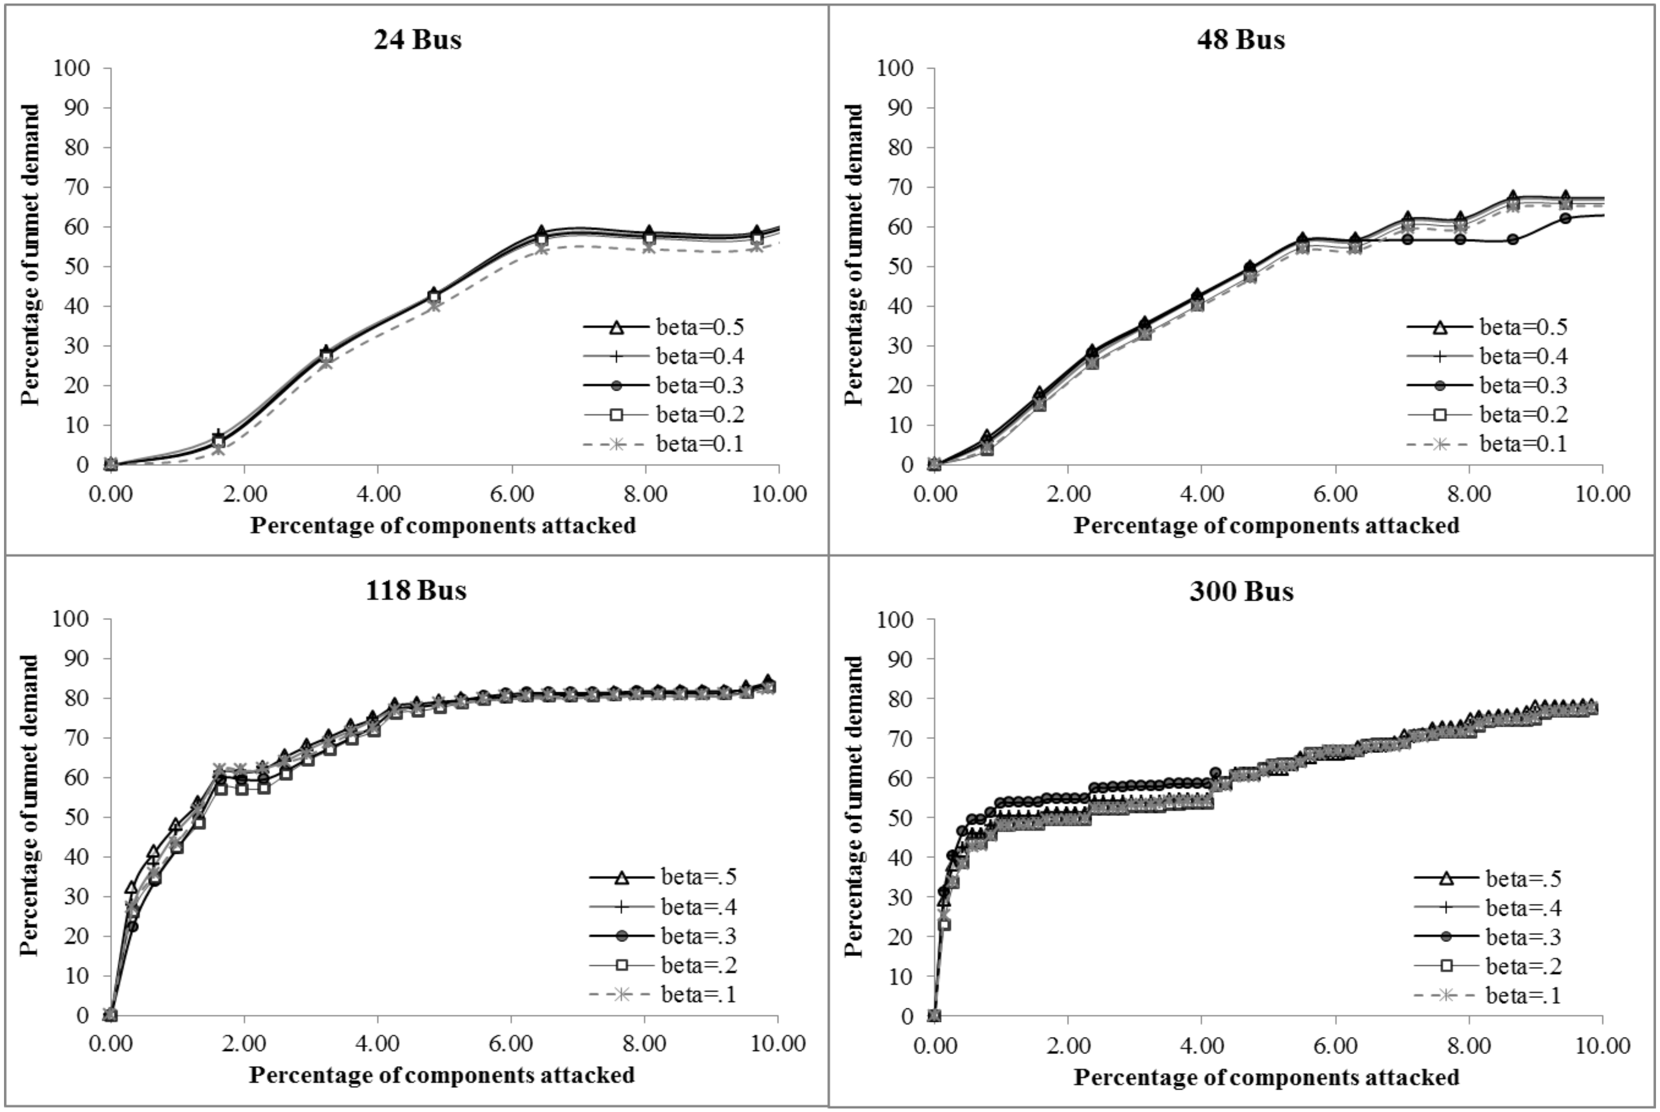

Appendix E. Various hardening strategies
(2% hardening against a static attacker with no cascading knowledge)

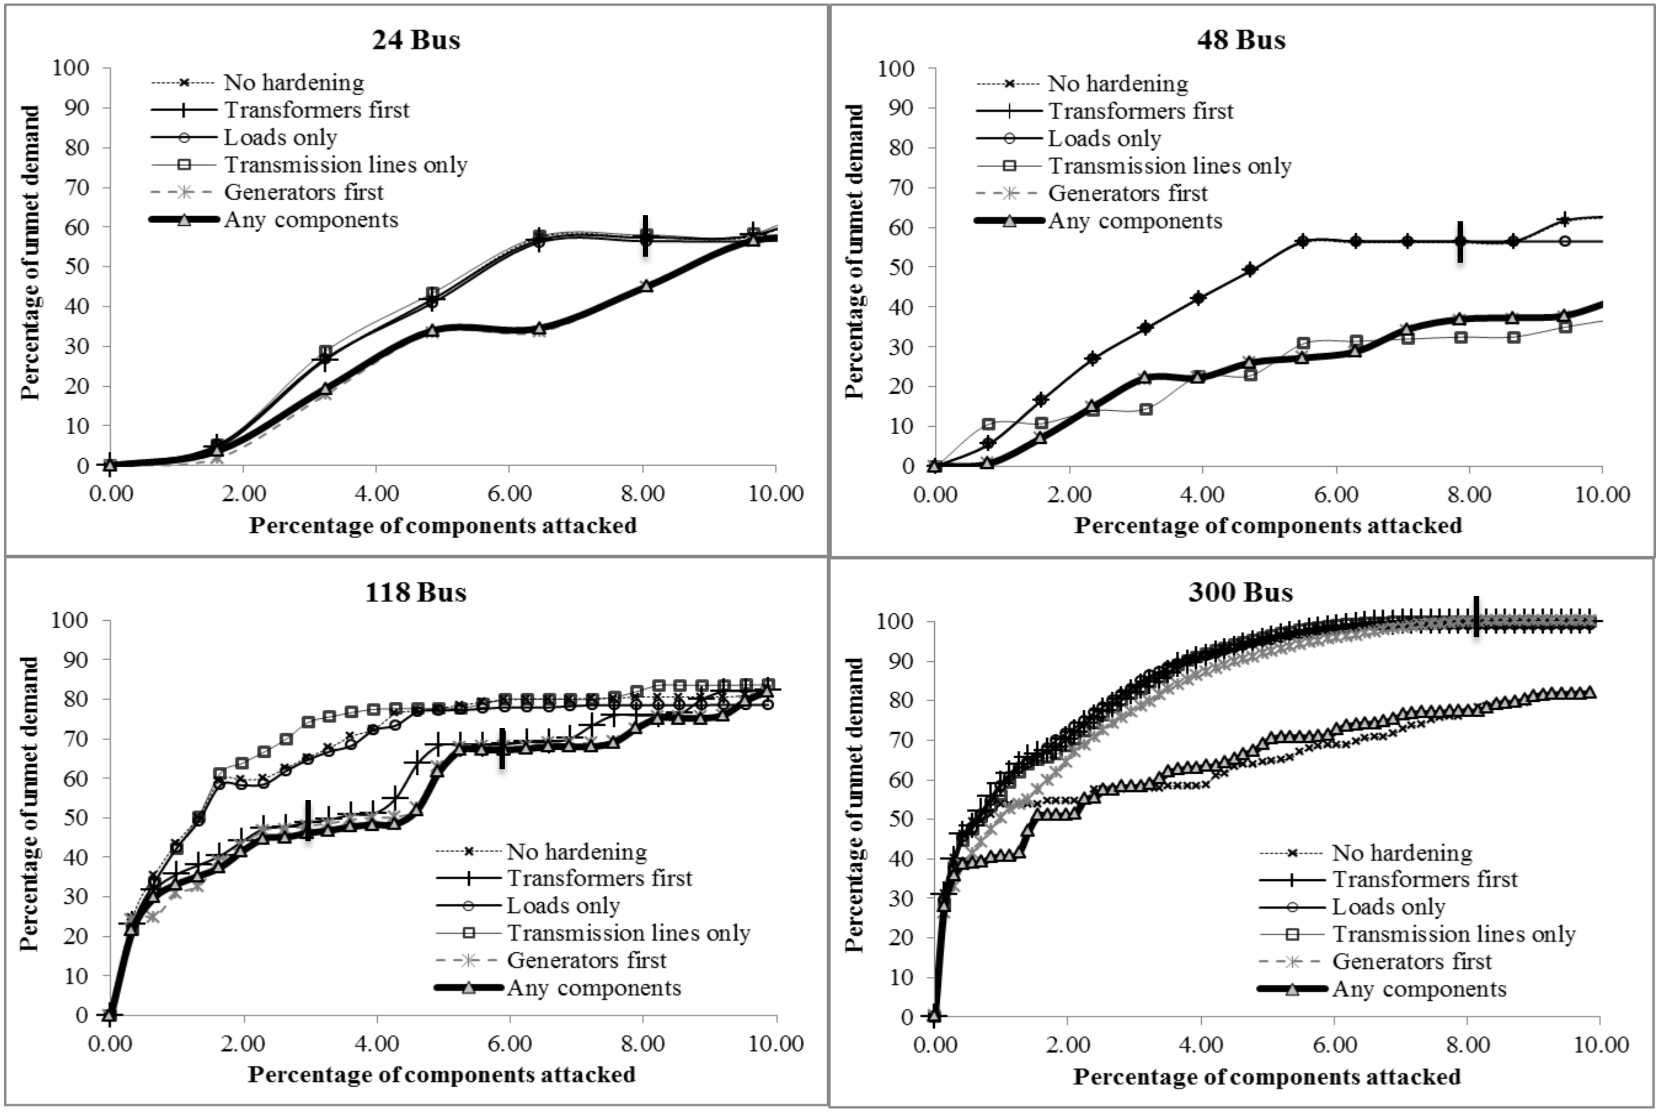

Appendix E. Various hardening strategies
(5% hardening against a static attacker with cascading knowledge)

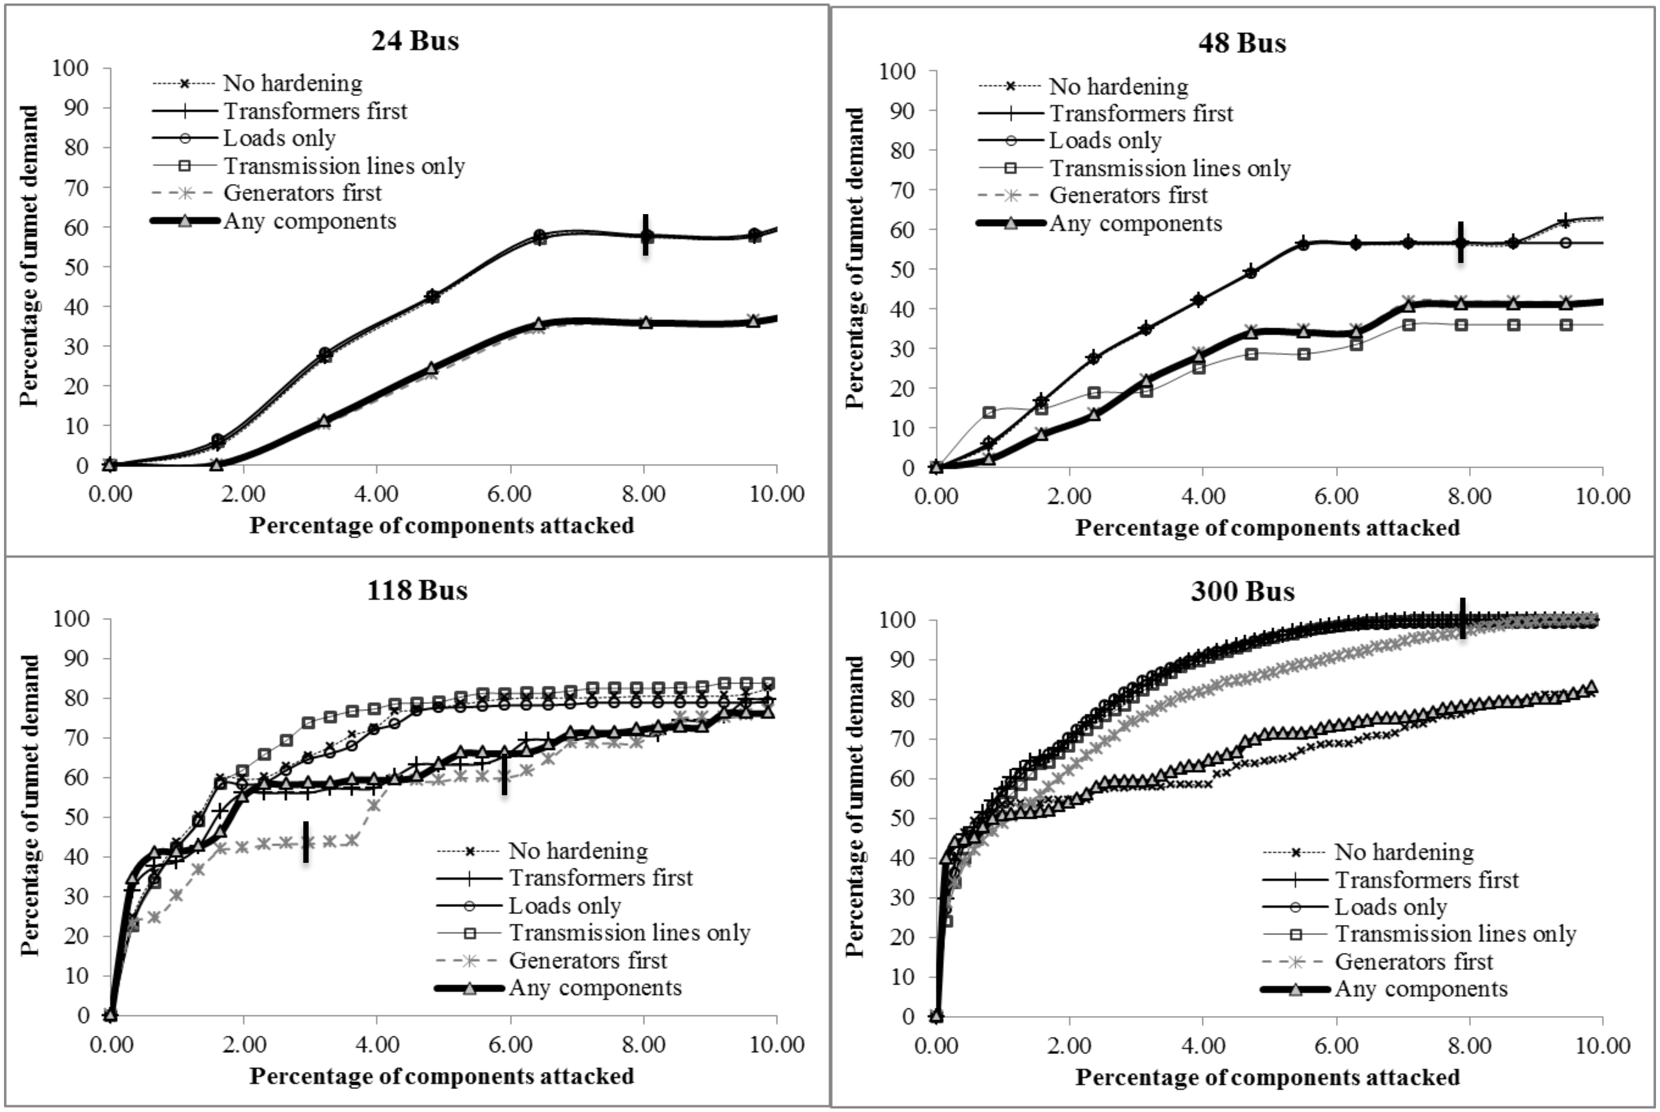

Appendix E. Various hardening strategies
(10% hardening against a static attacker with cascading knowledge)

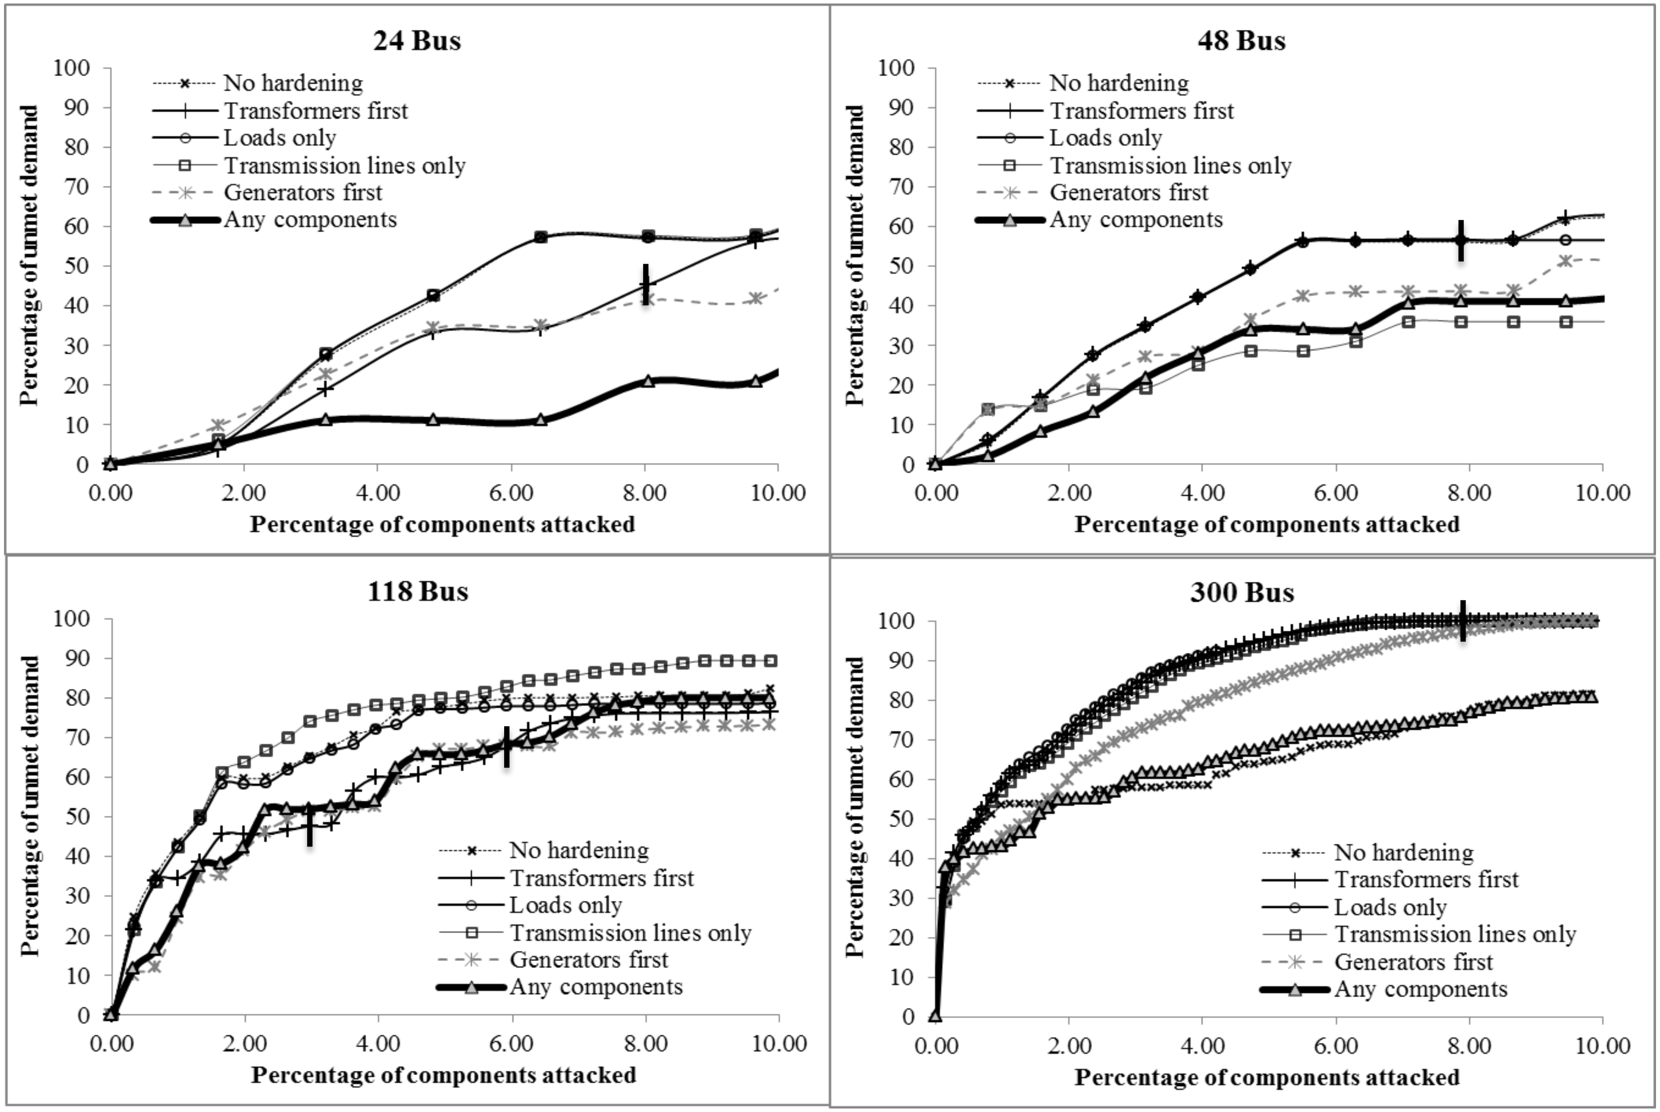

Appendix E. Various hardening strategies
(2% hardening against a static attacker with cascading knowledge)

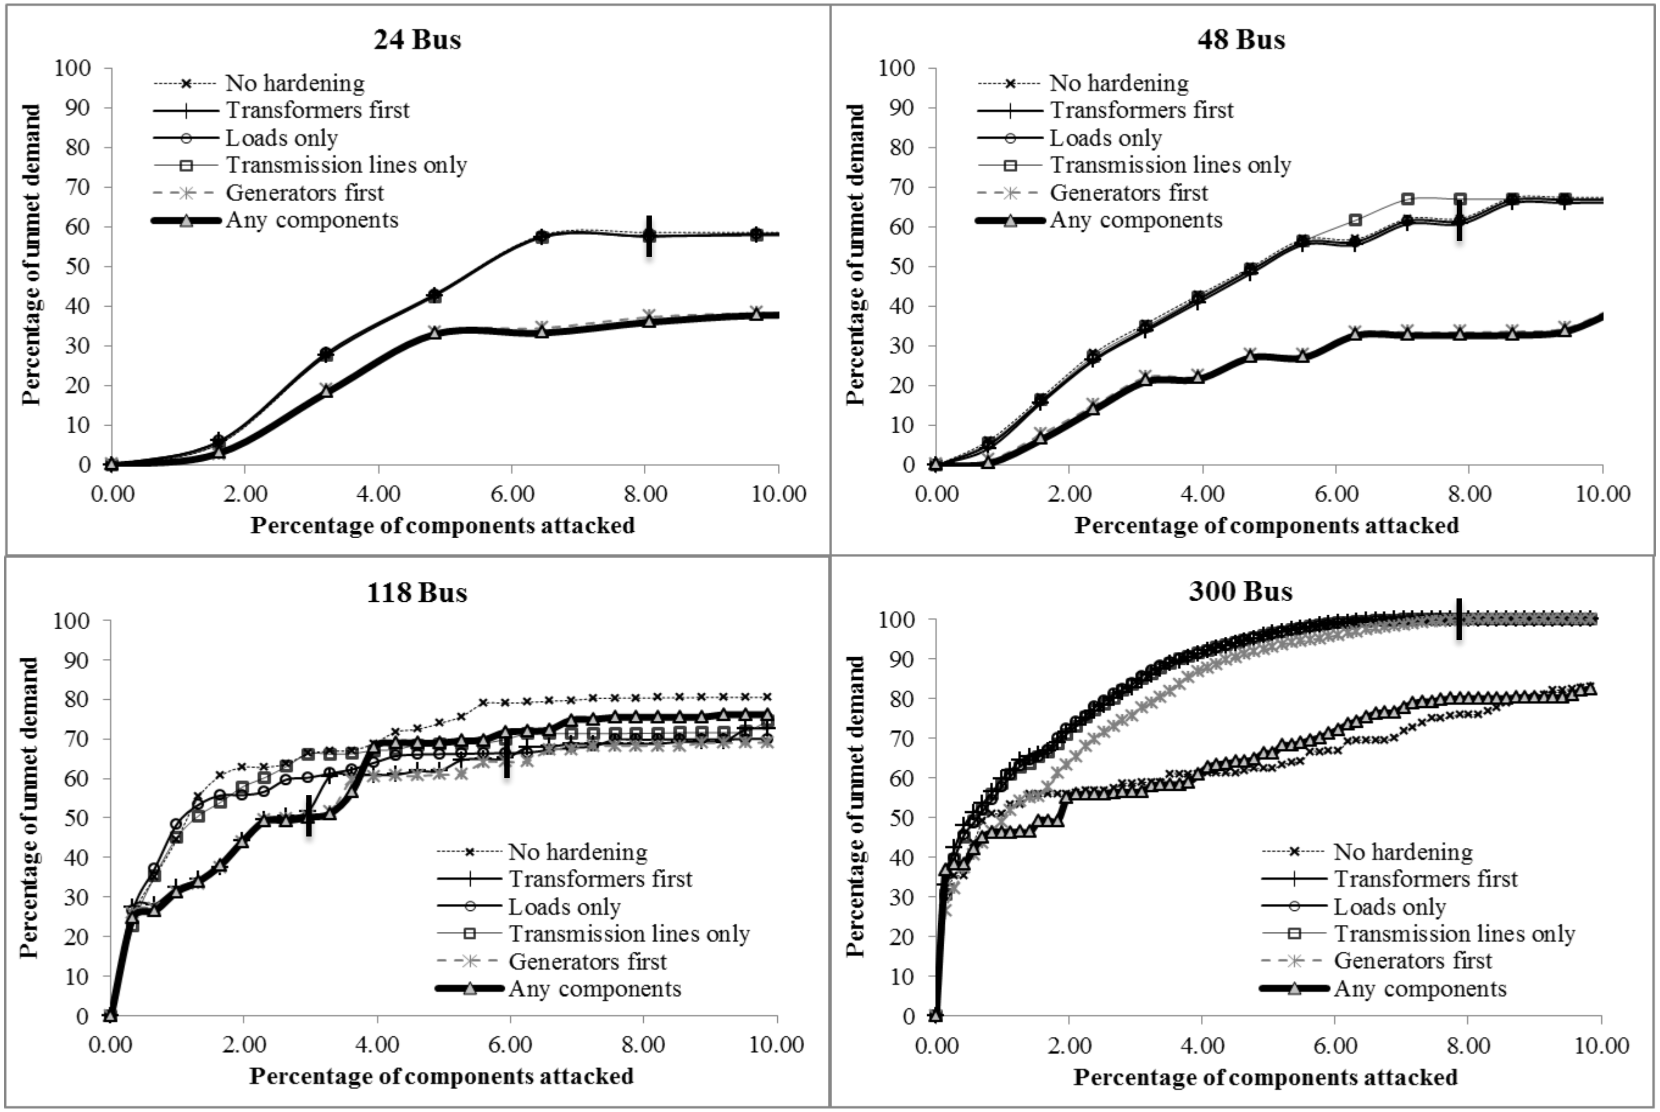

Appendix E. Various hardening strategies
(5% hardening against a static attacker with cascading knowledge)

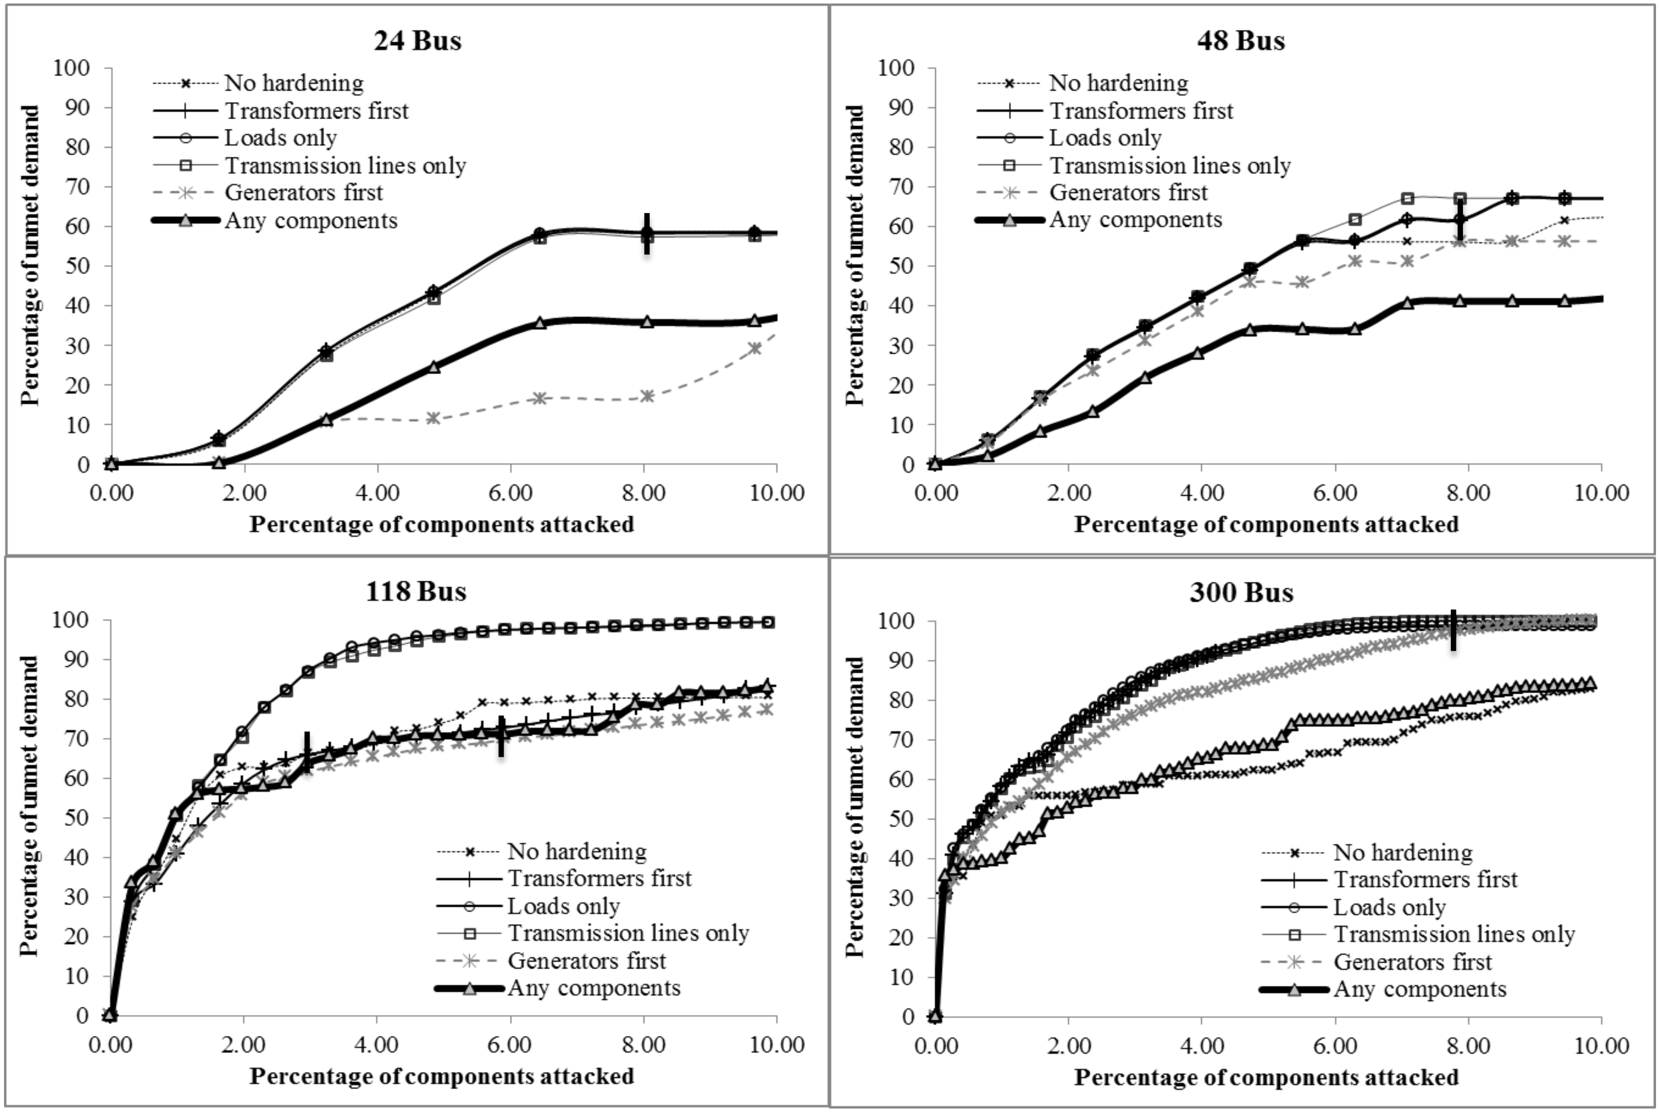

Appendix E. Various hardening strategies
(10% hardening against a static attacker with cascading knowledge)

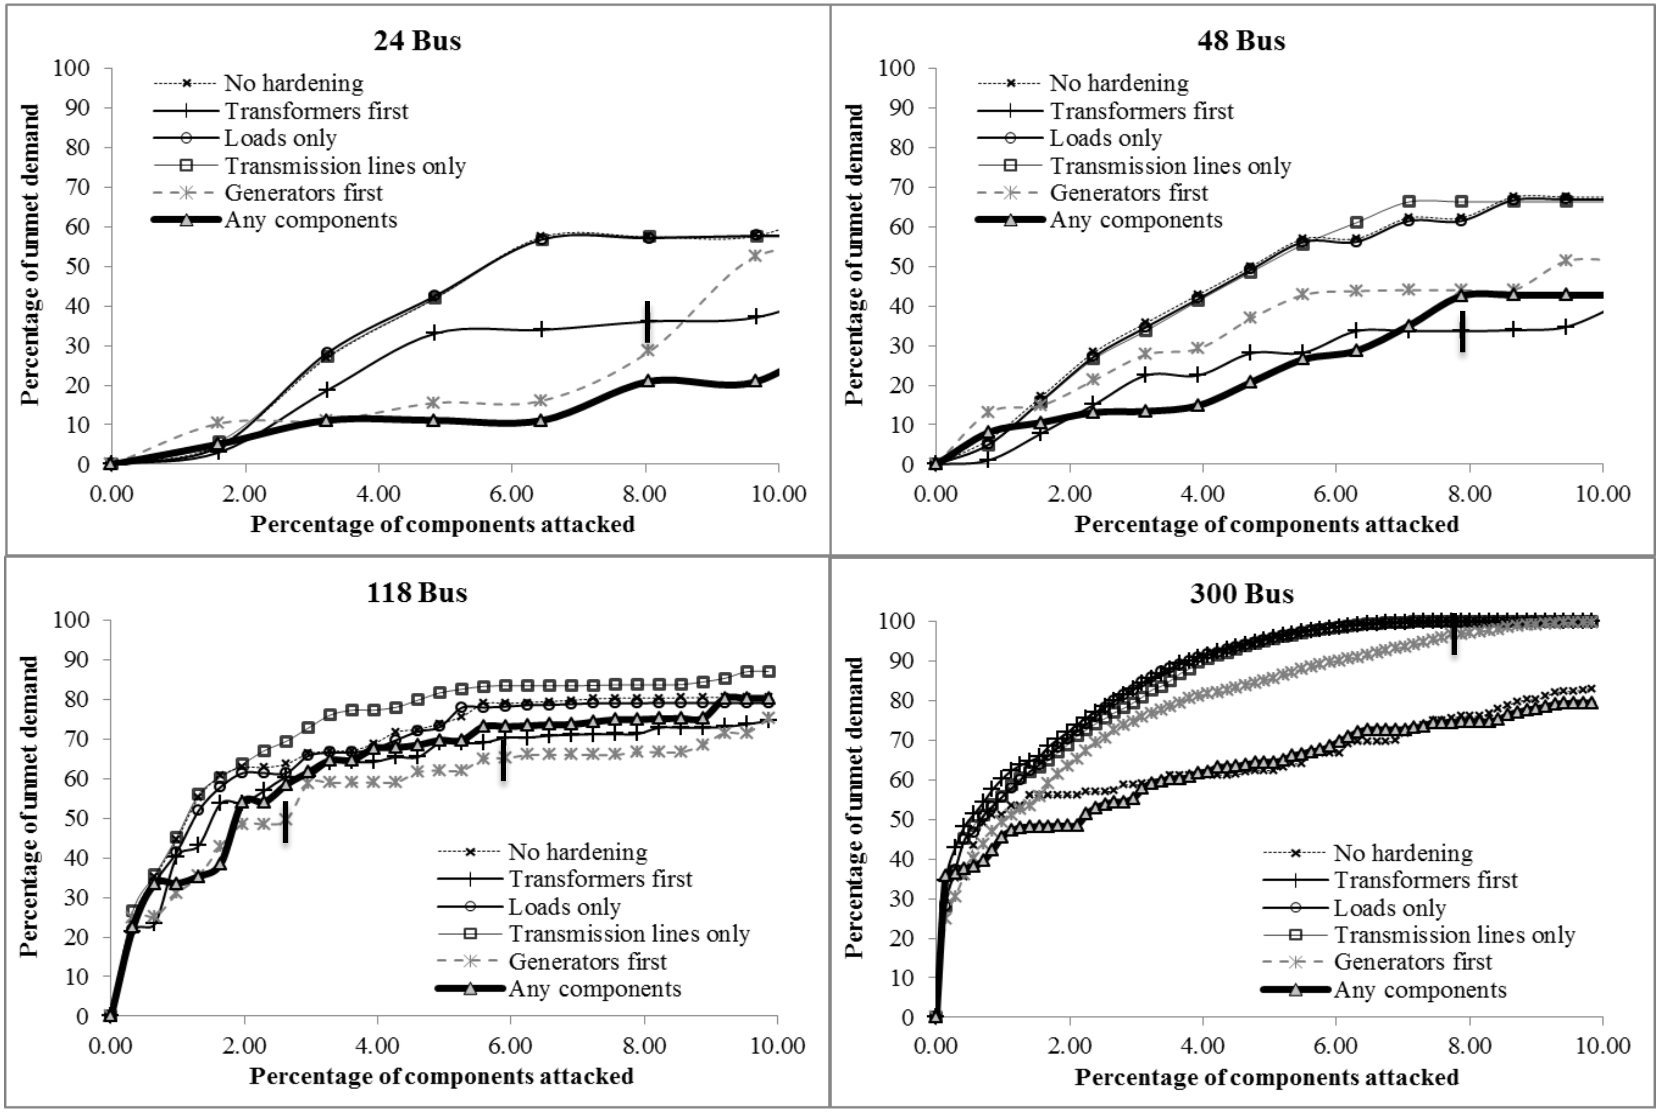

Appendix F. Static hardening versus degree-based hardening (static attacker with no cascading knowledge, 2% hardening scenario)

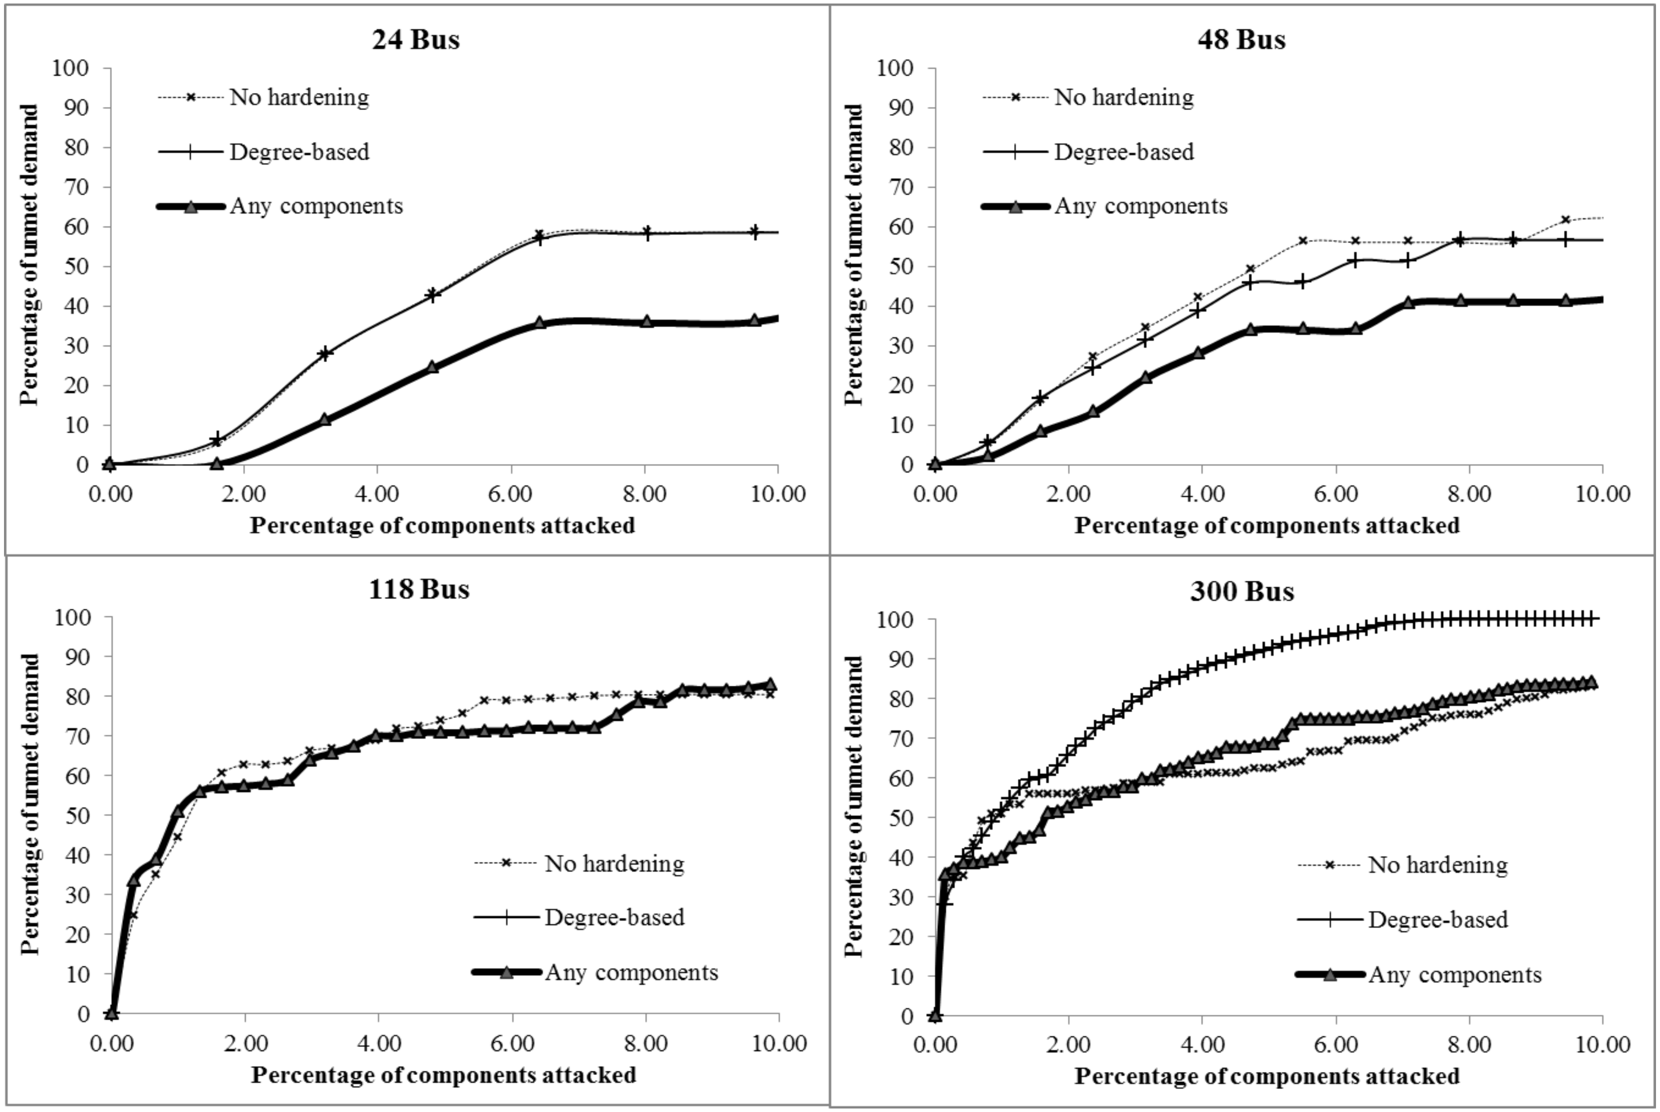

Appendix F. Static hardening versus degree-based hardening
(static attacker with cascading knowledge, 2% hardening scenario)

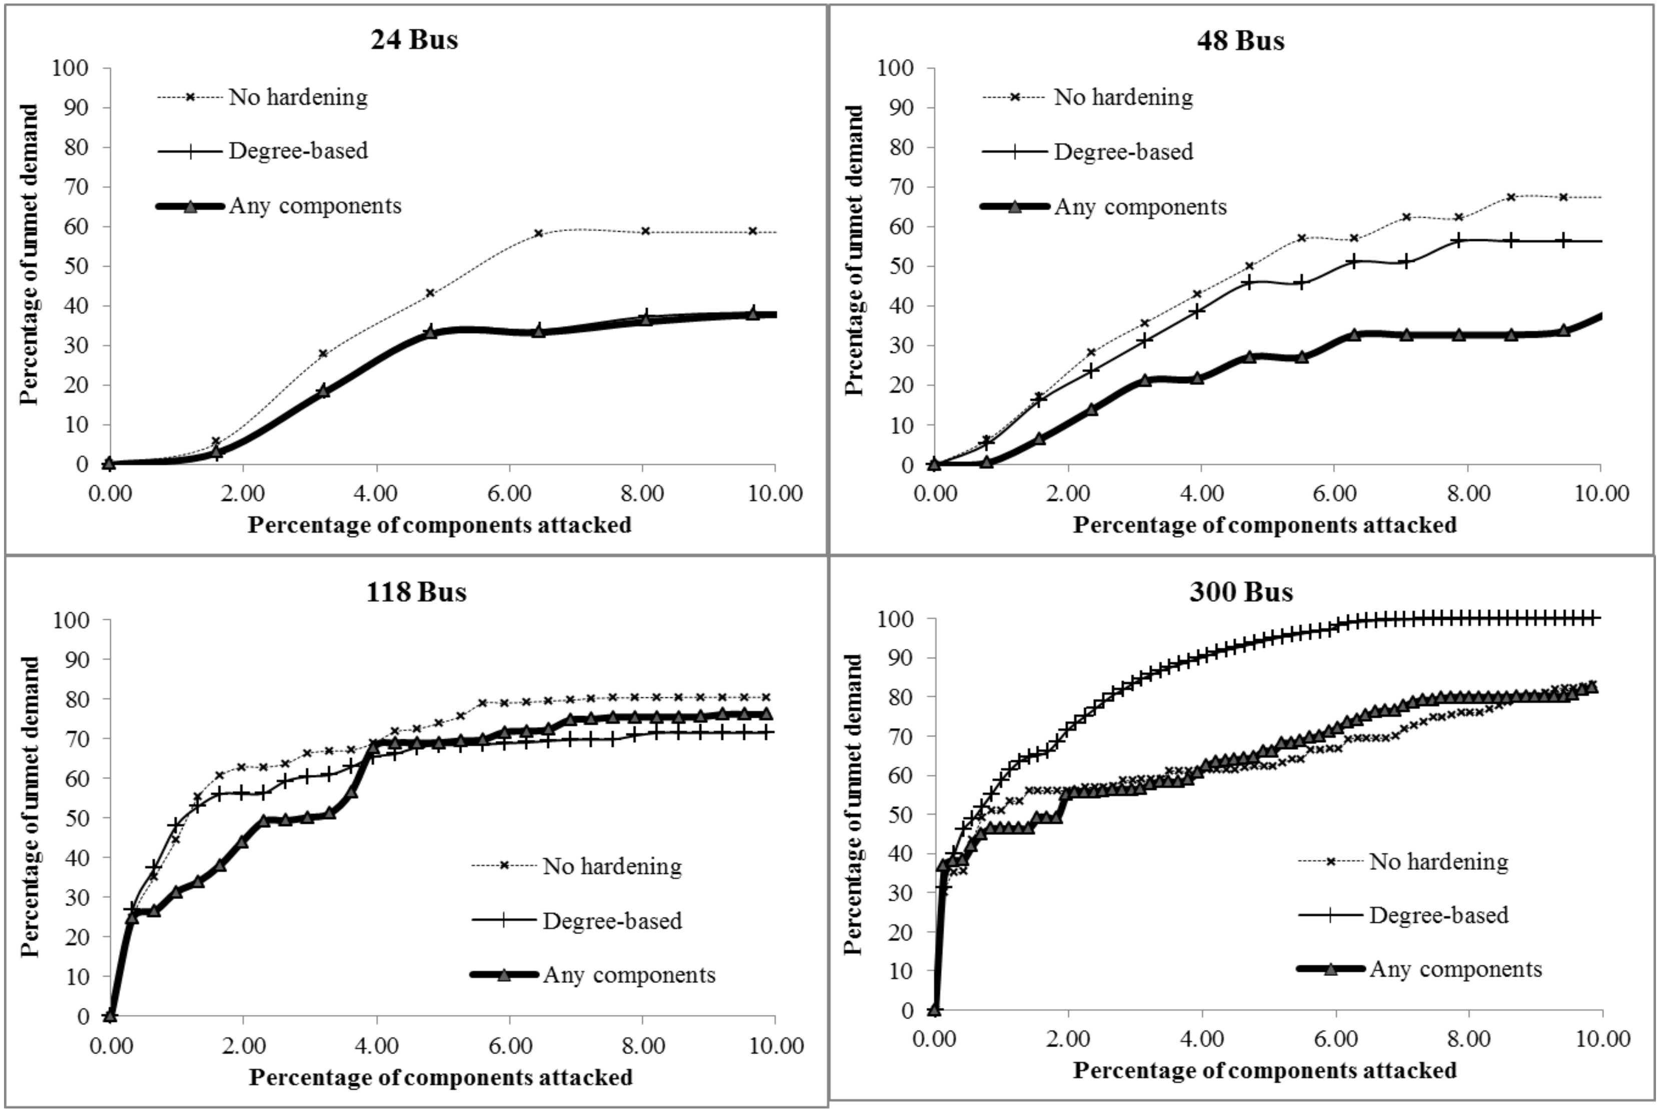

Appendix F. Static hardening versus degree-based hardening
(static attacker with no cascading knowledge, 5% hardening scenario)

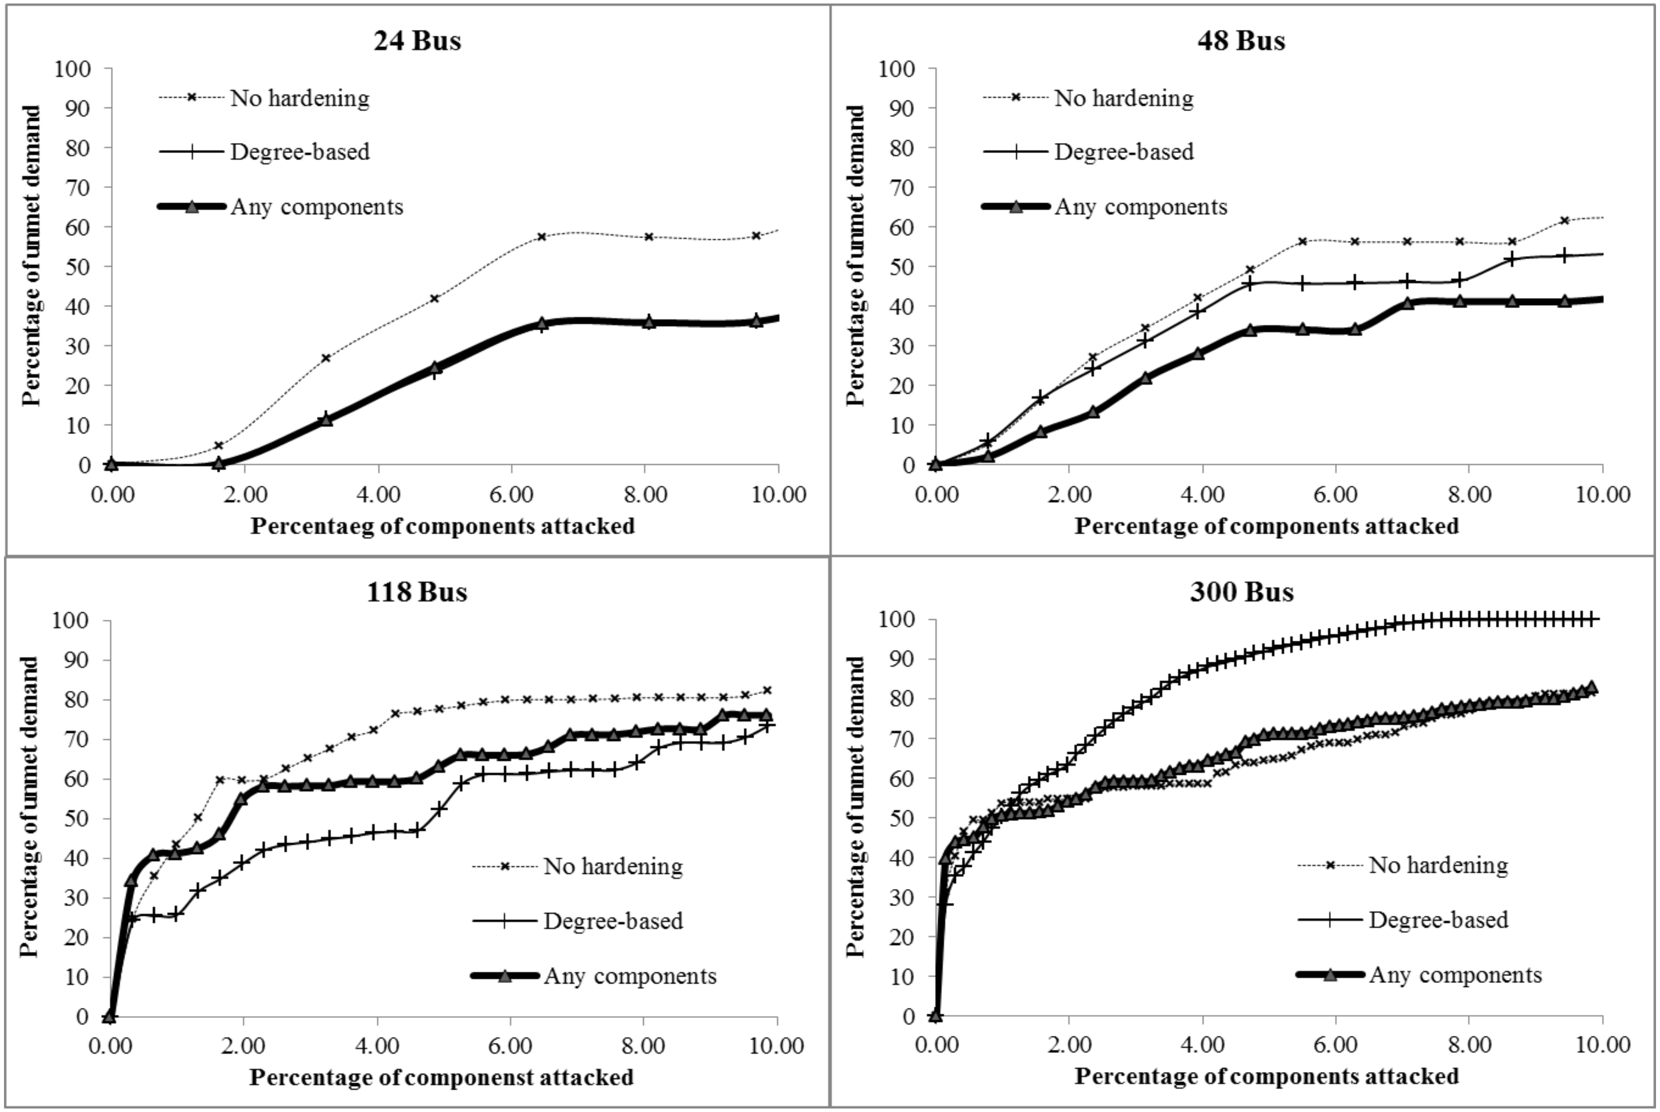

Appendix F. Static hardening versus degree-based hardening (static attacker with cascading knowledge, 5% hardening scenario)

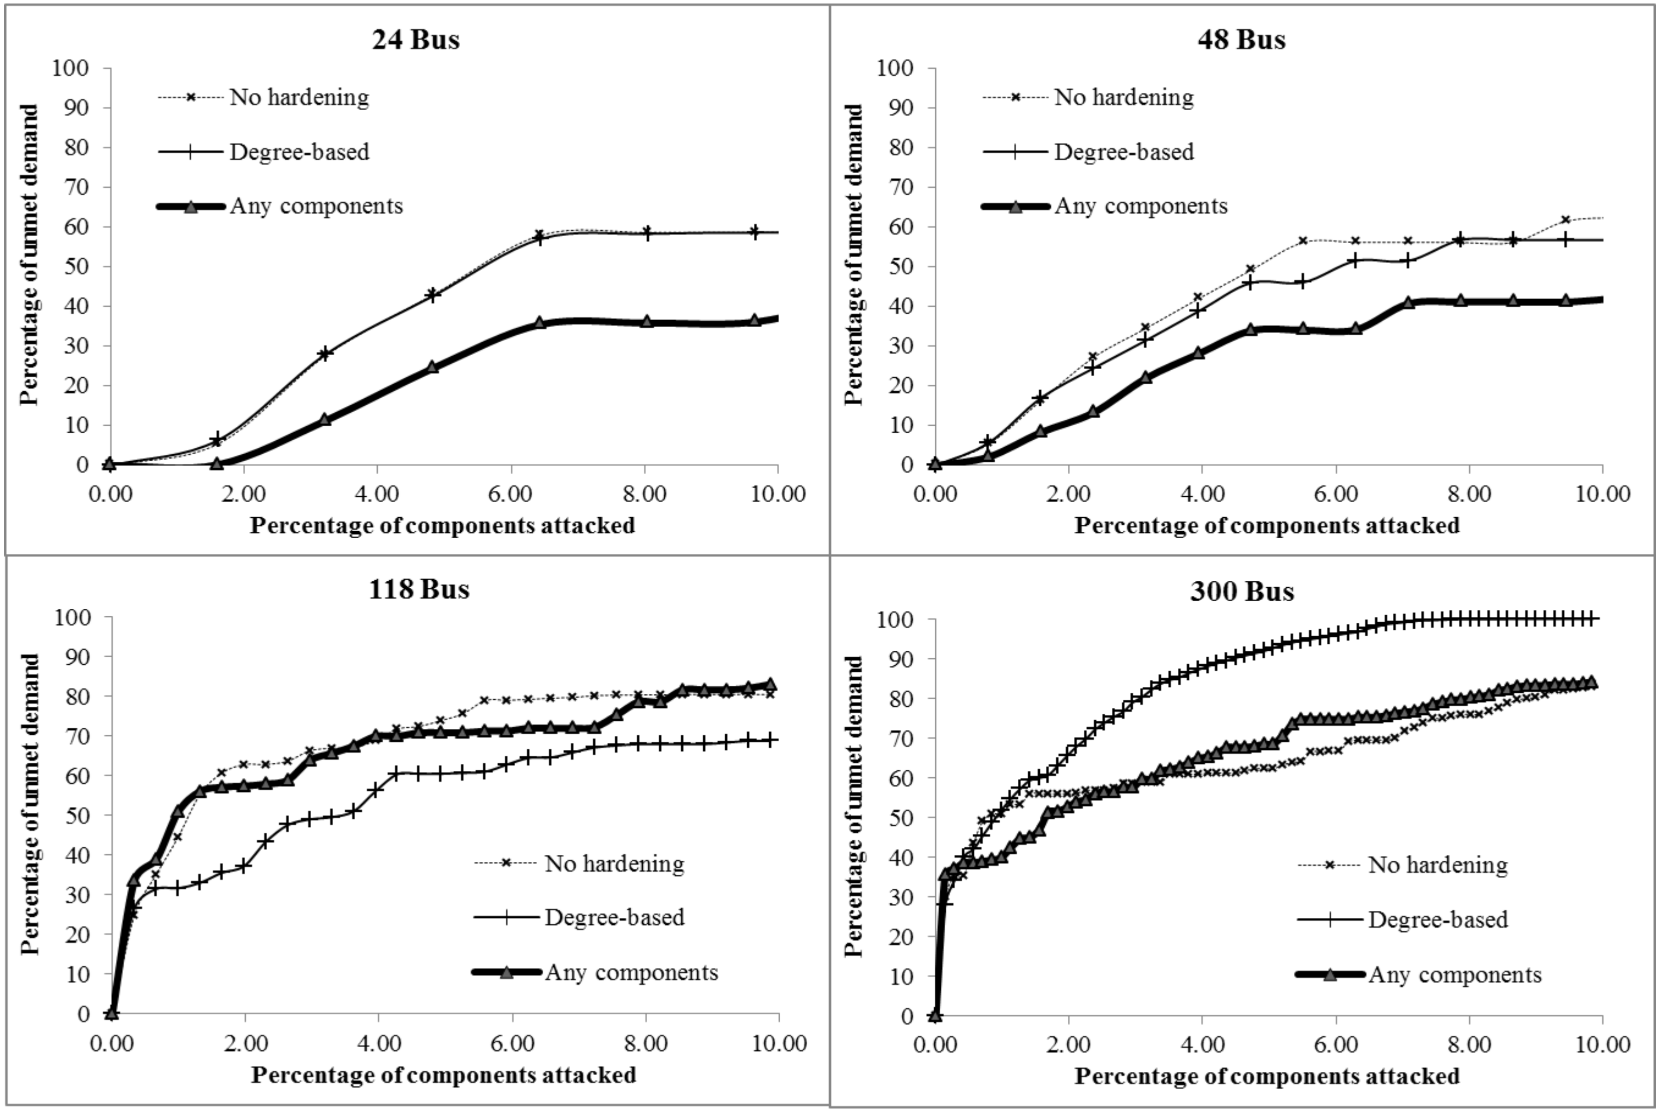

Appendix F. Static hardening versus degree-based hardening (static attacker with no cascading knowledge, 10% hardening scenario)

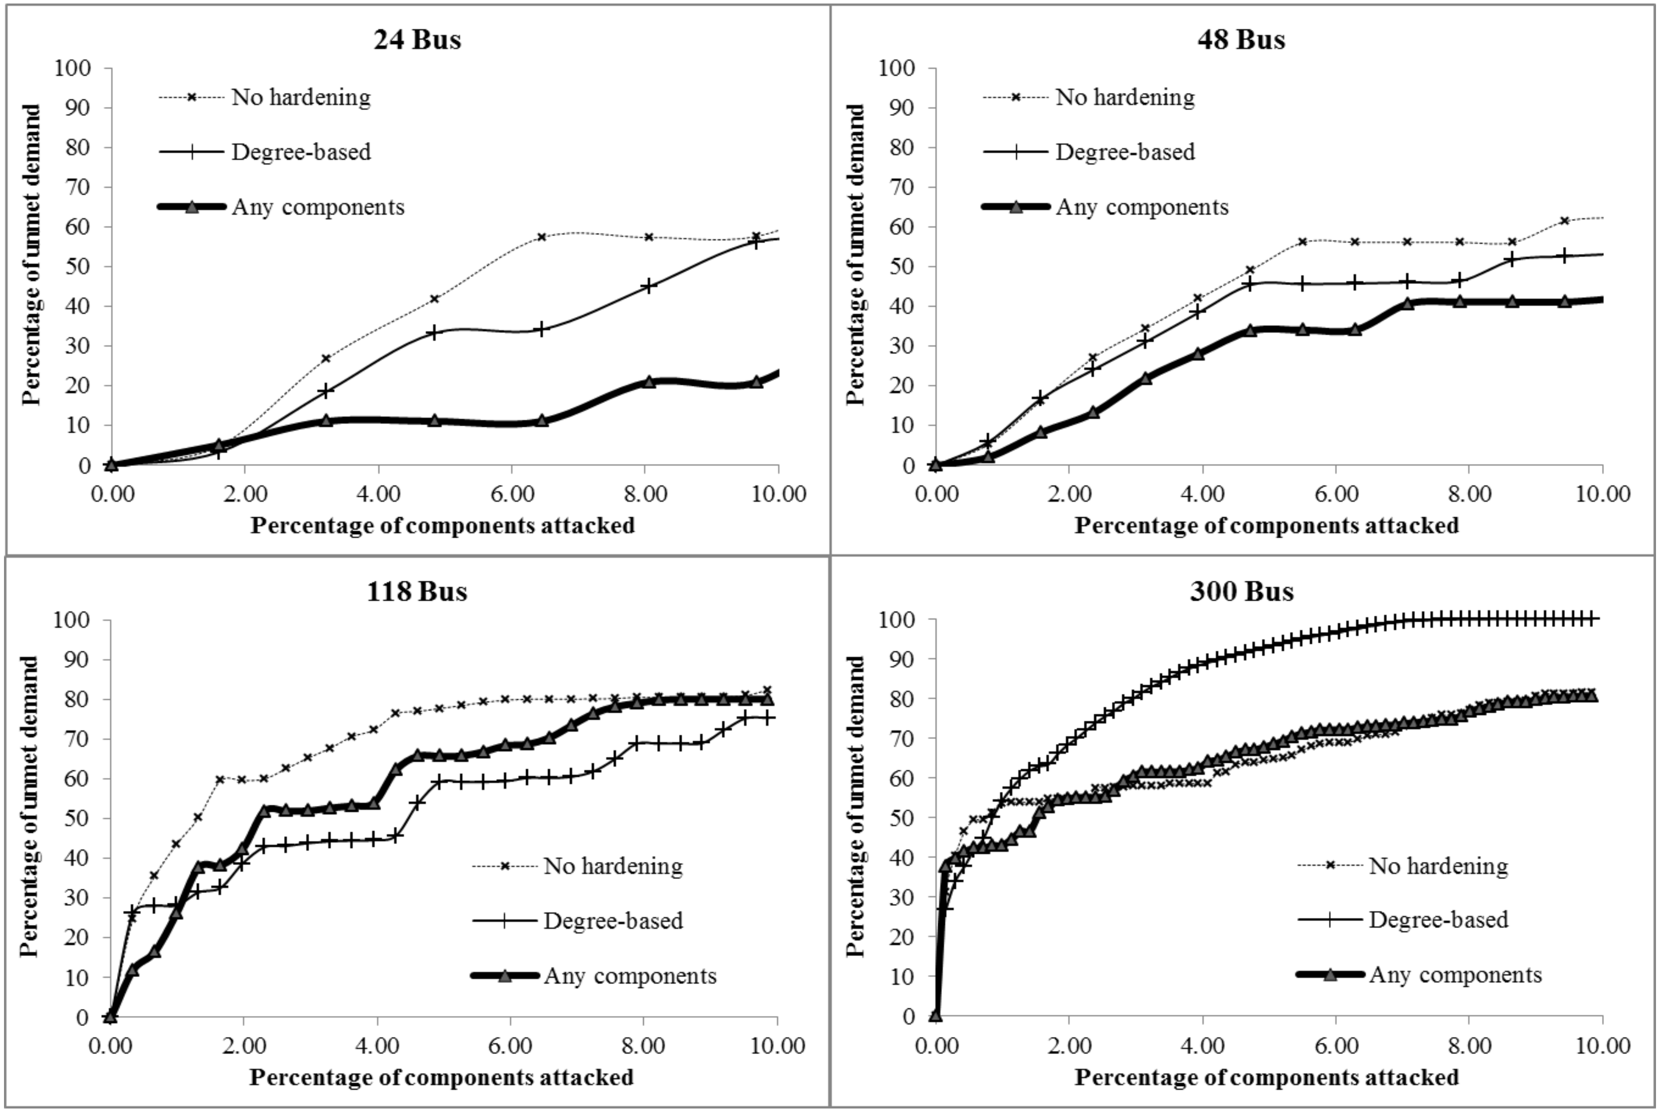

Appendix F. Static hardening versus degree-based hardening (static attacker with cascading knowledge, 10% hardening scenario)

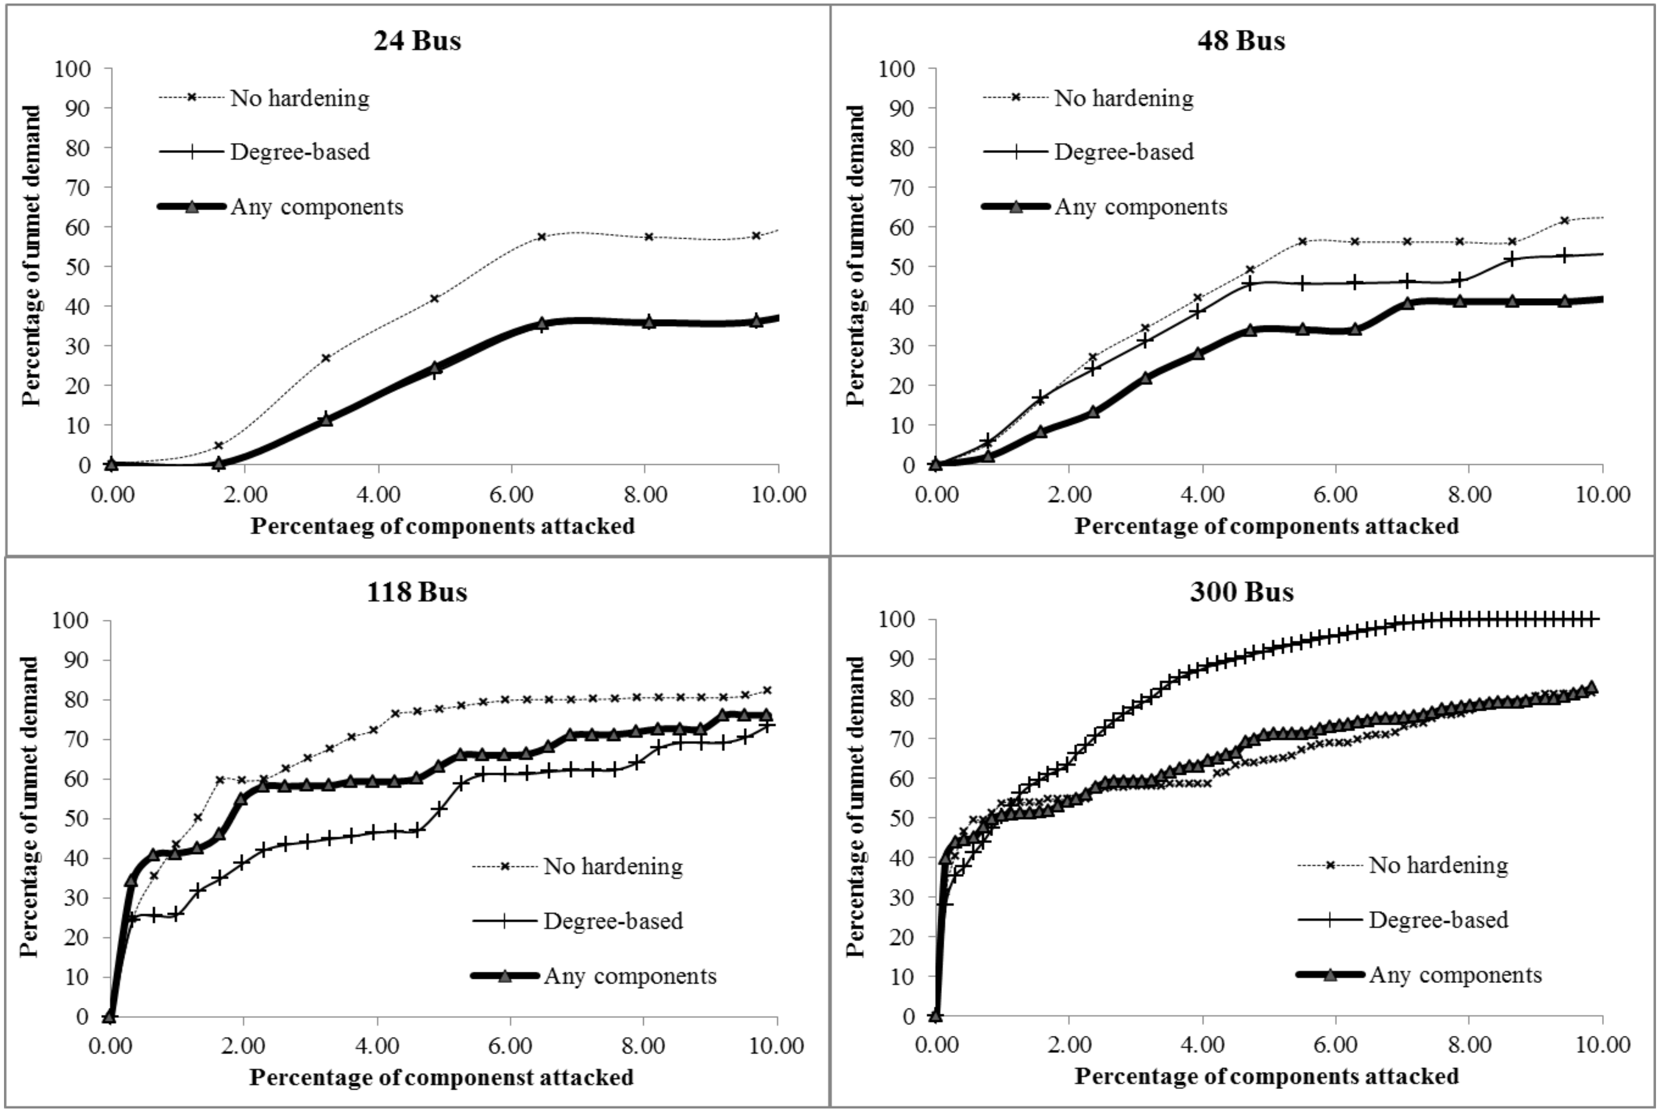

Appendix G. Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a small attack (static attacker with no cascading knowledge)

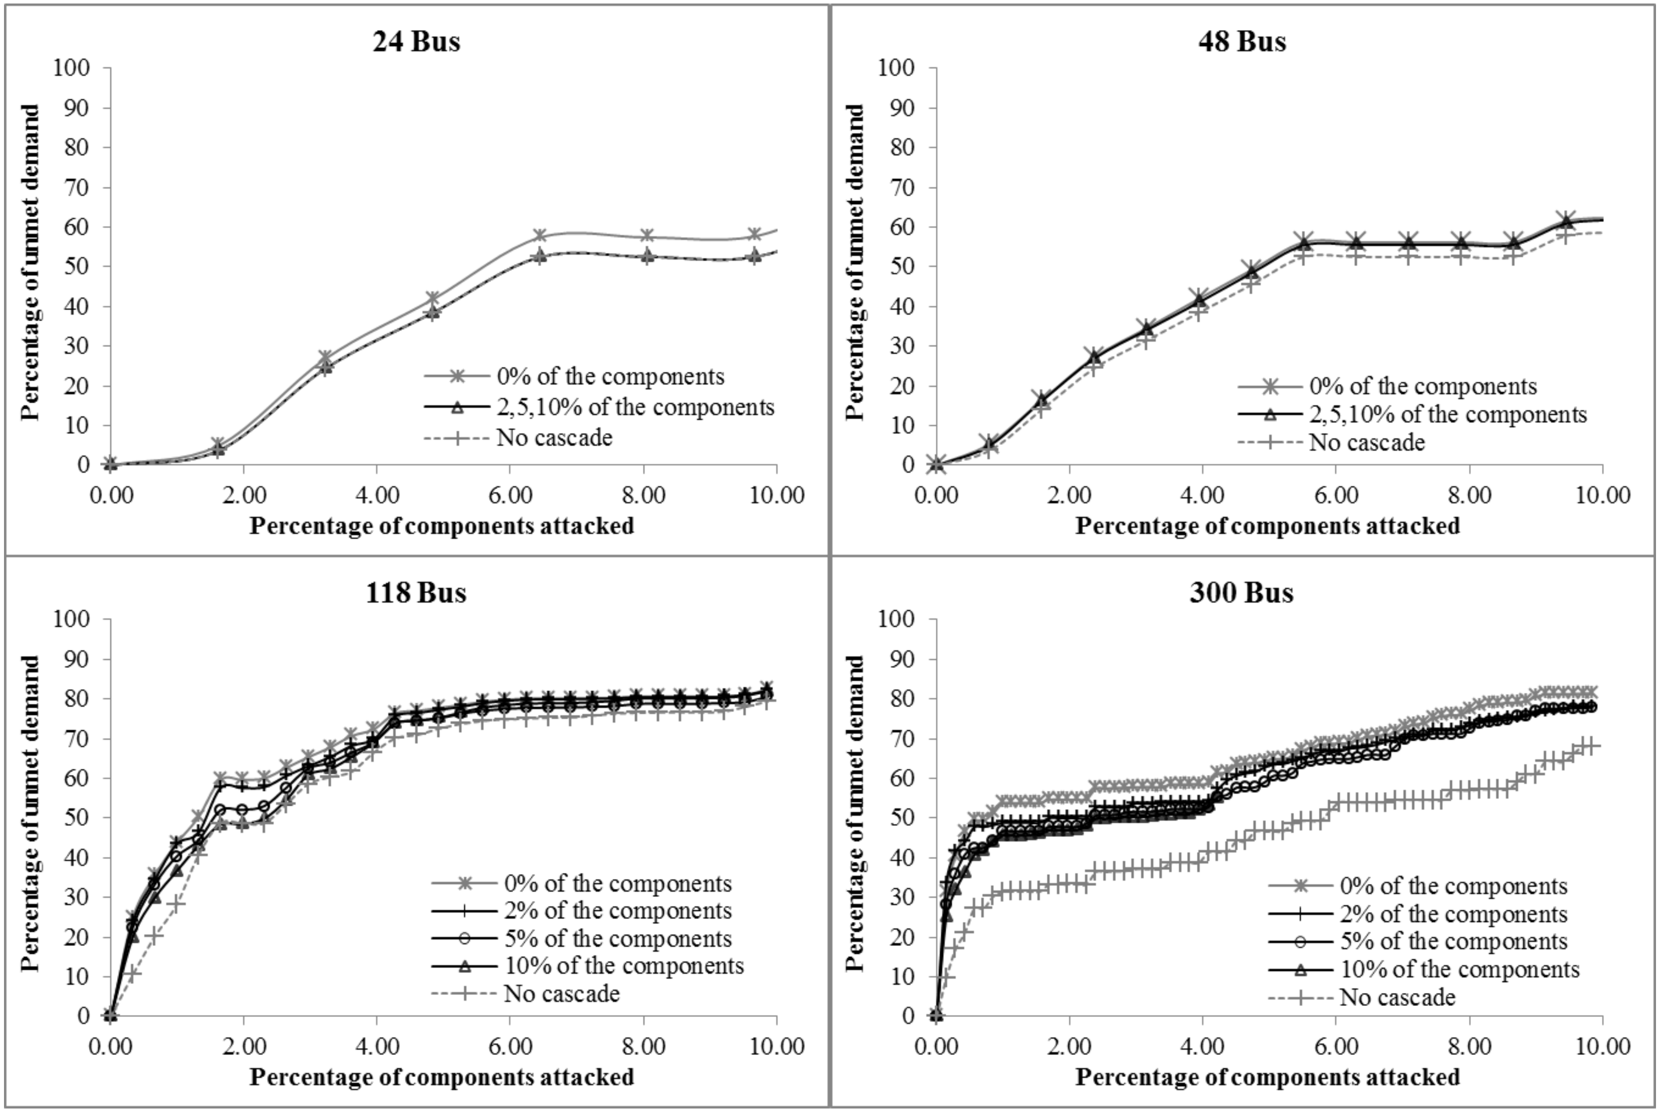

Appendix G. Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a large attack (static attacker with no cascading knowledge)

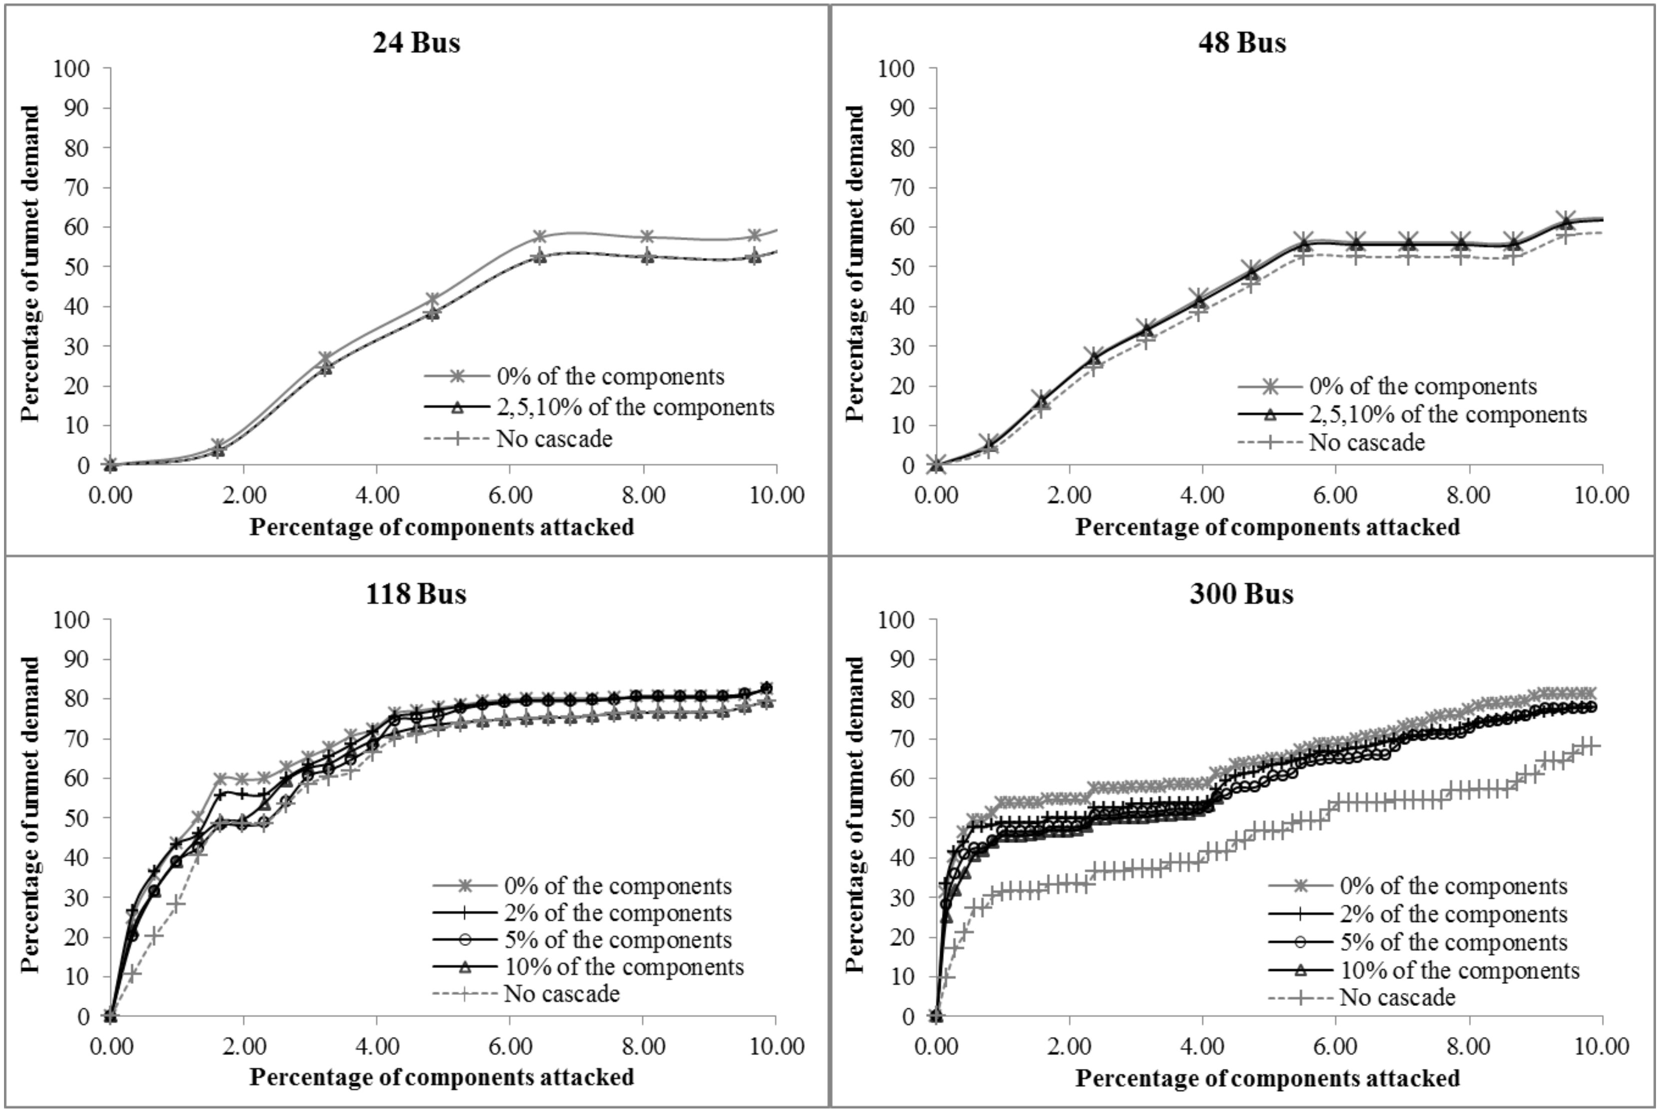

Appendix G. Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a small attack (static attacker with cascading knowledge).

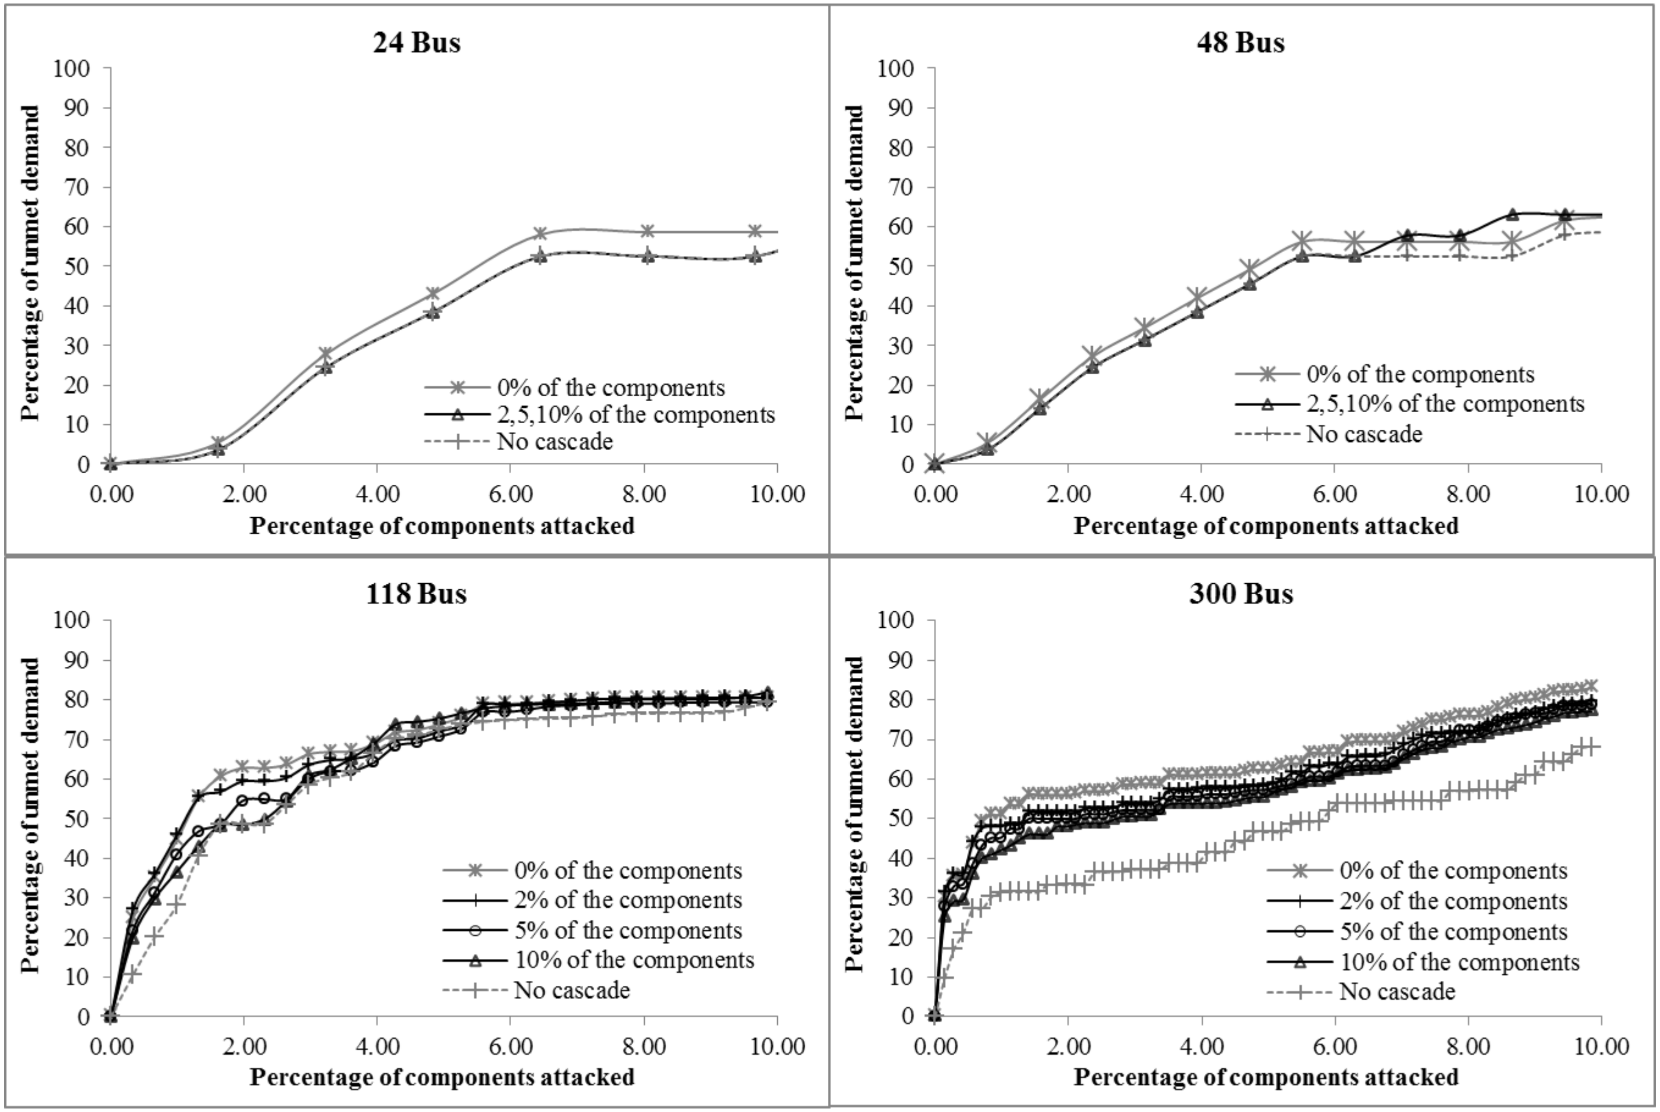

Appendix G. Various levels of investments in doubling capacities of the transmission lines that are mostly likely to cascade in a large attack (static attacker with cascading knowledge)

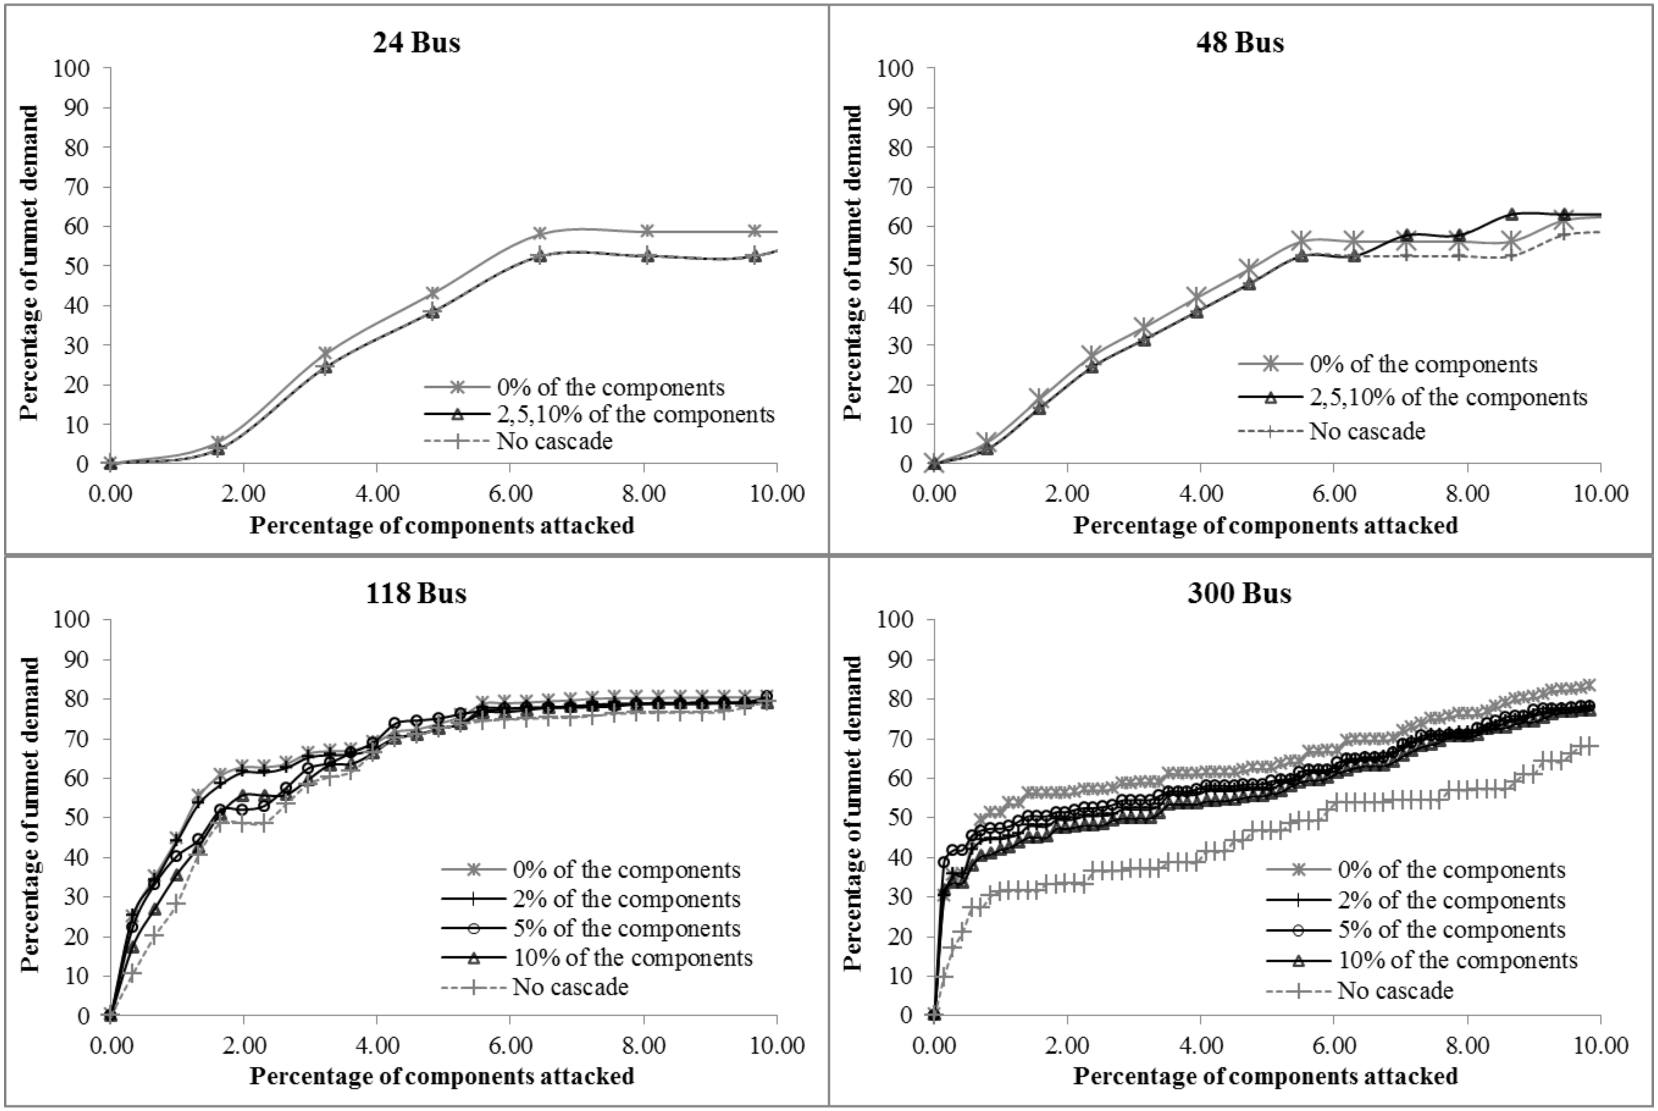

Appendix H. Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a small attack (static attacker with no cascading knowledge)

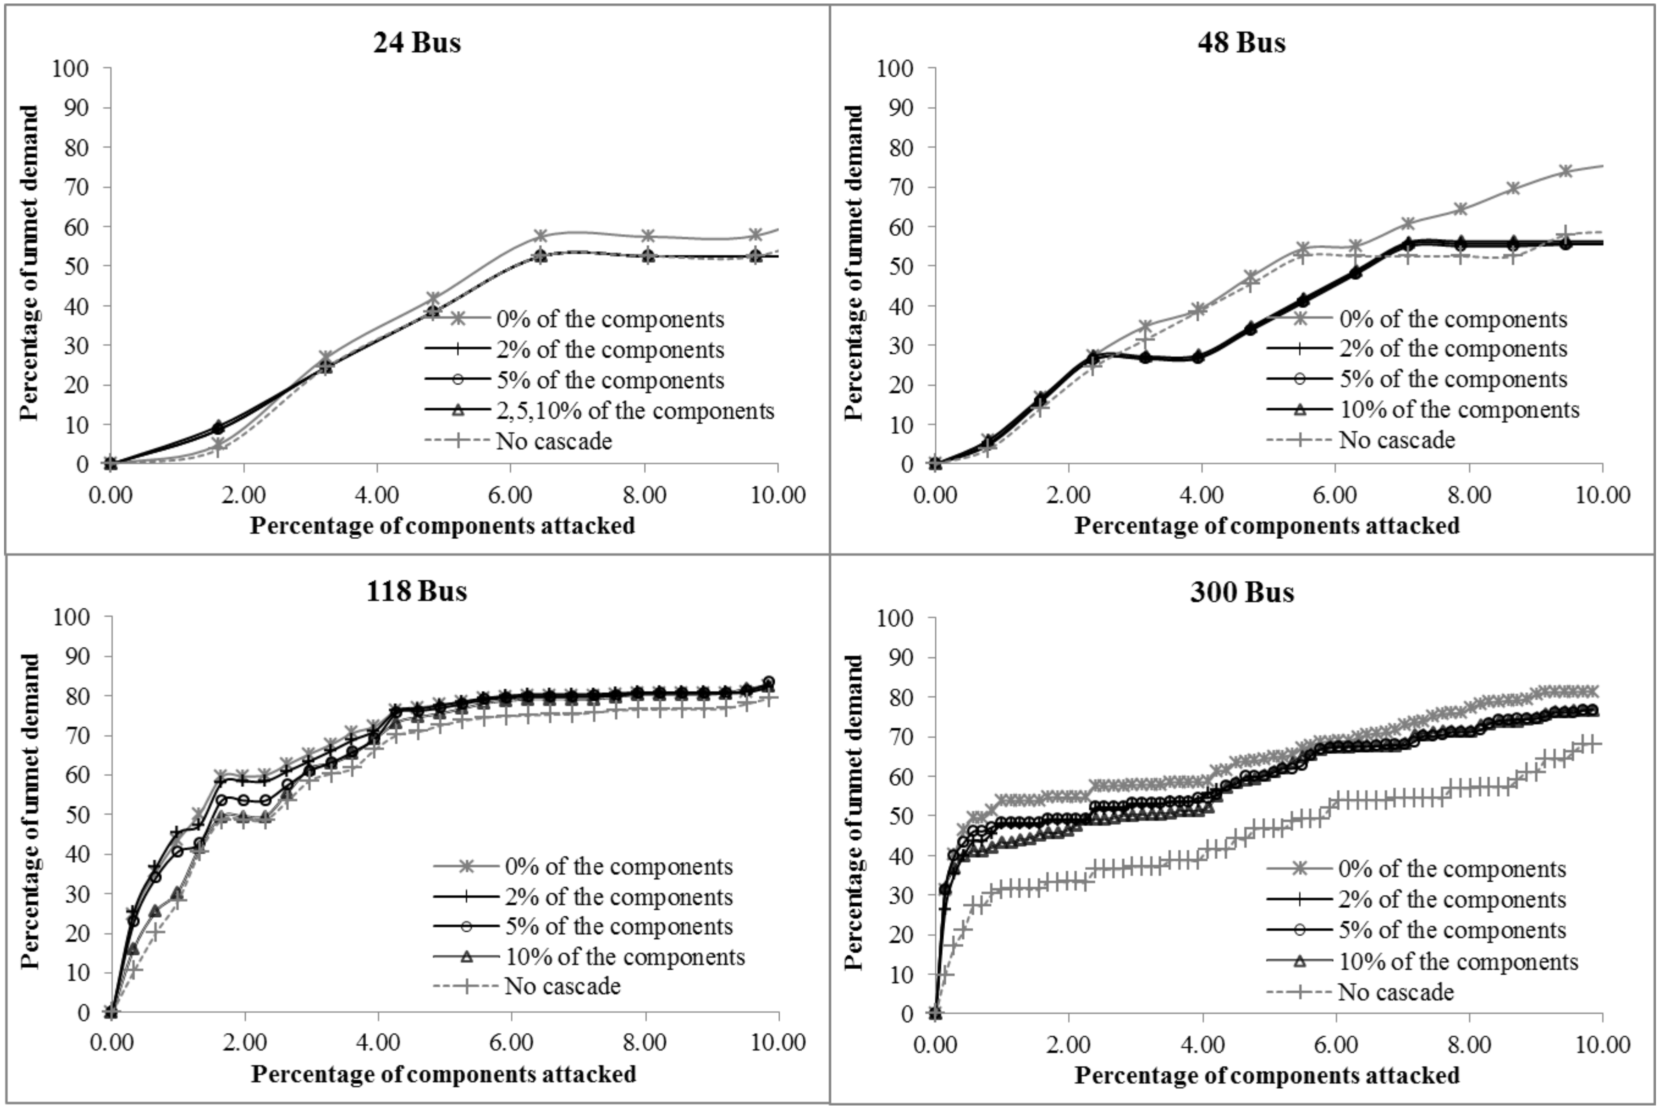

Appendix H. Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a large attack (static attacker with no cascading knowledge)

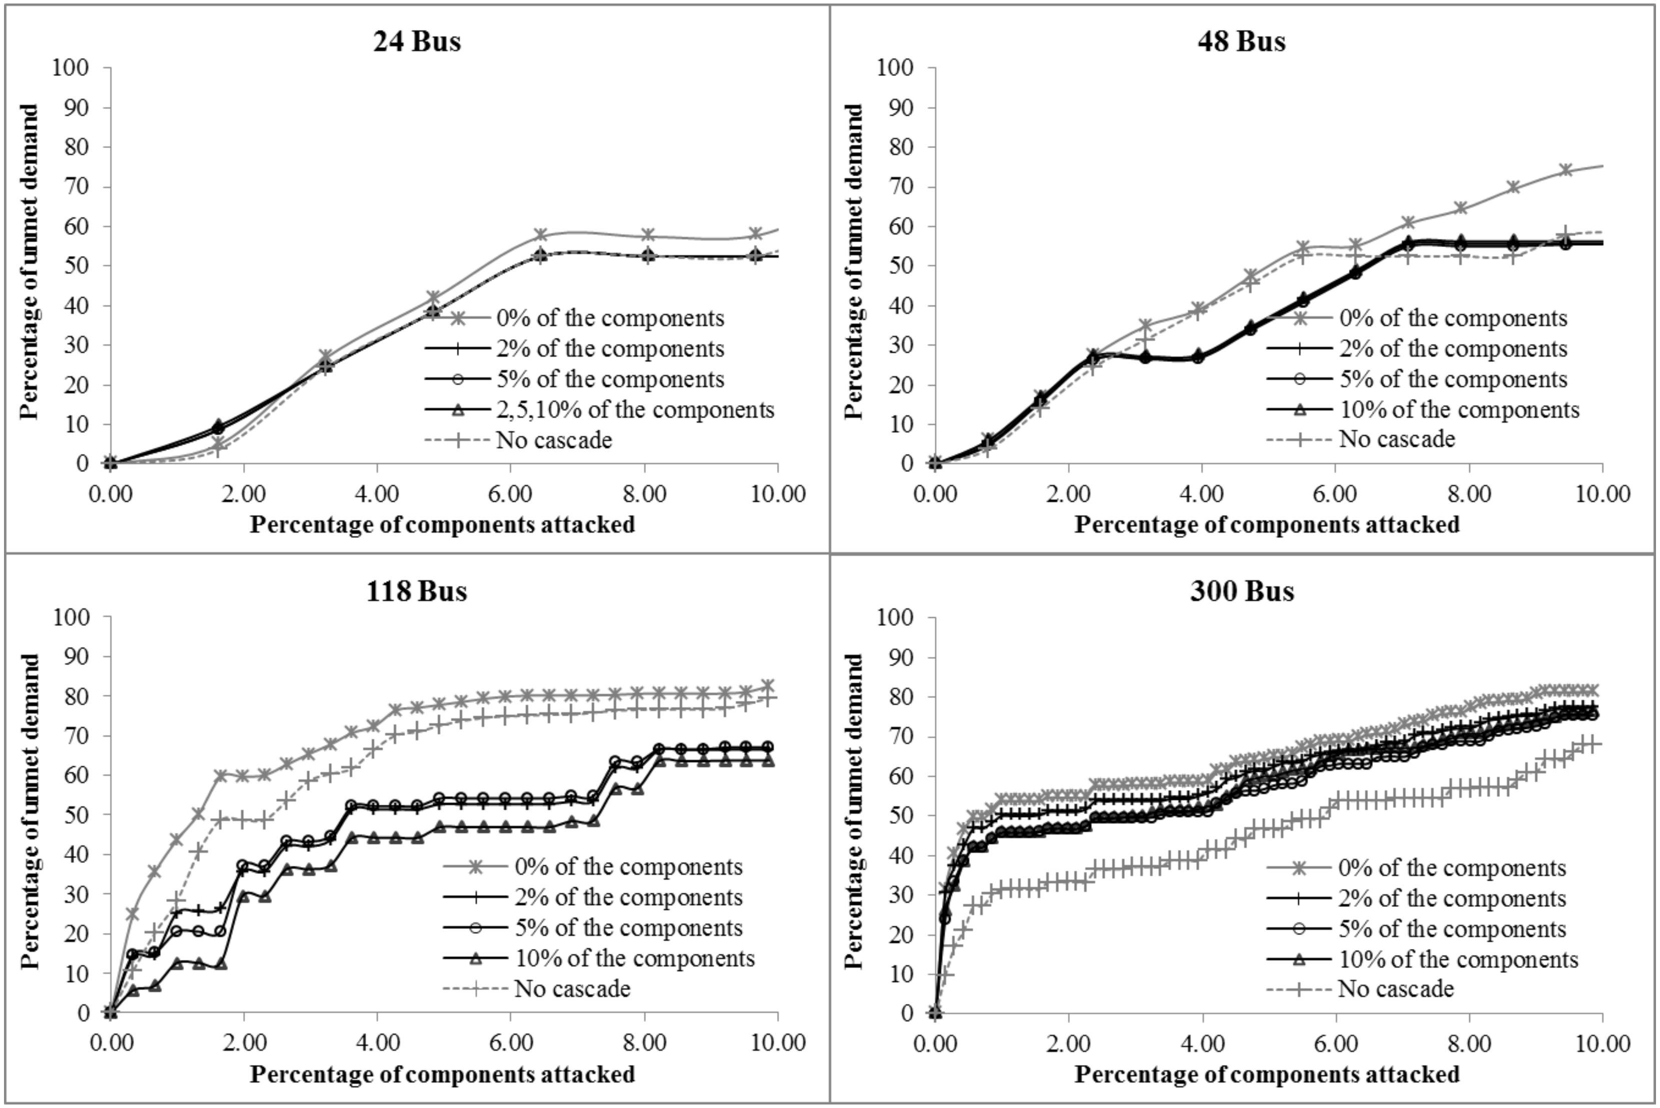

Appendix H. Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a small attack (static attacker with cascading knowledge)

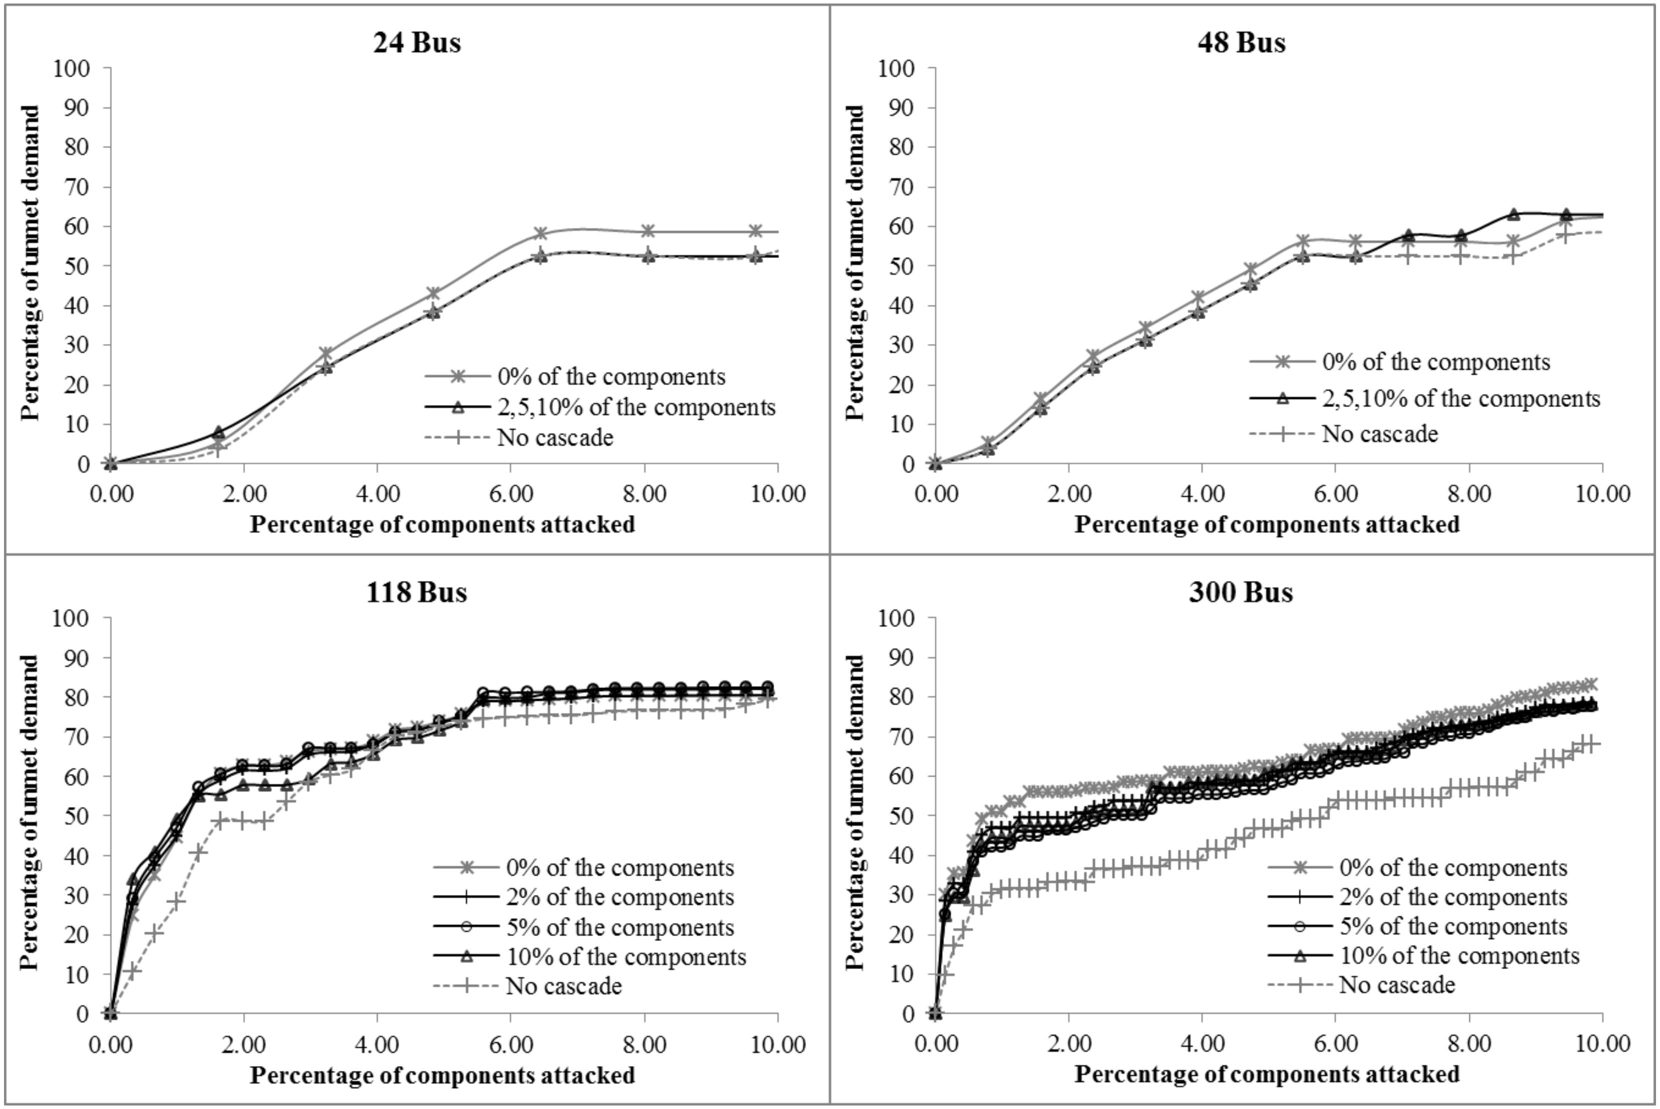

Appendix H. Various levels of investments in adding new parallel lines to the transmission lines that are mostly likely to cascade in a large attack (static attacker with cascading knowledge)

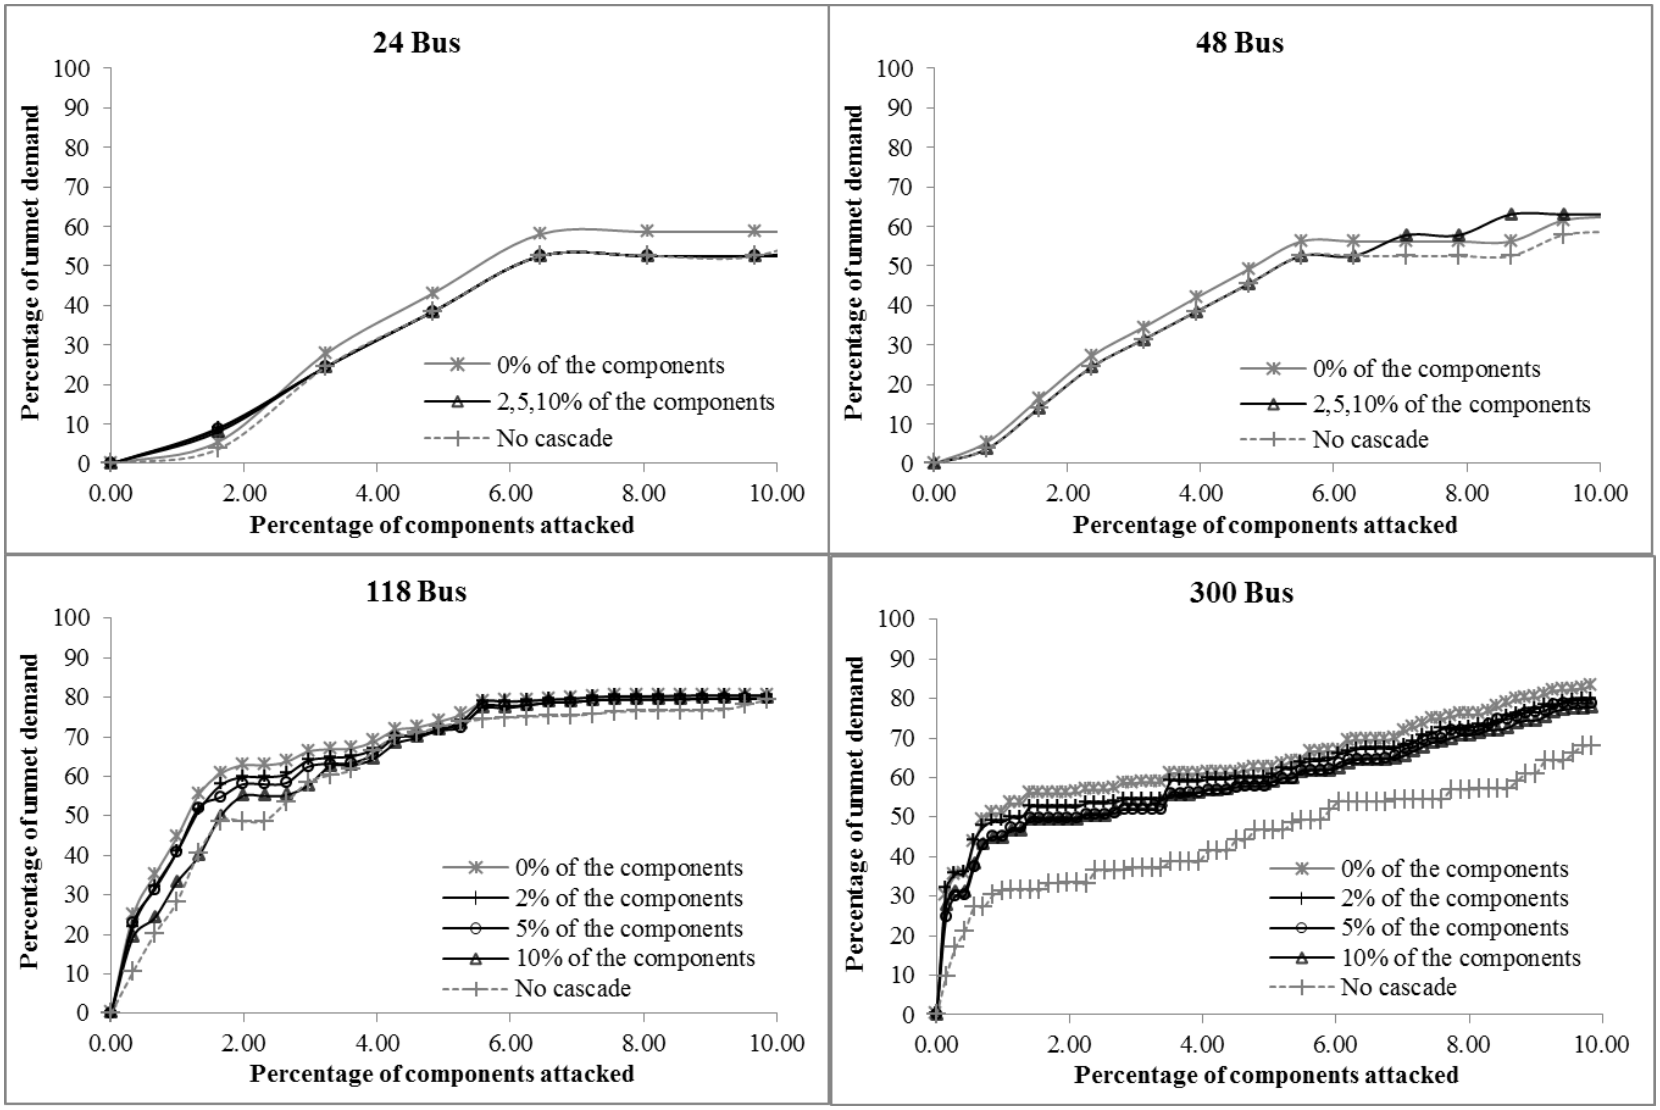

Appendix I. Comparing defensive investments against a static attacker (static attacker with no cascading knowledge (2% investment scenario)

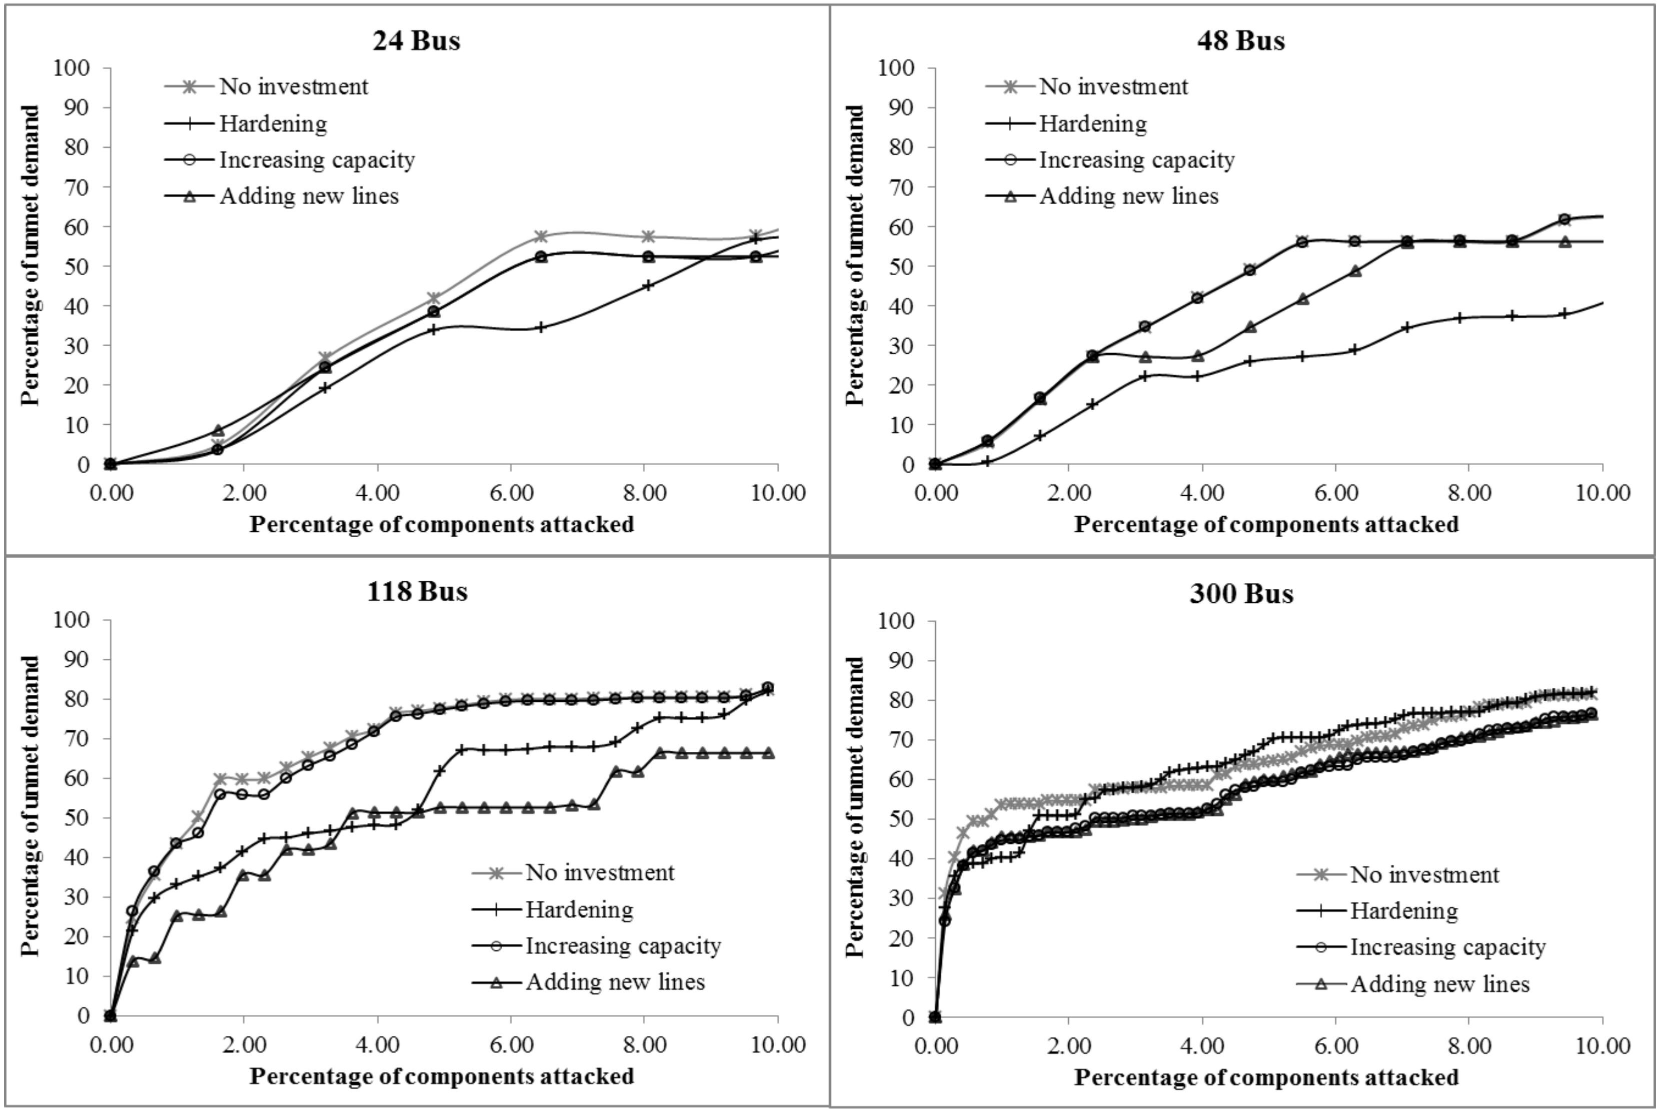

Appendix I. Comparing defensive investments against a static attacker (static attacker with no cascading knowledge, 5% investment scenario)

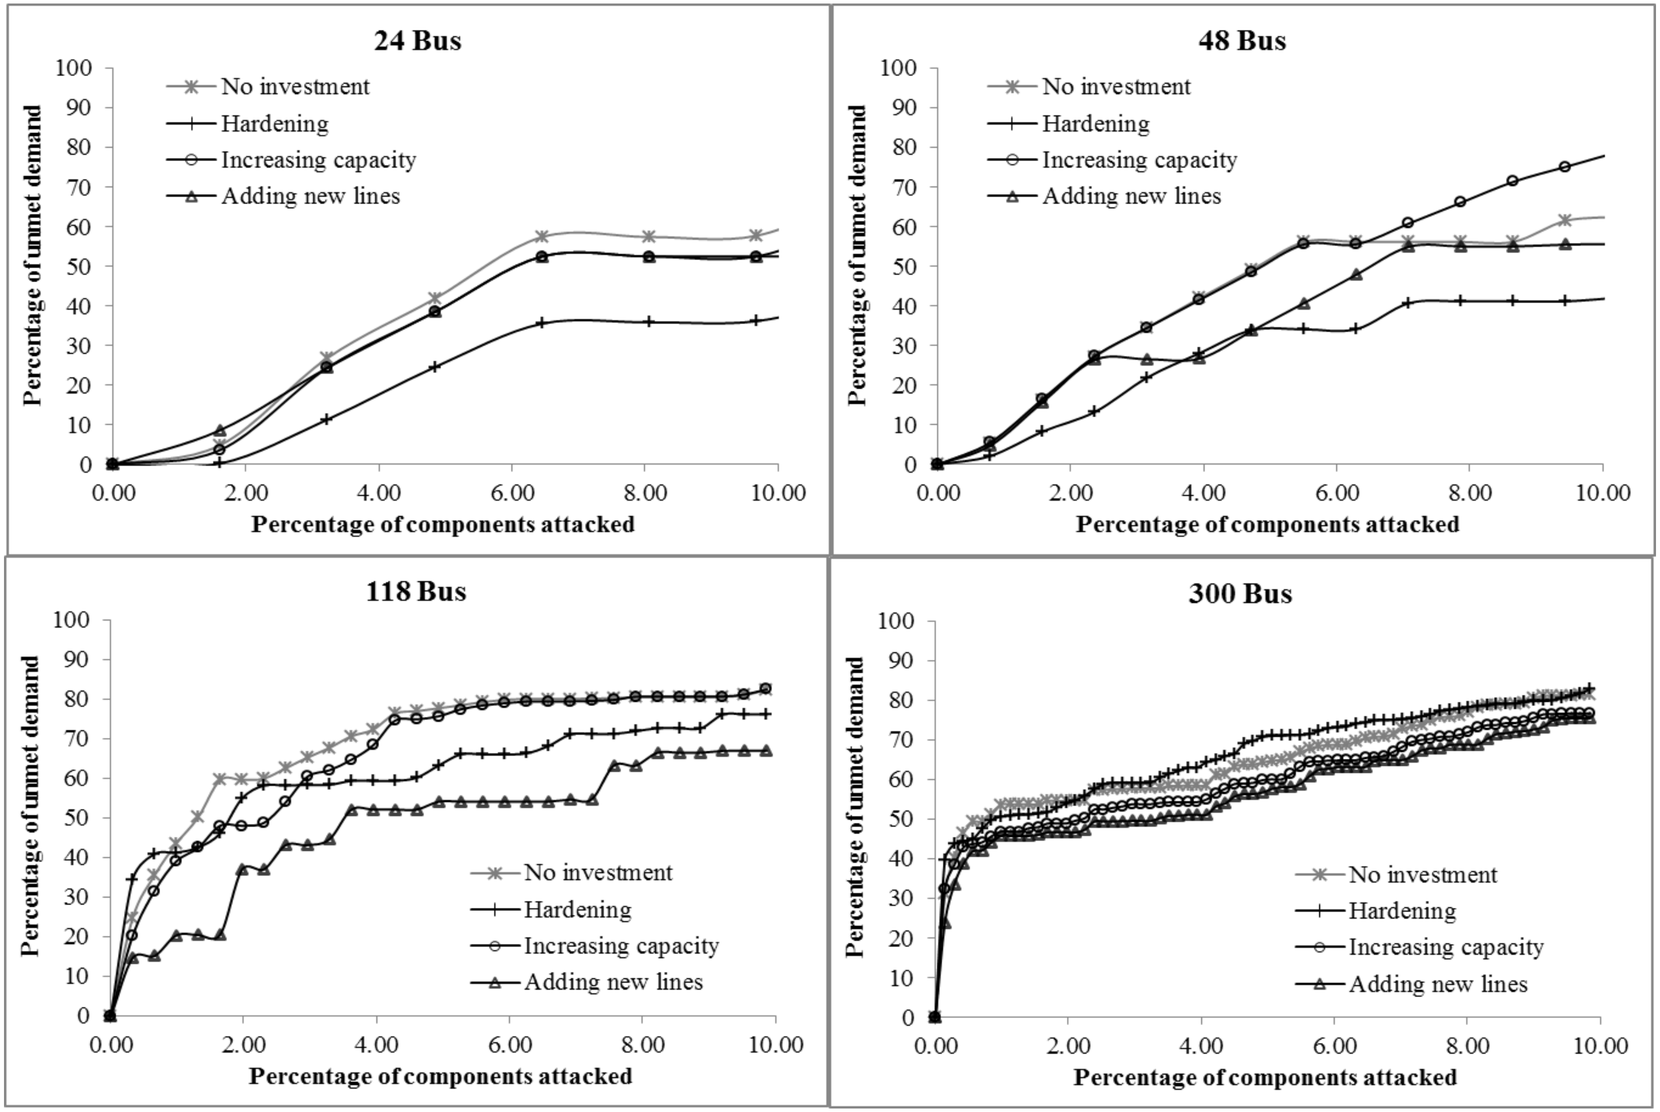

Appendix I. Comparing defensive investments against a static attacker (static attacker with no cascading knowledge, 10 % investment scenario)

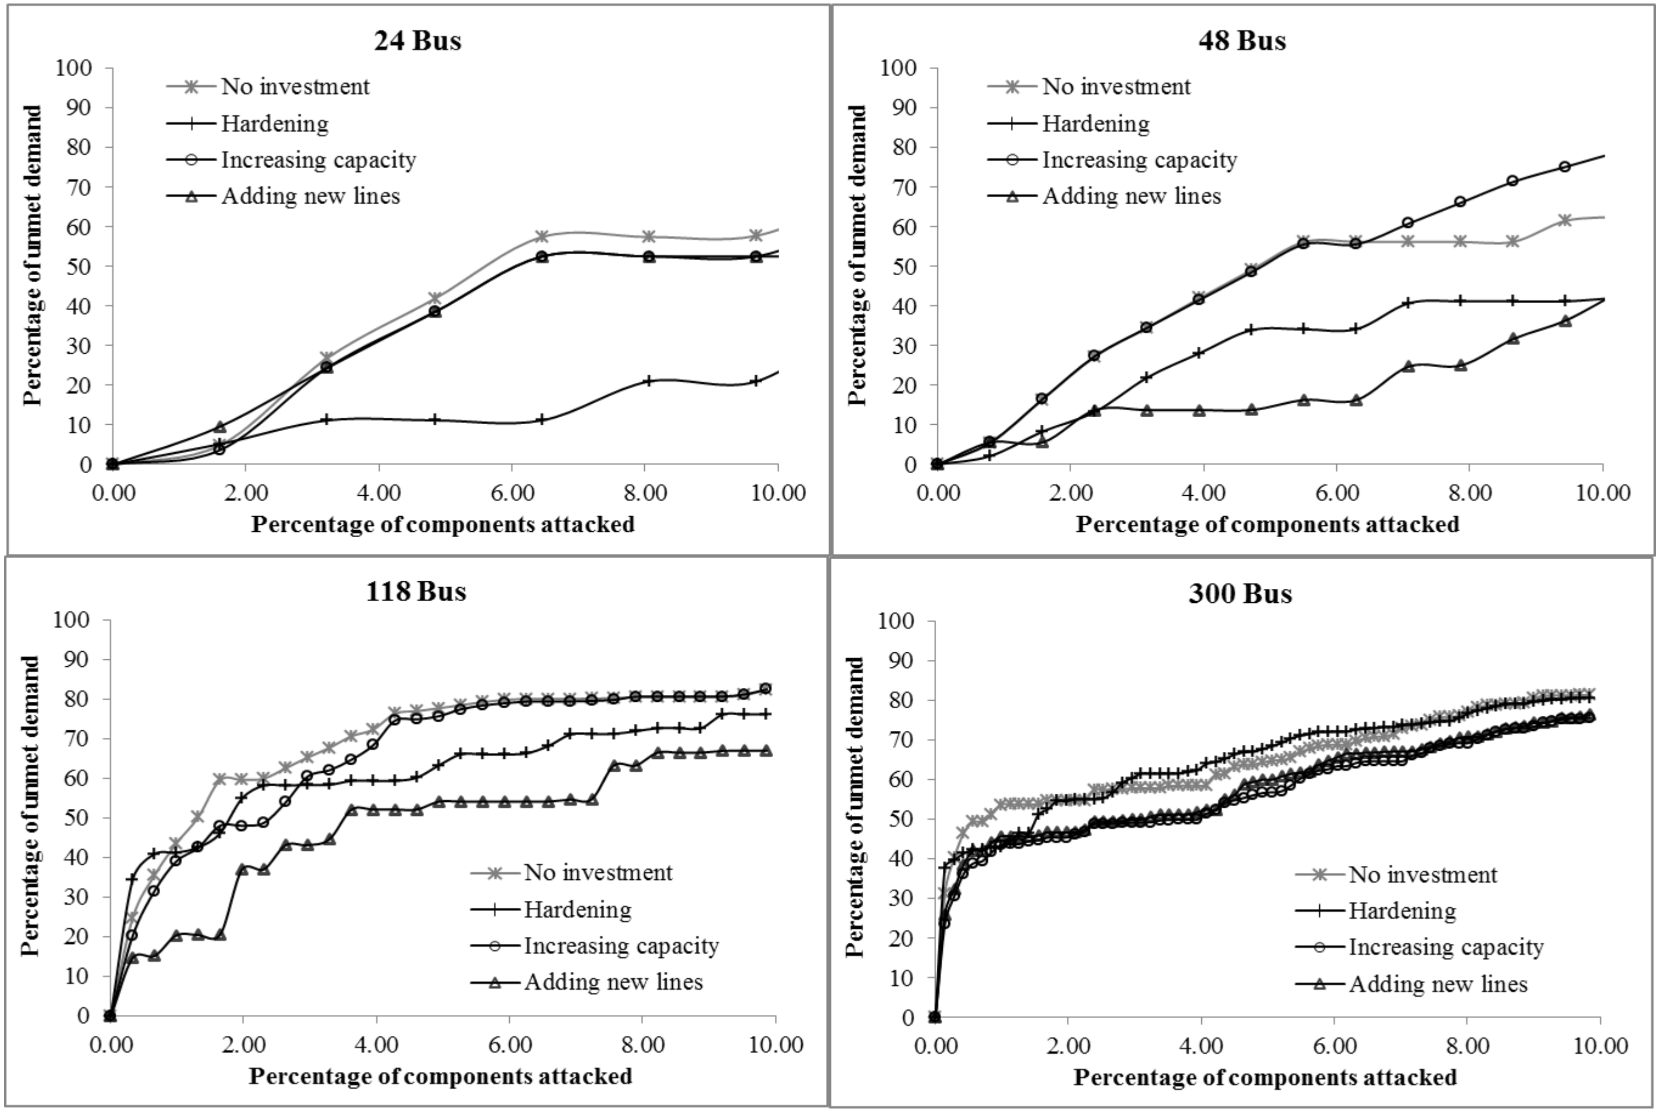

Appendix I. Comparing defensive investments against a static attacker (static attacker with cascading knowledge, 2% investment scenario).

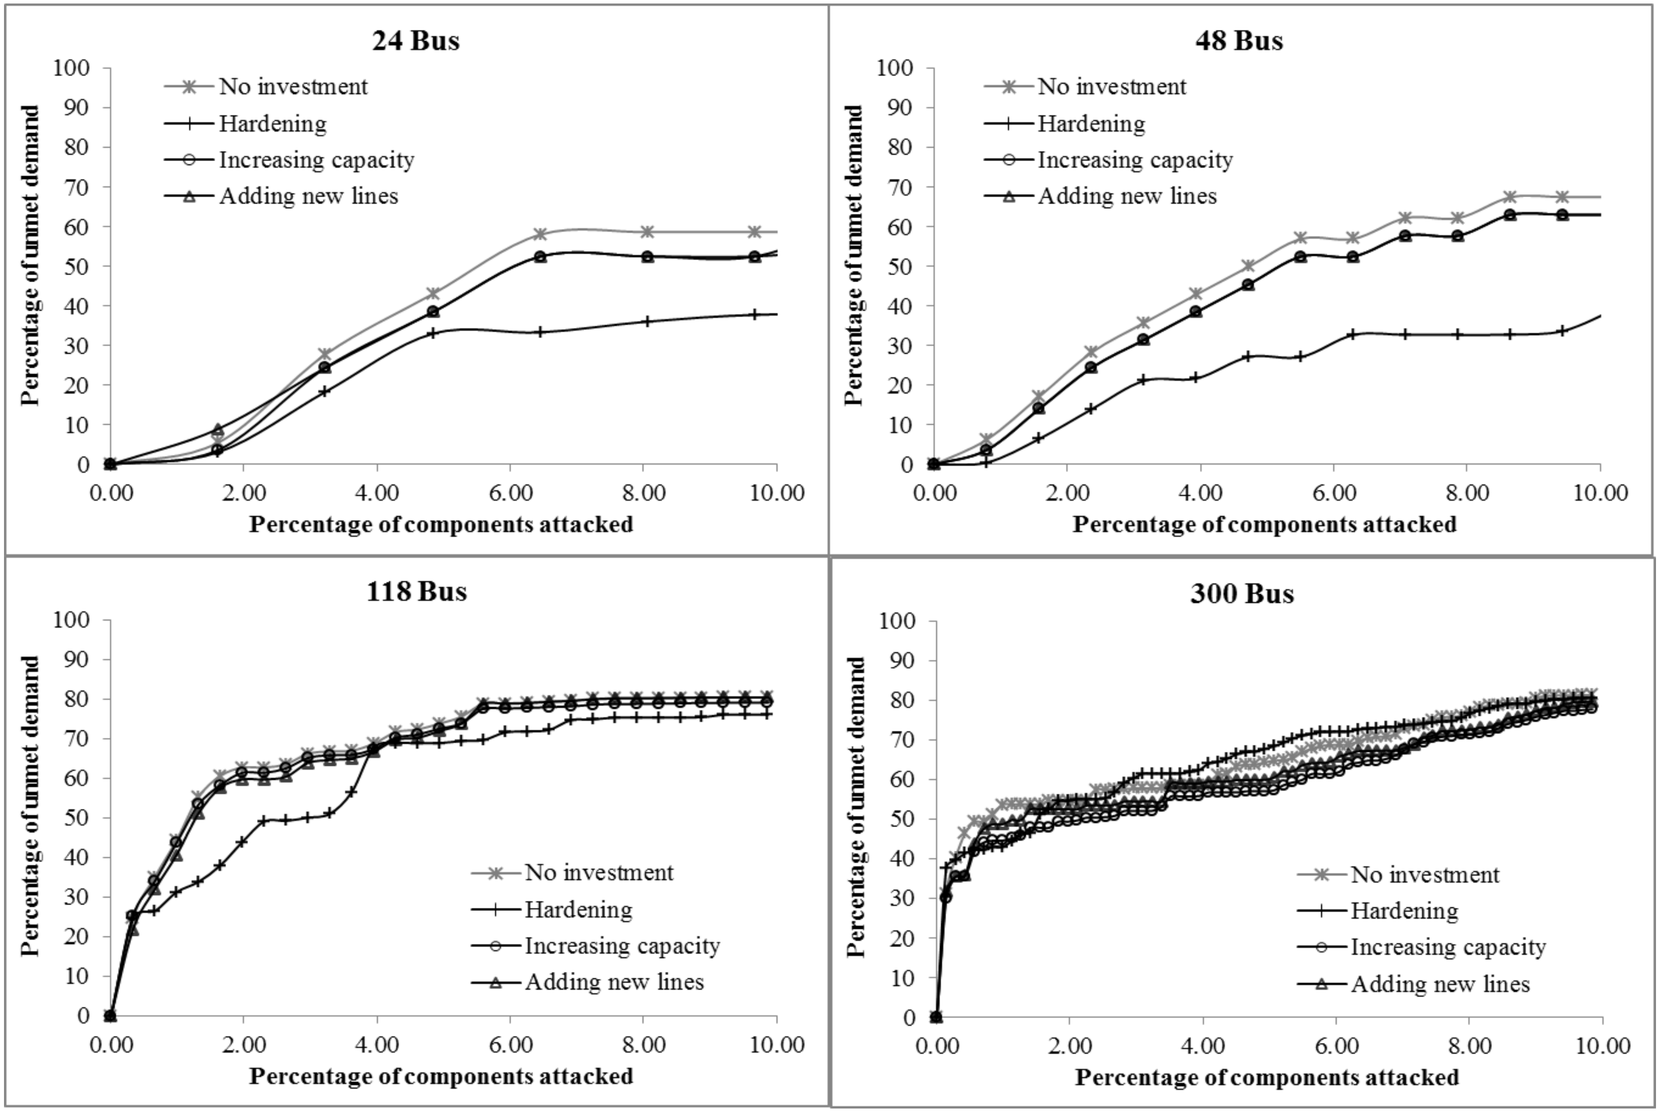
24 Bus

48 Bus

118 Bus

300 Bus

Appendix I. Comparing defensive investments against a static attacker (static attacker with cascading knowledge, 5% investment scenario).

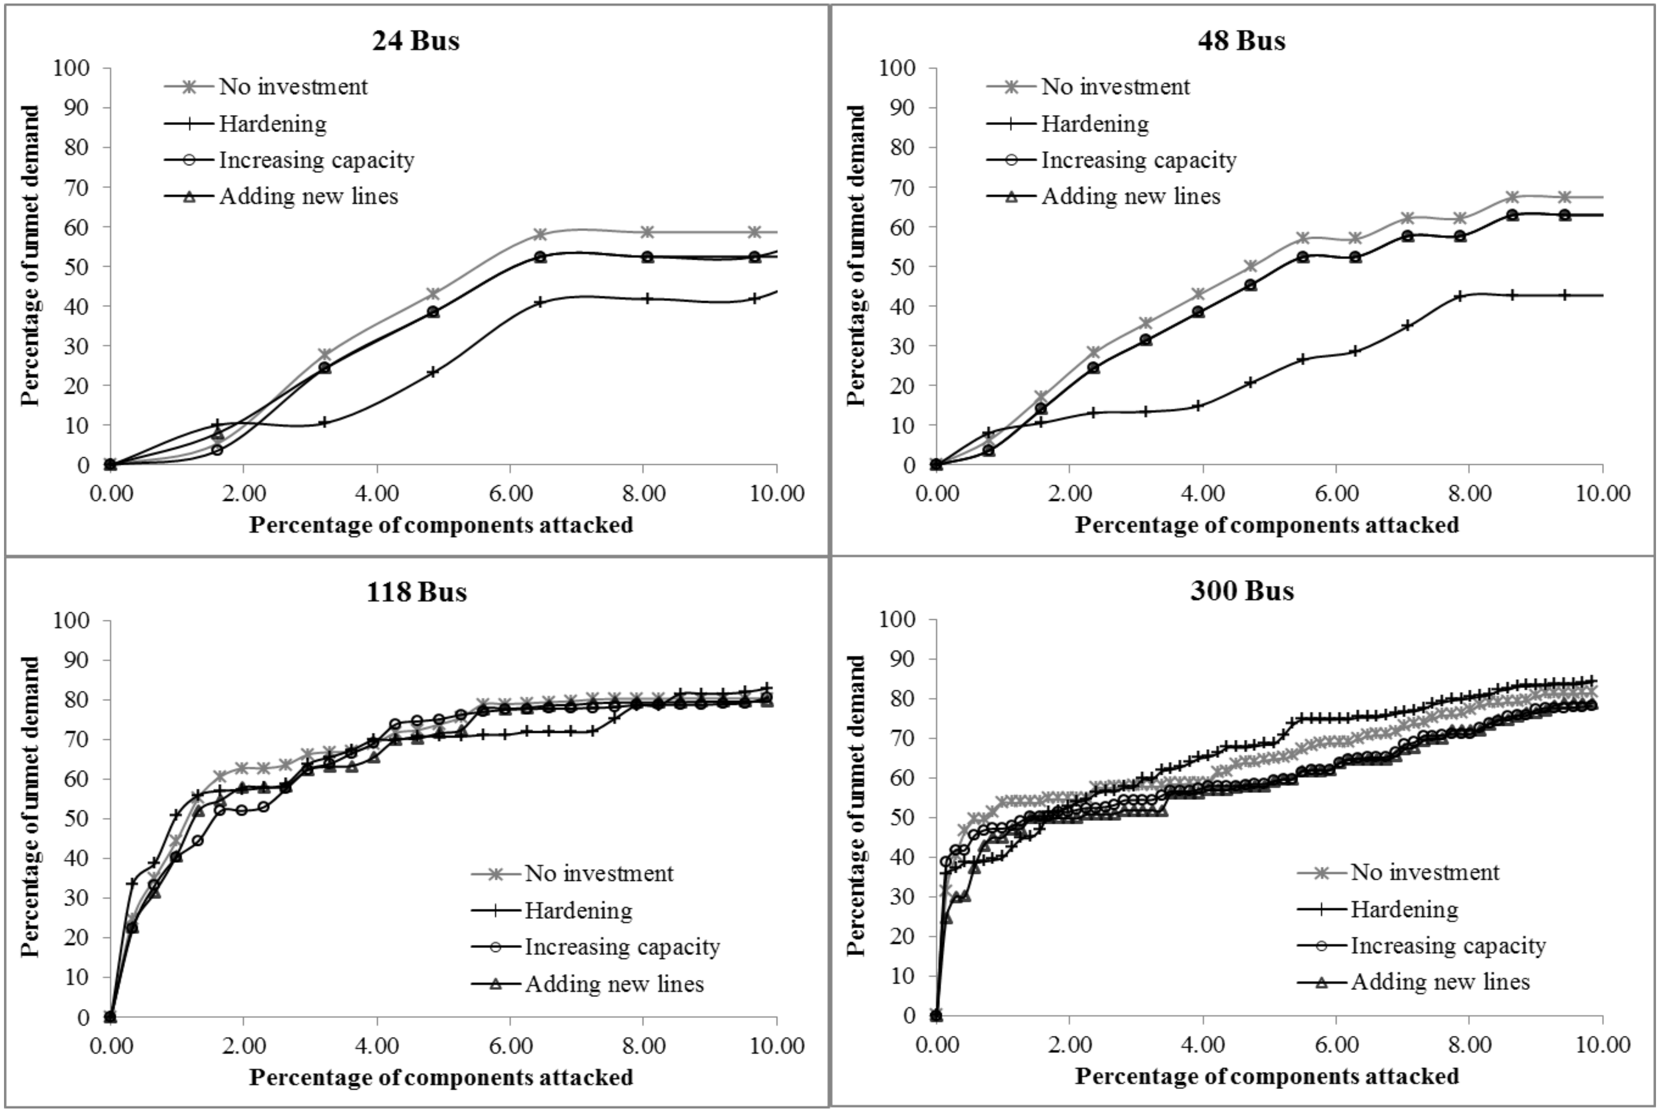
24 Bus

48 Bus

118 Bus

300 Bus

Appendix I. Comparing defensive investments against a static attacker (static attacker with cascading knowledge (10% investment scenario).

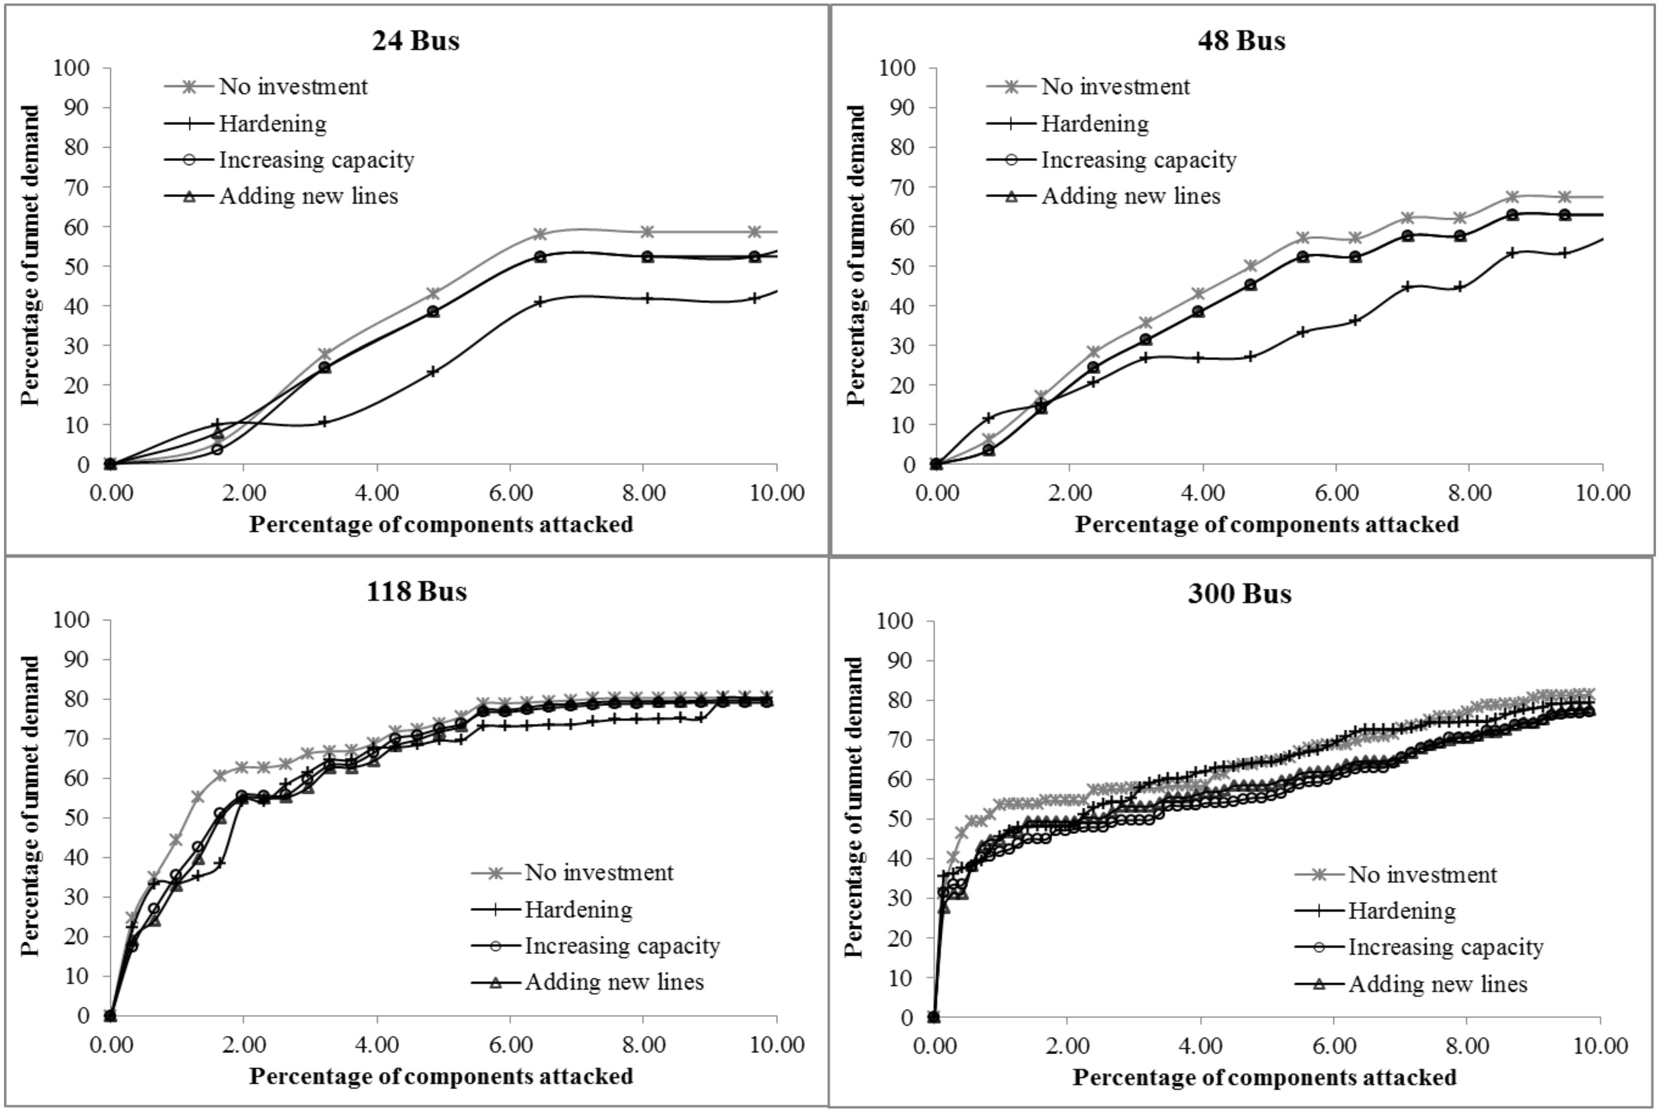

Appendix J. Impact of considering restoration times against a static attacker (static attacker with no cascading knowledge, 24-bus system)

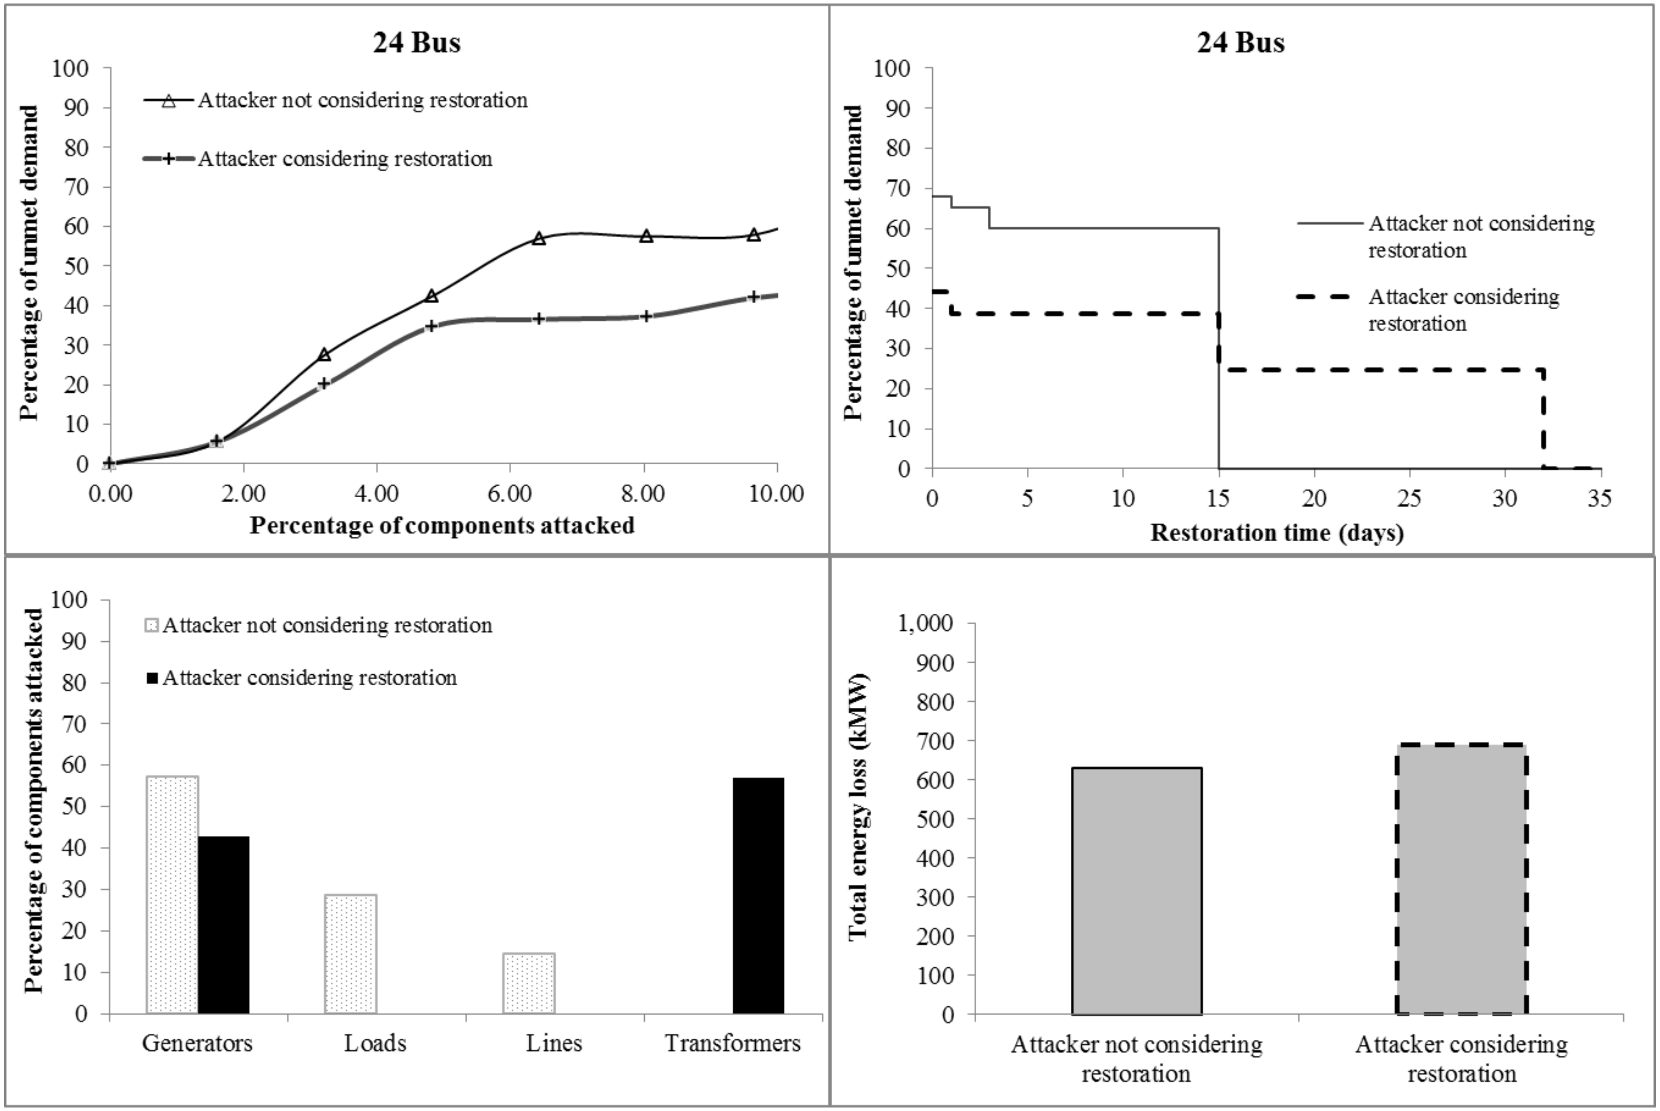

Appendix J. Impact of considering restoration times against a static attacker (static attacker with no cascading knowledge, 48-bus system)

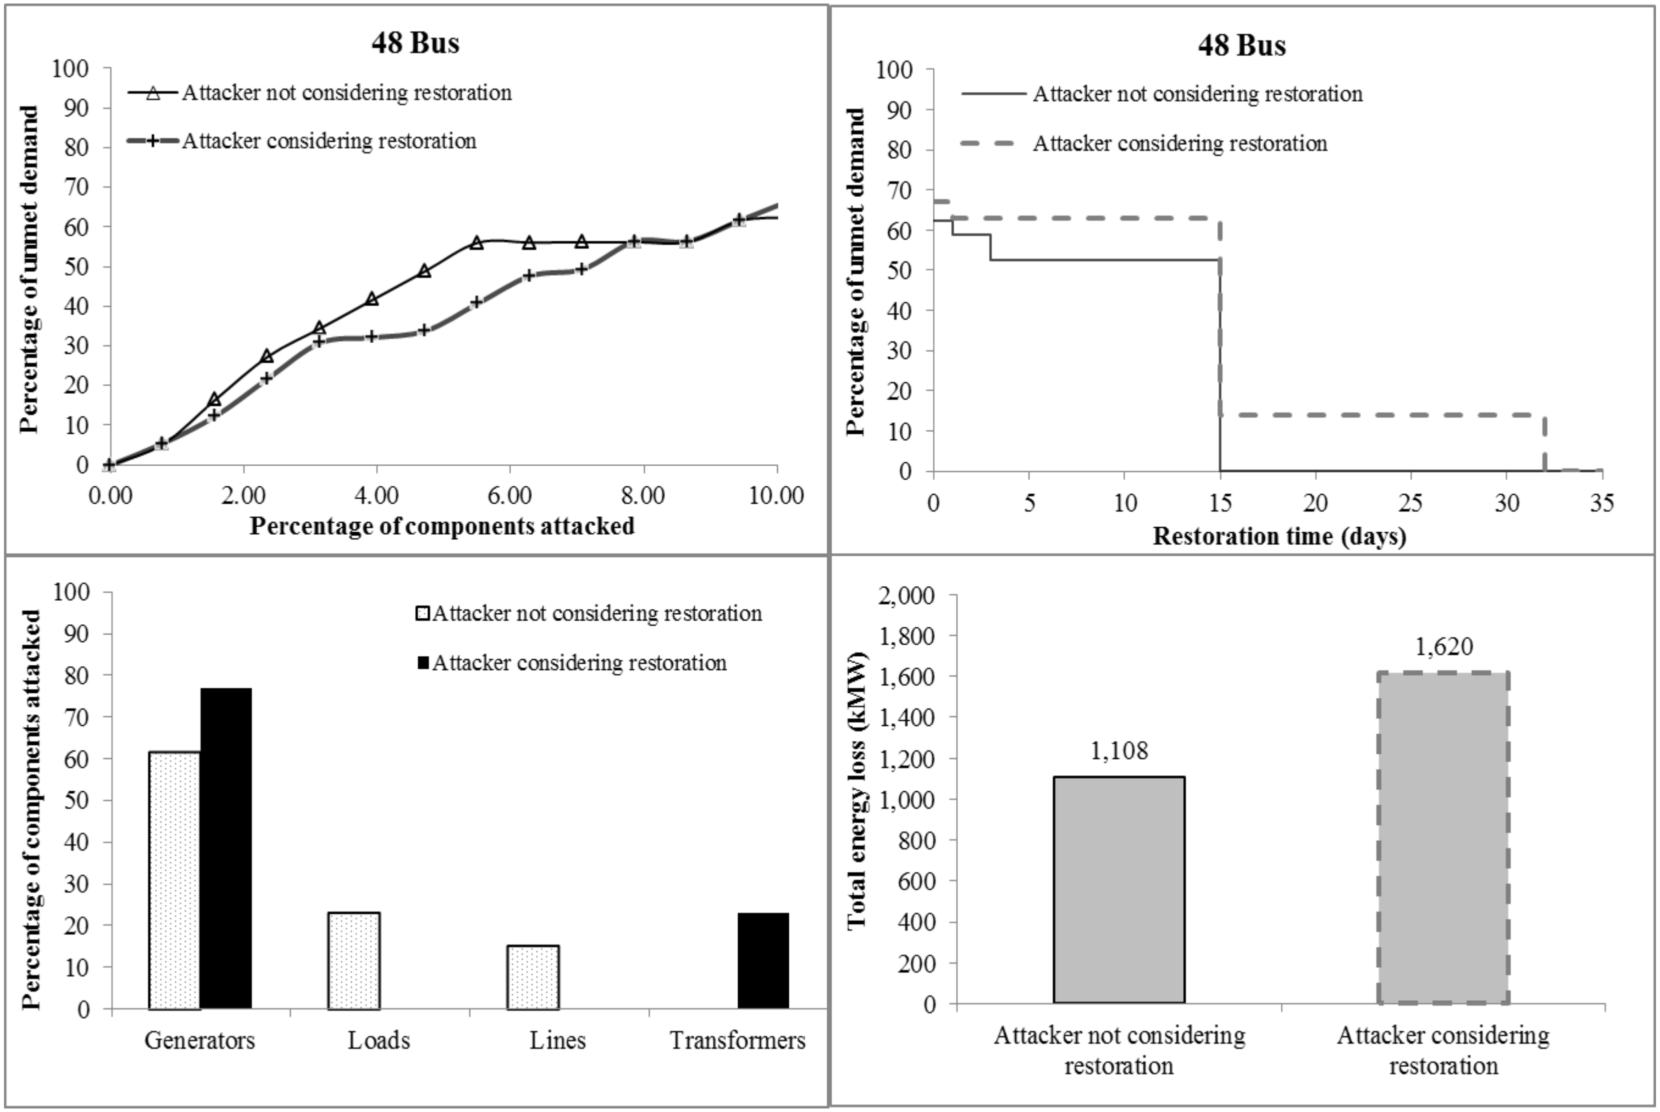

Appendix J. Impact of considering restoration times against a static attacker (static attacker with no cascading knowledge, 118-bus system)

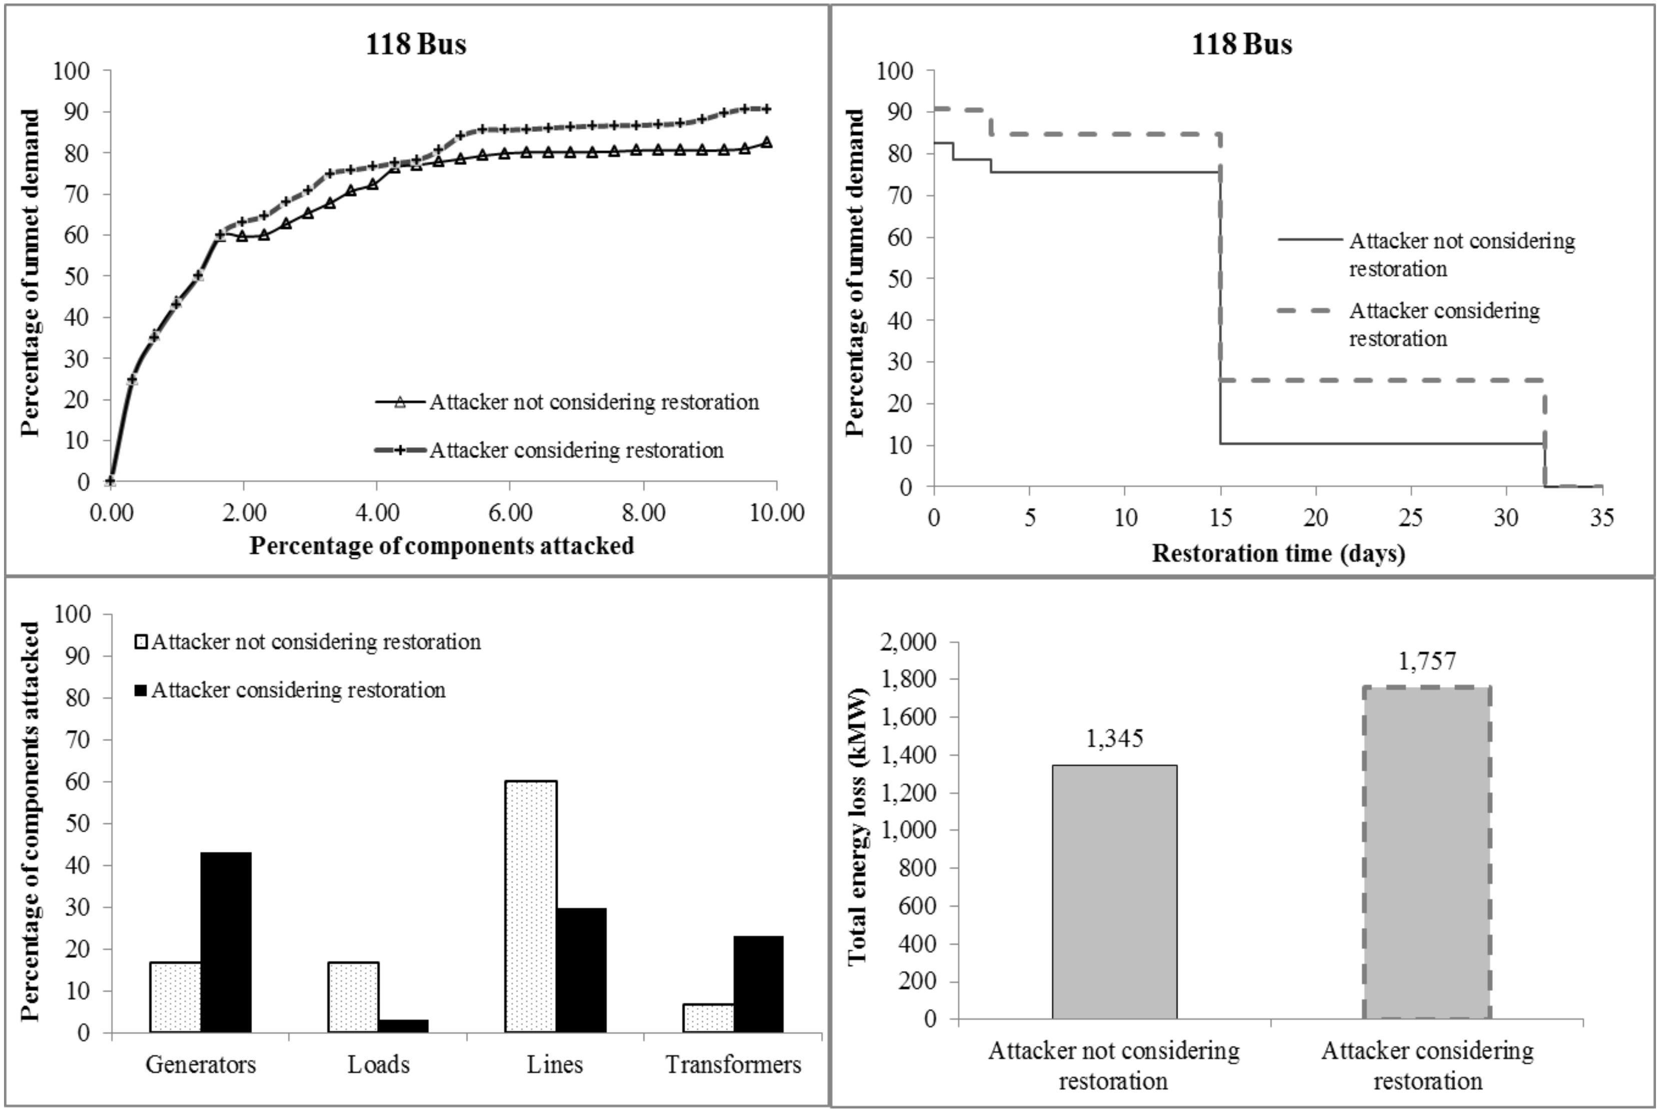

Appendix J. Impact of considering restoration times against a static attacker (static attacker with no cascading knowledge, 300-bus system)

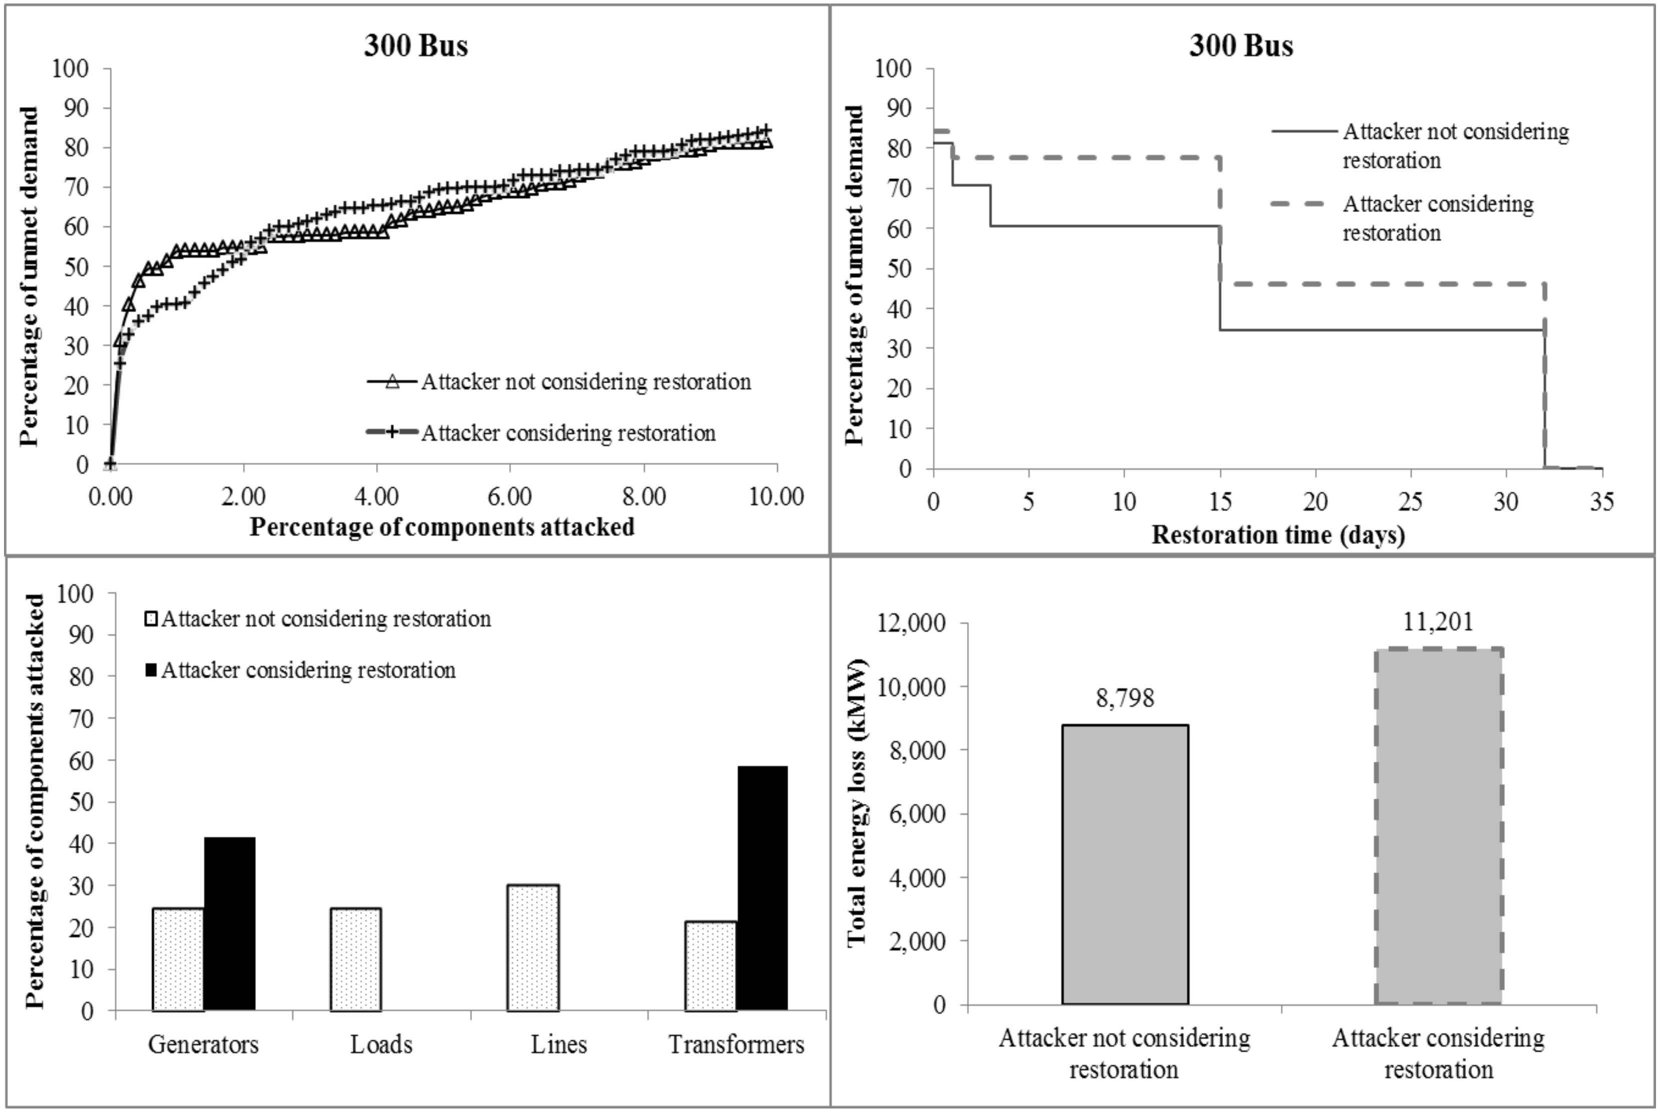

Appendix J. Impact of considering restoration times against a static attacker (static attacker with cascading knowledge, 24-bus system)

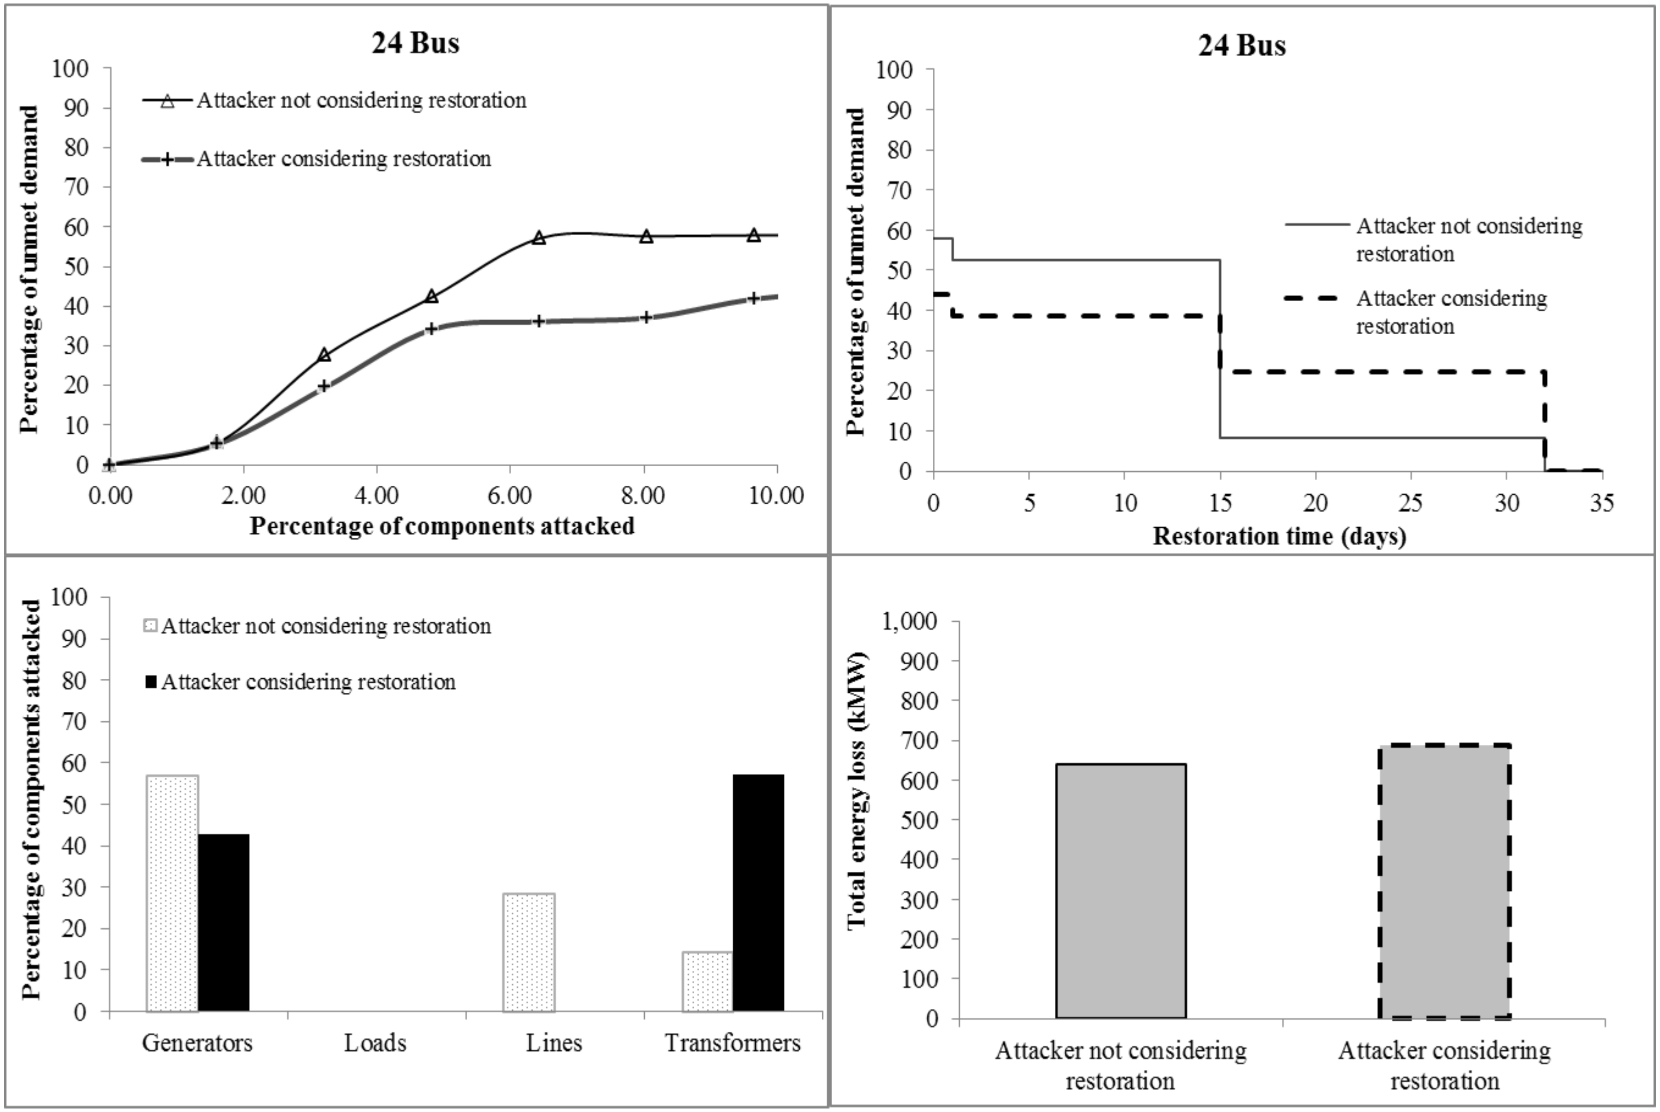

Appendix J. Impact of considering restoration times against a static attacker (static attacker with cascading knowledge, 48-bus system)

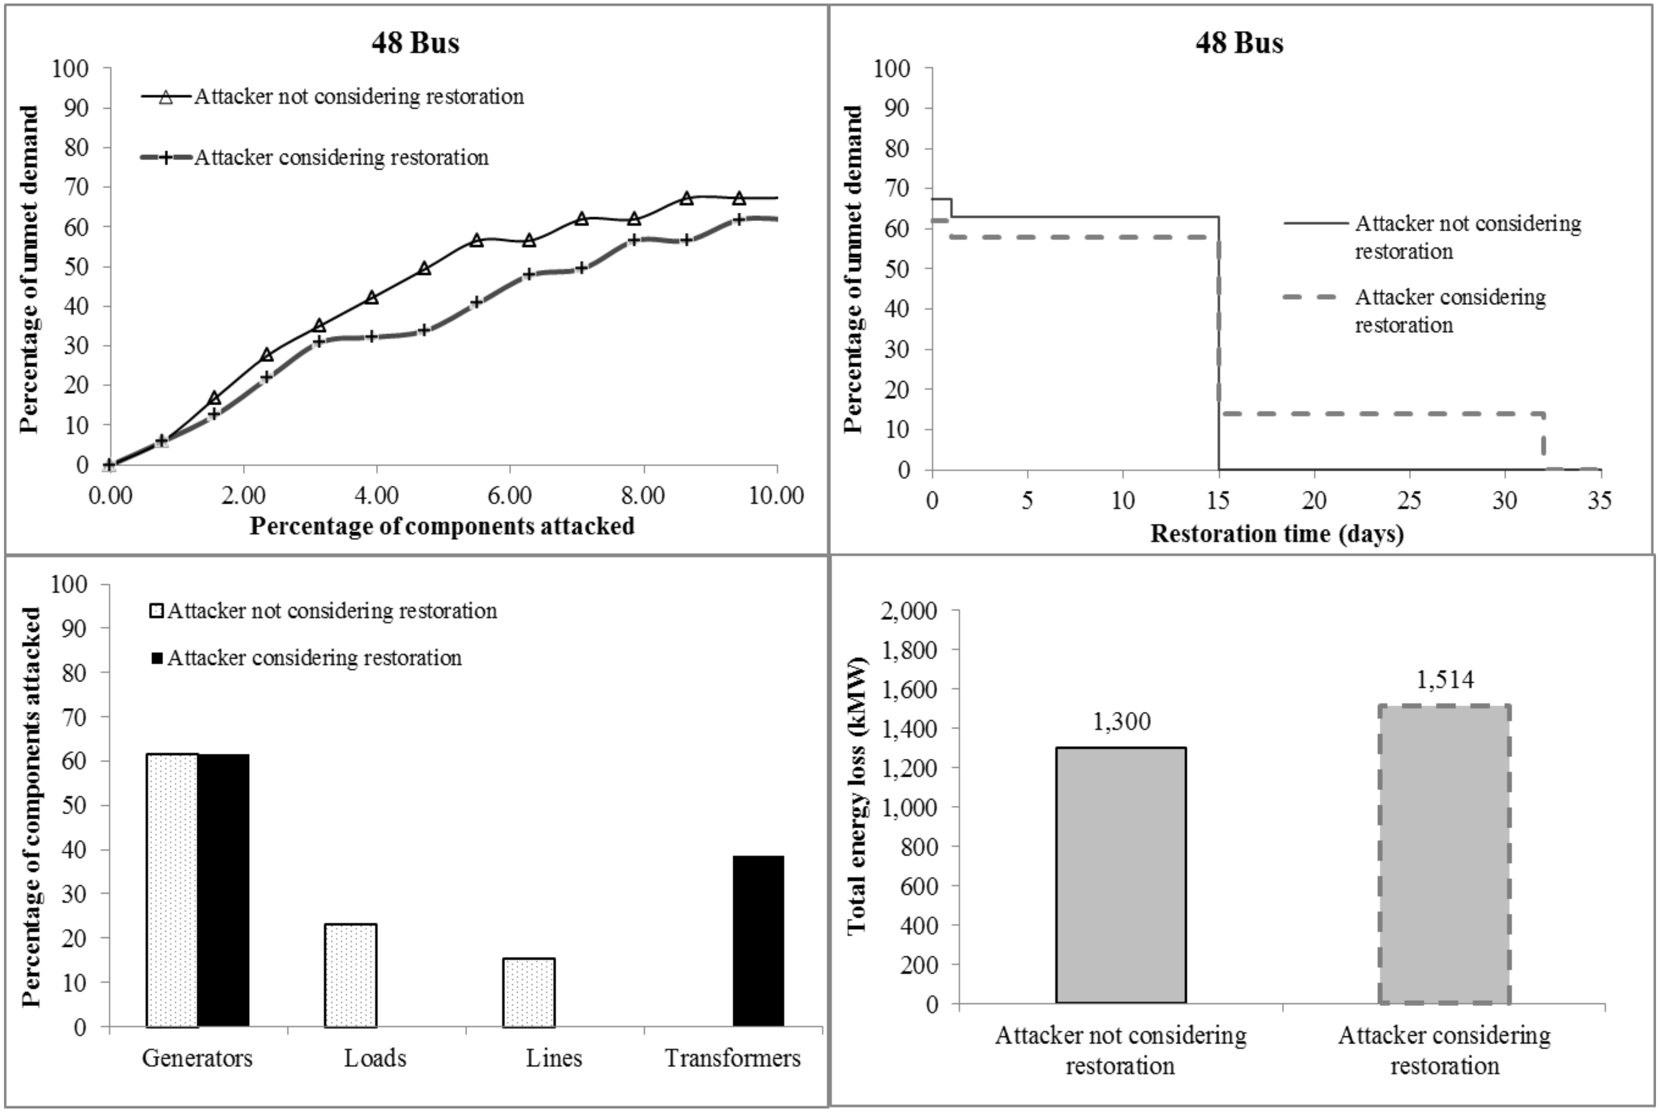

Appendix J. Impact of considering restoration times against a static attacker (static attacker with cascading knowledge, 118-bus system)

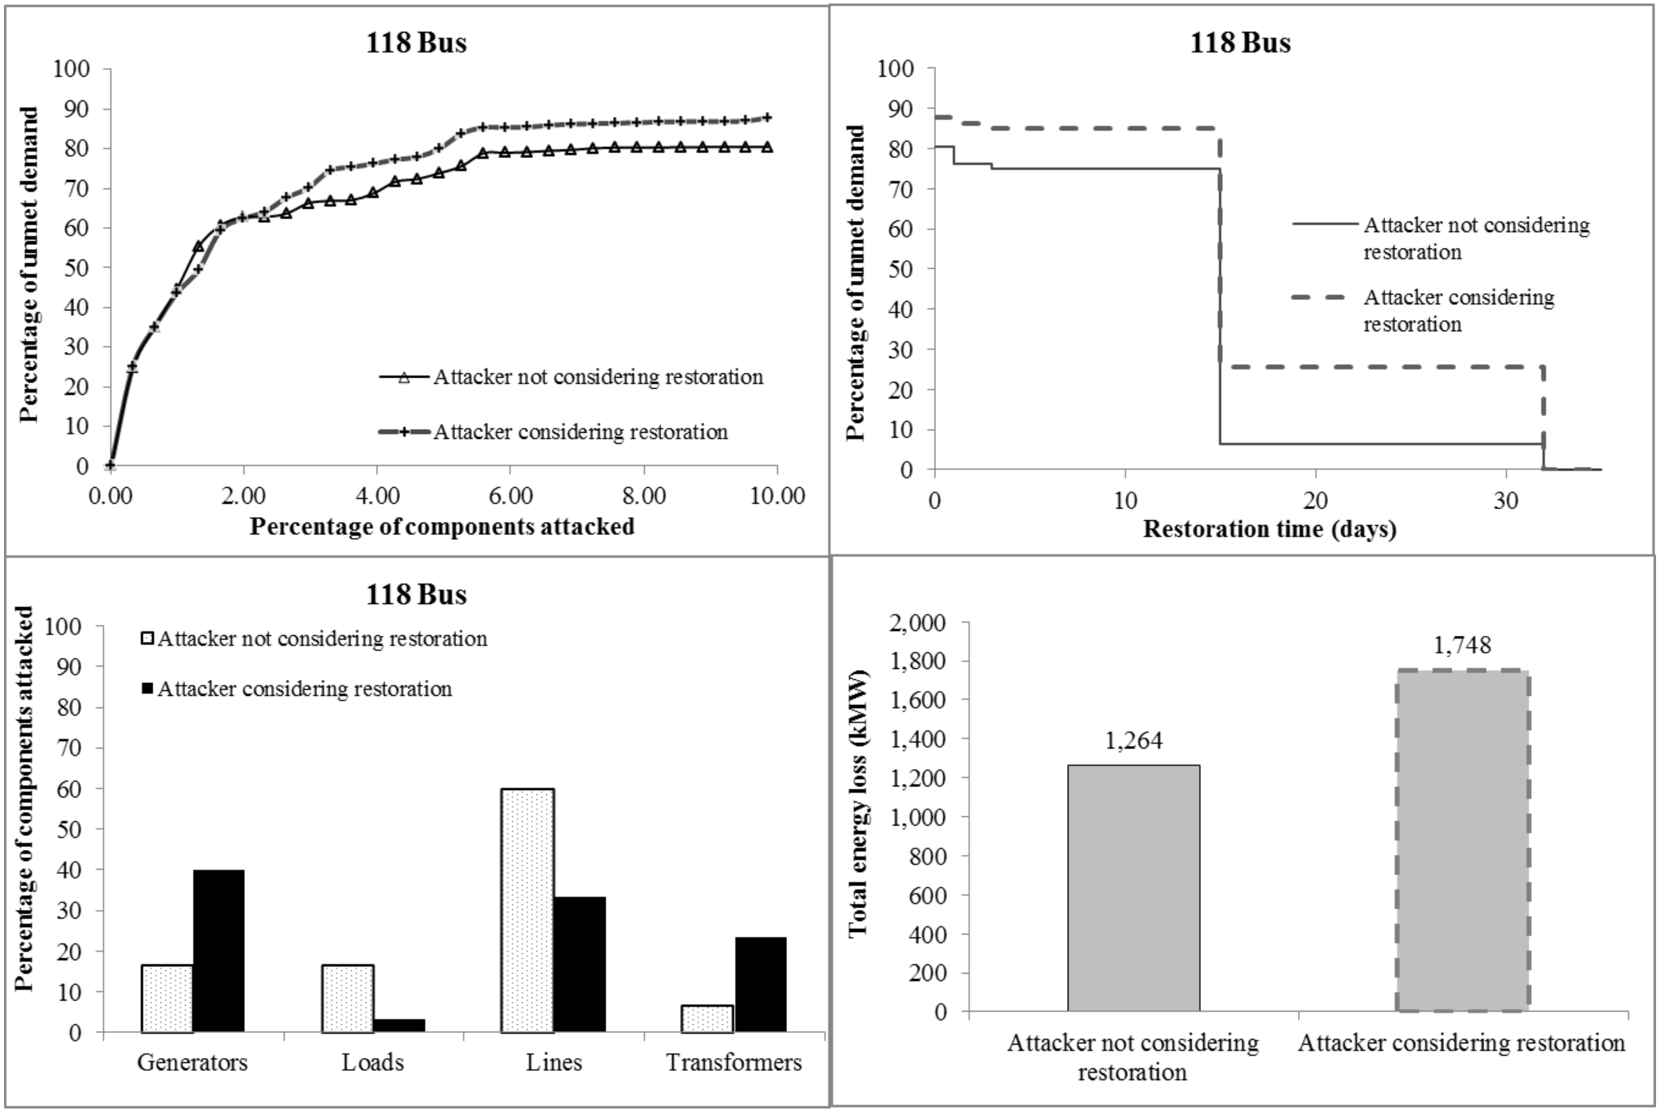

Appendix J. Impact of considering restoration times (static attacker with cascading knowledge, 300-bus system)

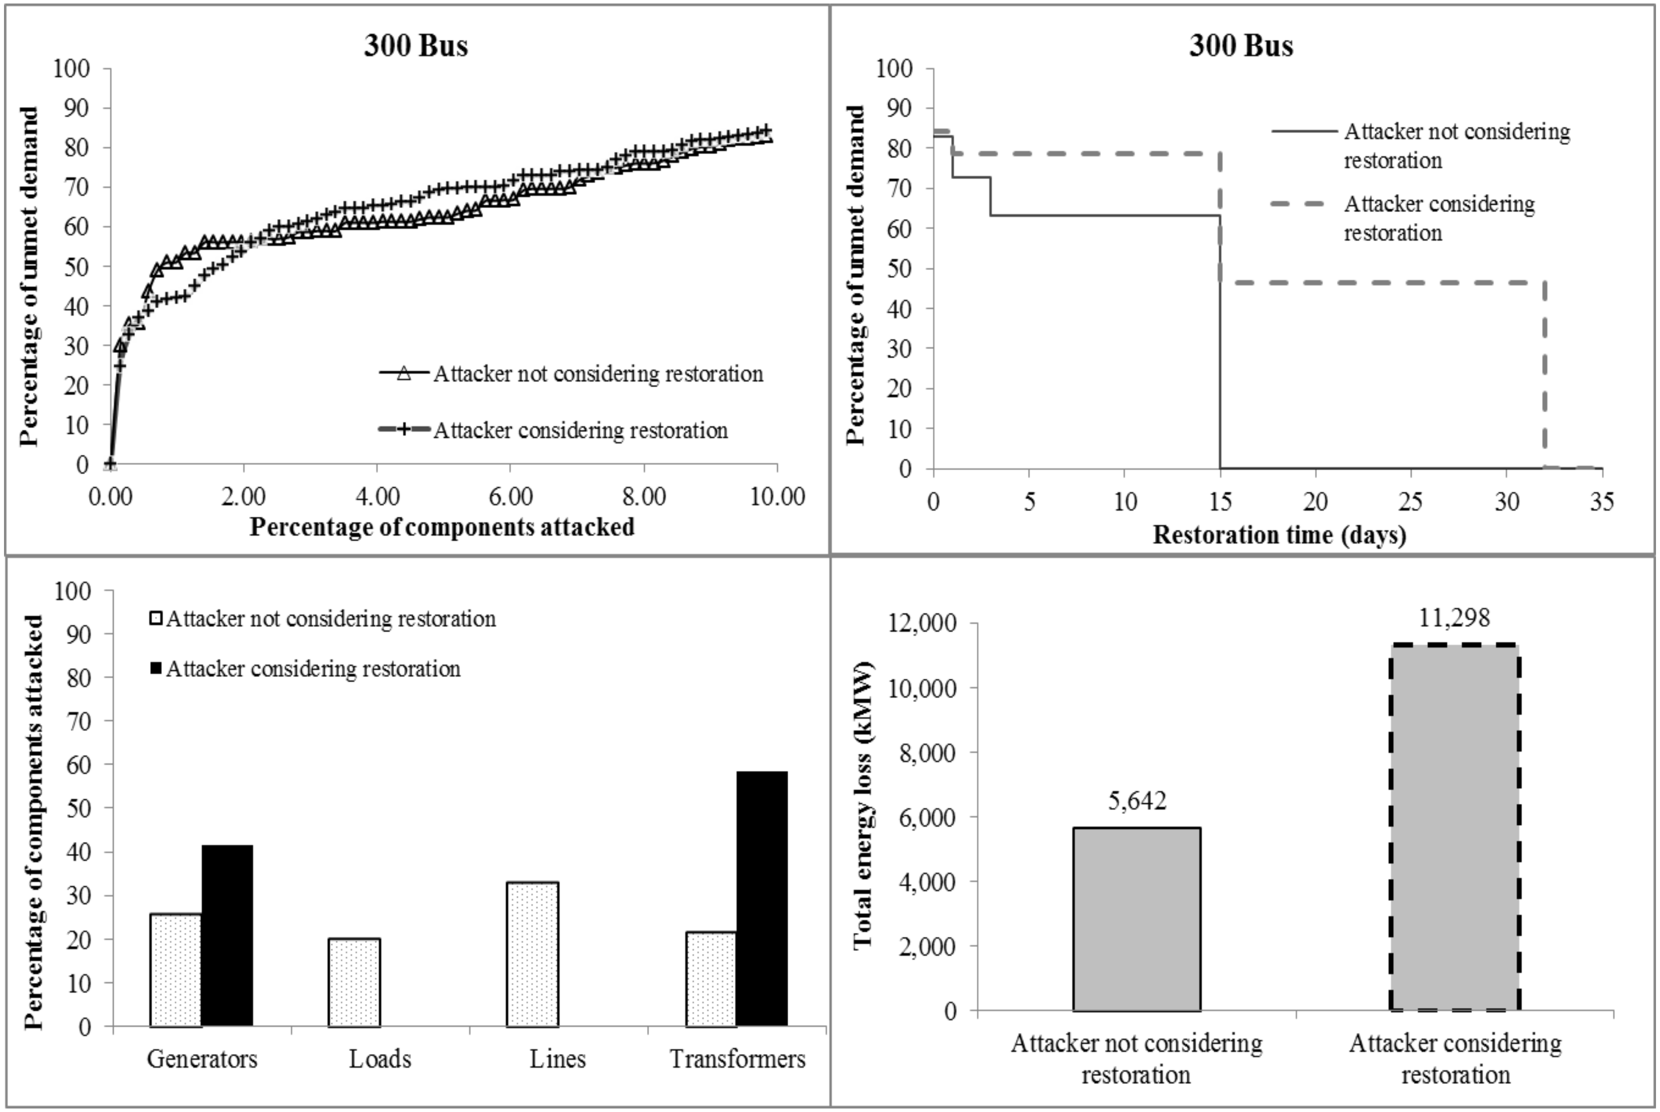

Appendix K. Comparison of three attacker assumptions with cascading knowledge when they consider restoration times (24-bus system)

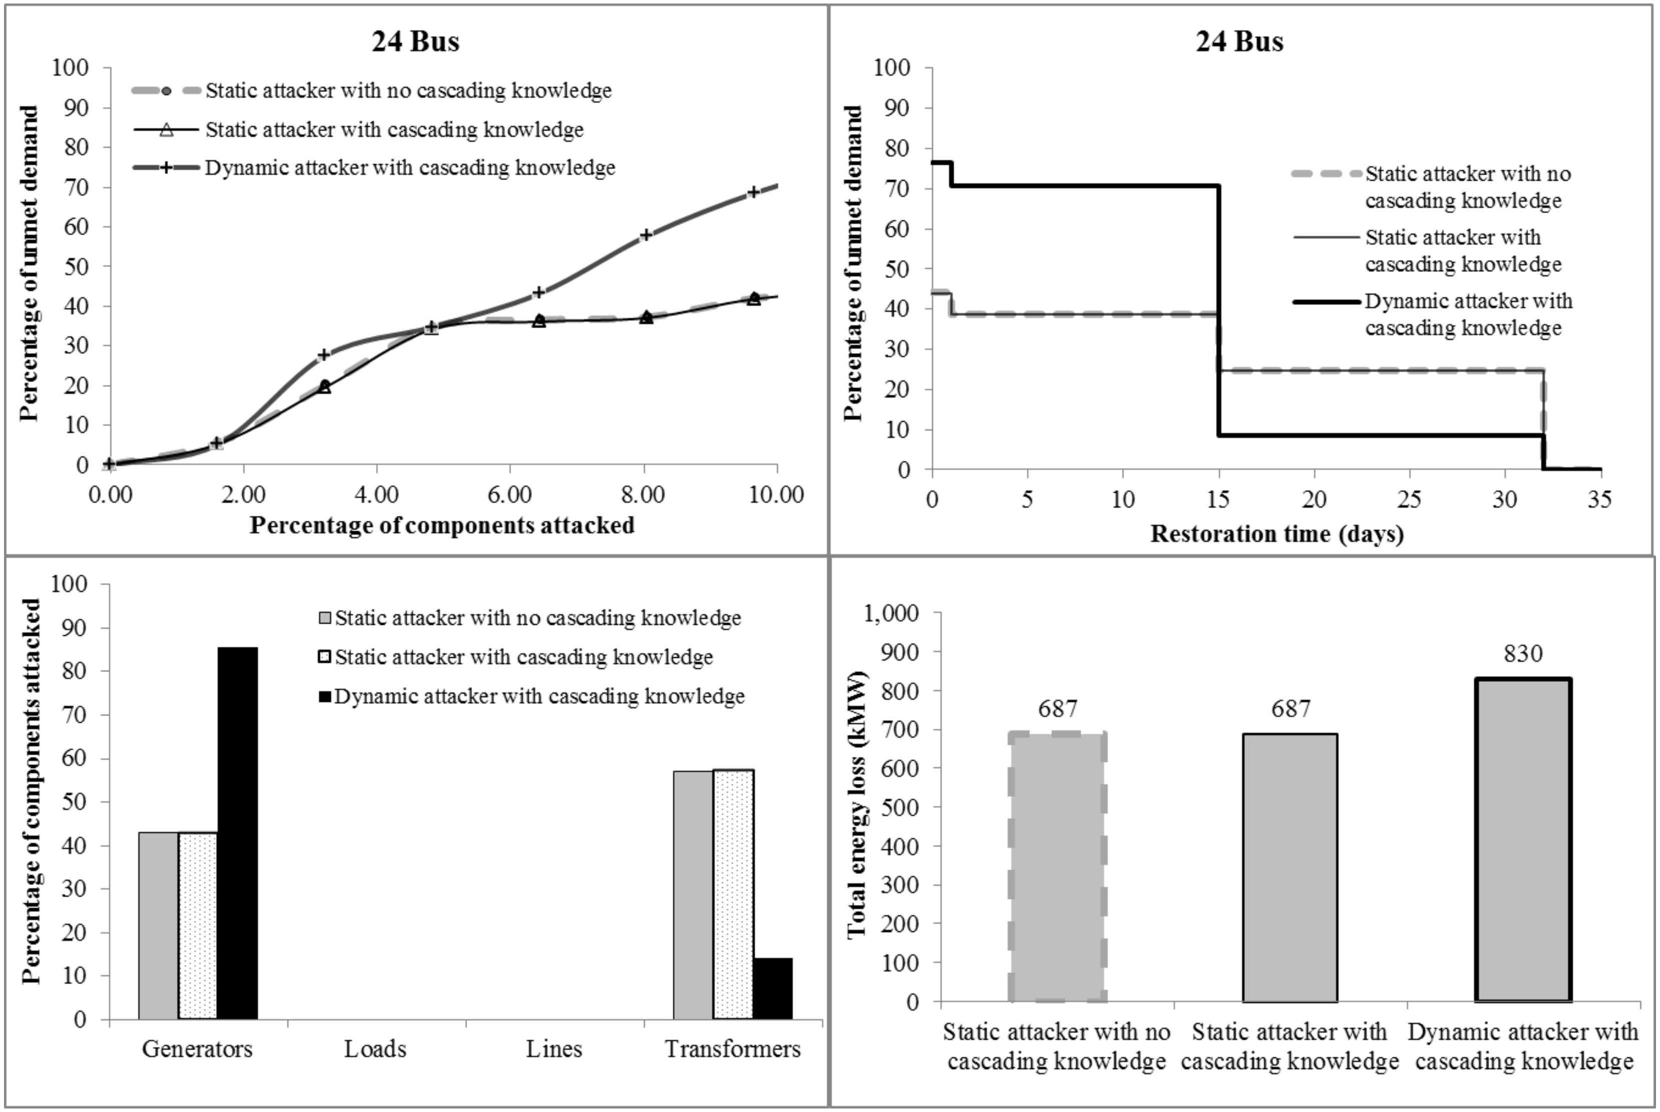

Appendix K. Comparison of three attacker assumptions with cascading knowledge when they consider restoration times (48-bus system)

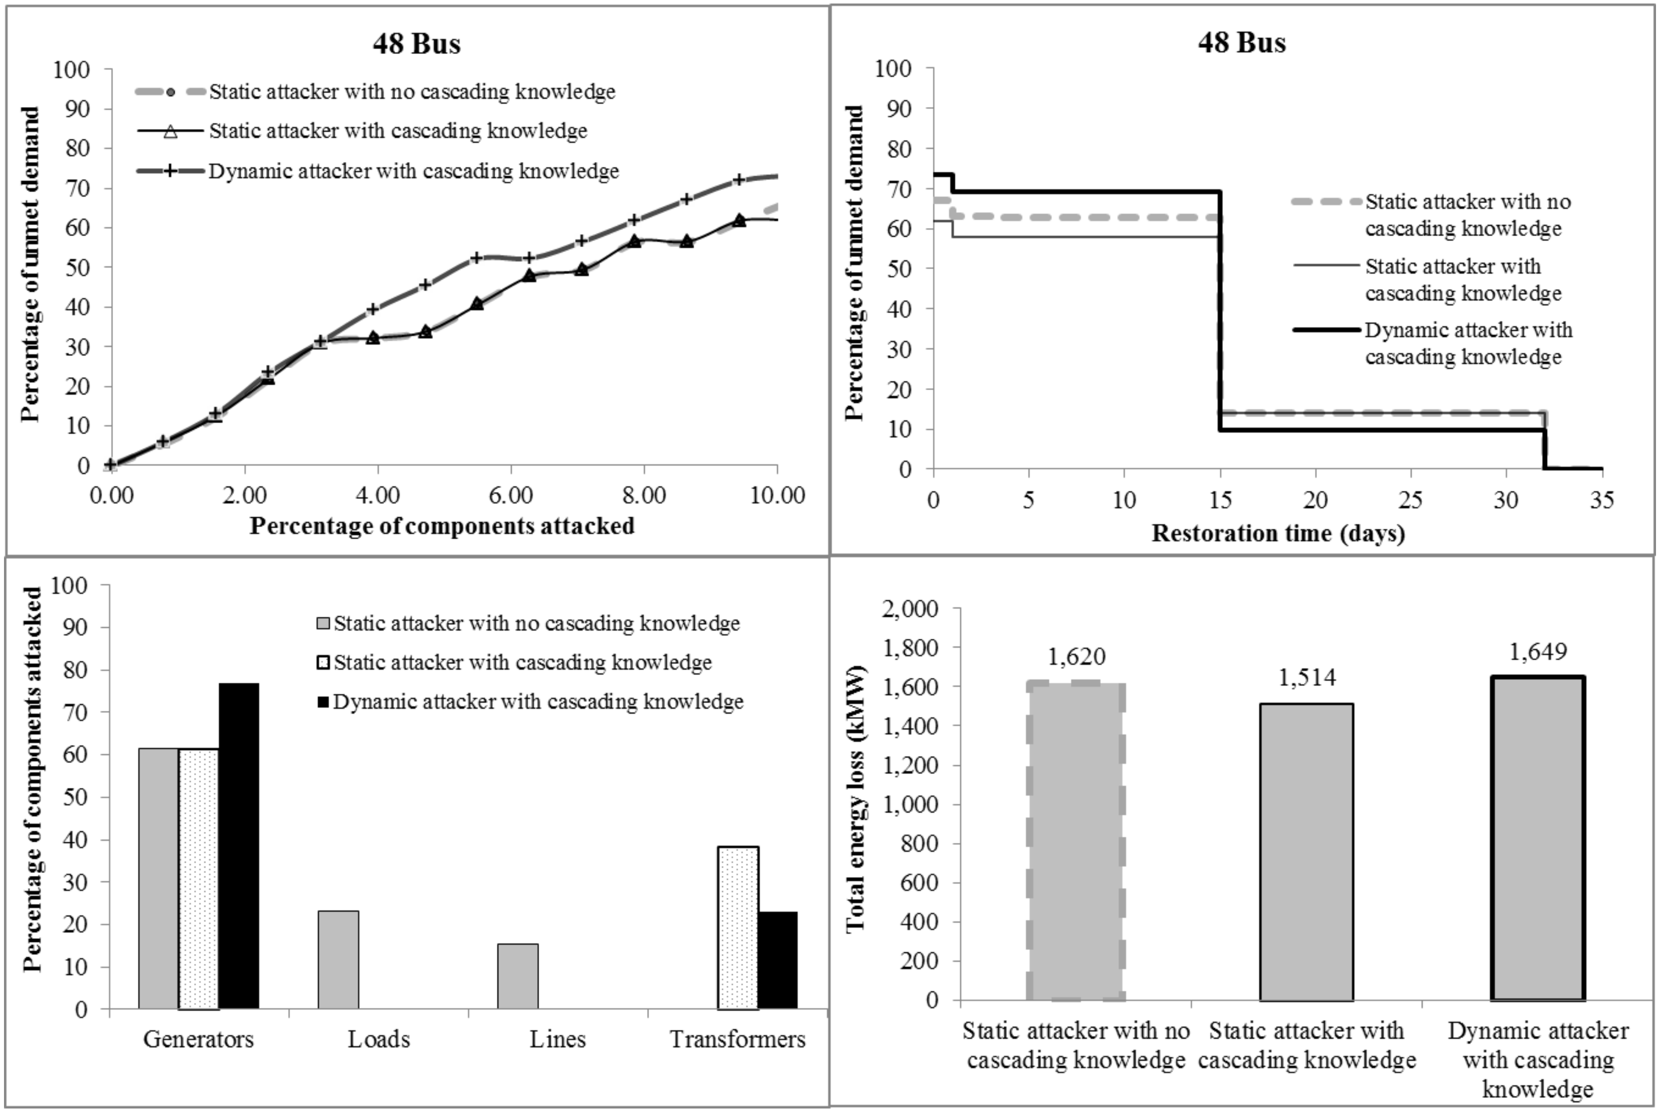

Appendix K. Comparison of three attacker assumptions with cascading knowledge when they consider restoration times (118-bus system)

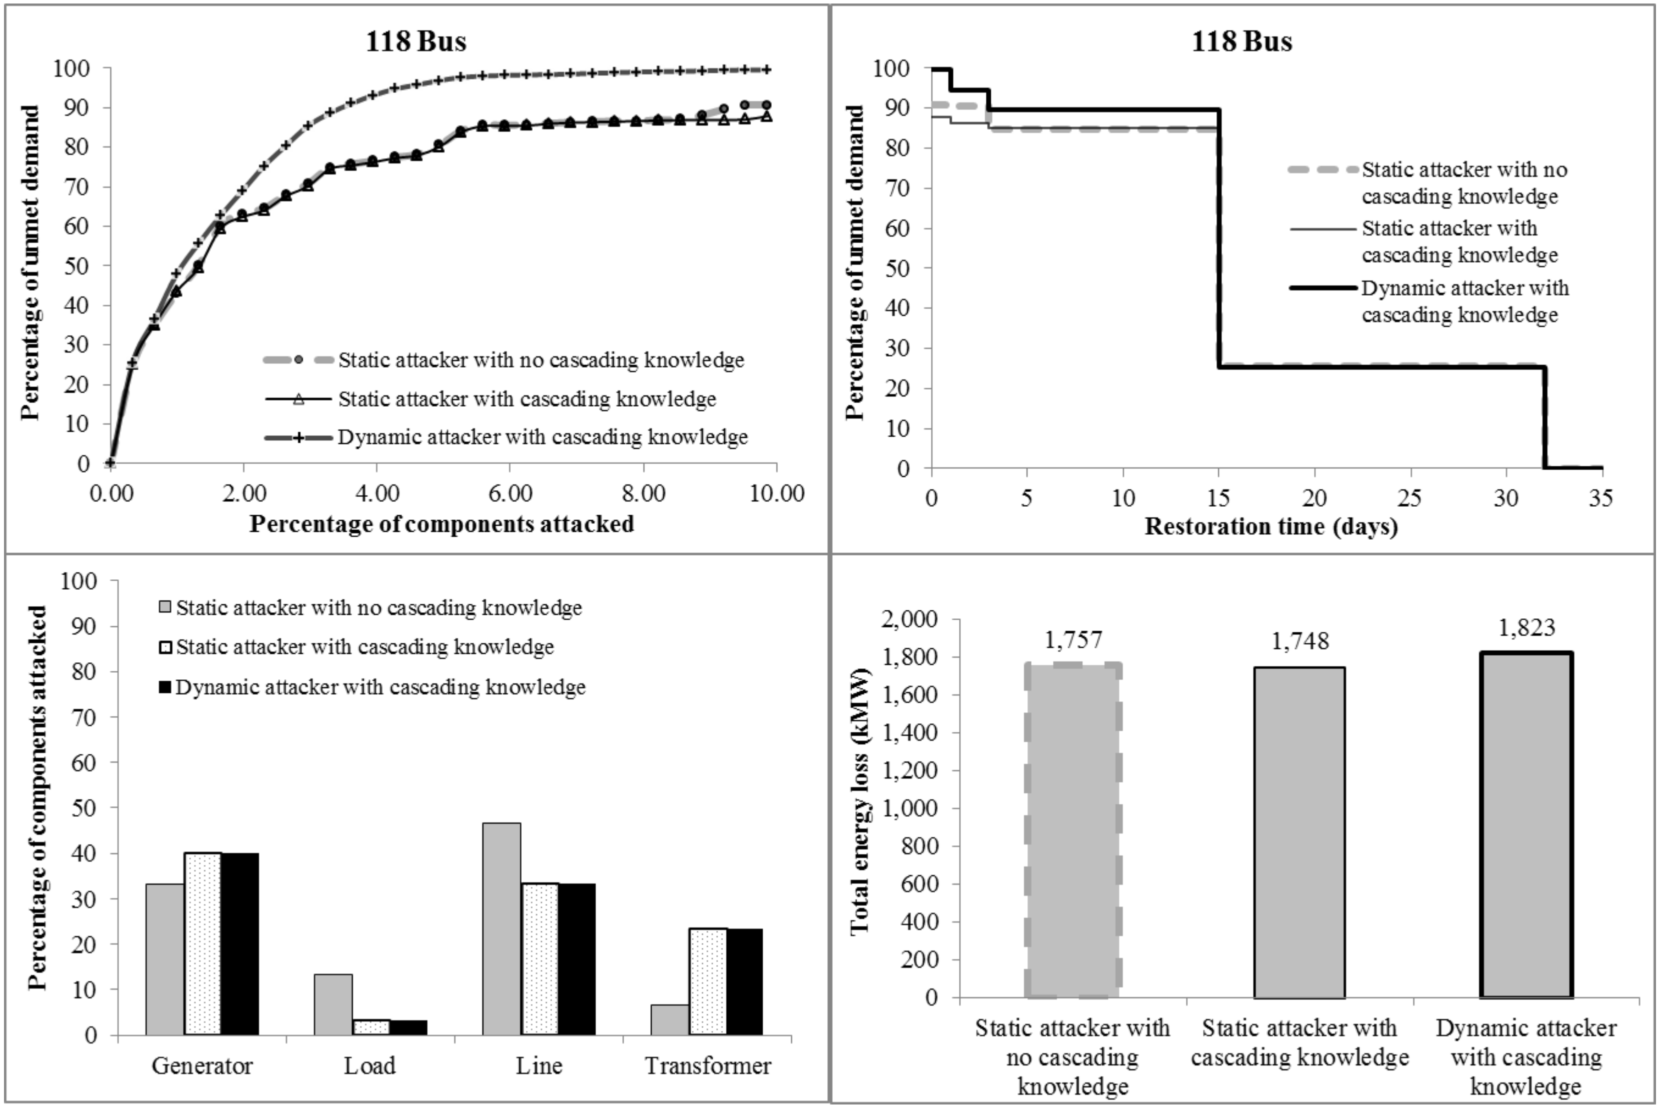

Appendix K. Comparison of three attacker assumptions with cascading knowledge when they consider restoration times (300-bus system)

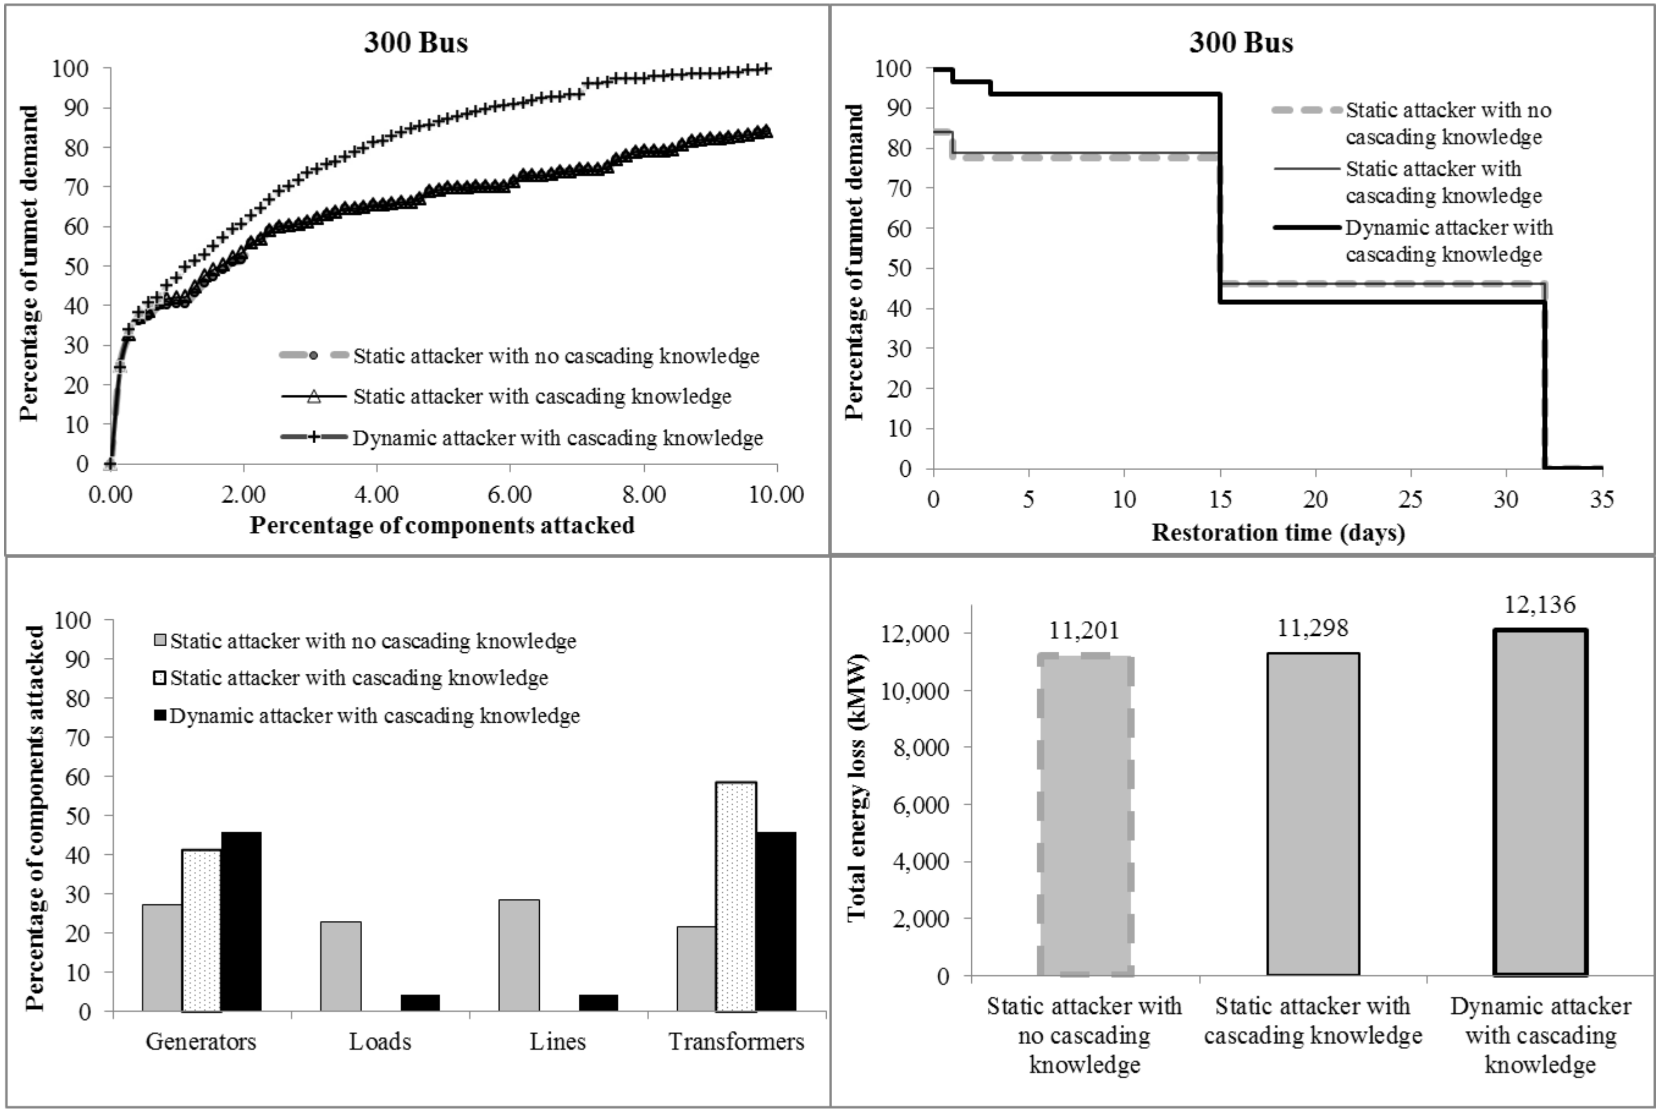

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with no cascading knowledge, 24-bus system)

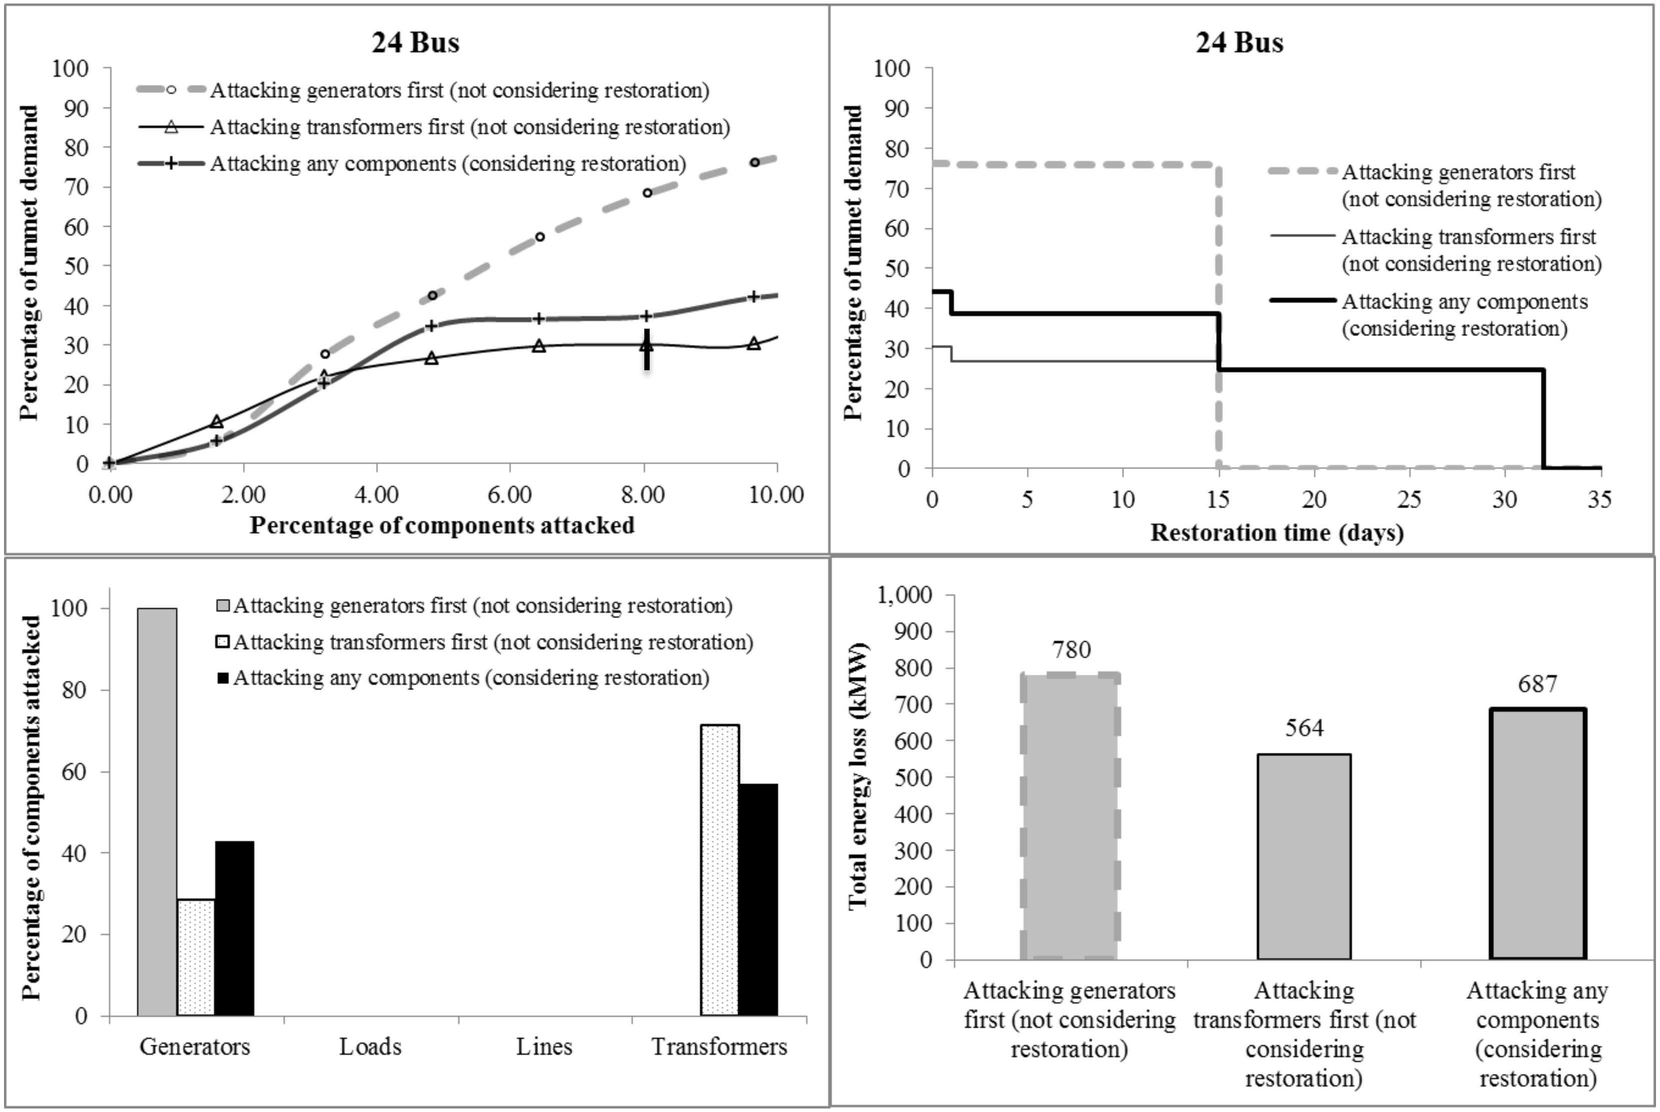

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with no cascading knowledge, 48-bus system)

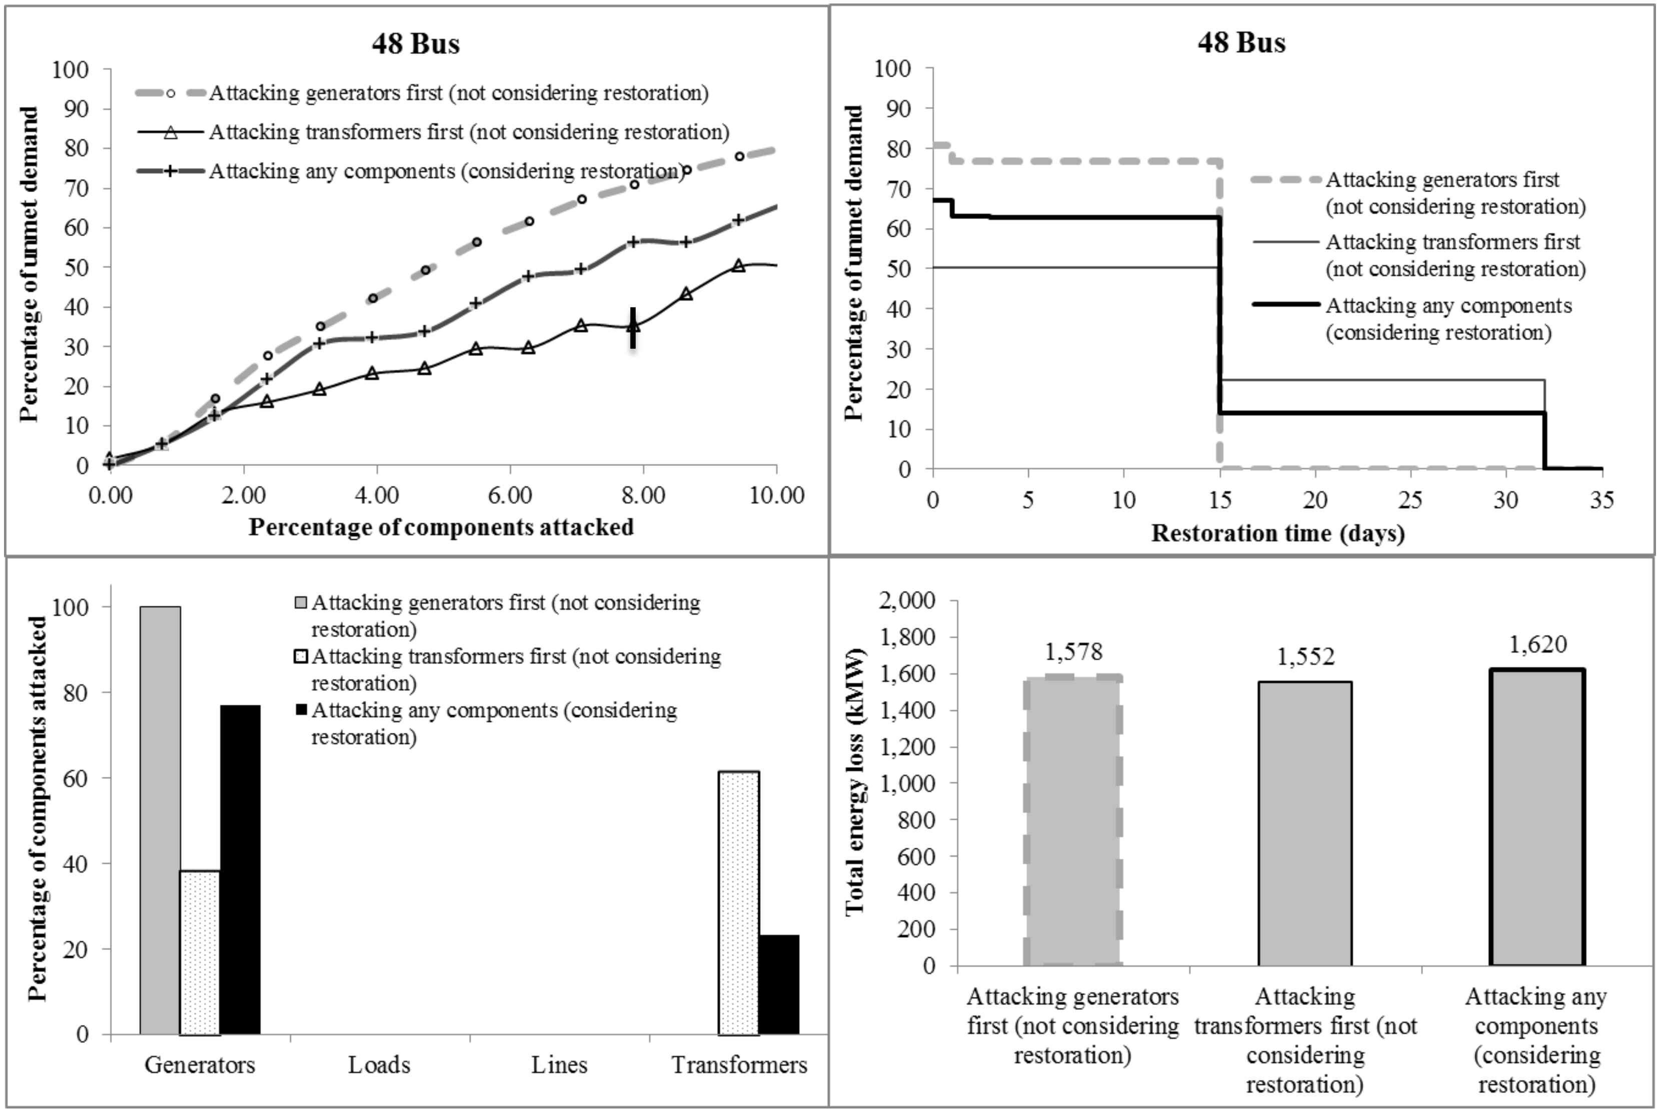

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with no cascading knowledge, 118-bus system)

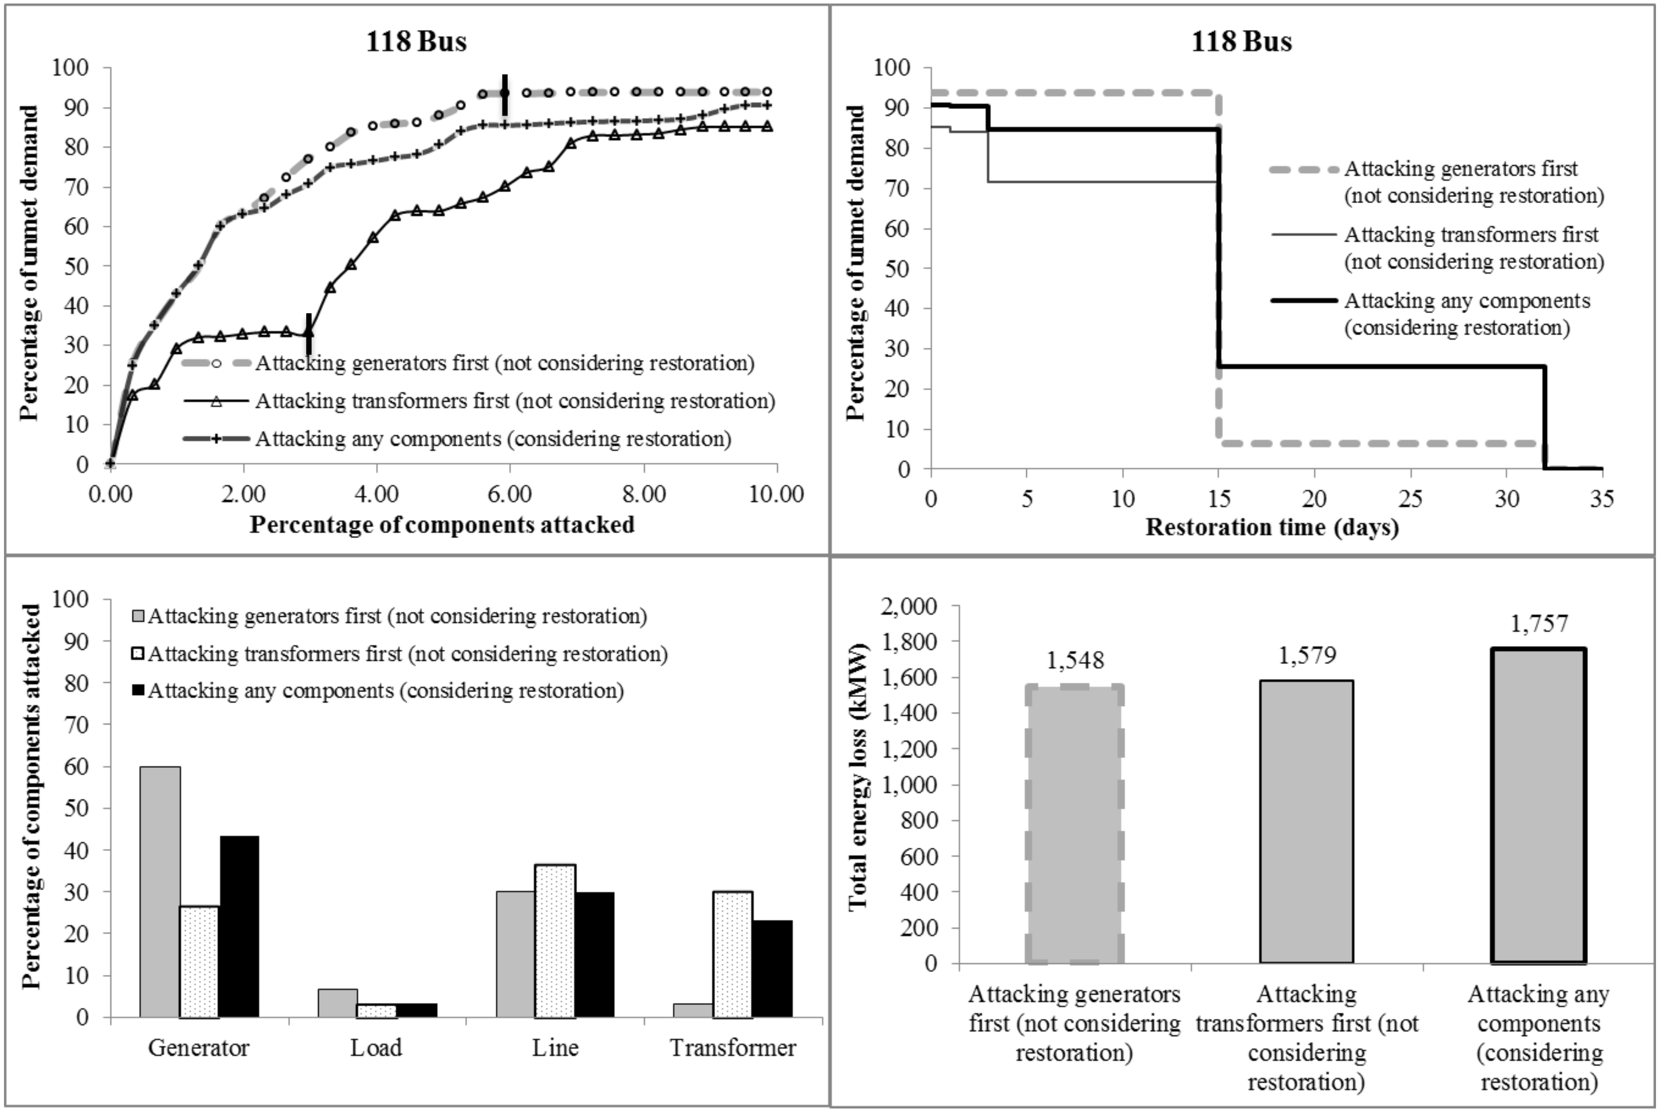

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with no cascading knowledge, 300-bus system)

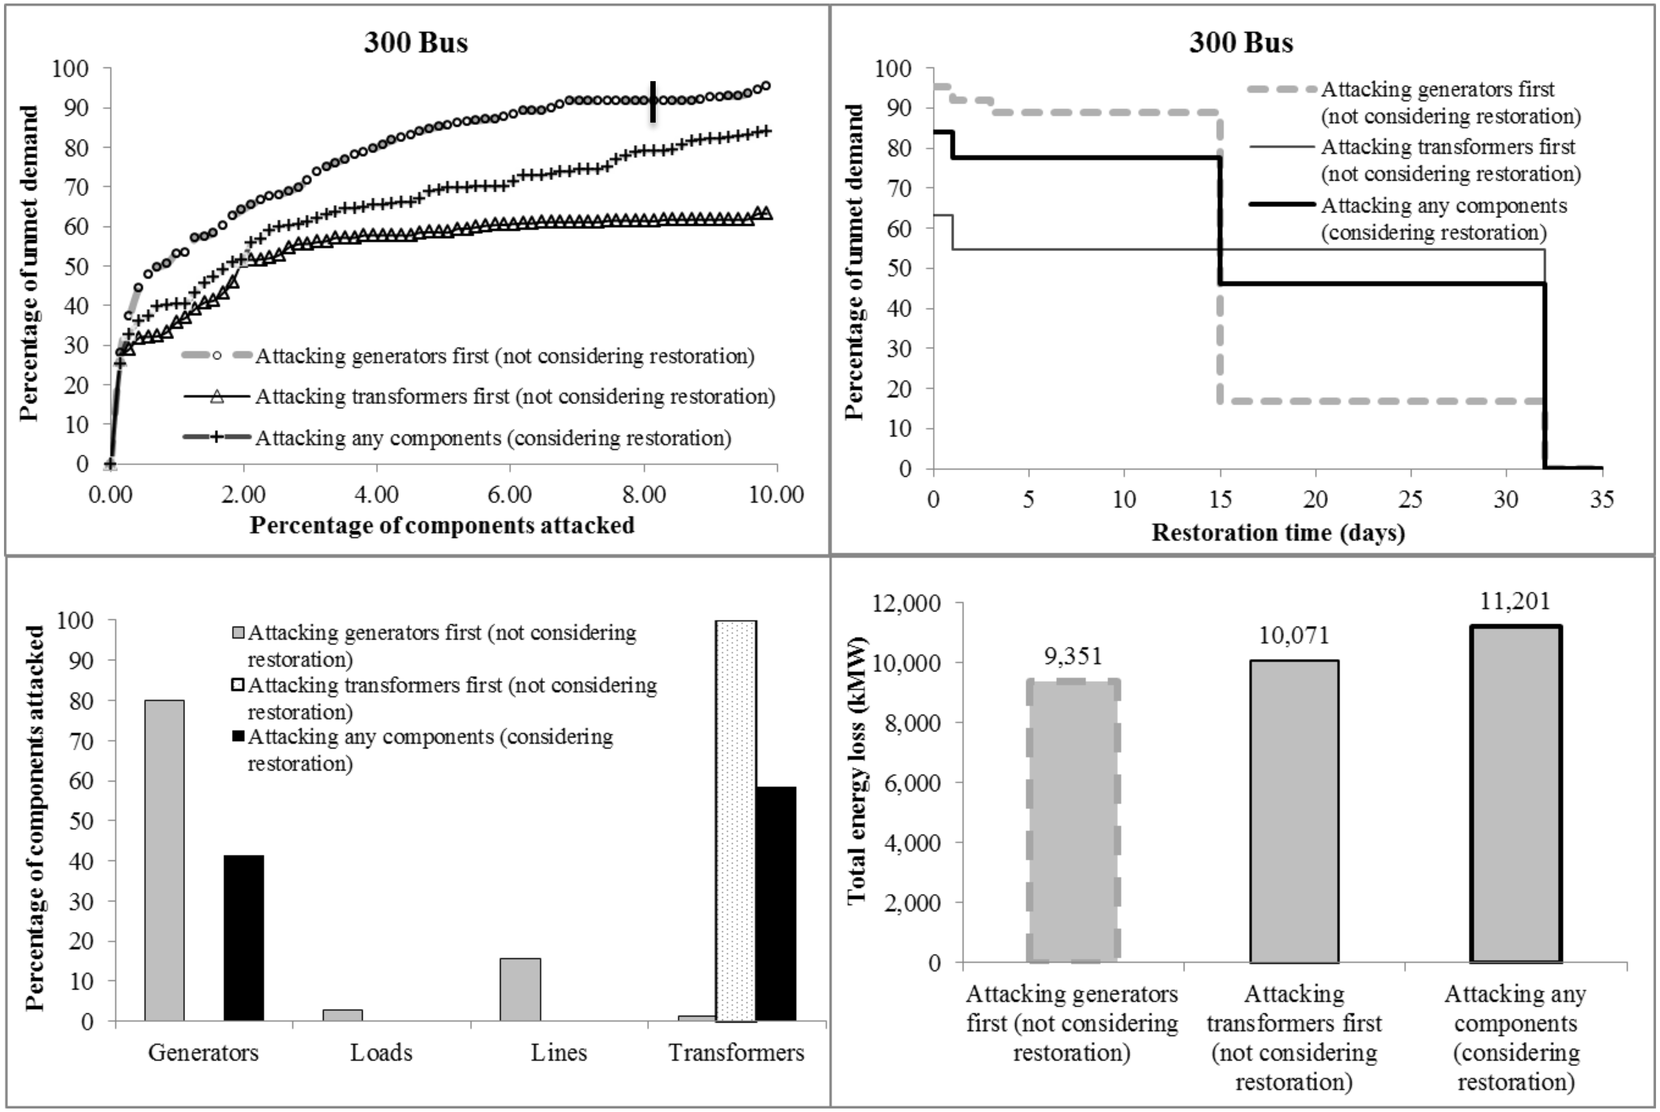

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with cascading knowledge, 24-bus system)

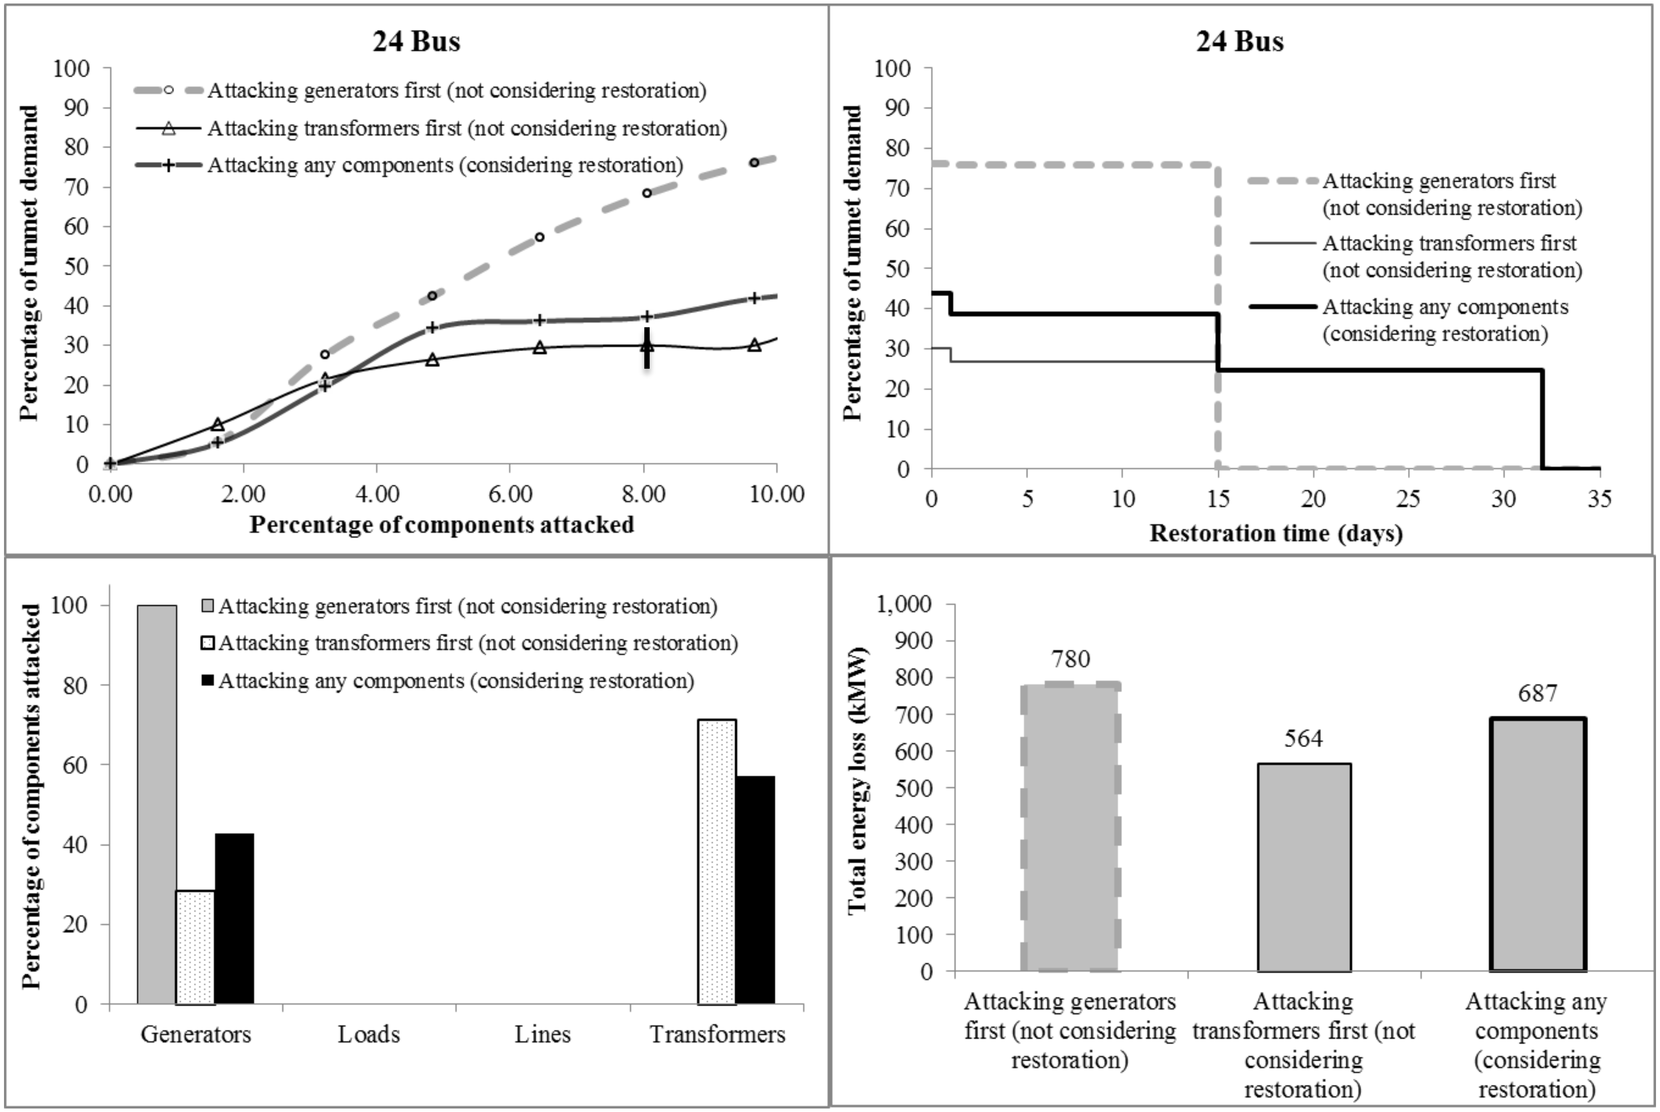

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with cascading knowledge, 48-bus system)

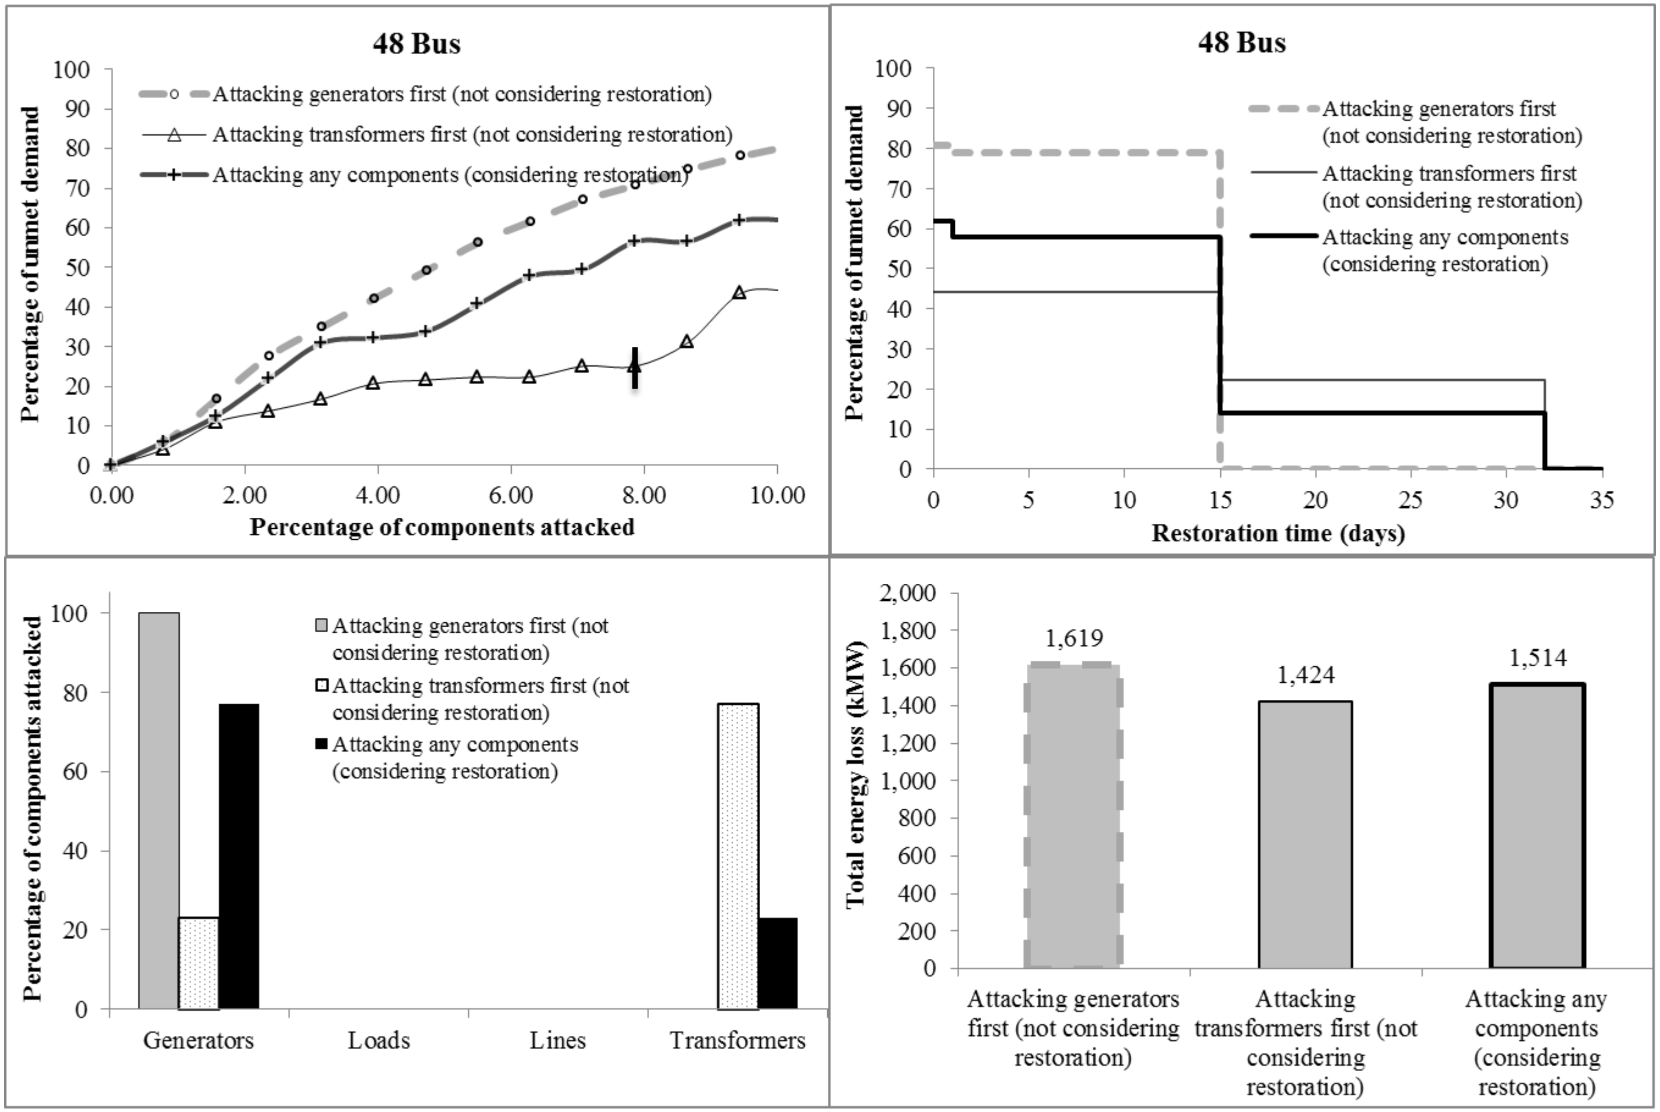

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with cascading knowledge, 118-bus system)

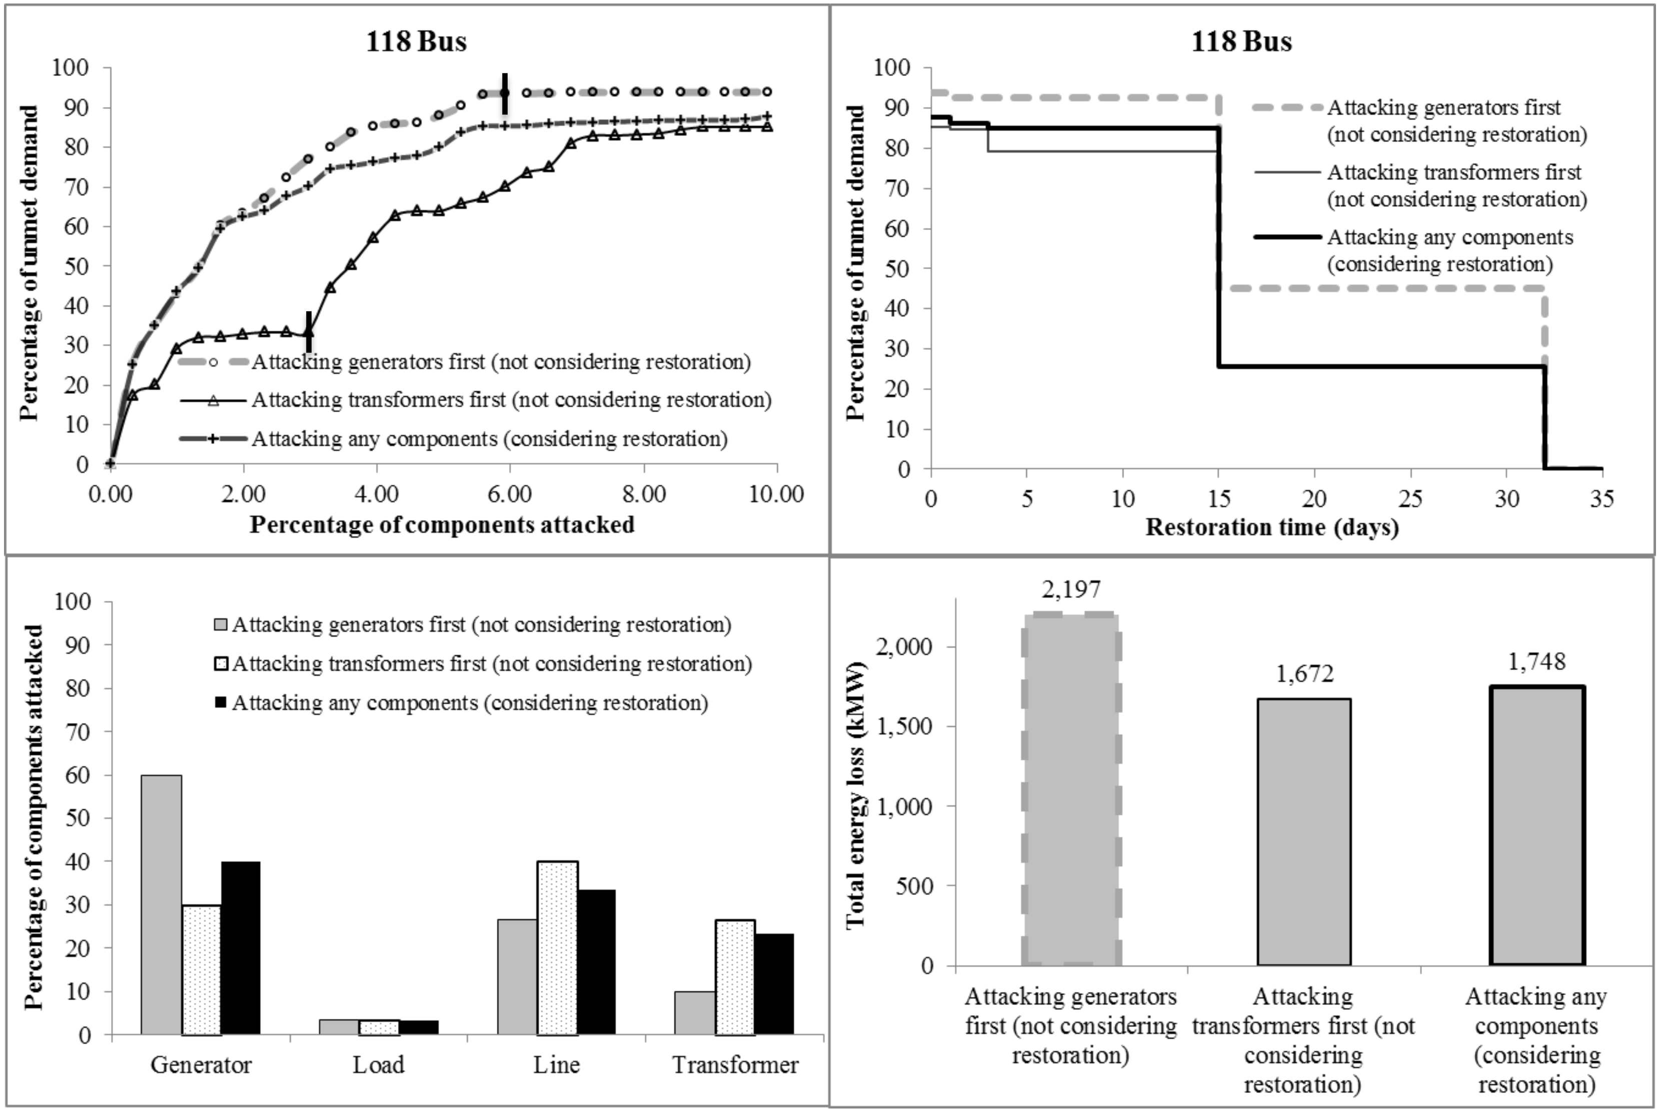

Appendix L. Comparison of an attacker attacking generators/transformers first, or an attacker considering restoration times (static attacker with cascading knowledge, 300-bus system)

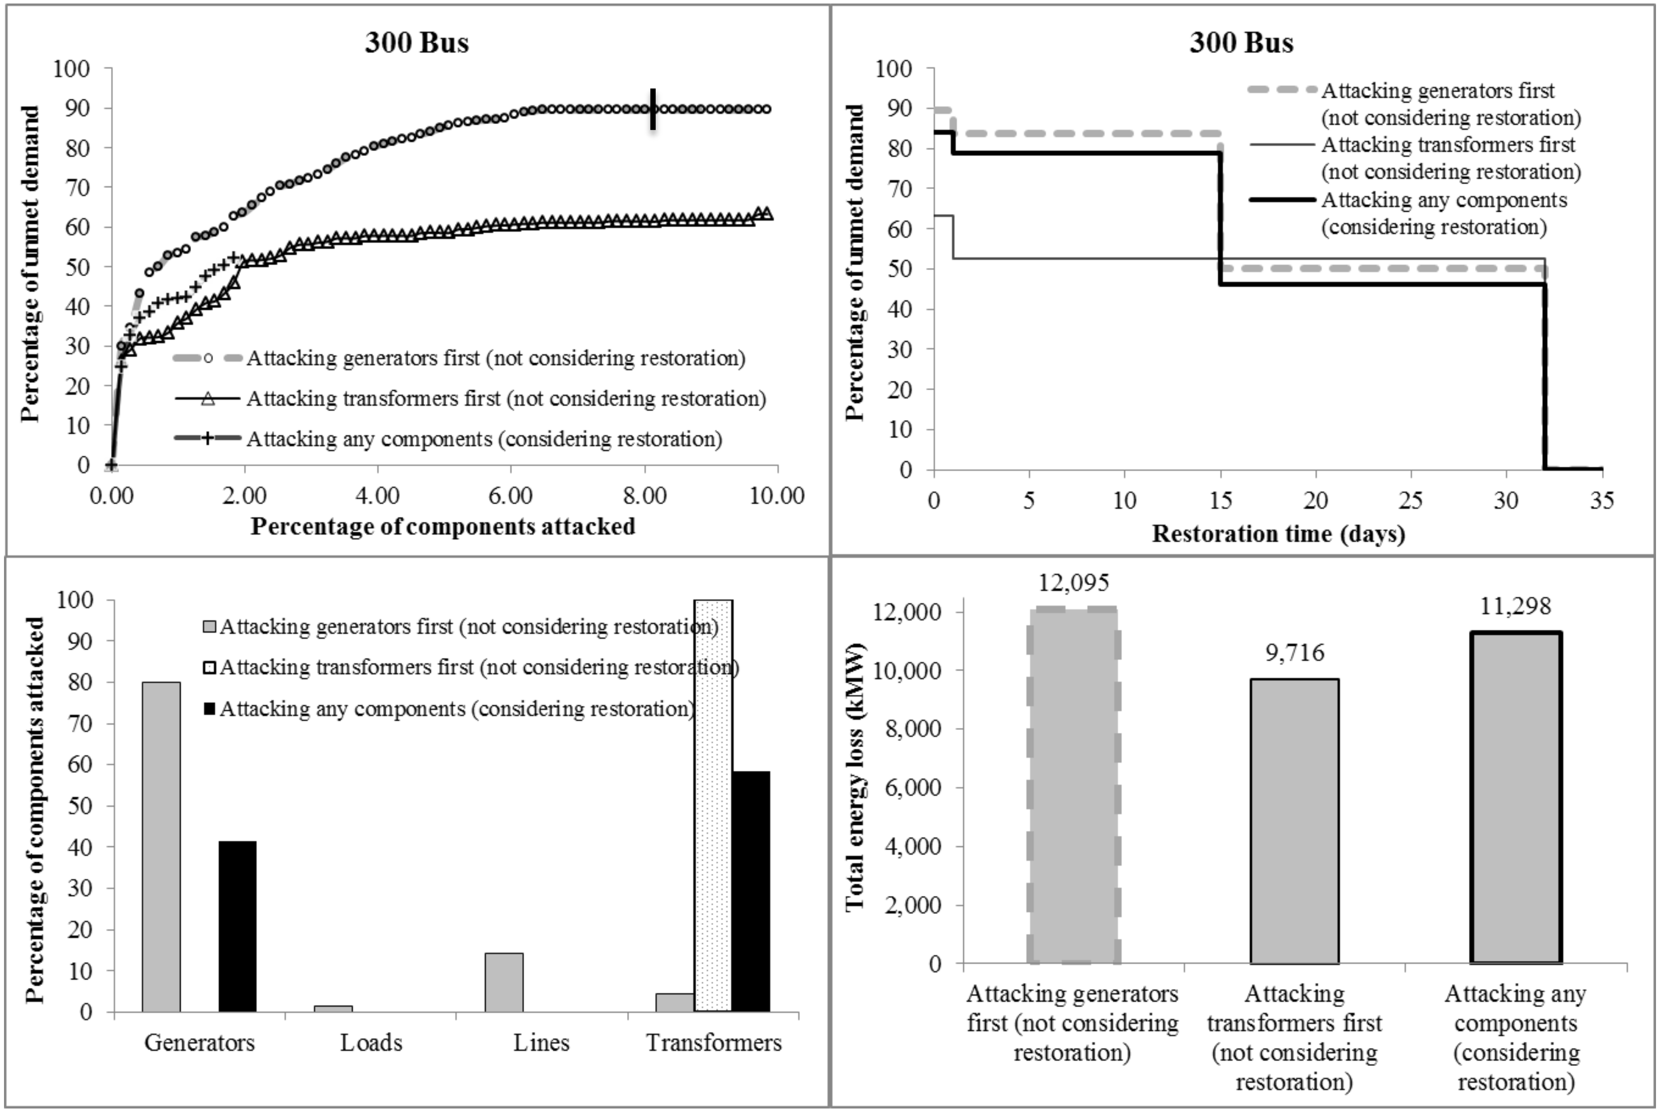

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a dynamic attacker who does not consider restoration times, 24-bus system)

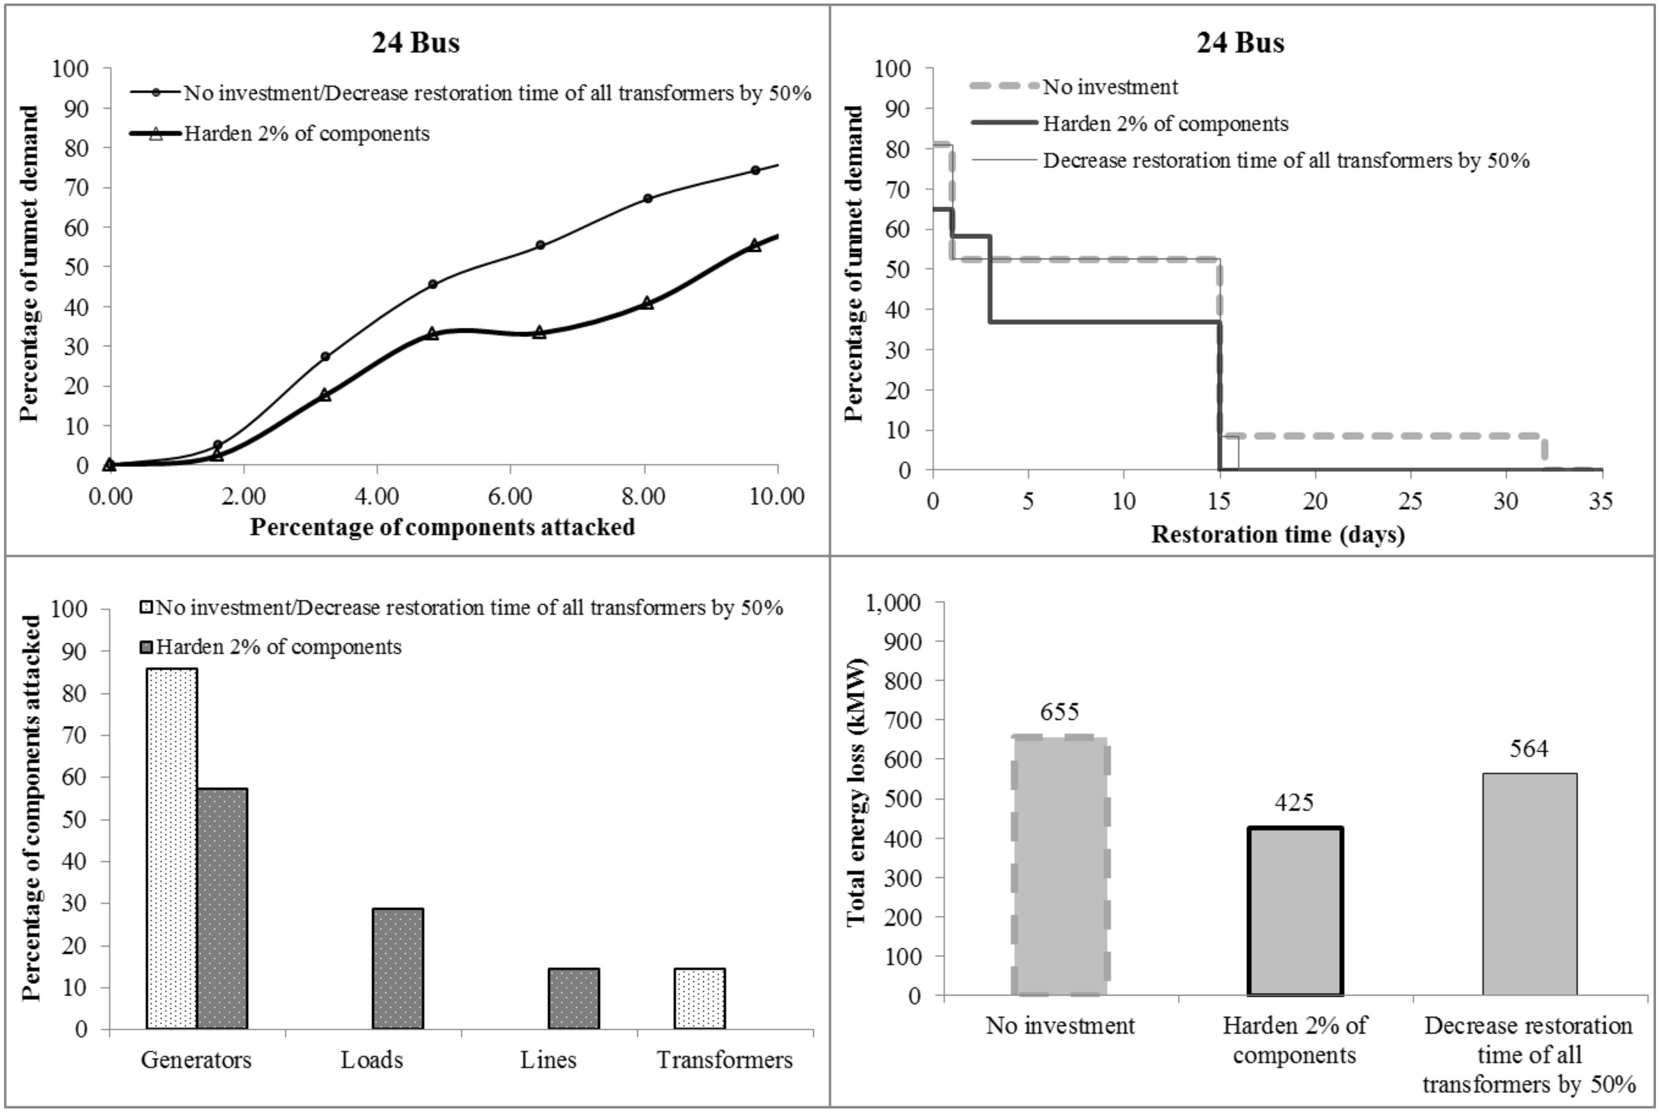

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 2% of the components against a dynamic attacker who does not consider restoration times, 48-bus system)

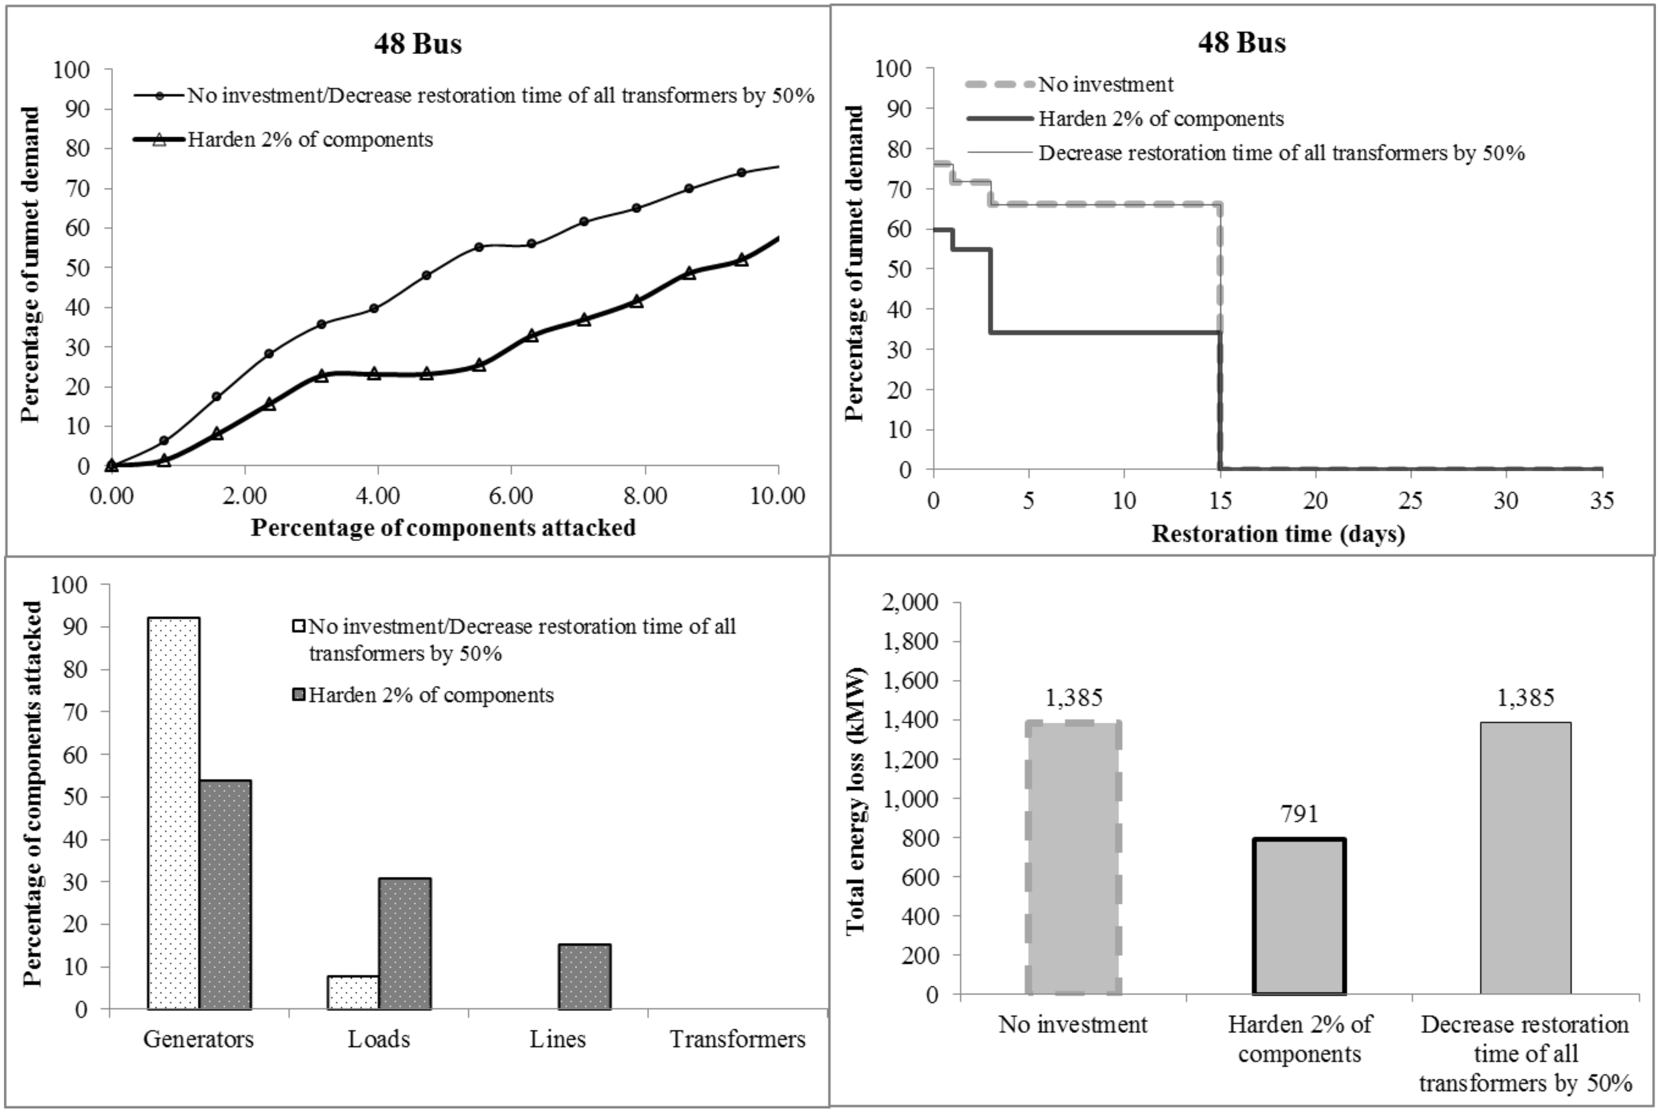

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a dynamic attacker who does not consider restoration times, 118-bus system)

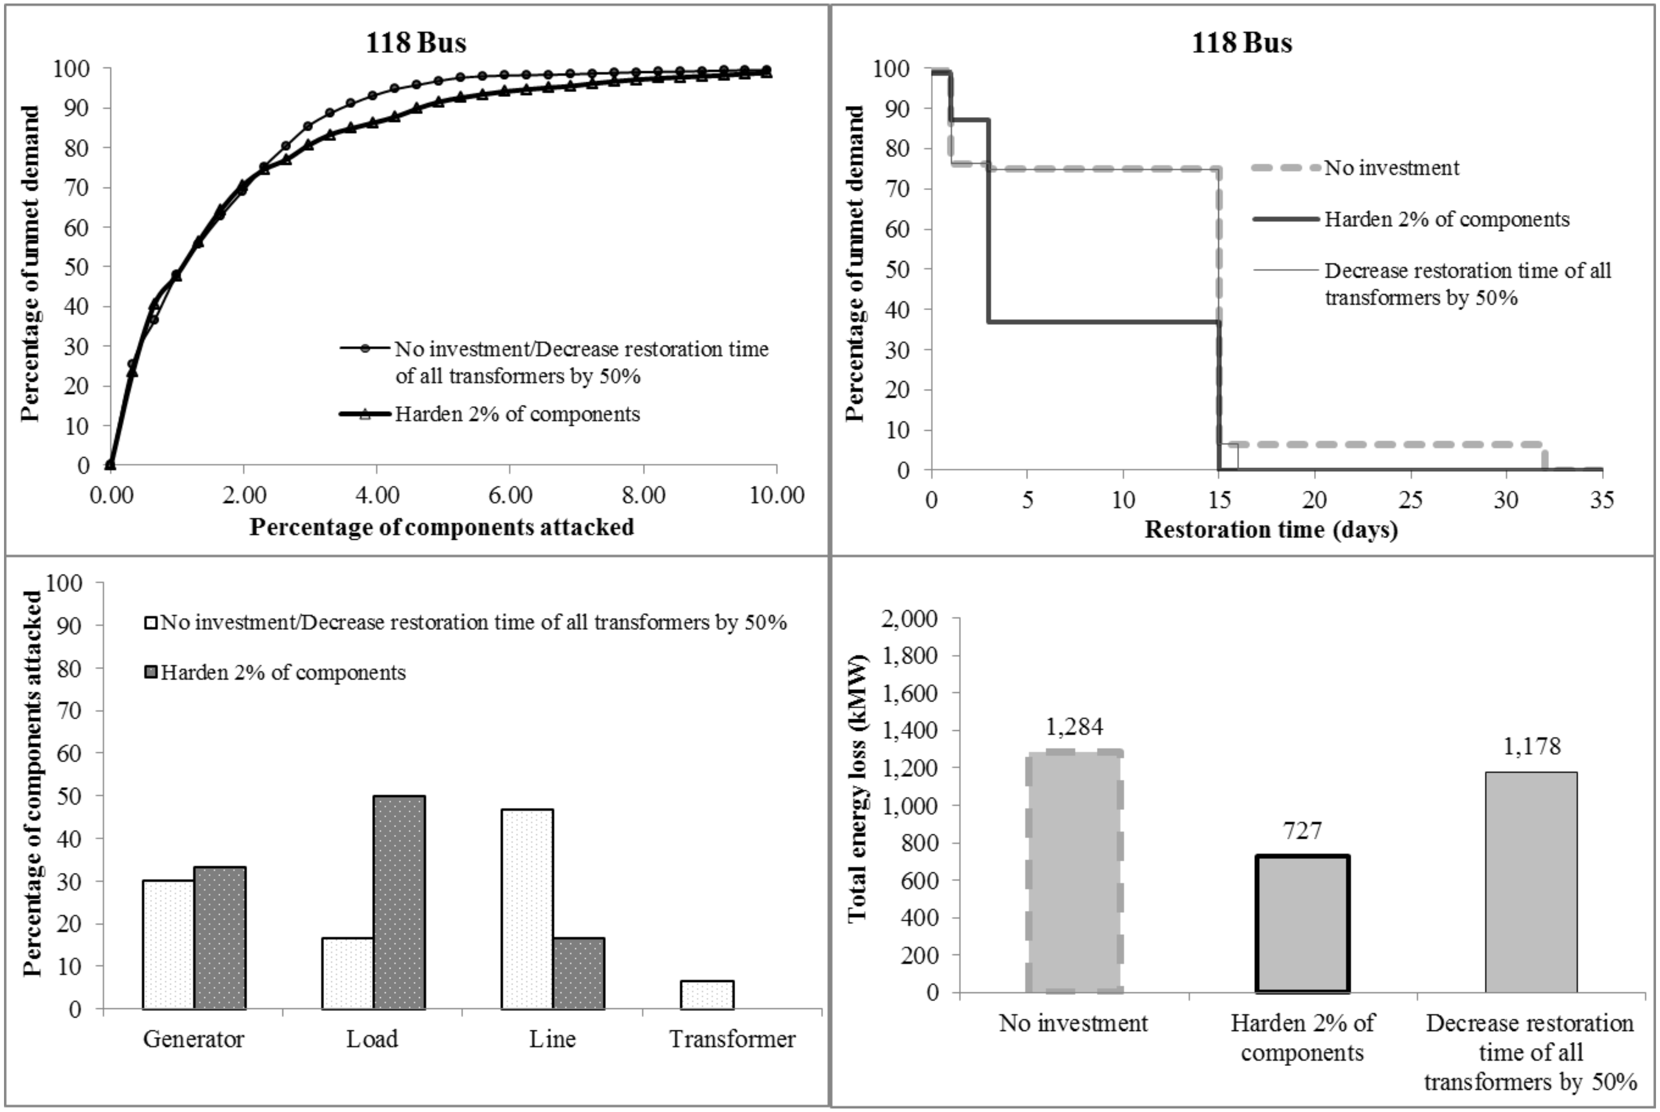

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a dynamic attacker who does not consider restoration times, 300-bus system)

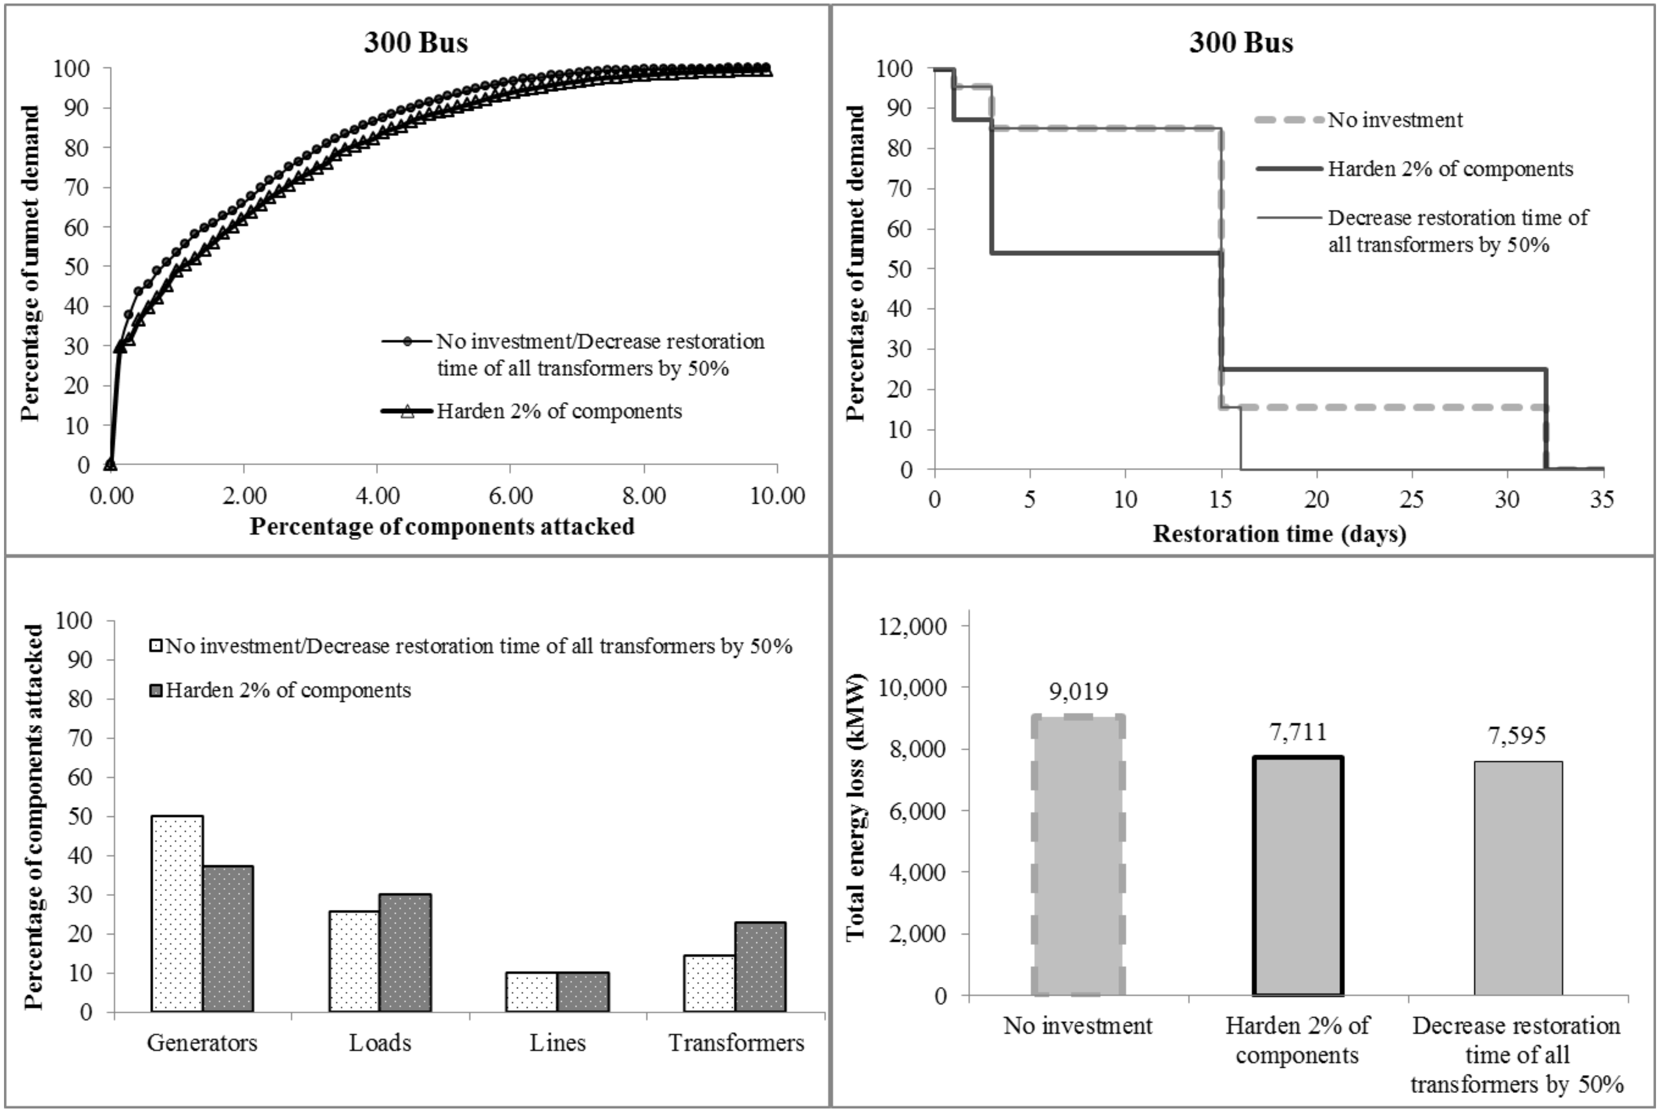

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 5% of the components against a dynamic attacker who does not consider restoration times, 24-bus system)

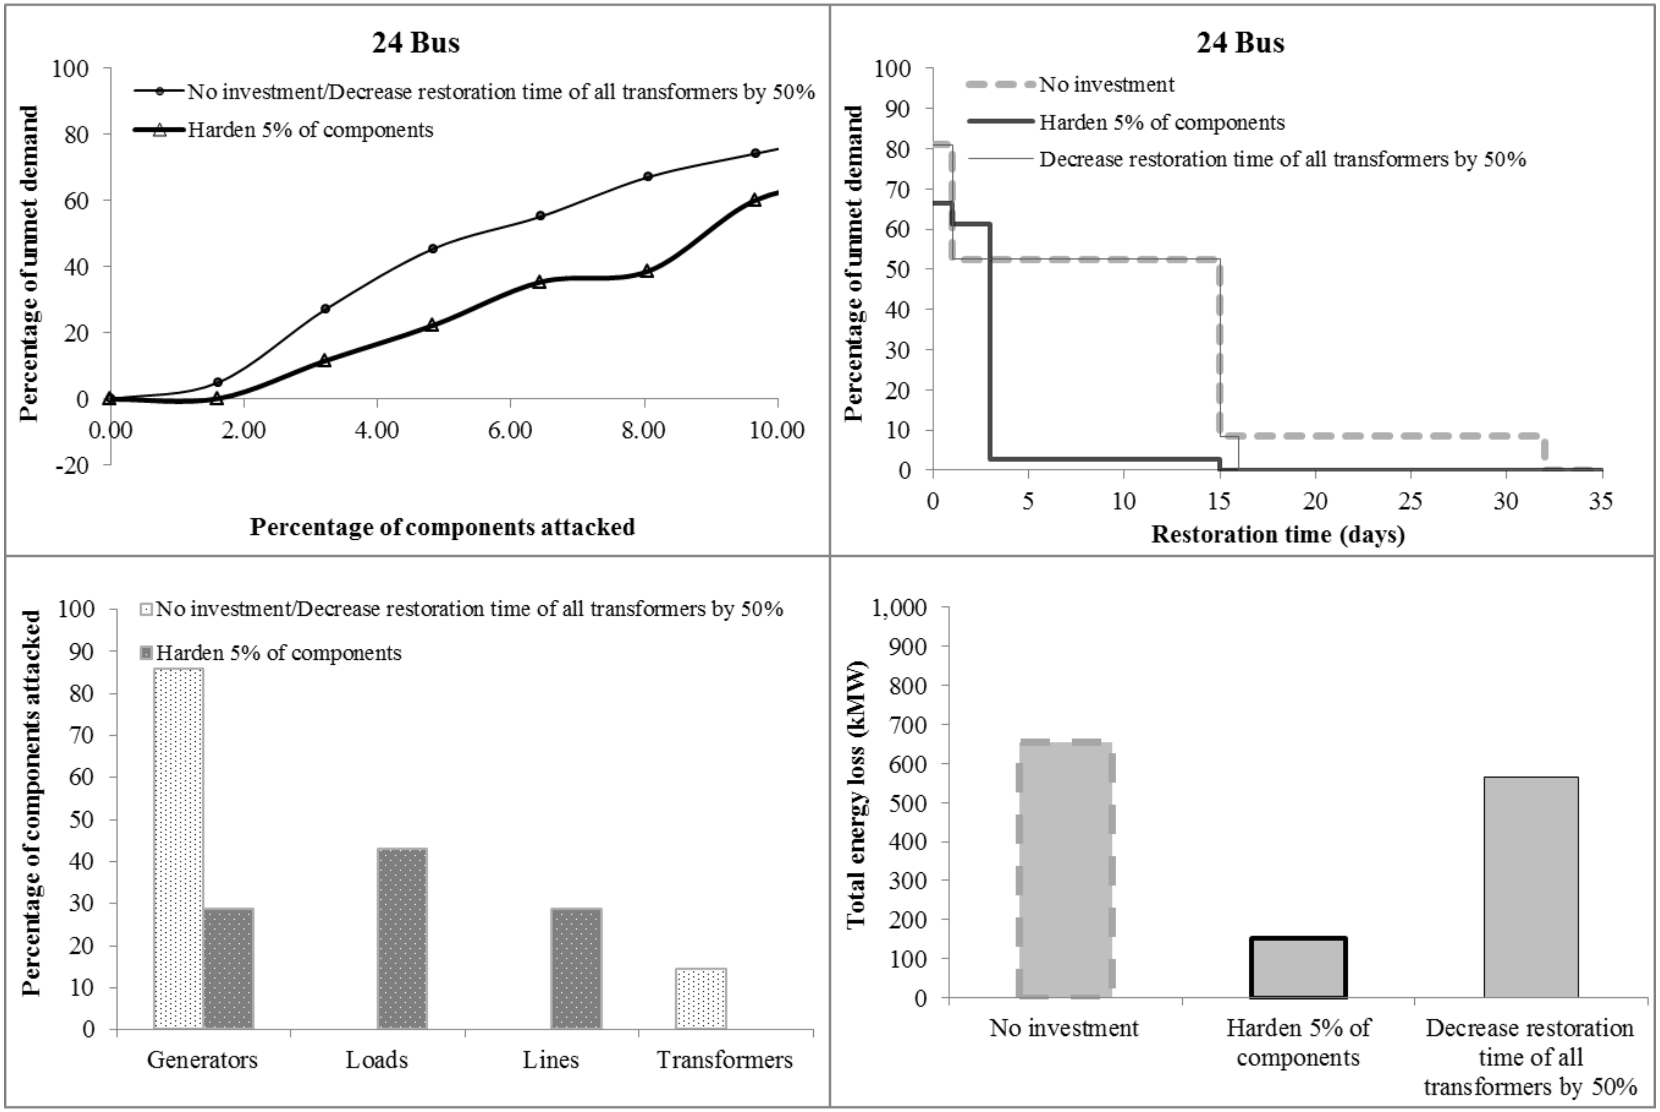

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 5% of the components against a dynamic attacker who does not consider restoration times, 48-bus system)

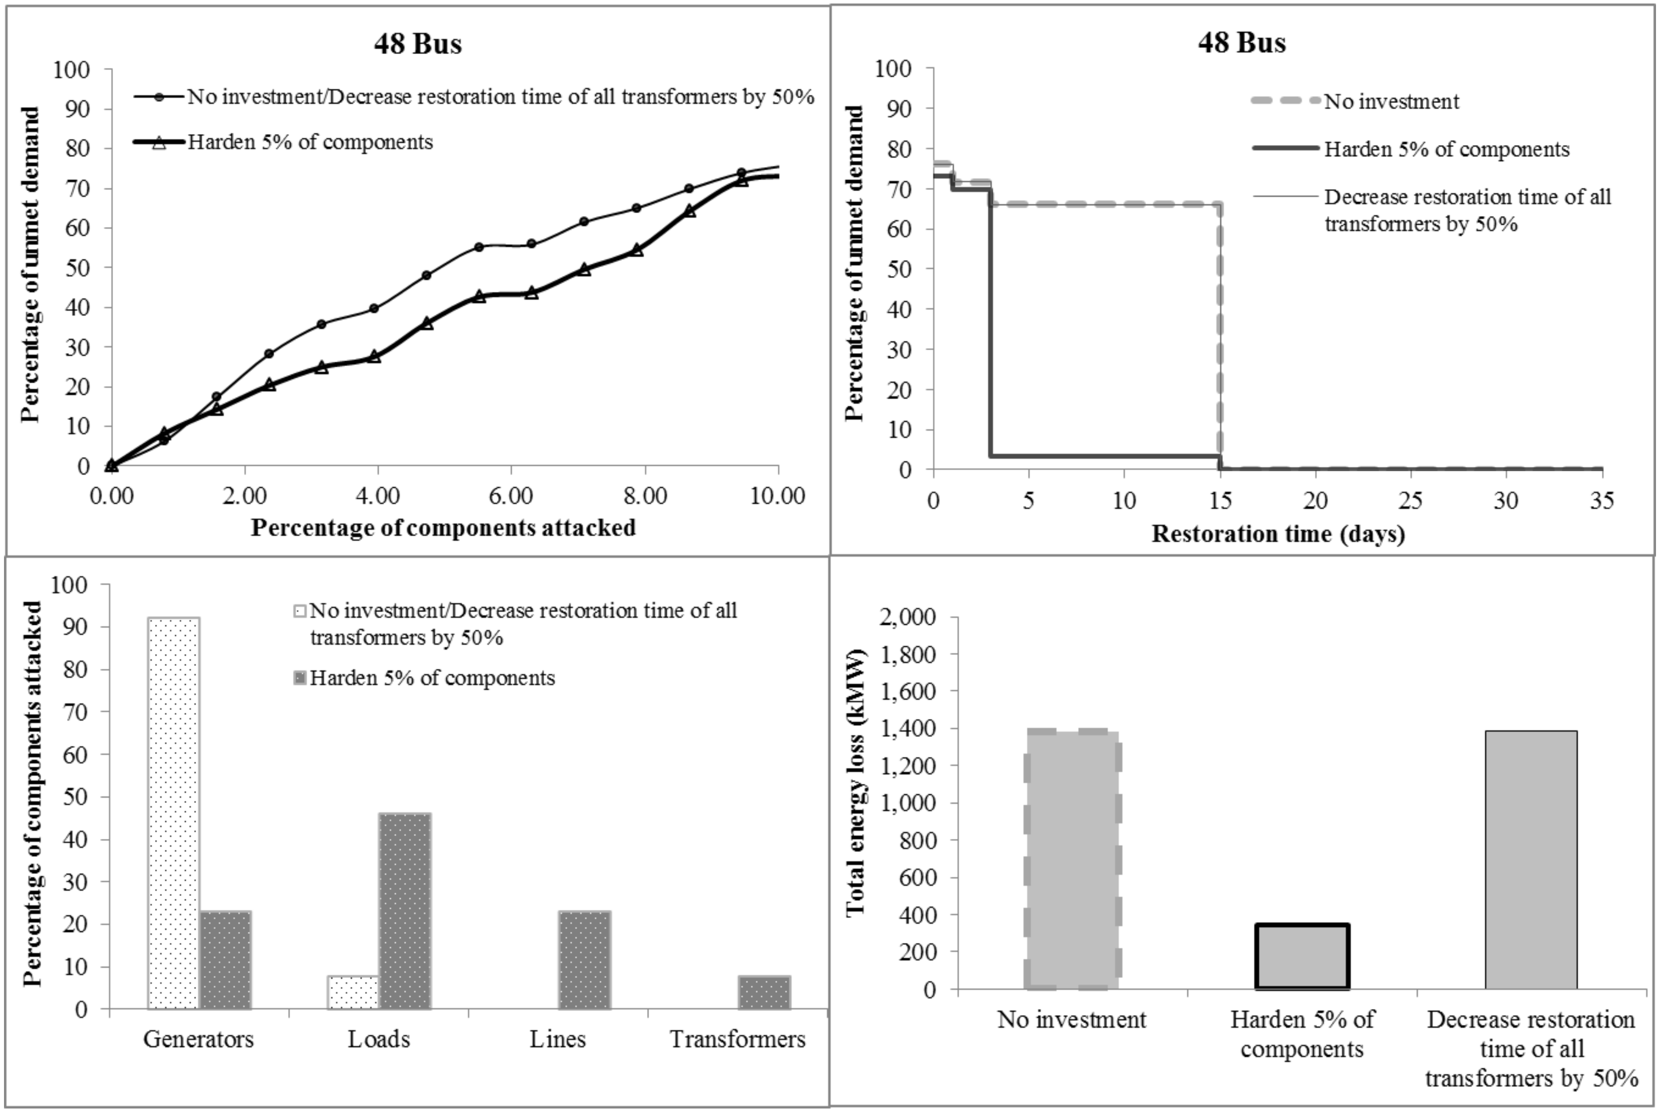

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a dynamic attacker who does not consider restoration times, 118-bus system)

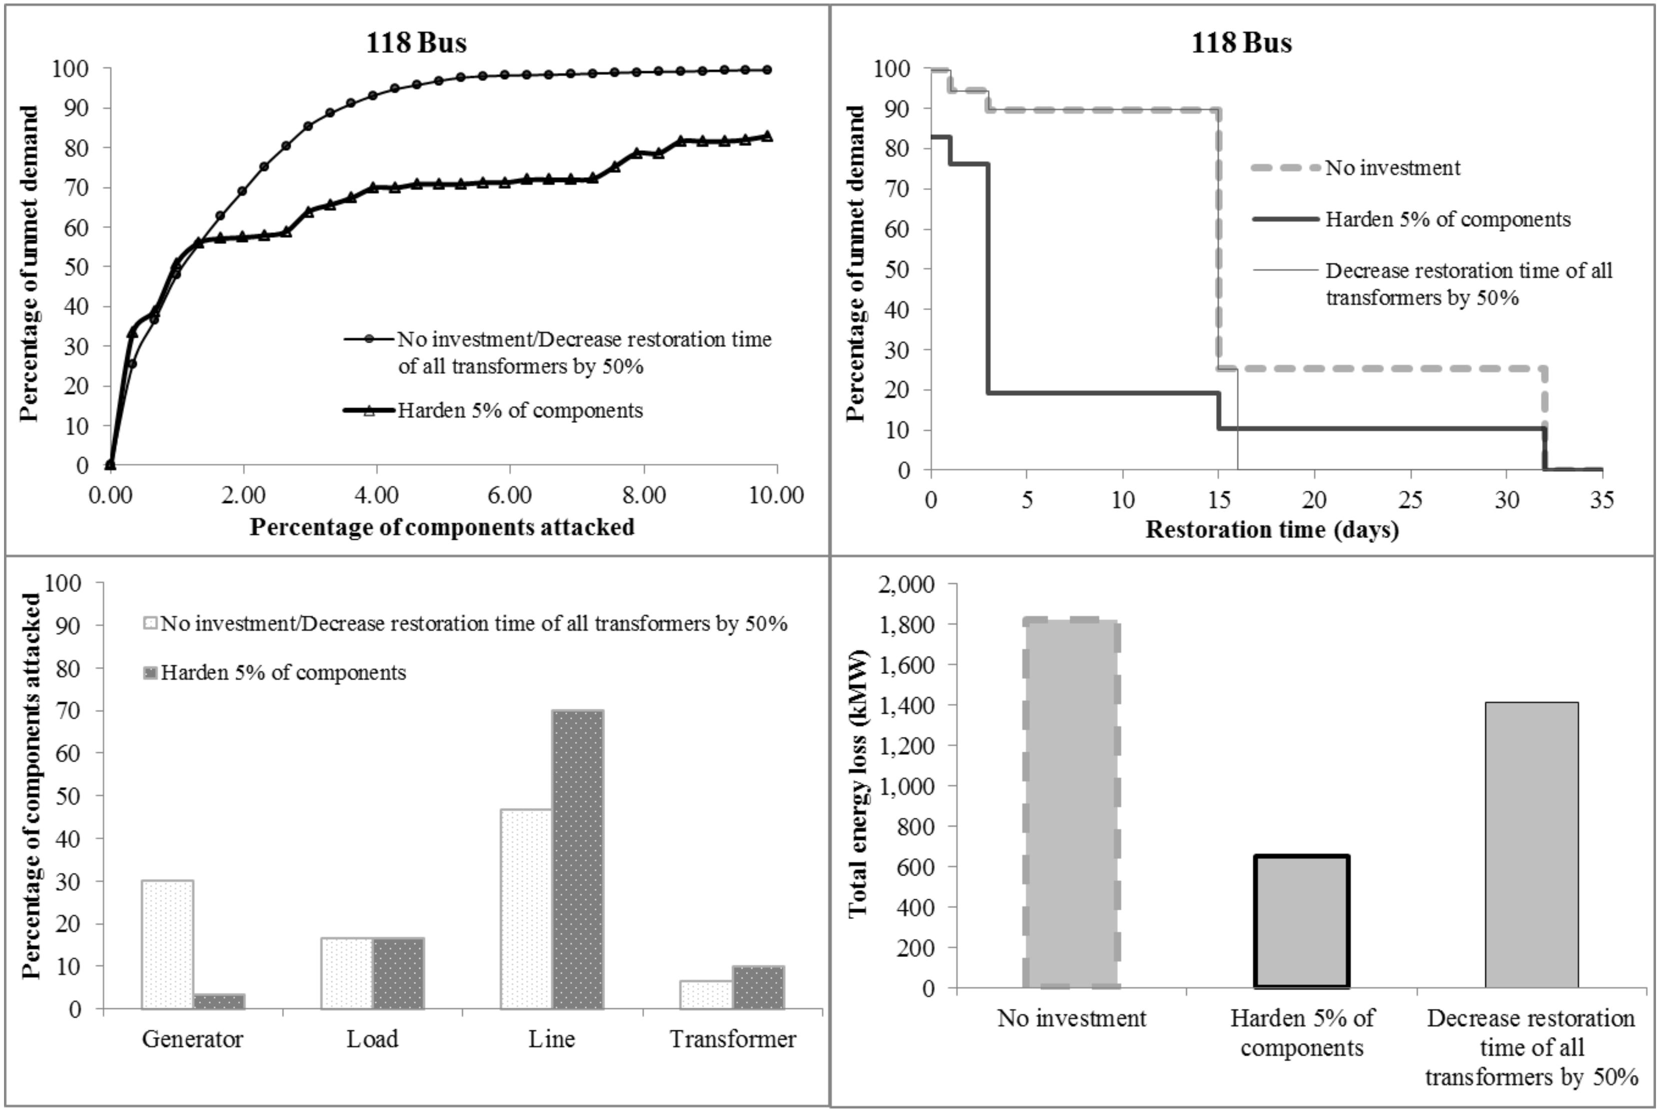

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a dynamic attacker who does not consider restoration times, 300-bus system)

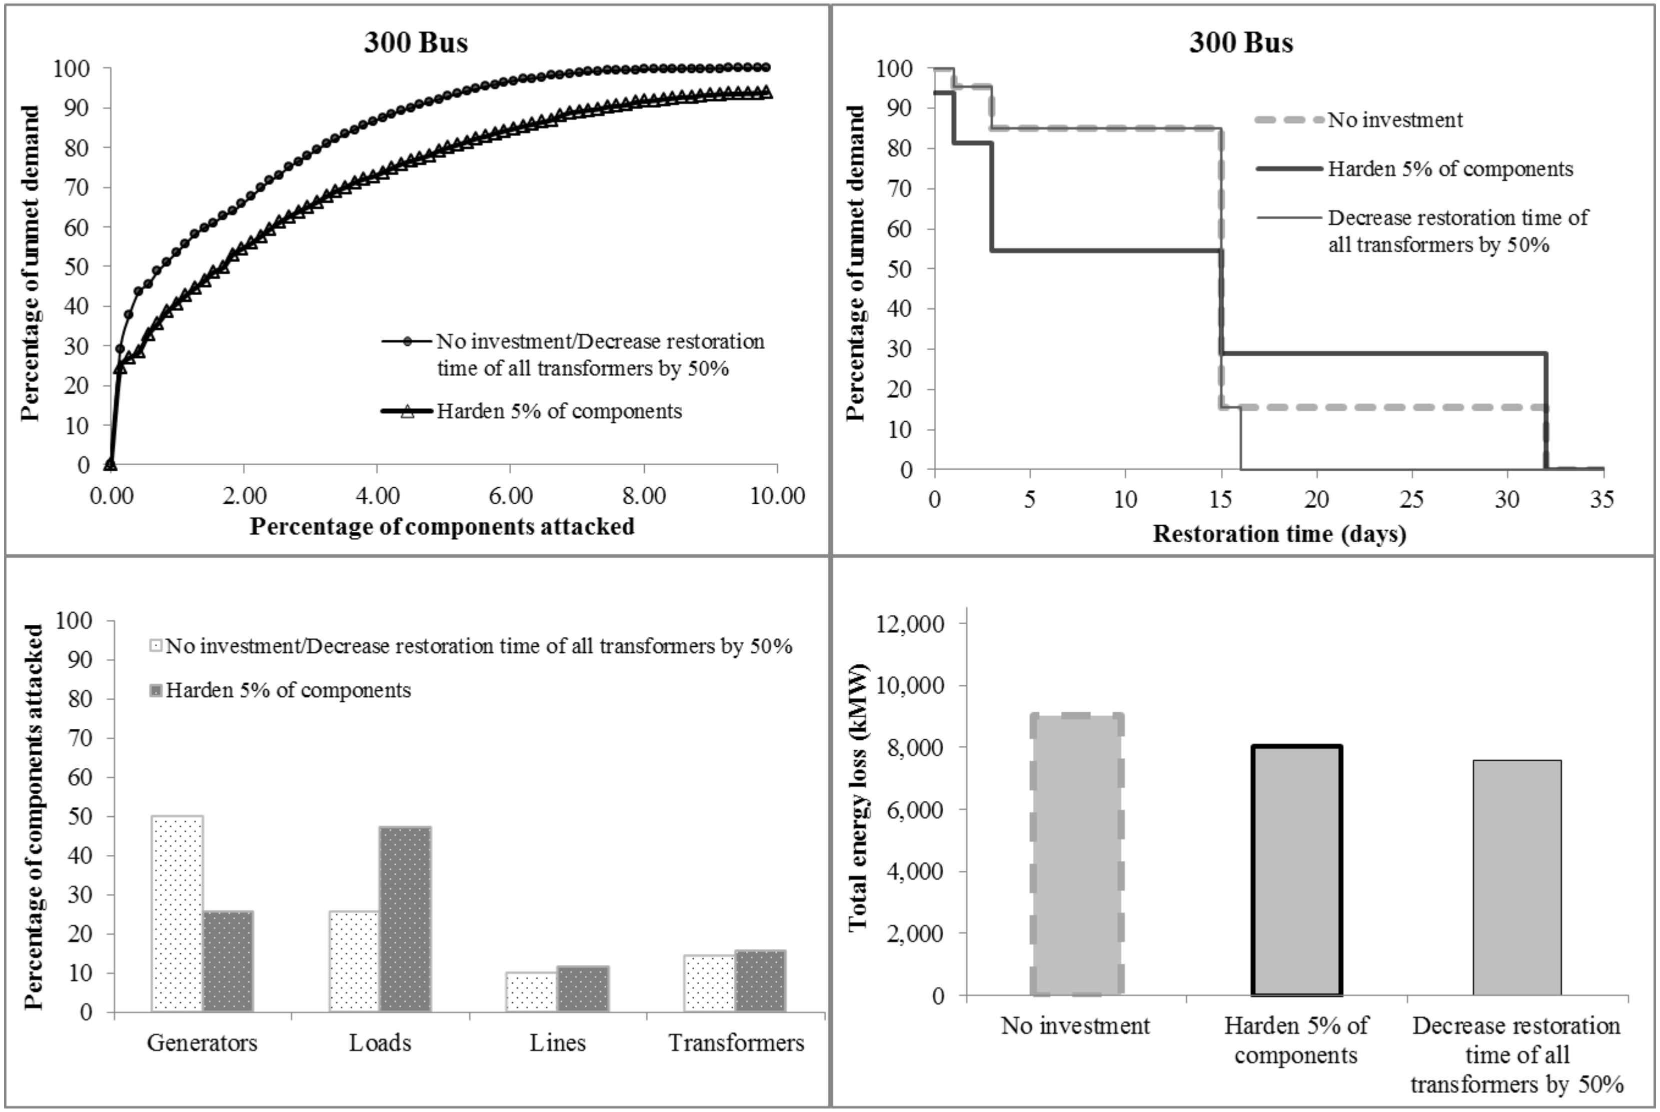

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who does not consider restoration times, 24-bus system)

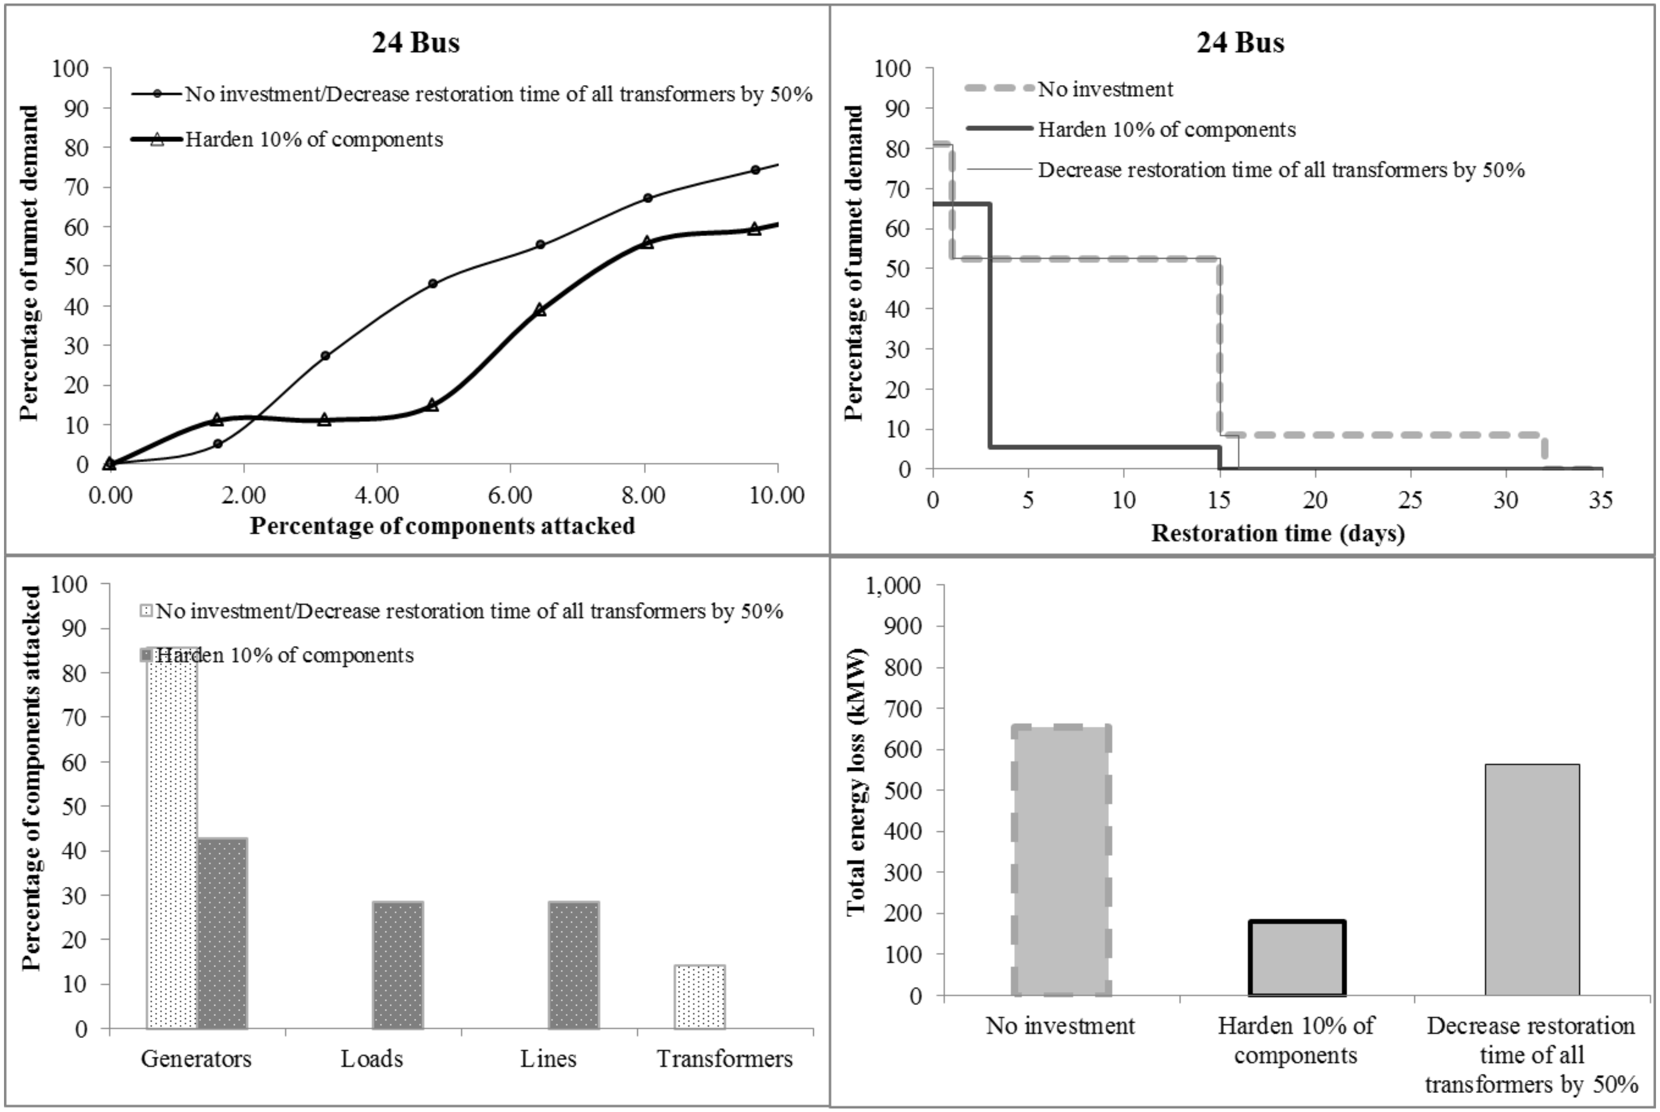

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who does not consider restoration times, 48-bus system)

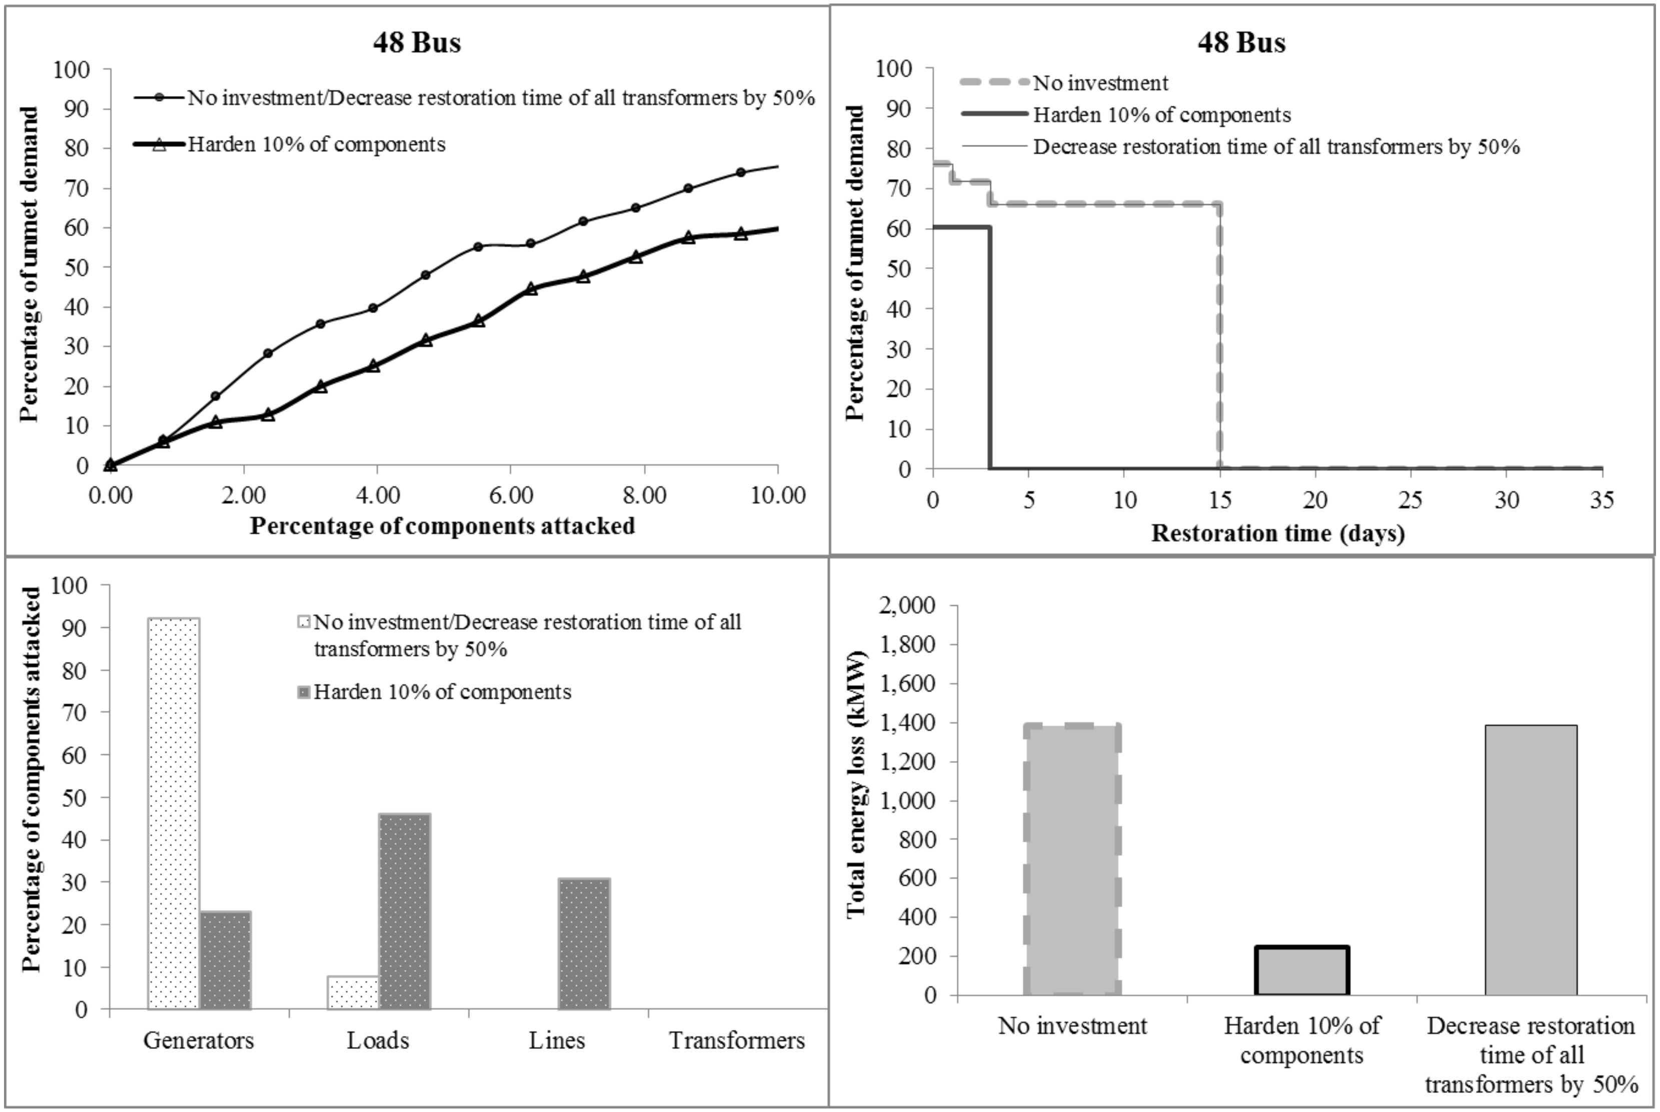

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 10% of the components against a dynamic attacker who does not consider restoration times, 118-bus system)

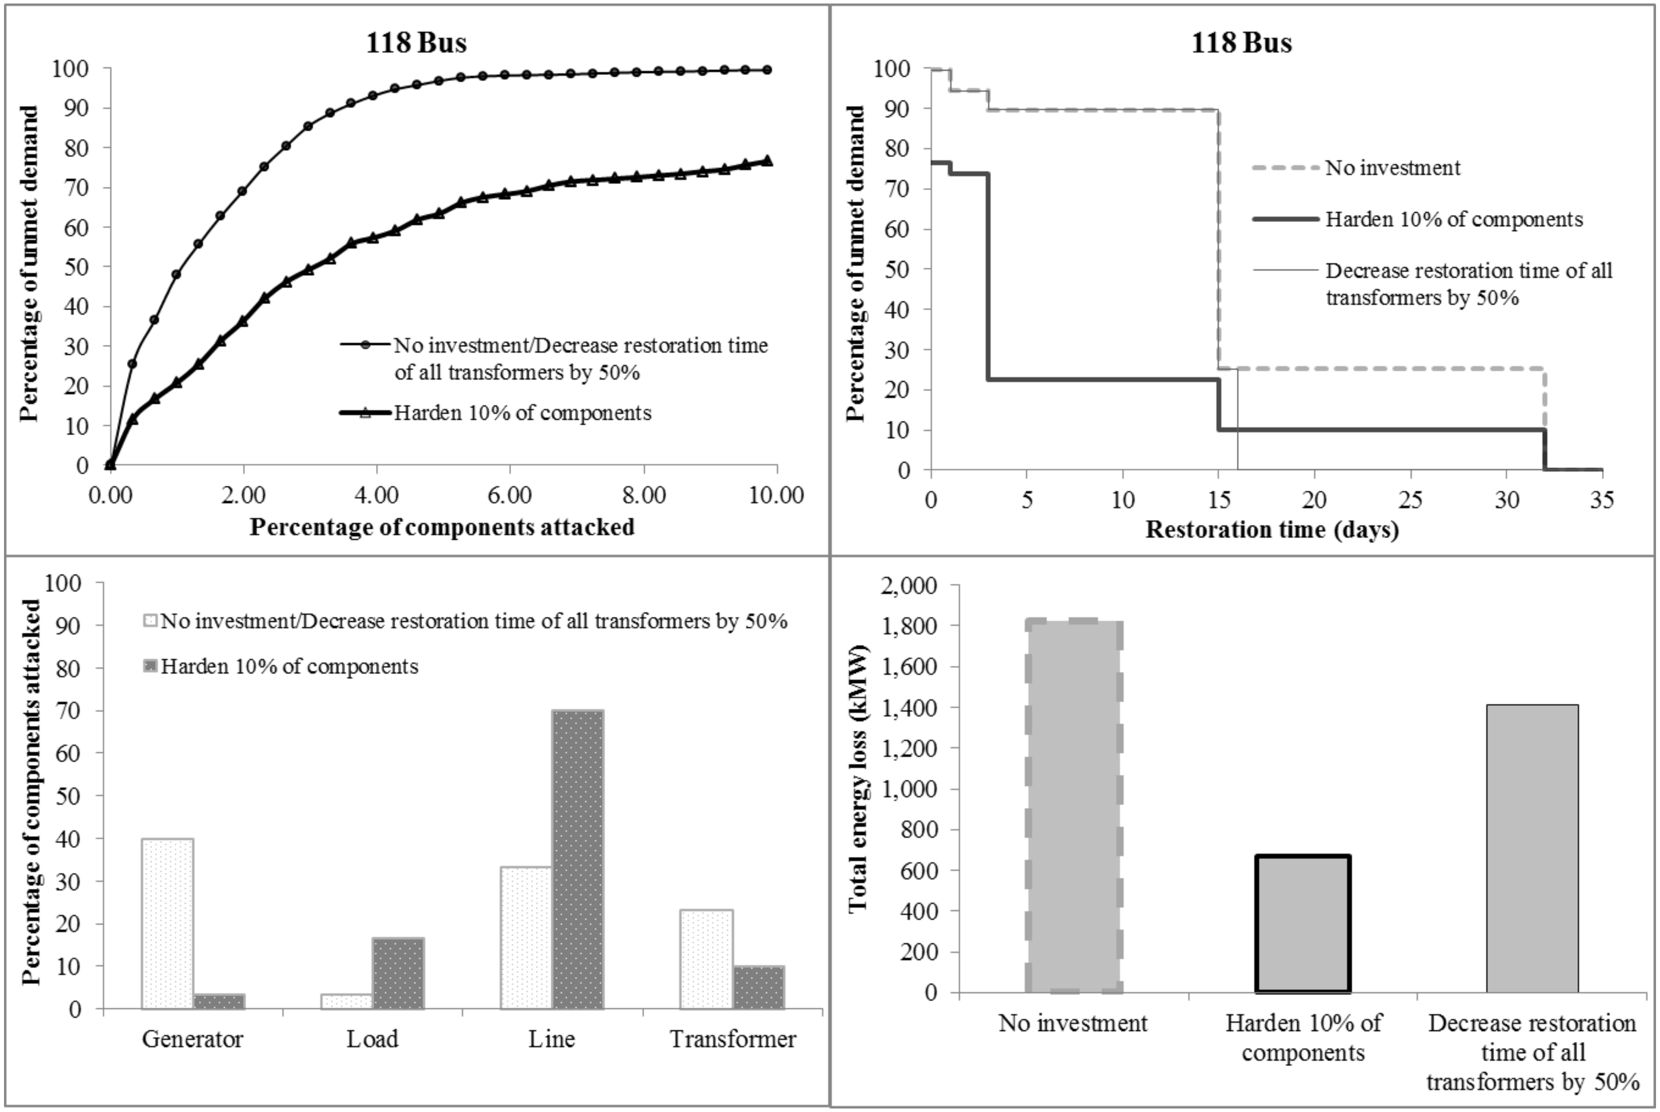

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 10% of the components against a dynamic attacker who does not consider restoration times, 300-bus system)

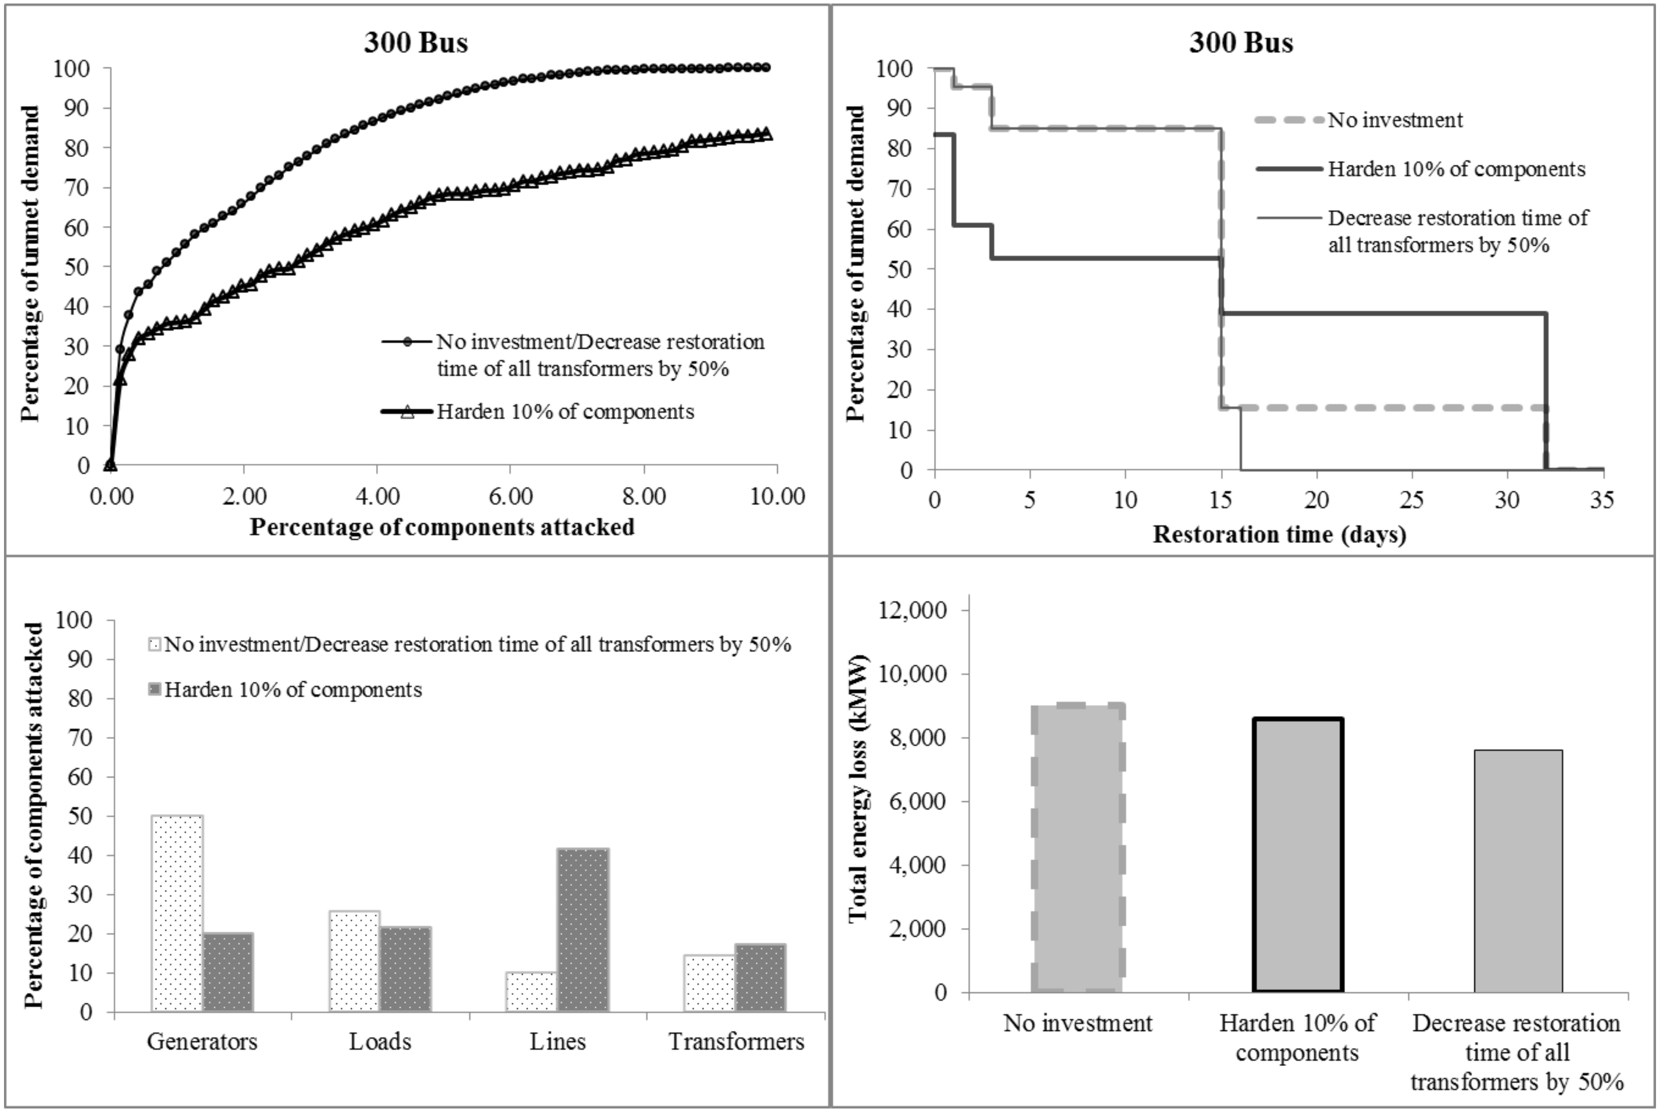

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a dynamic attacker who considers restoration times, 24-bus system)

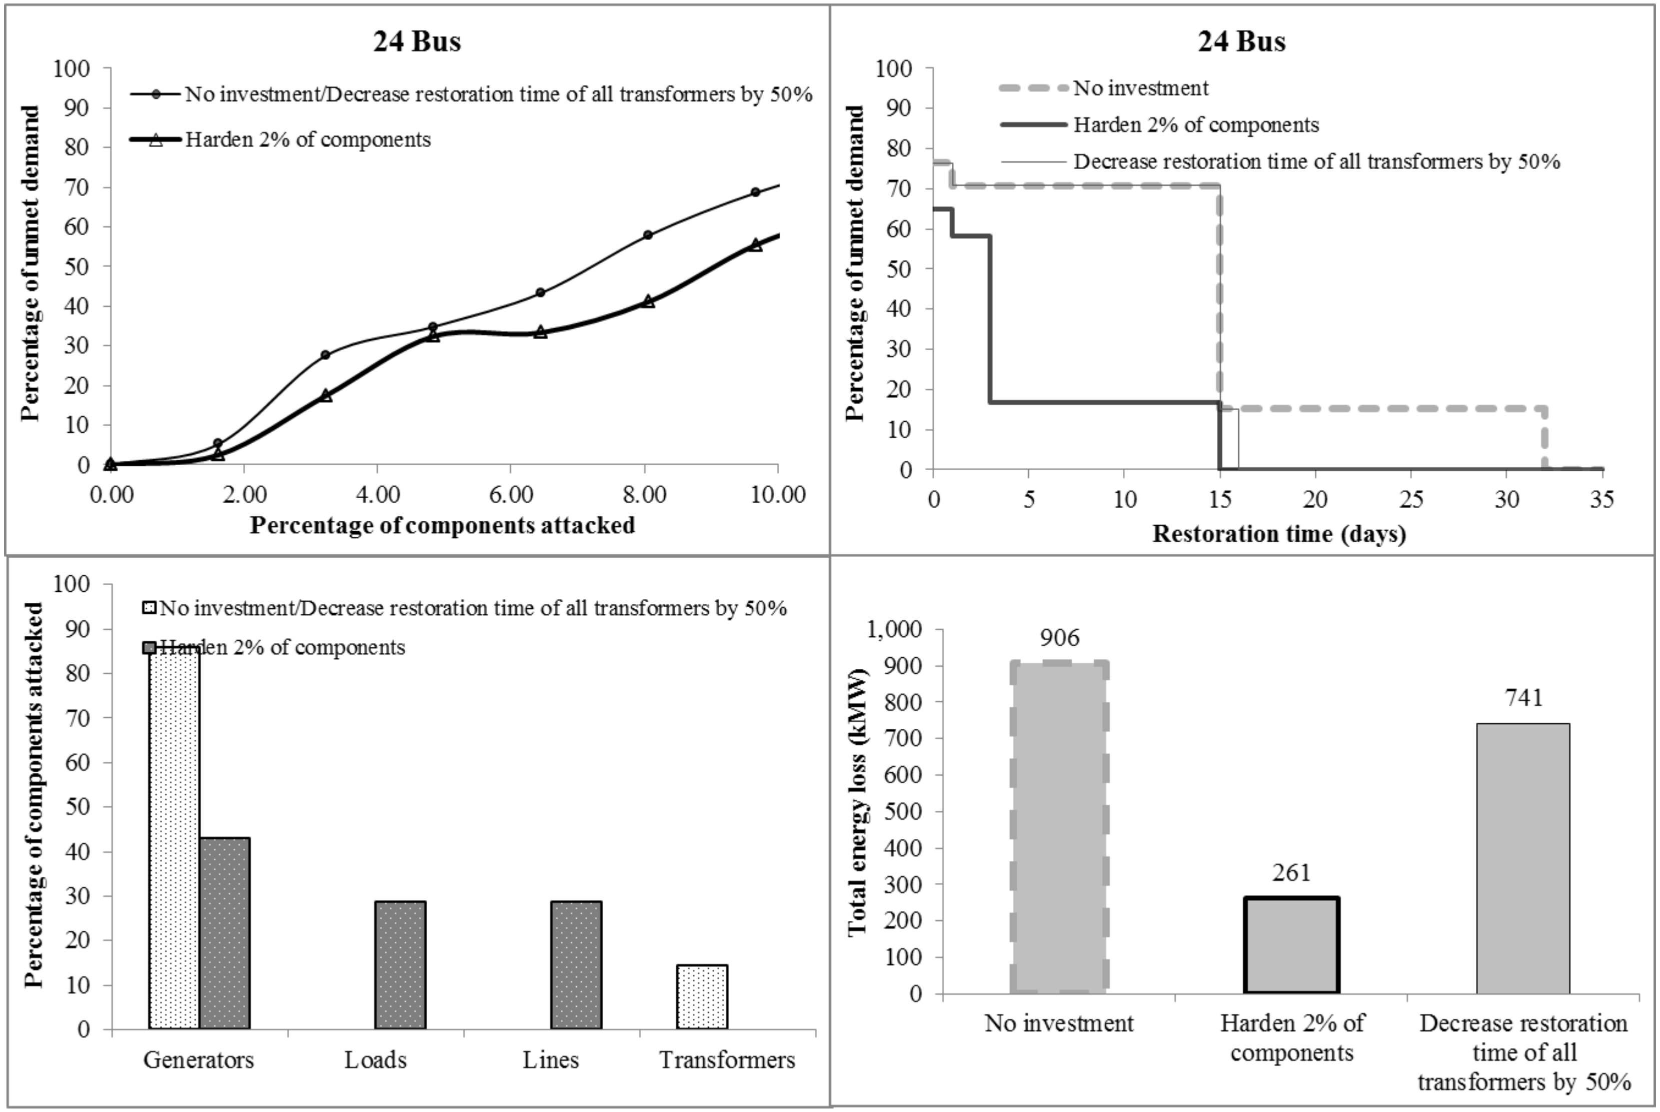

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 2% of the components against a dynamic attacker who considers restoration times, 48-bus system)

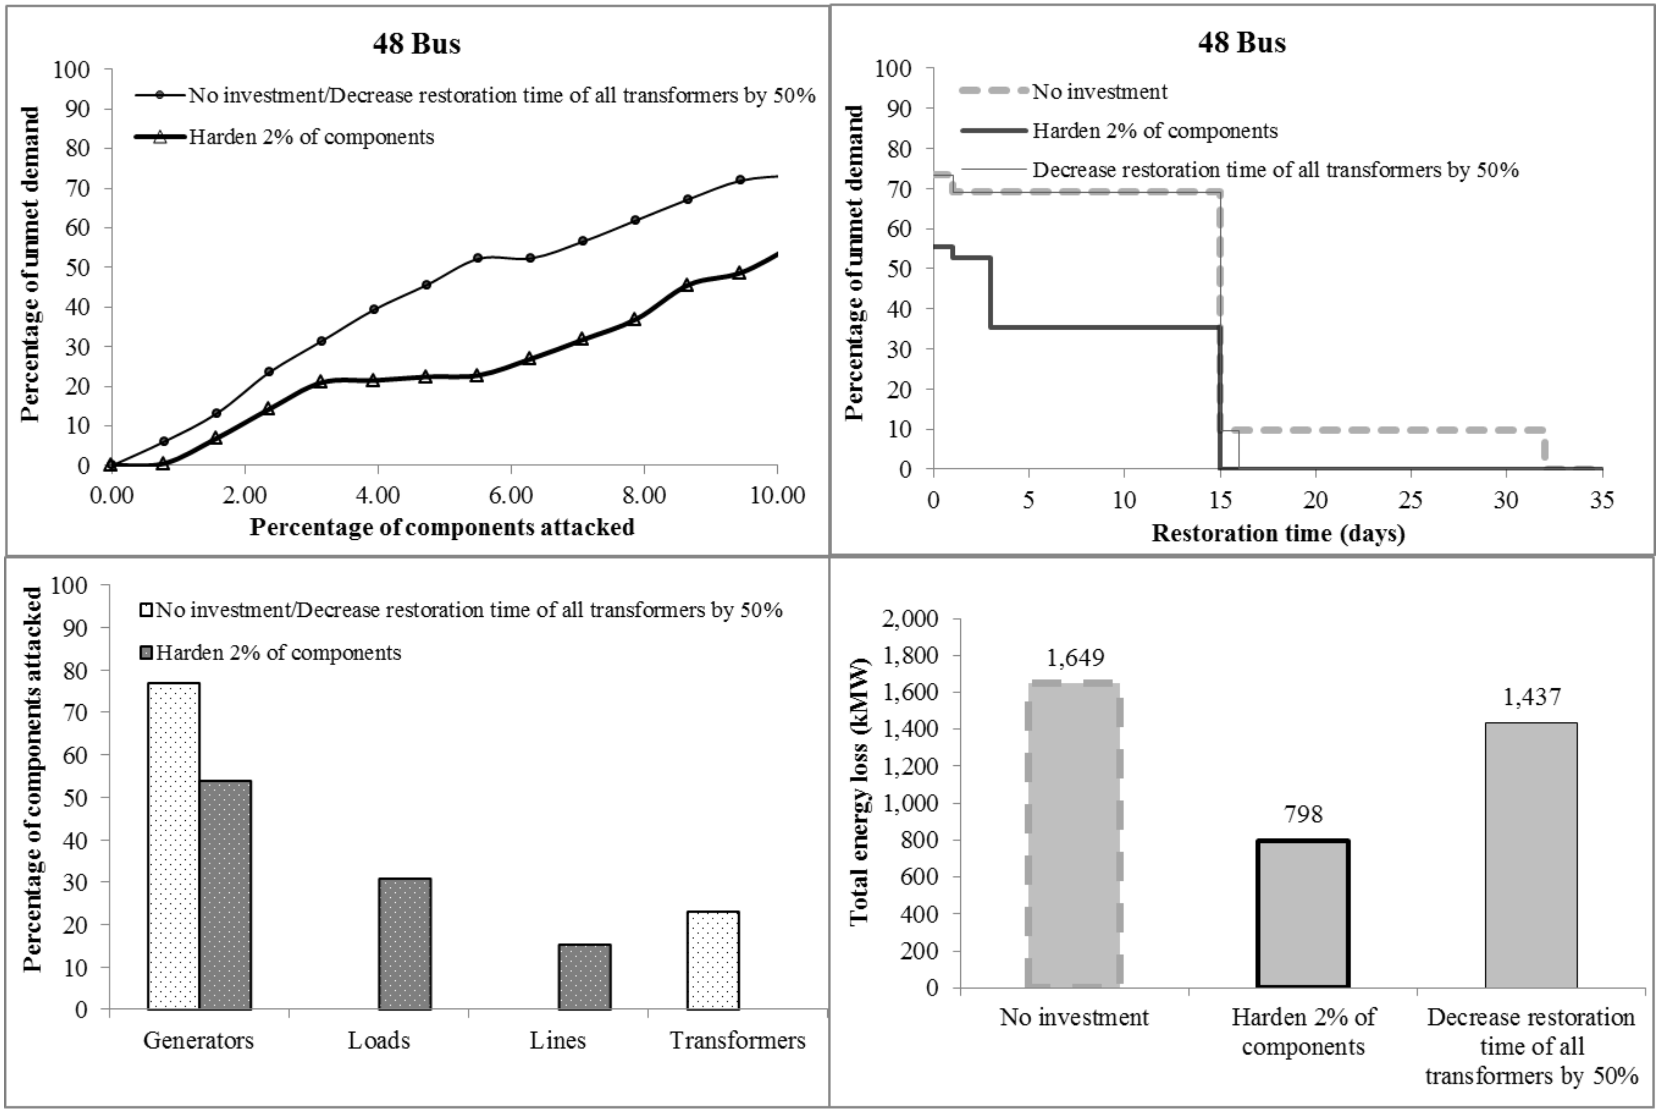

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 2% of the components against a dynamic attacker who considers restoration times, 118-bus system)

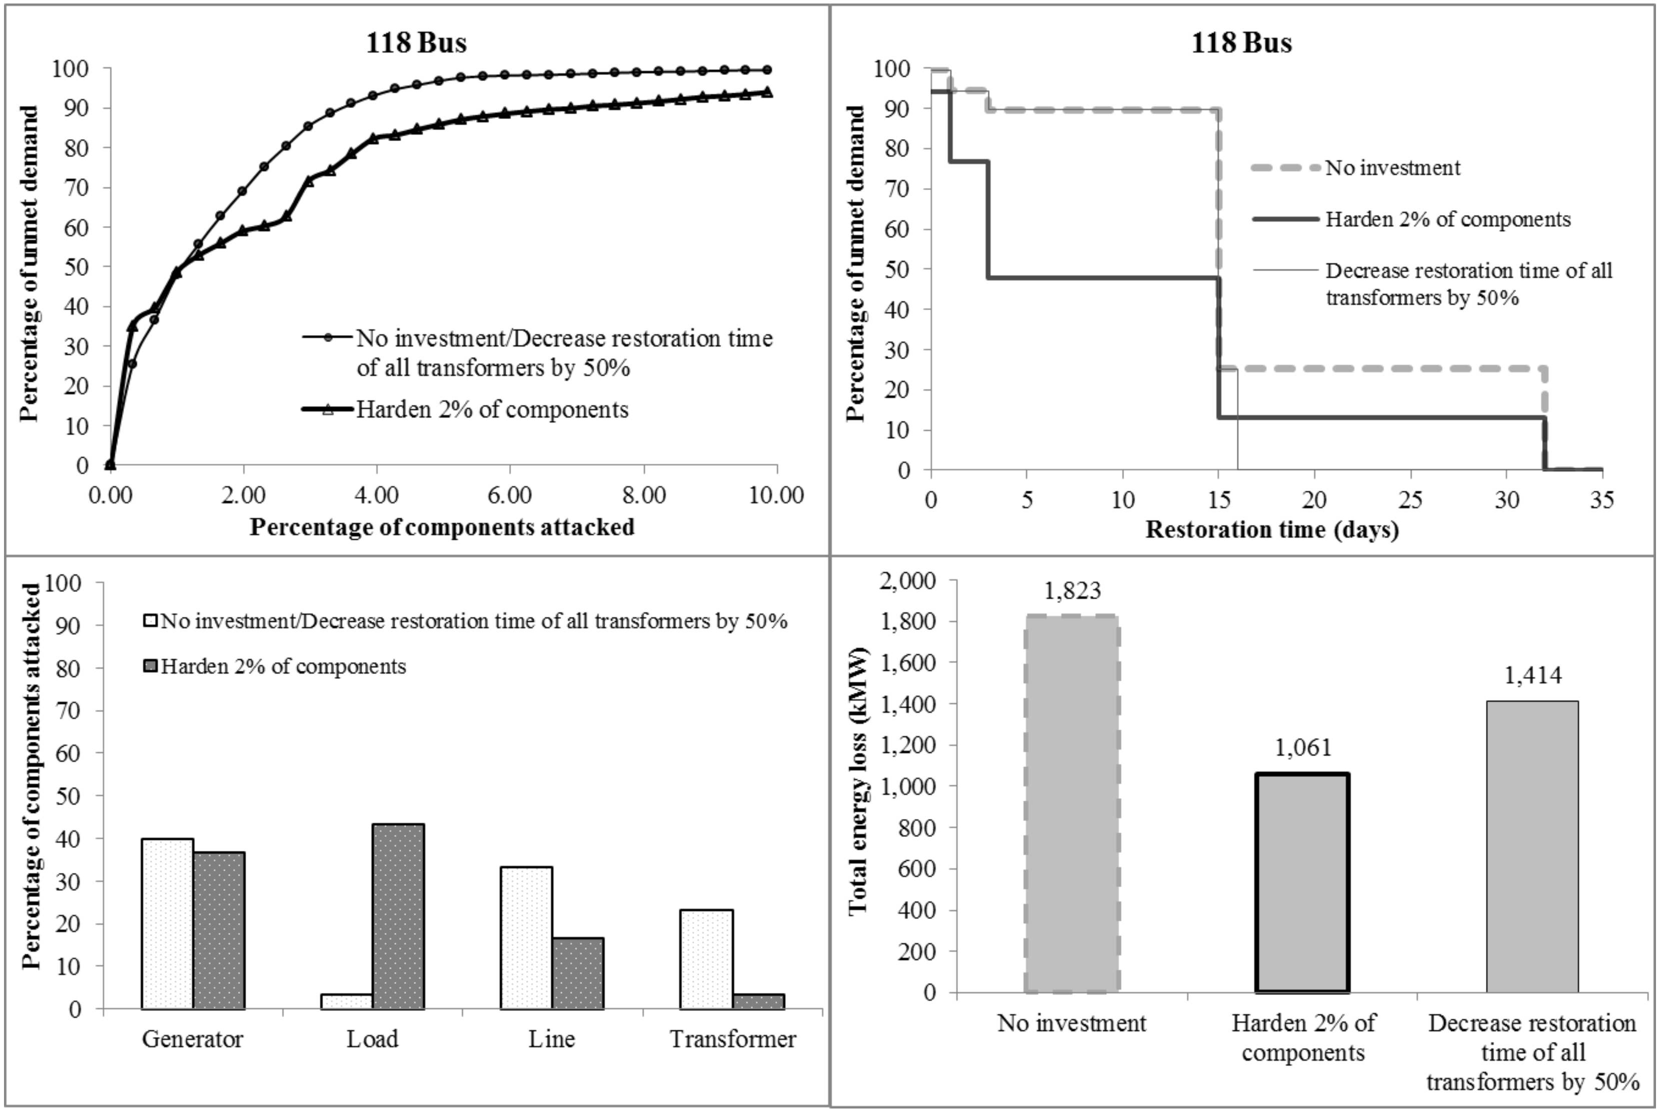

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a dynamic attacker who considers restoration times, 300-bus system)

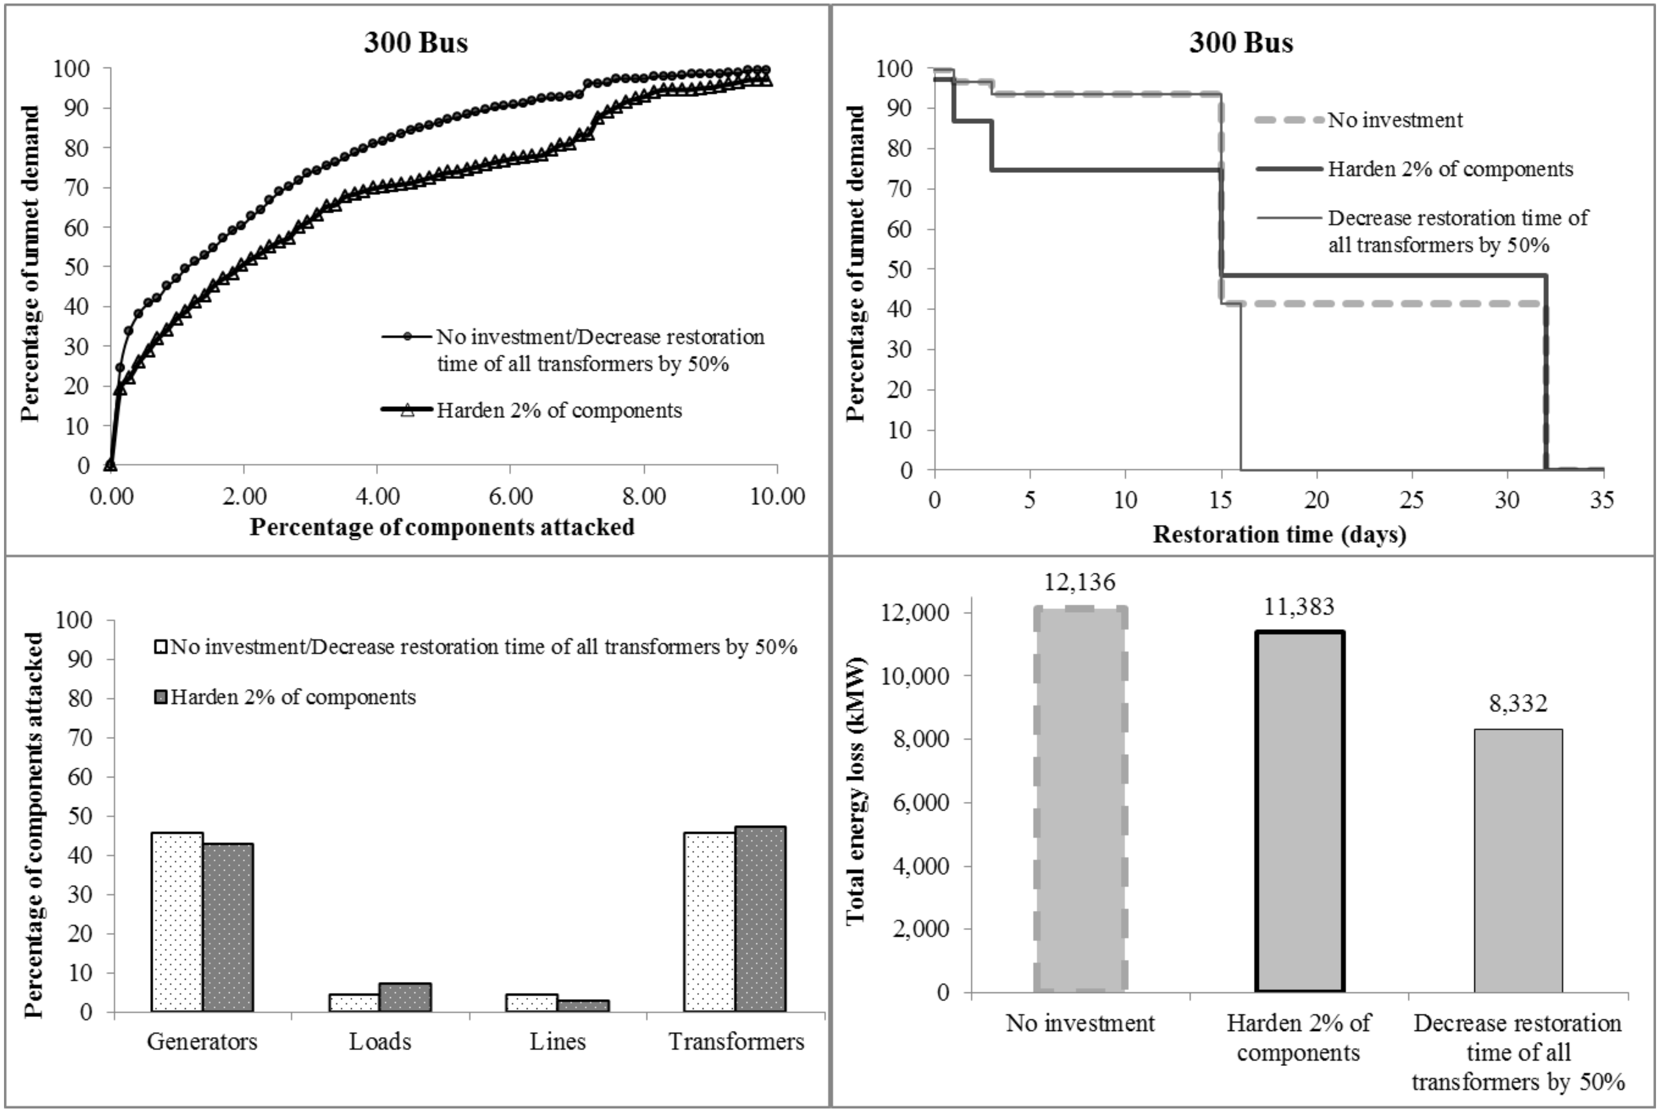

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a dynamic attacker who considers restoration times, 24-bus system)

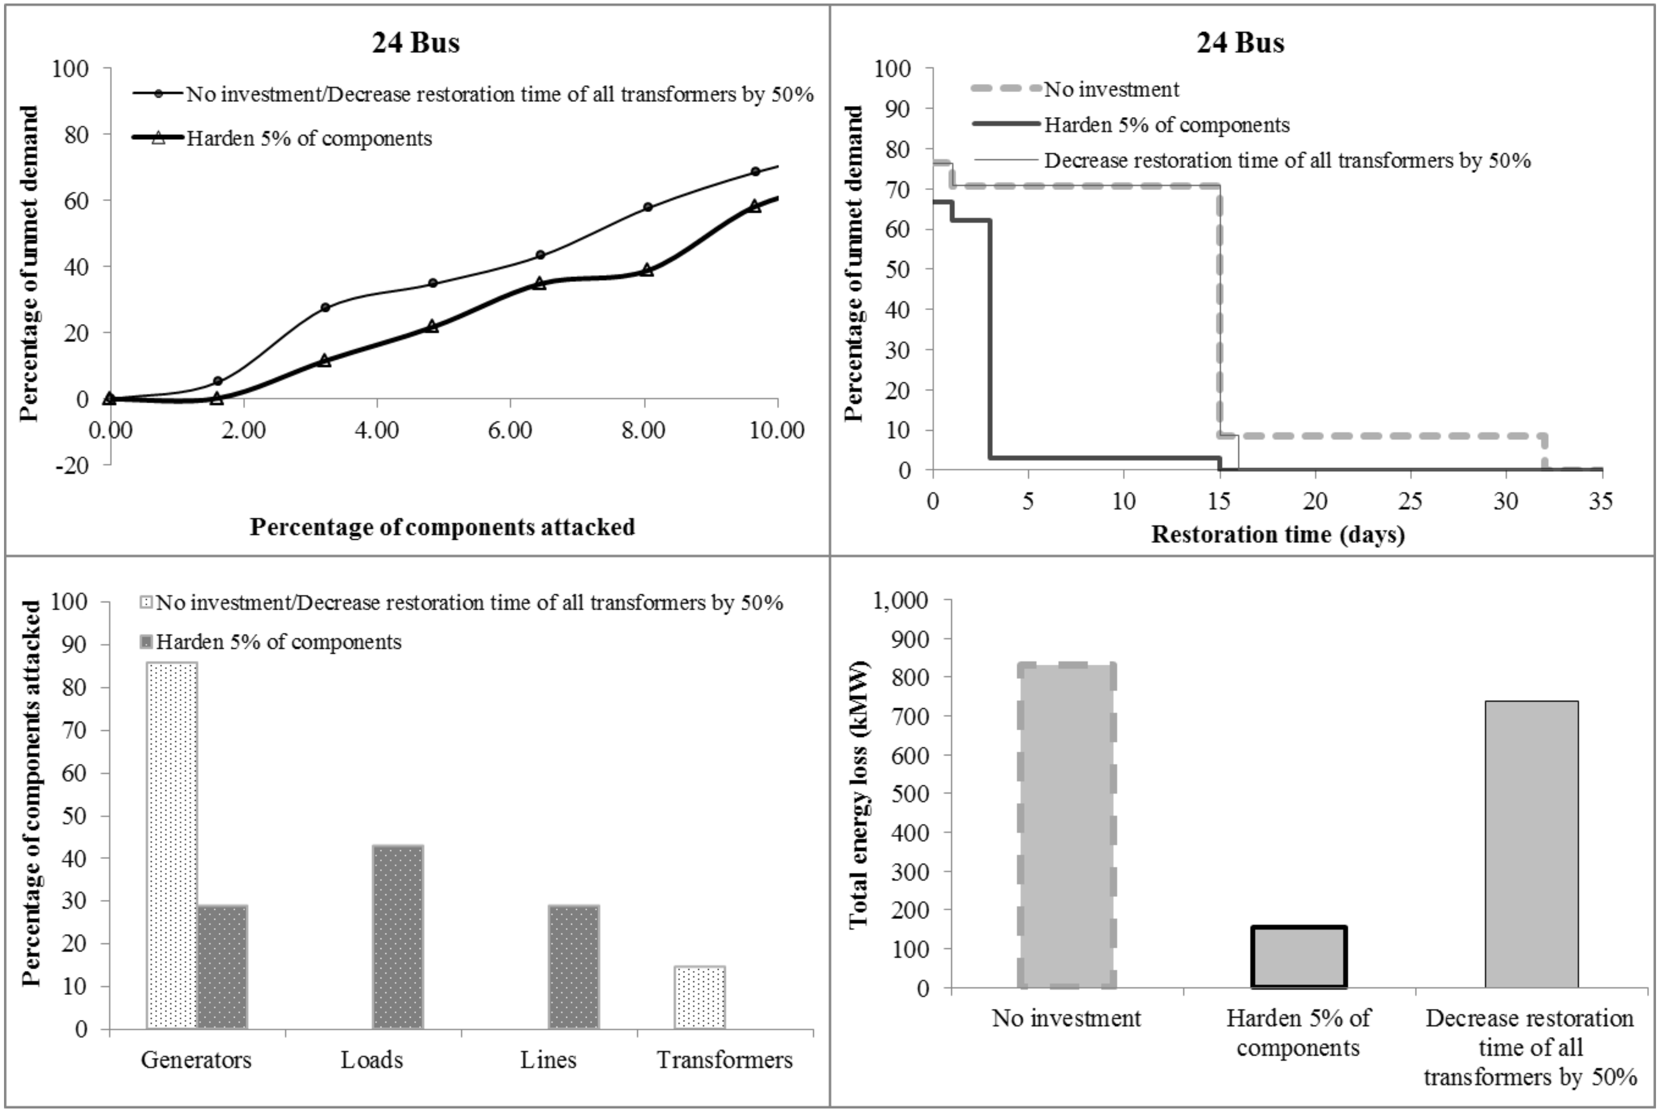

Appendix M. Hardening versus decreasing restoration times of transformers
(hardening 5% of the components against a dynamic attacker who considers restoration times, 48-bus system)

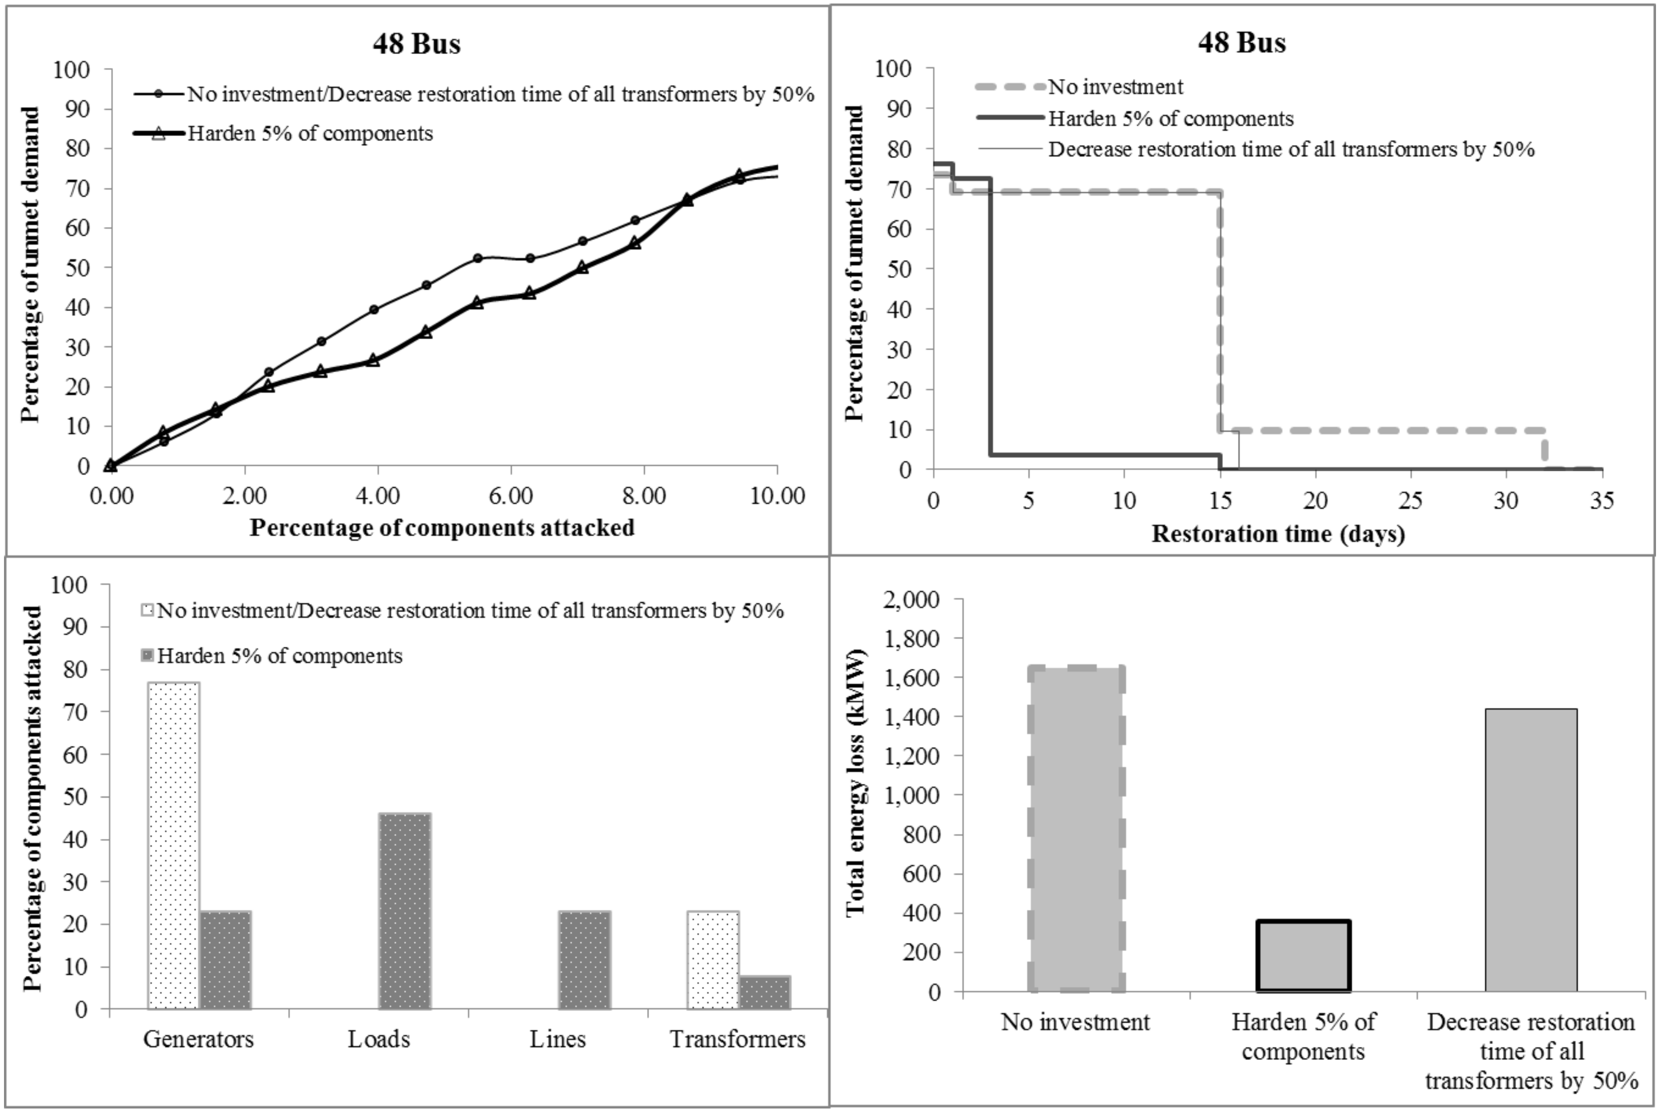

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a dynamic attacker who considers restoration times, 118-bus system)

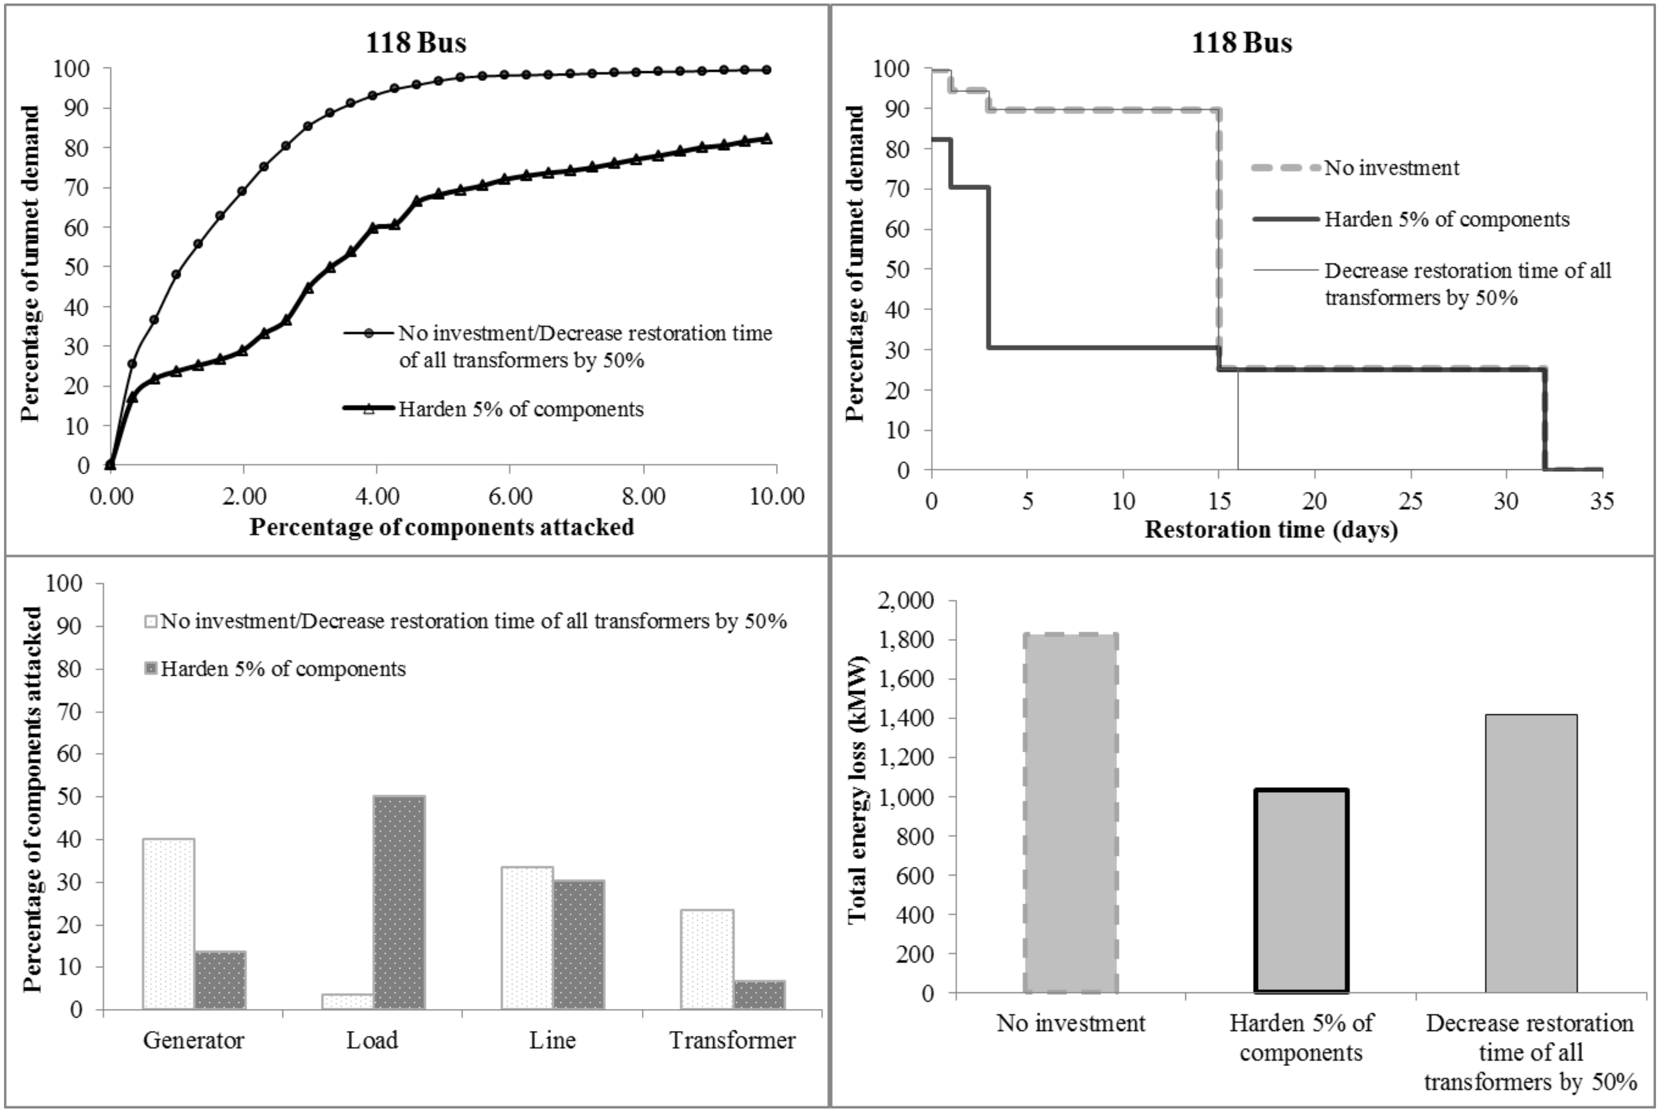

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a dynamic attacker who considers restoration times, 300-bus system)

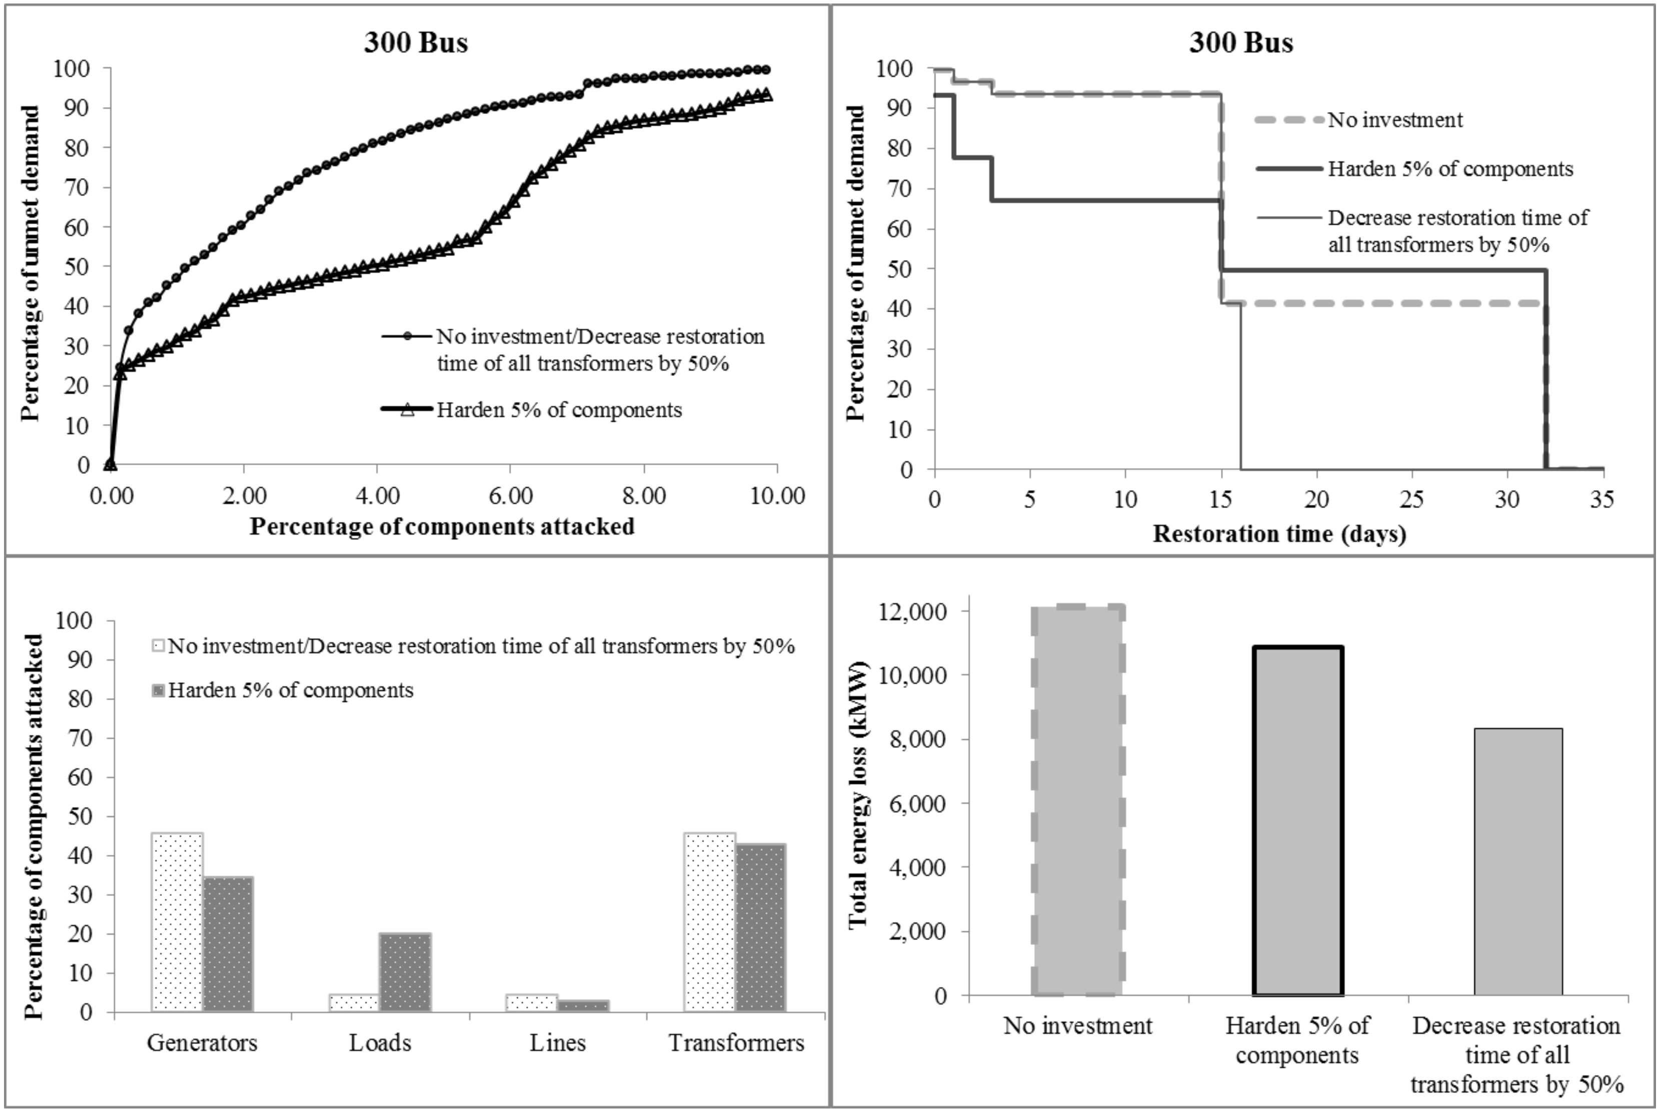

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who considers restoration times, 24-bus system)

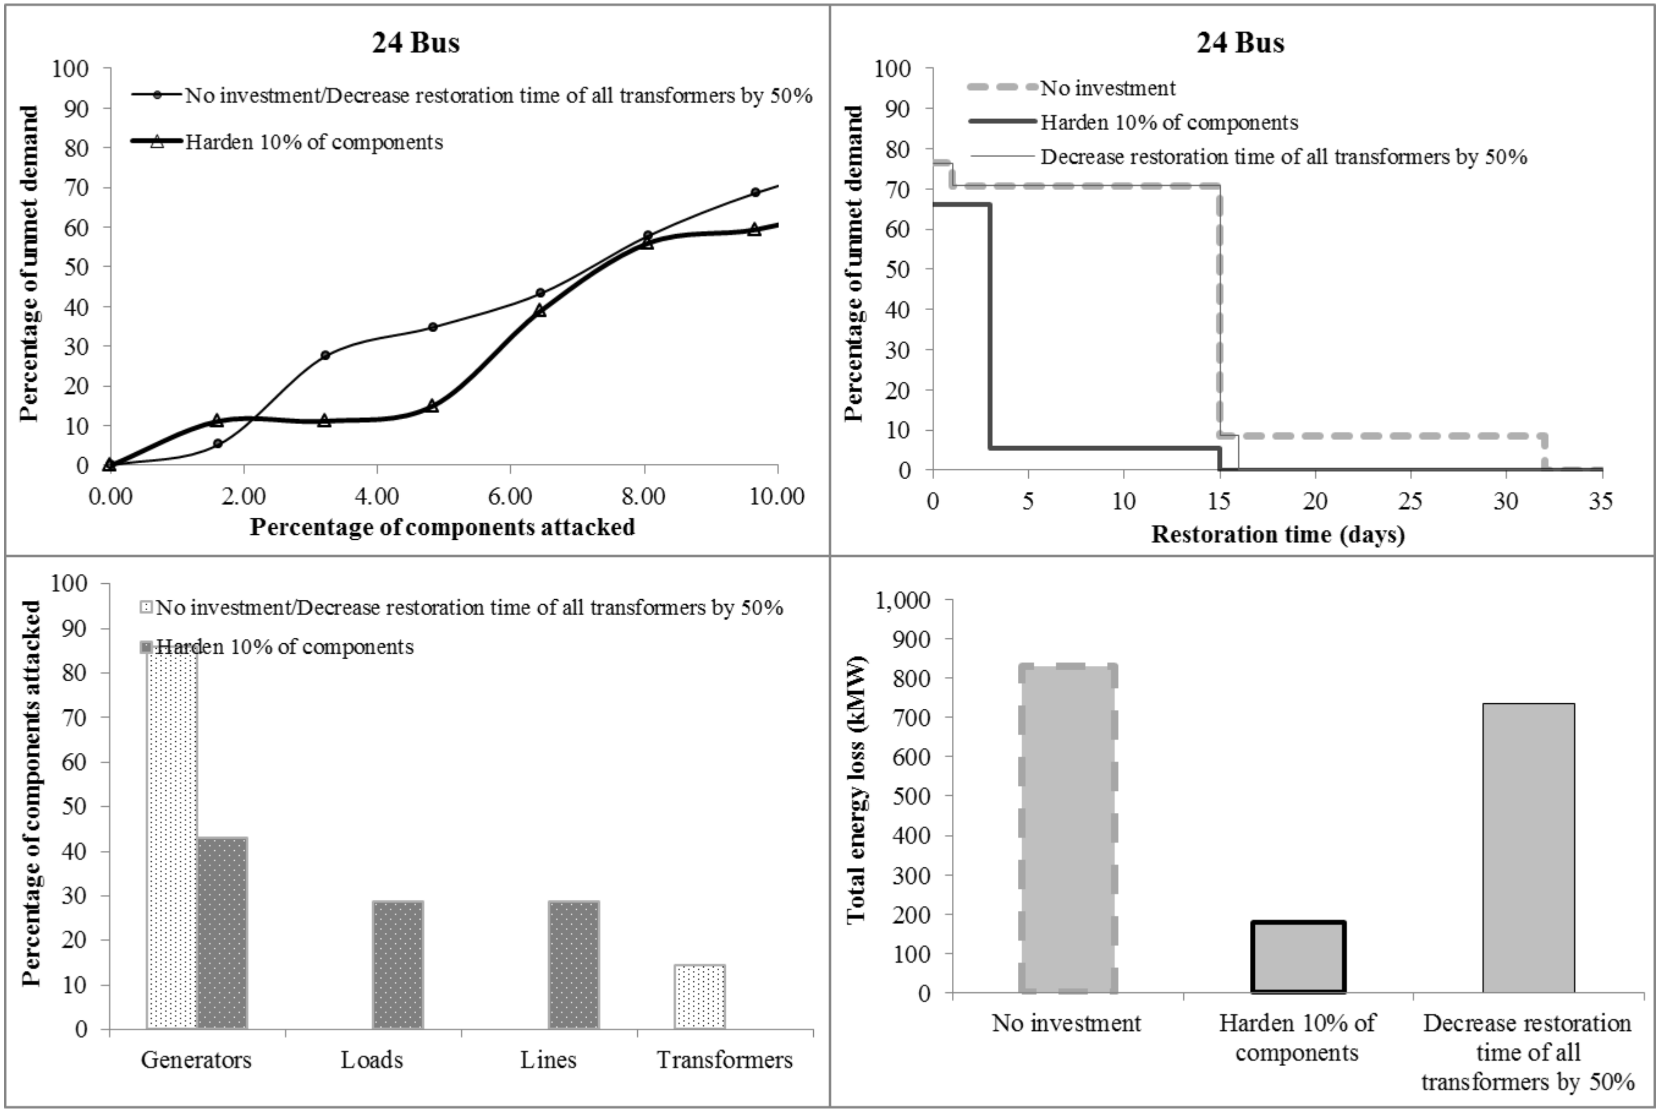

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who considers restoration times, 48-bus system)

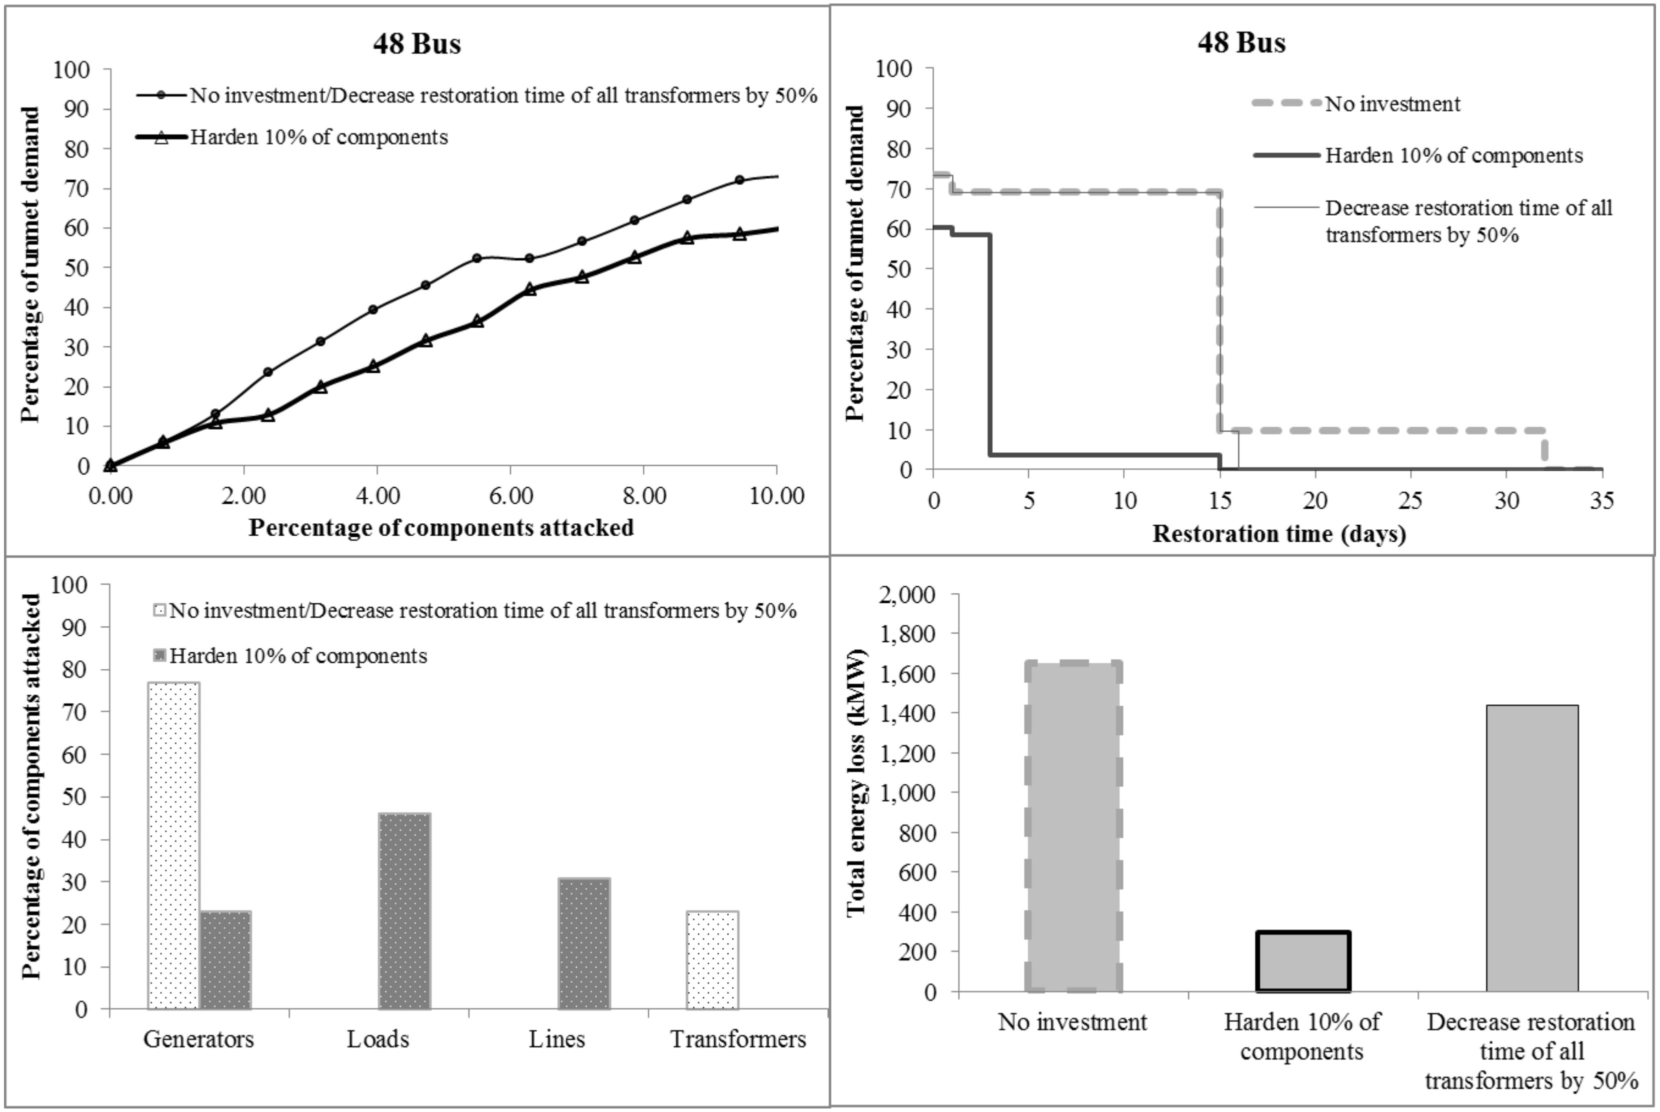

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who considers restoration times, 118-bus system)

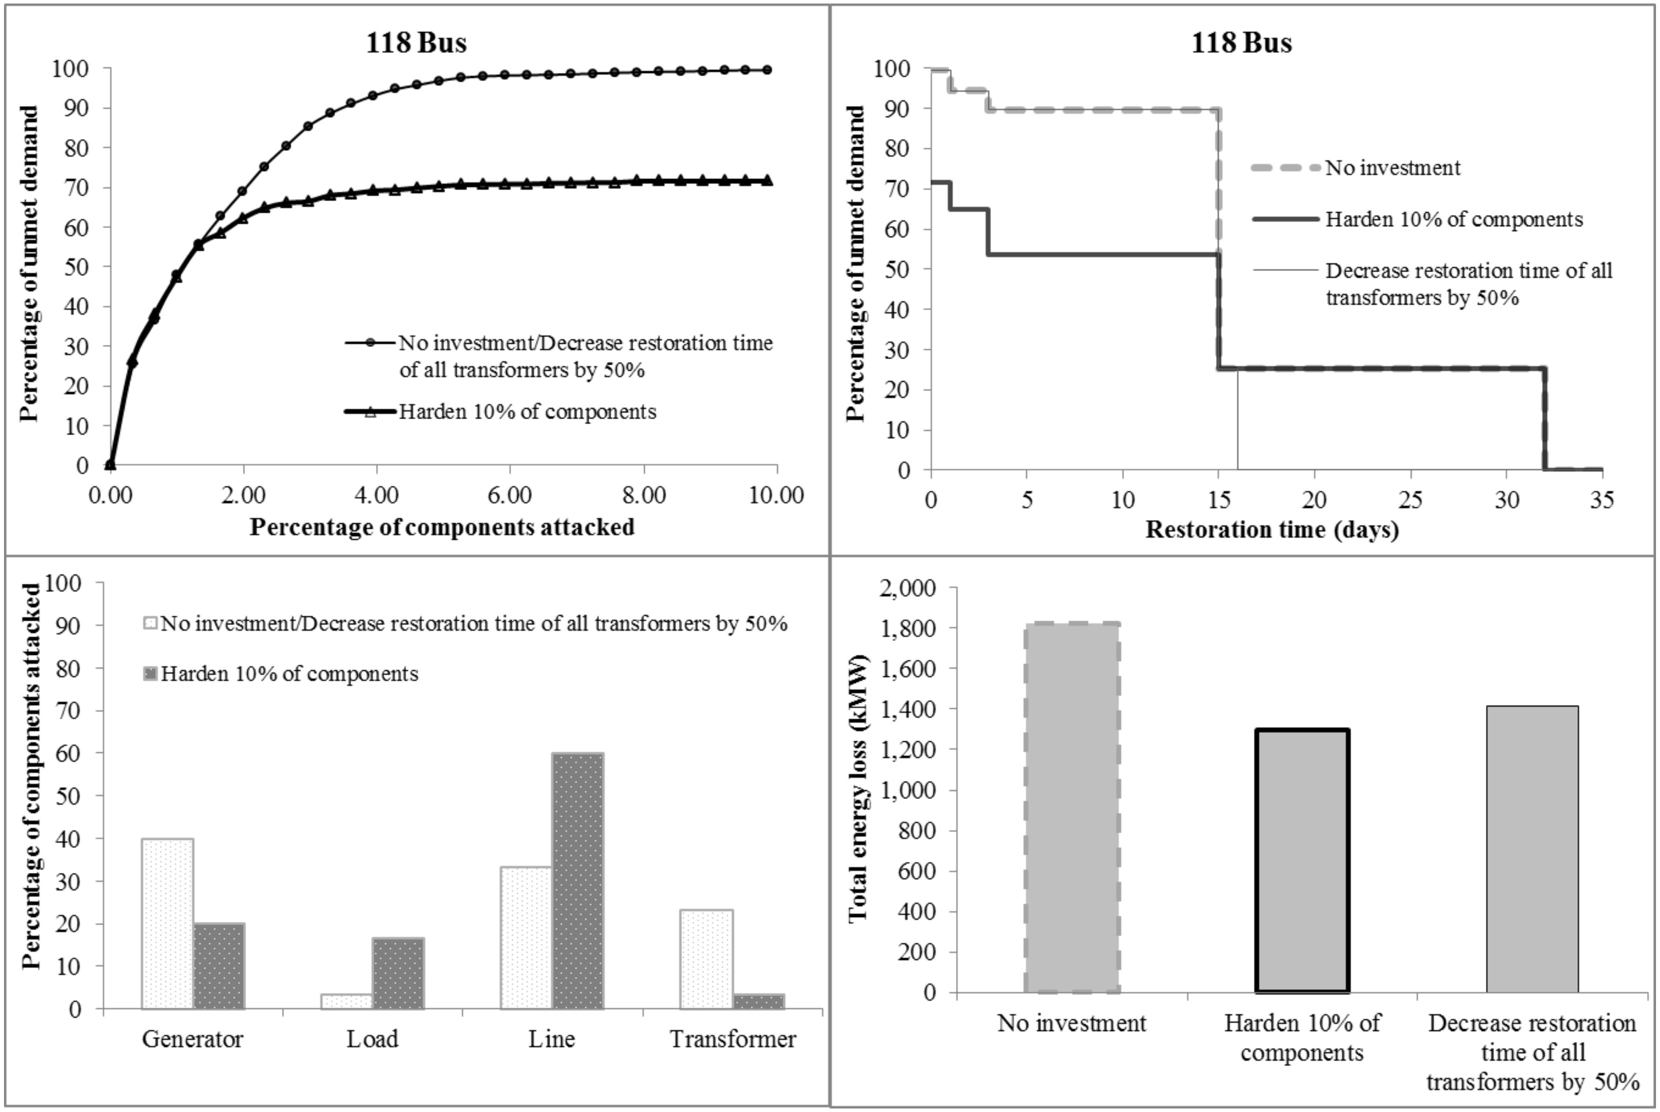

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a dynamic attacker who considers restoration times, 300-bus system)

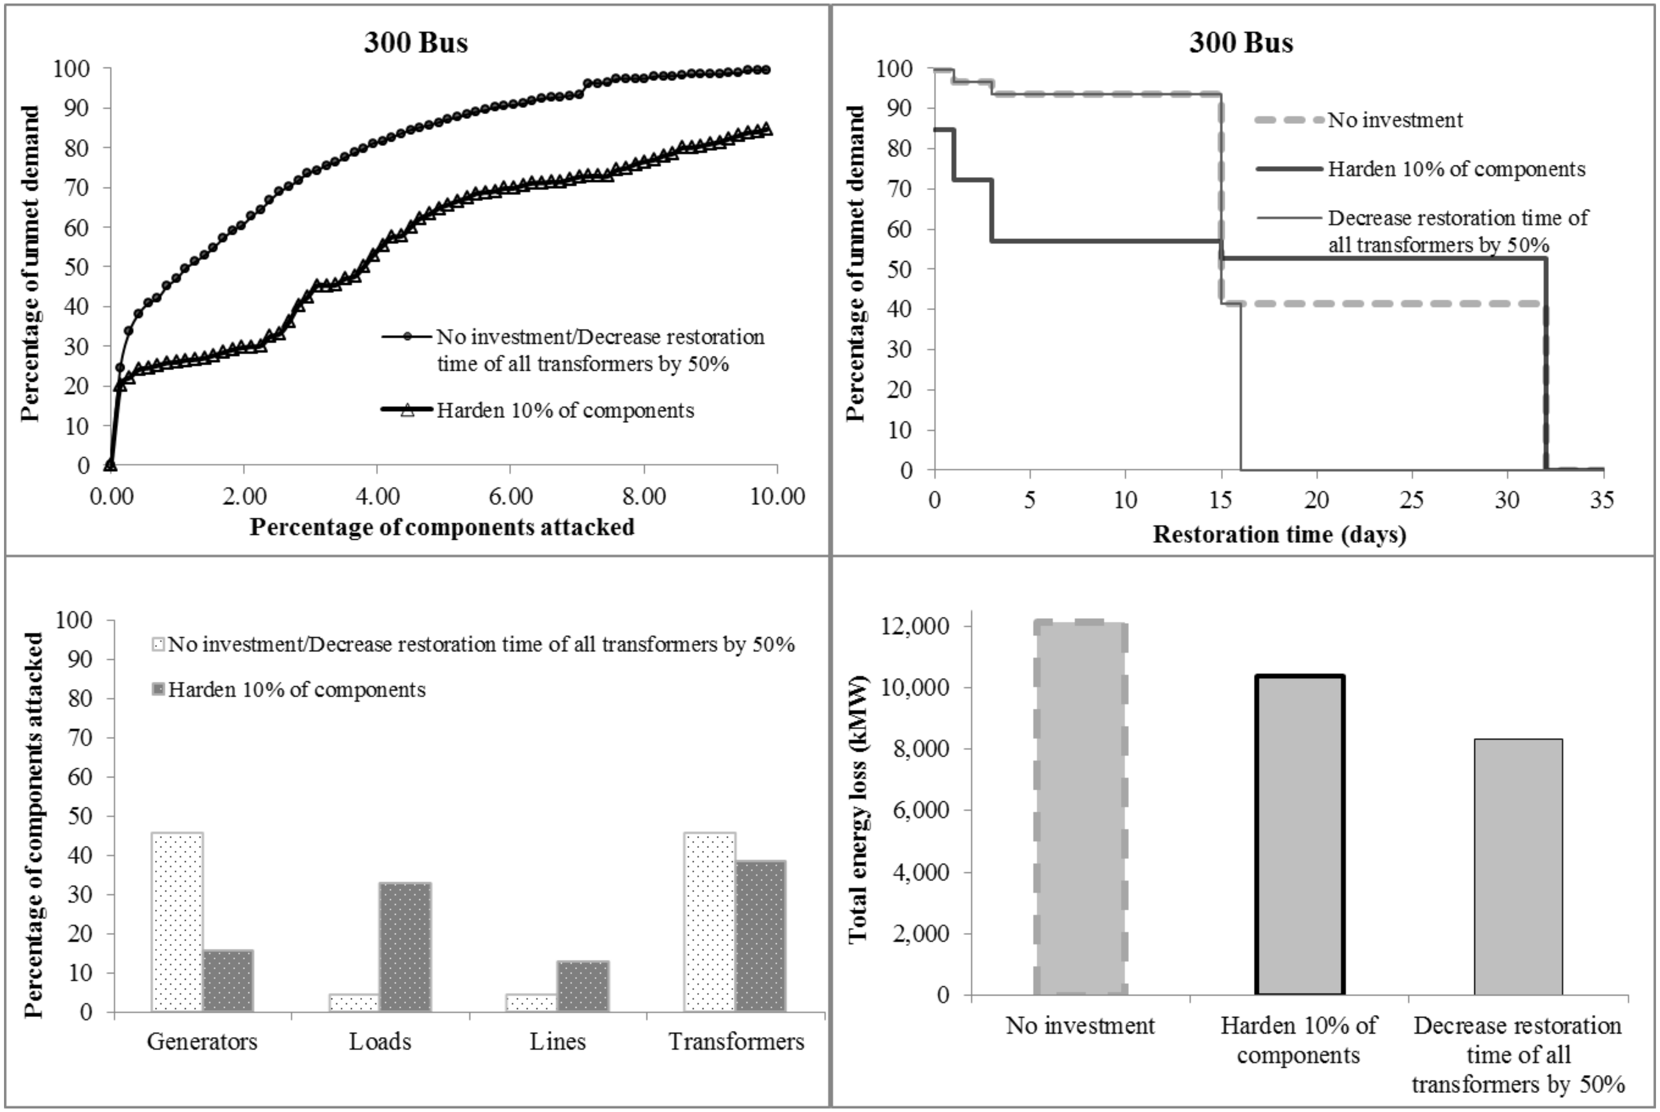

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and not considering restoration times, 24-bus system)

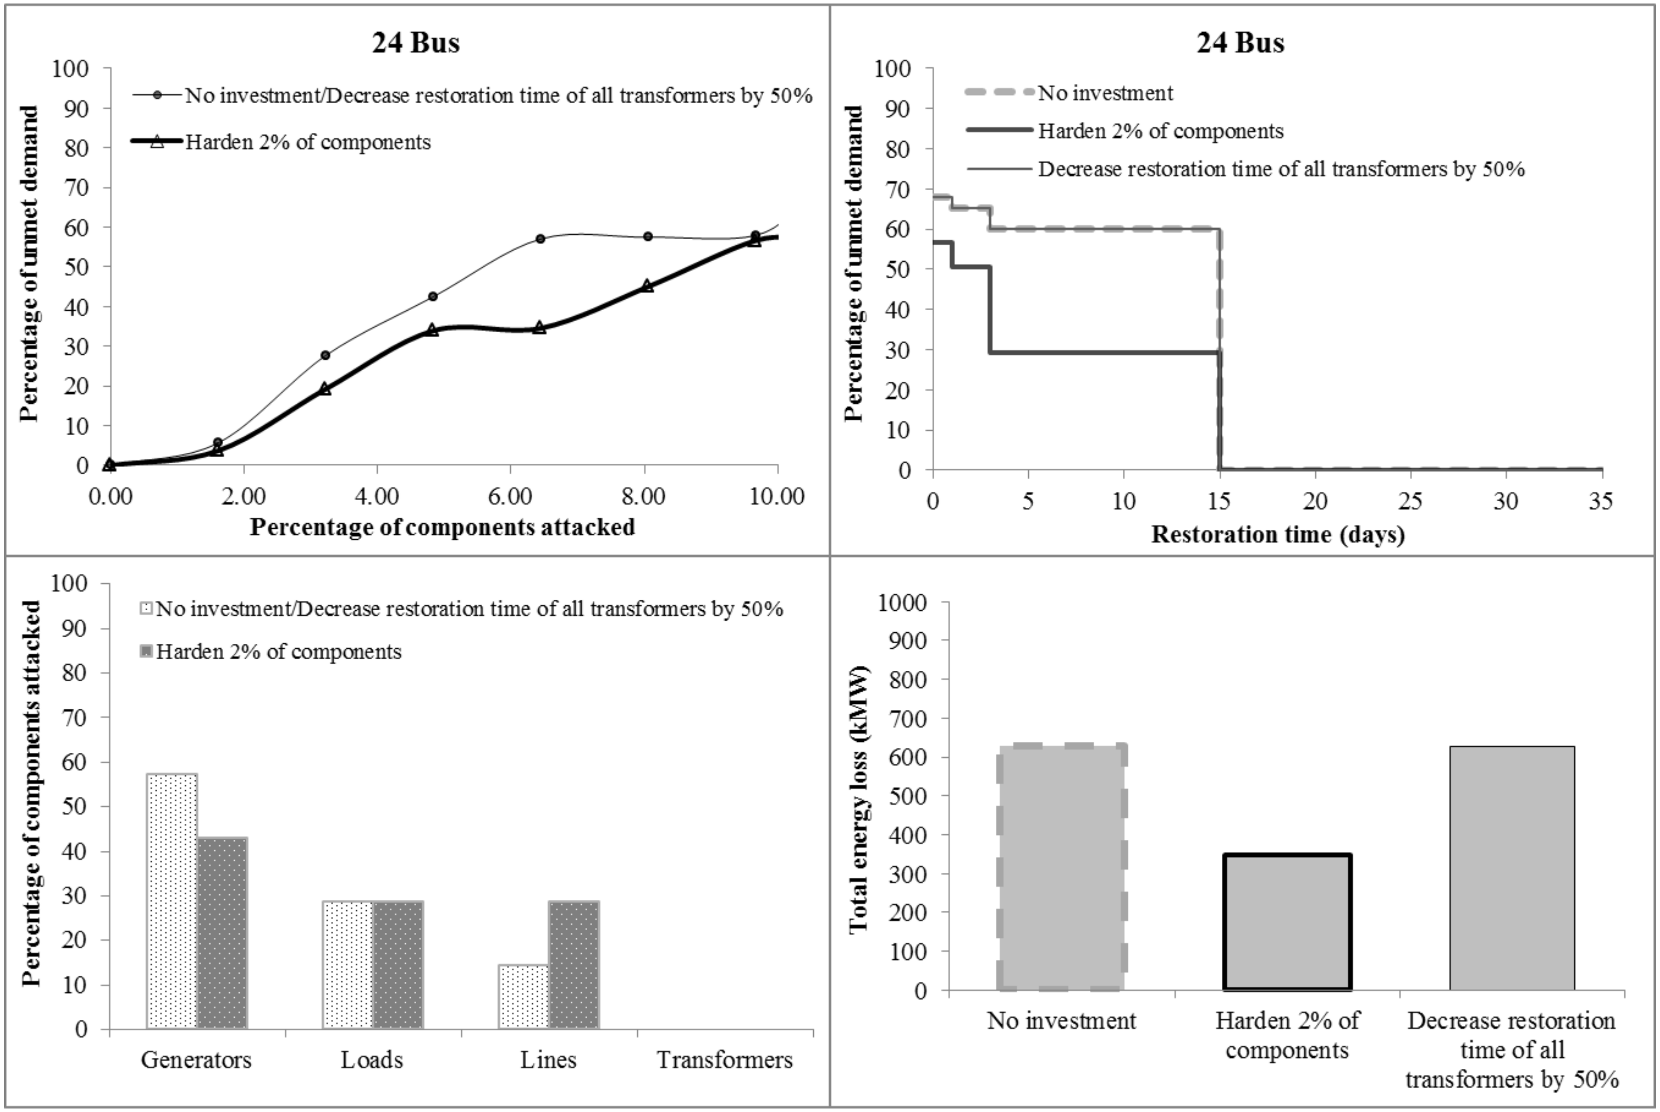

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and not considering restoration times, 48-bus system)

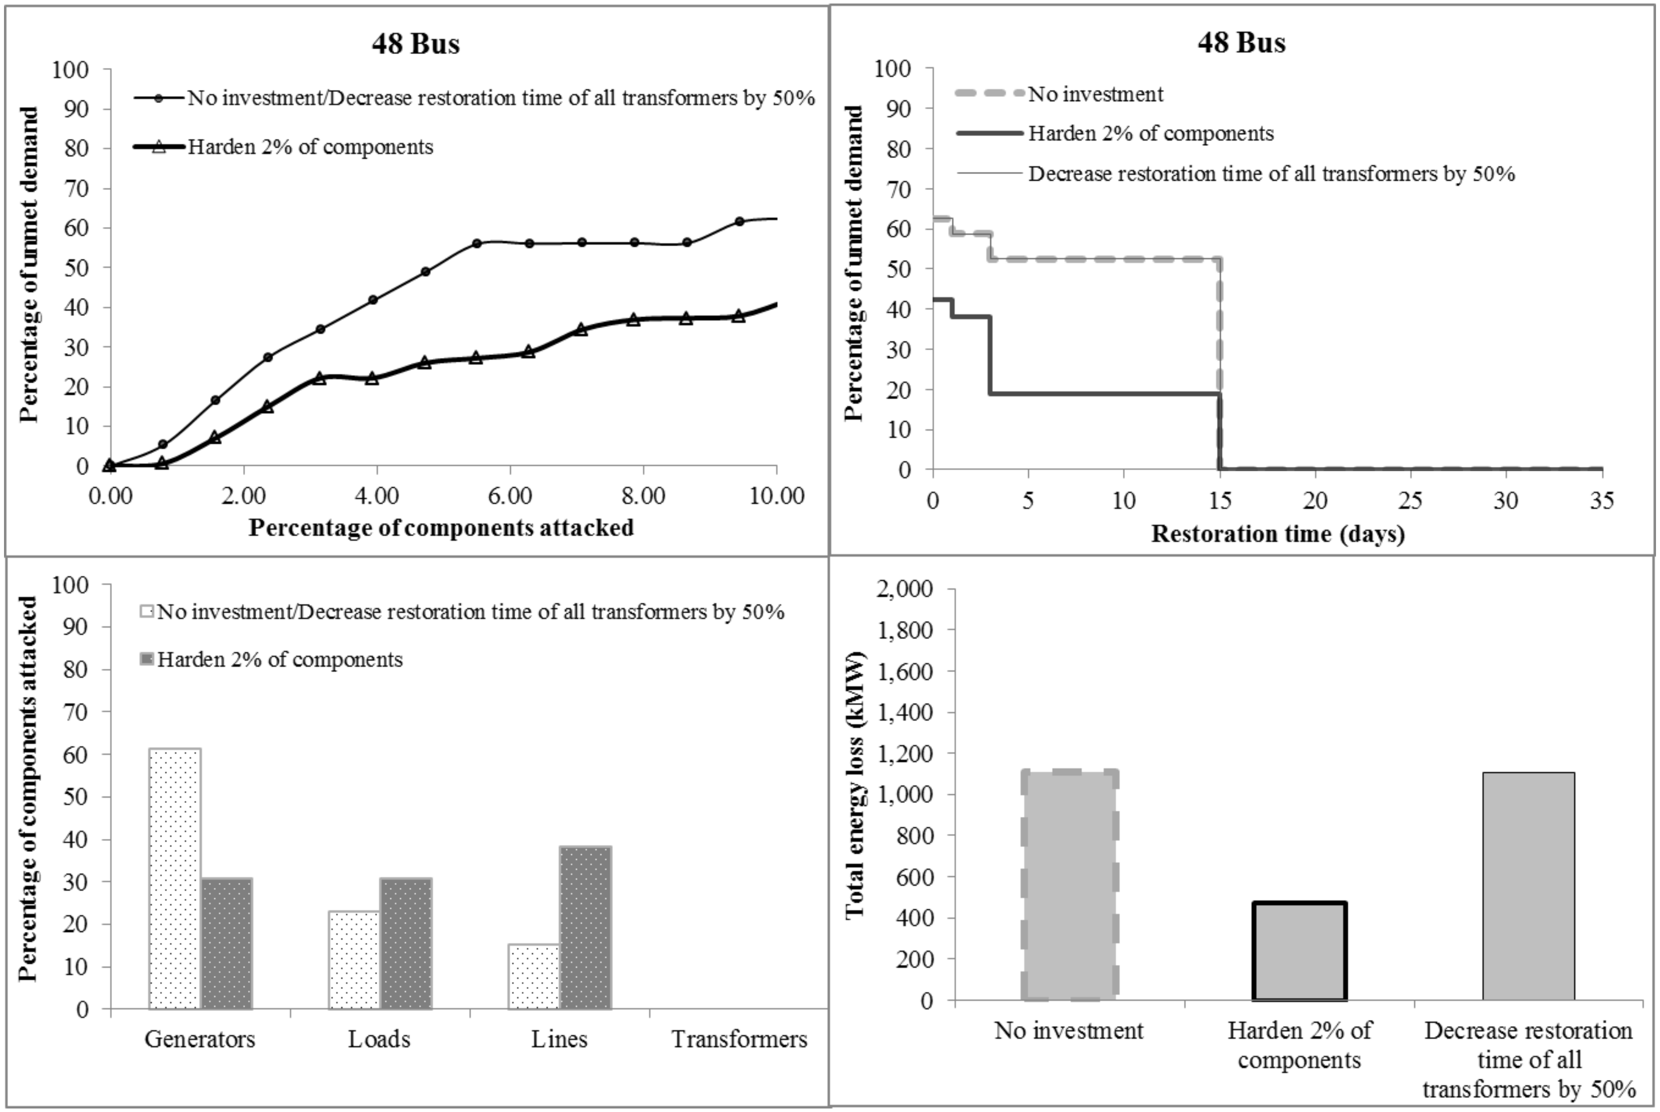

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and not considering restoration times, 118-bus system)

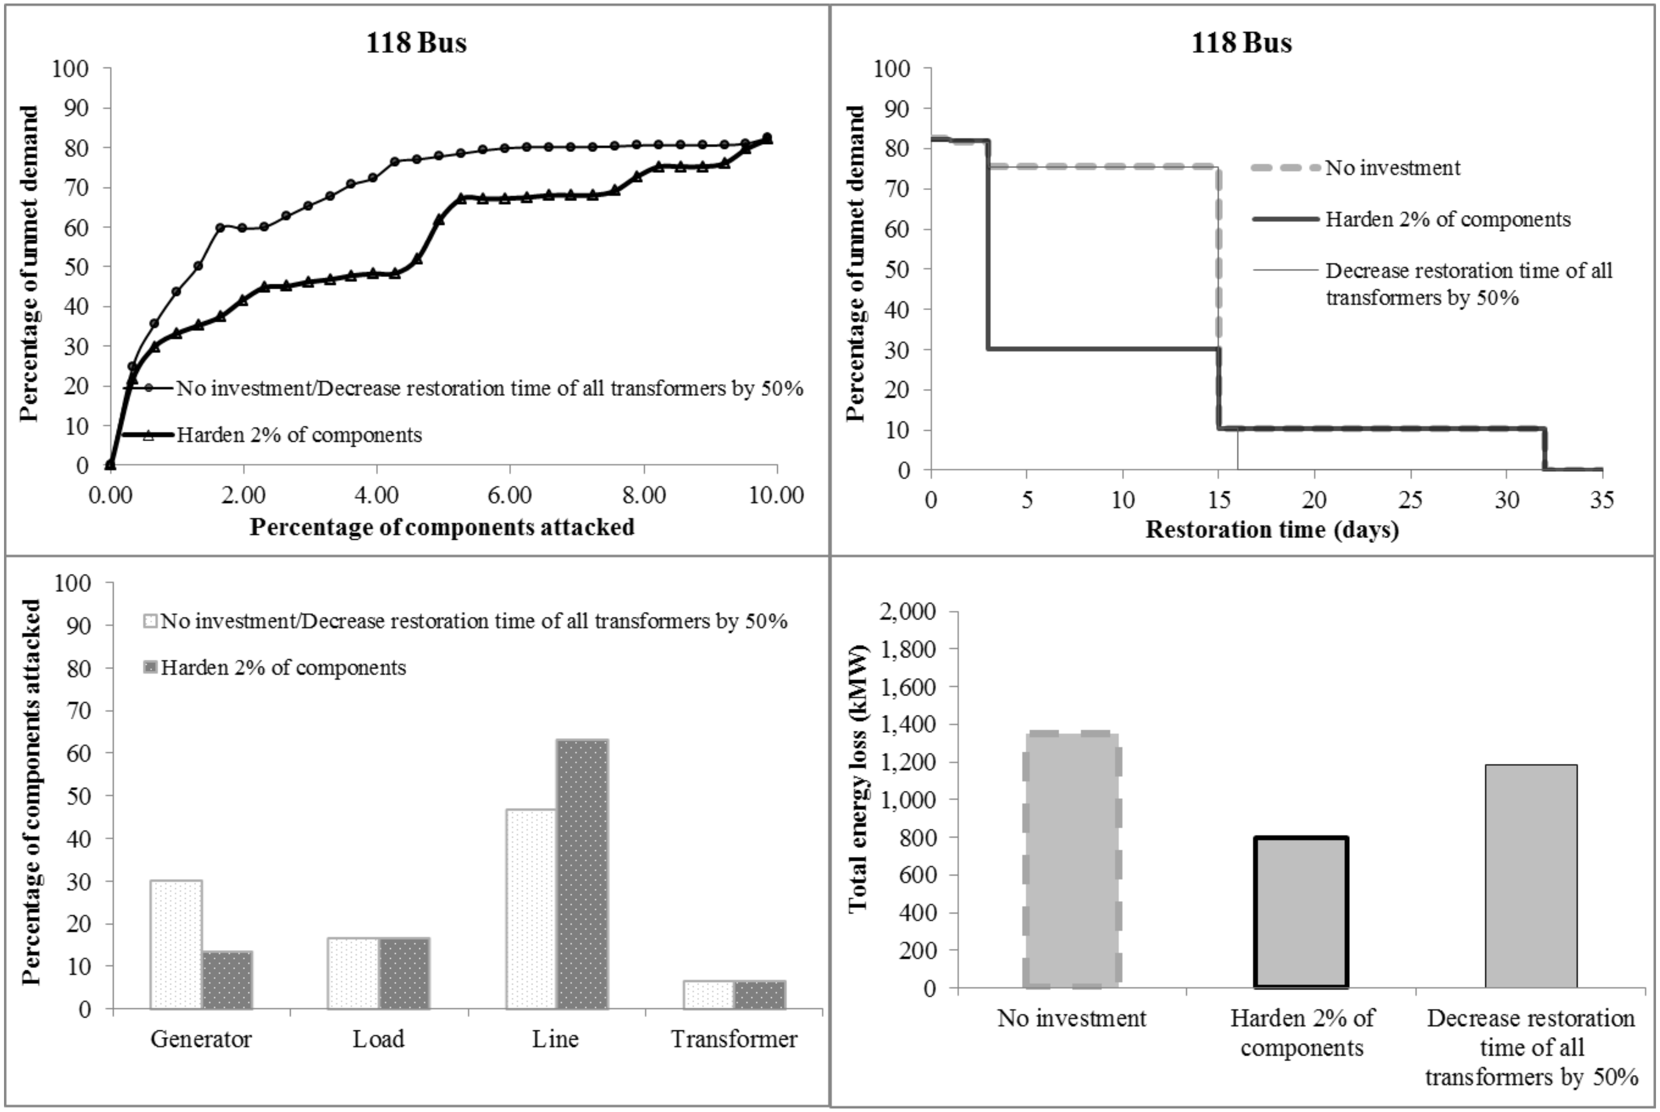

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and not considering restoration times, 300-bus system)

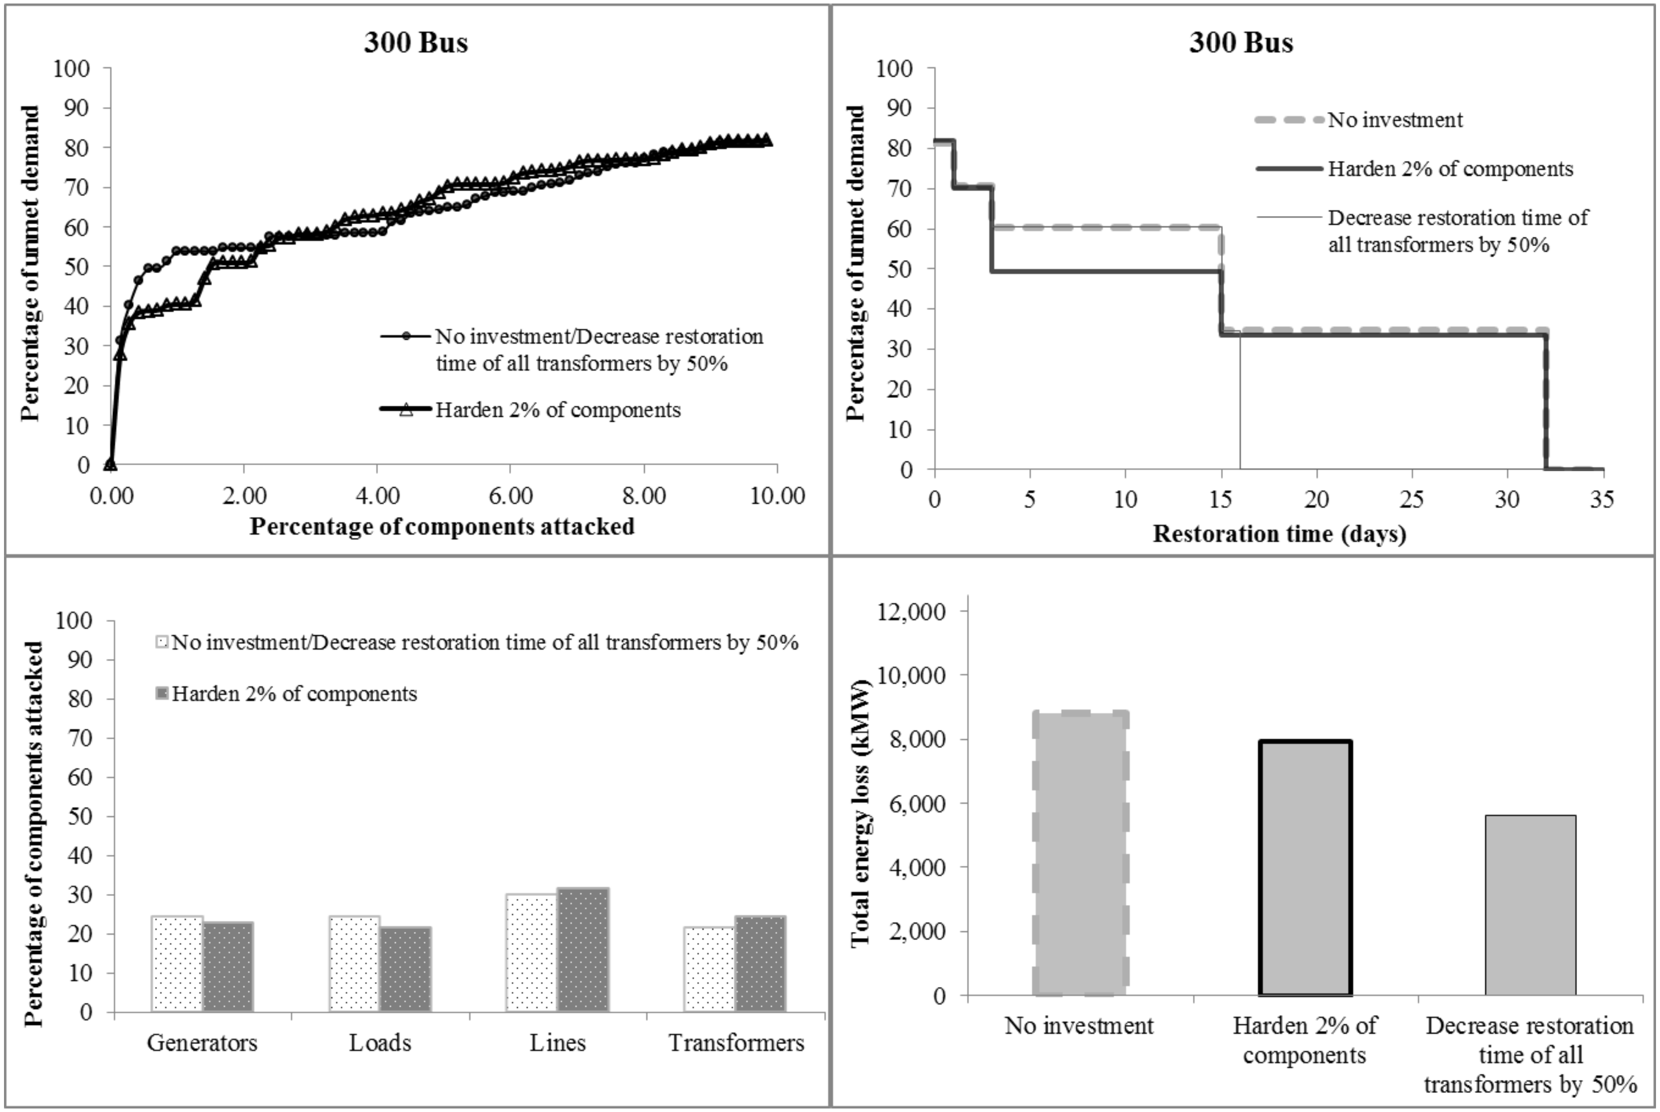

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and not considering restoration times, 24-bus system)

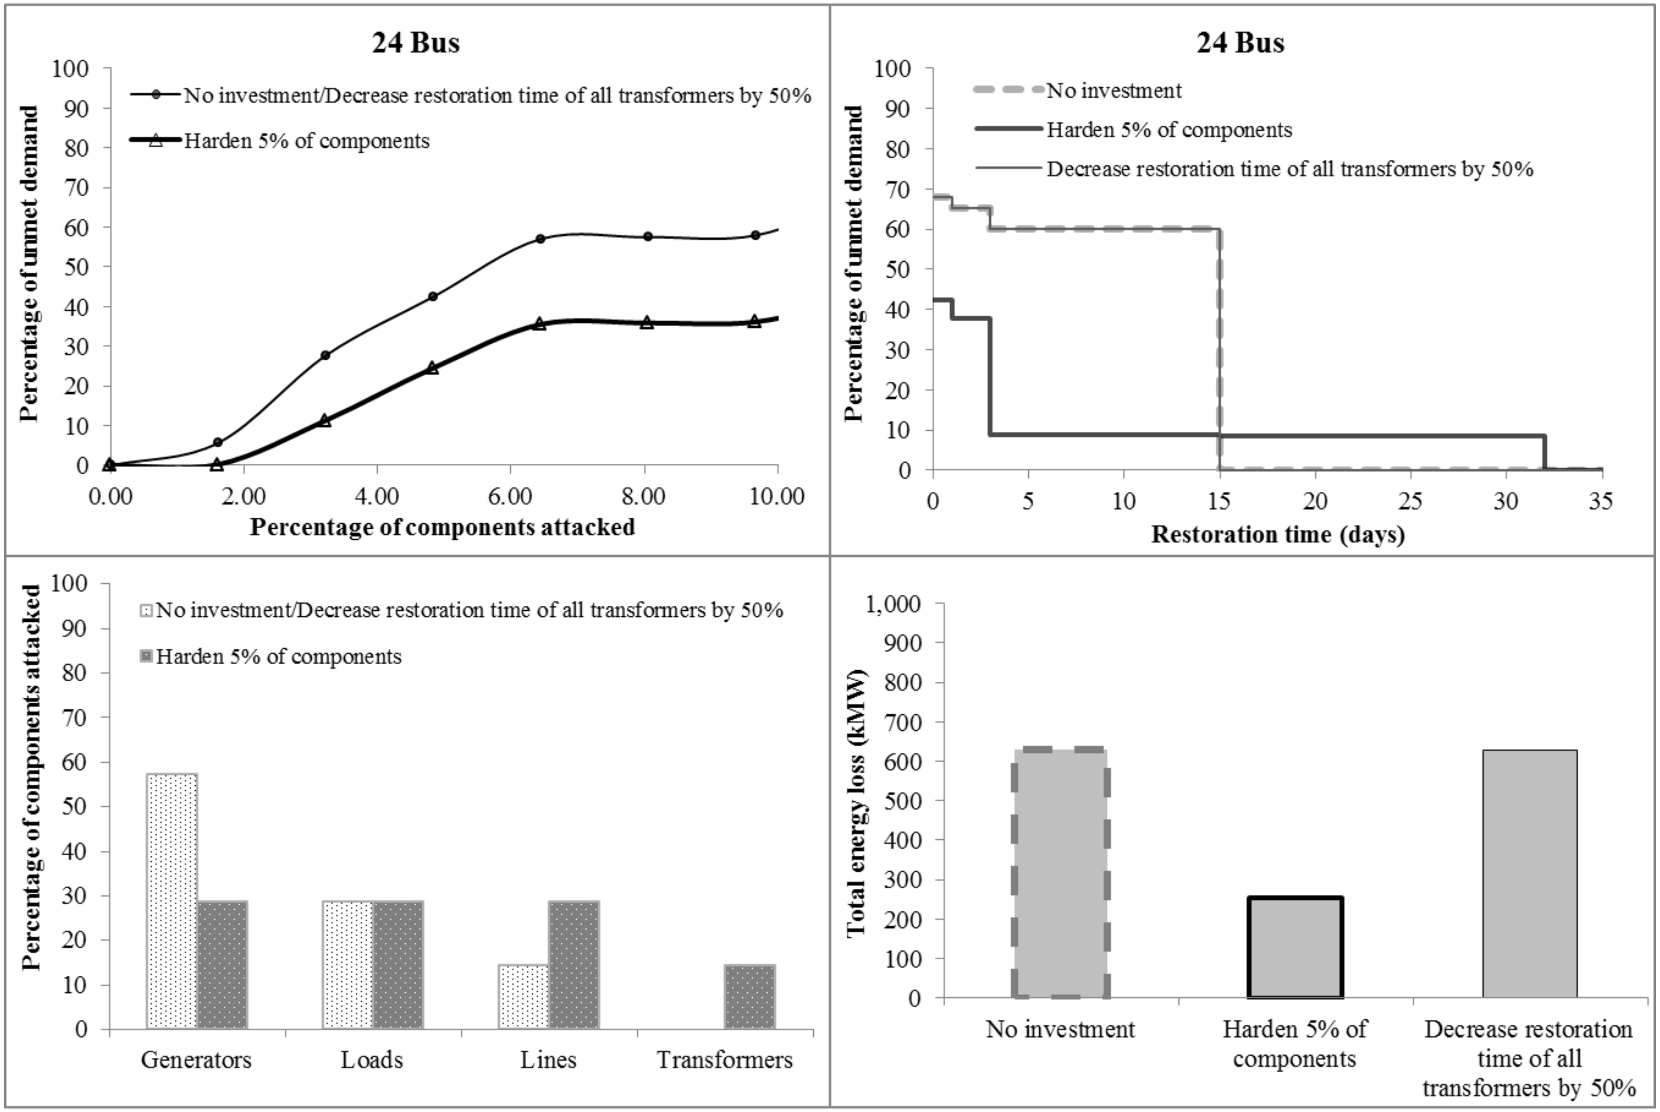

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and not considering restoration times, 48-bus system)

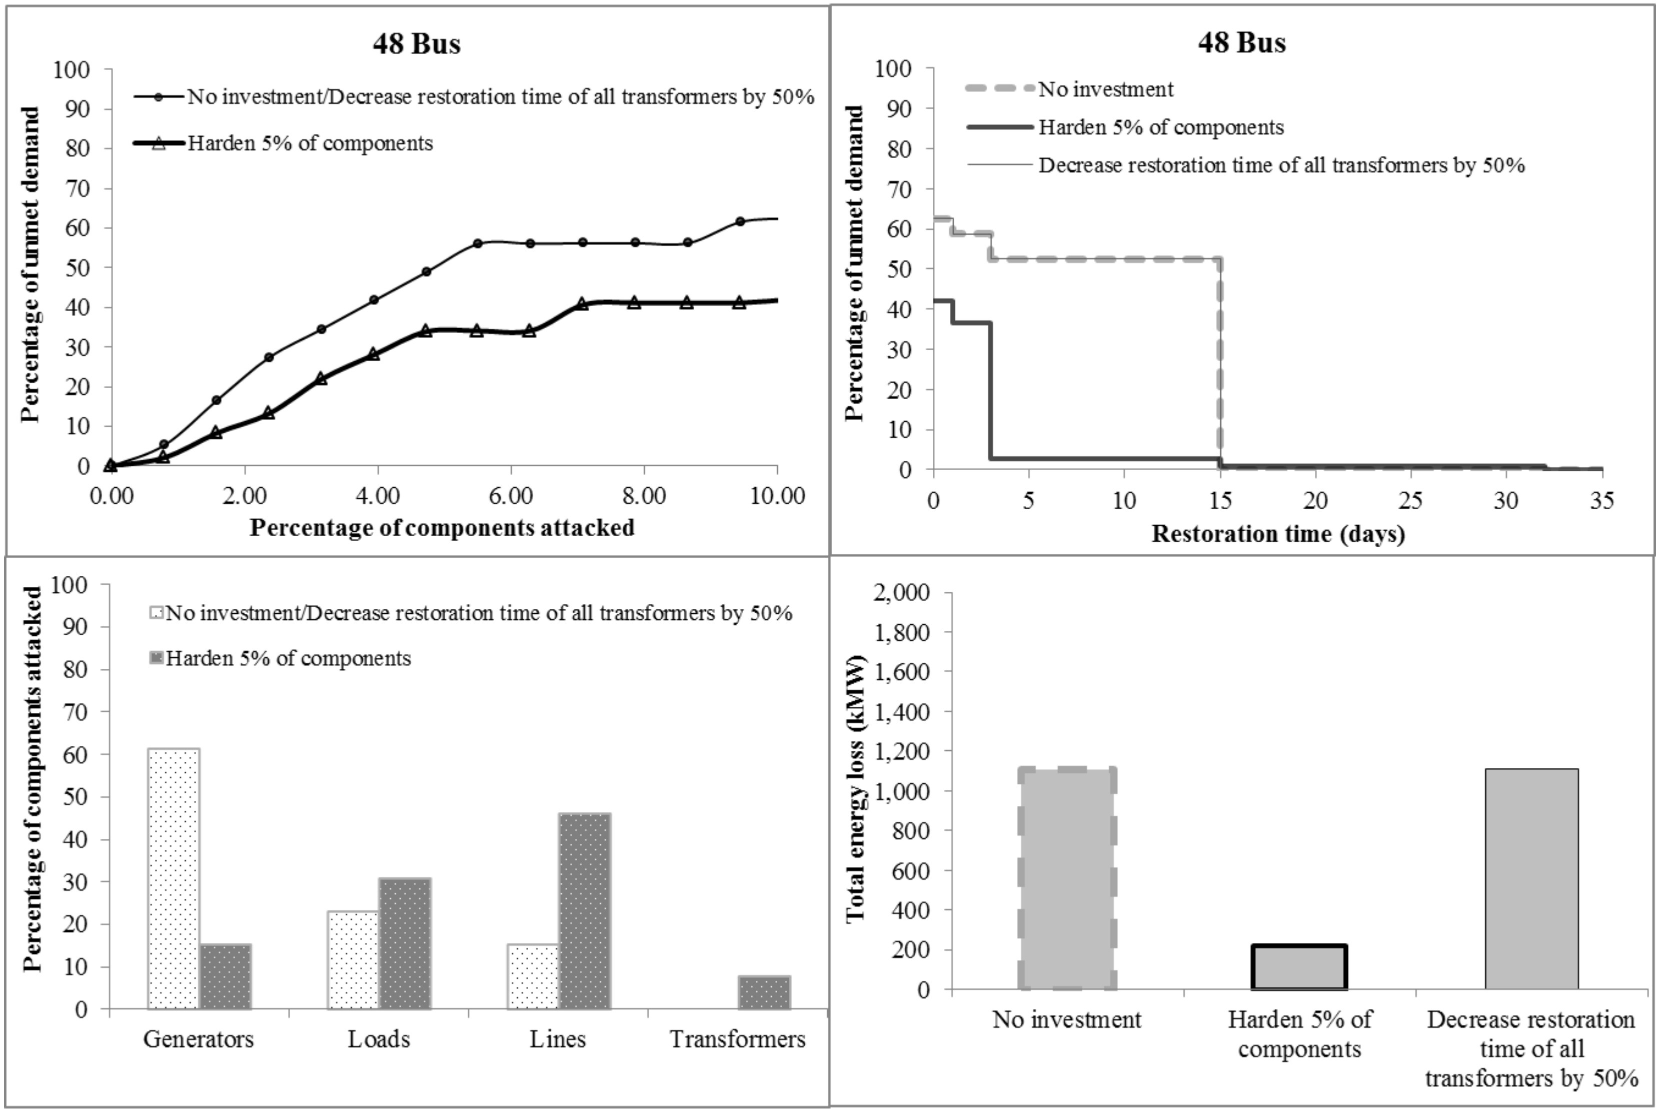

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and not considering restoration times, 118-bus system)

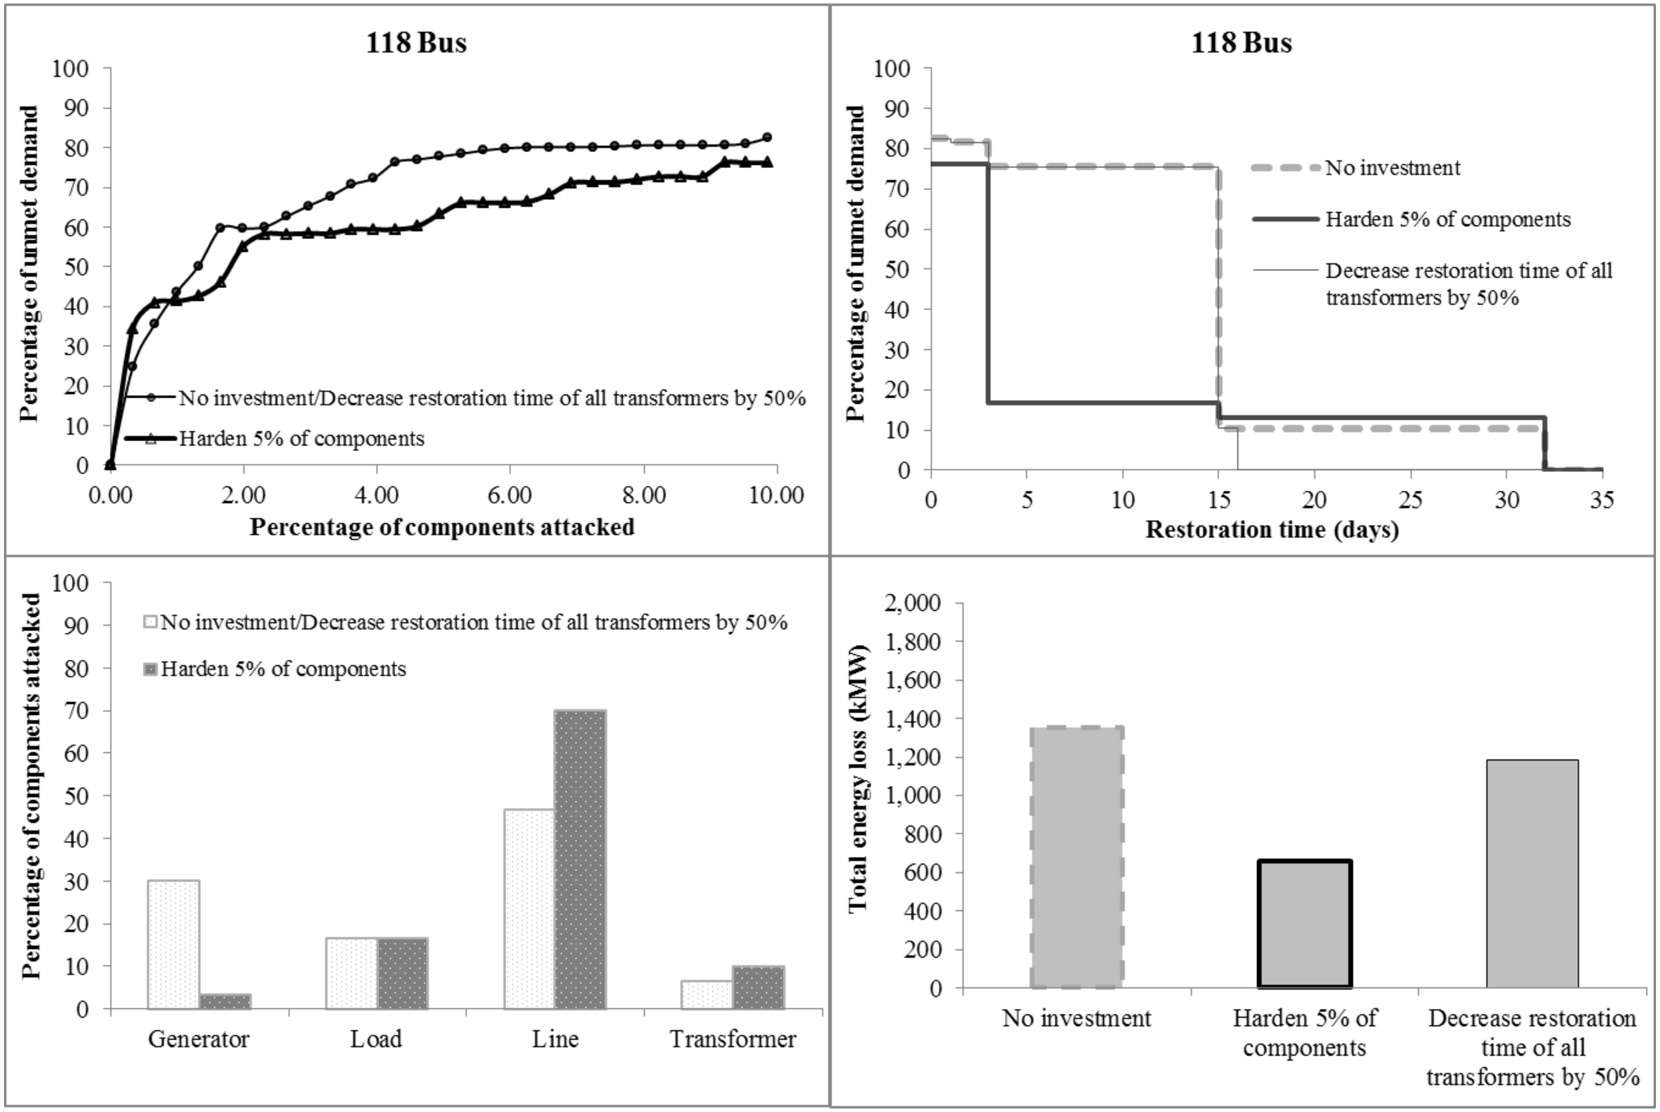

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and not considering restoration times, 300-bus system)

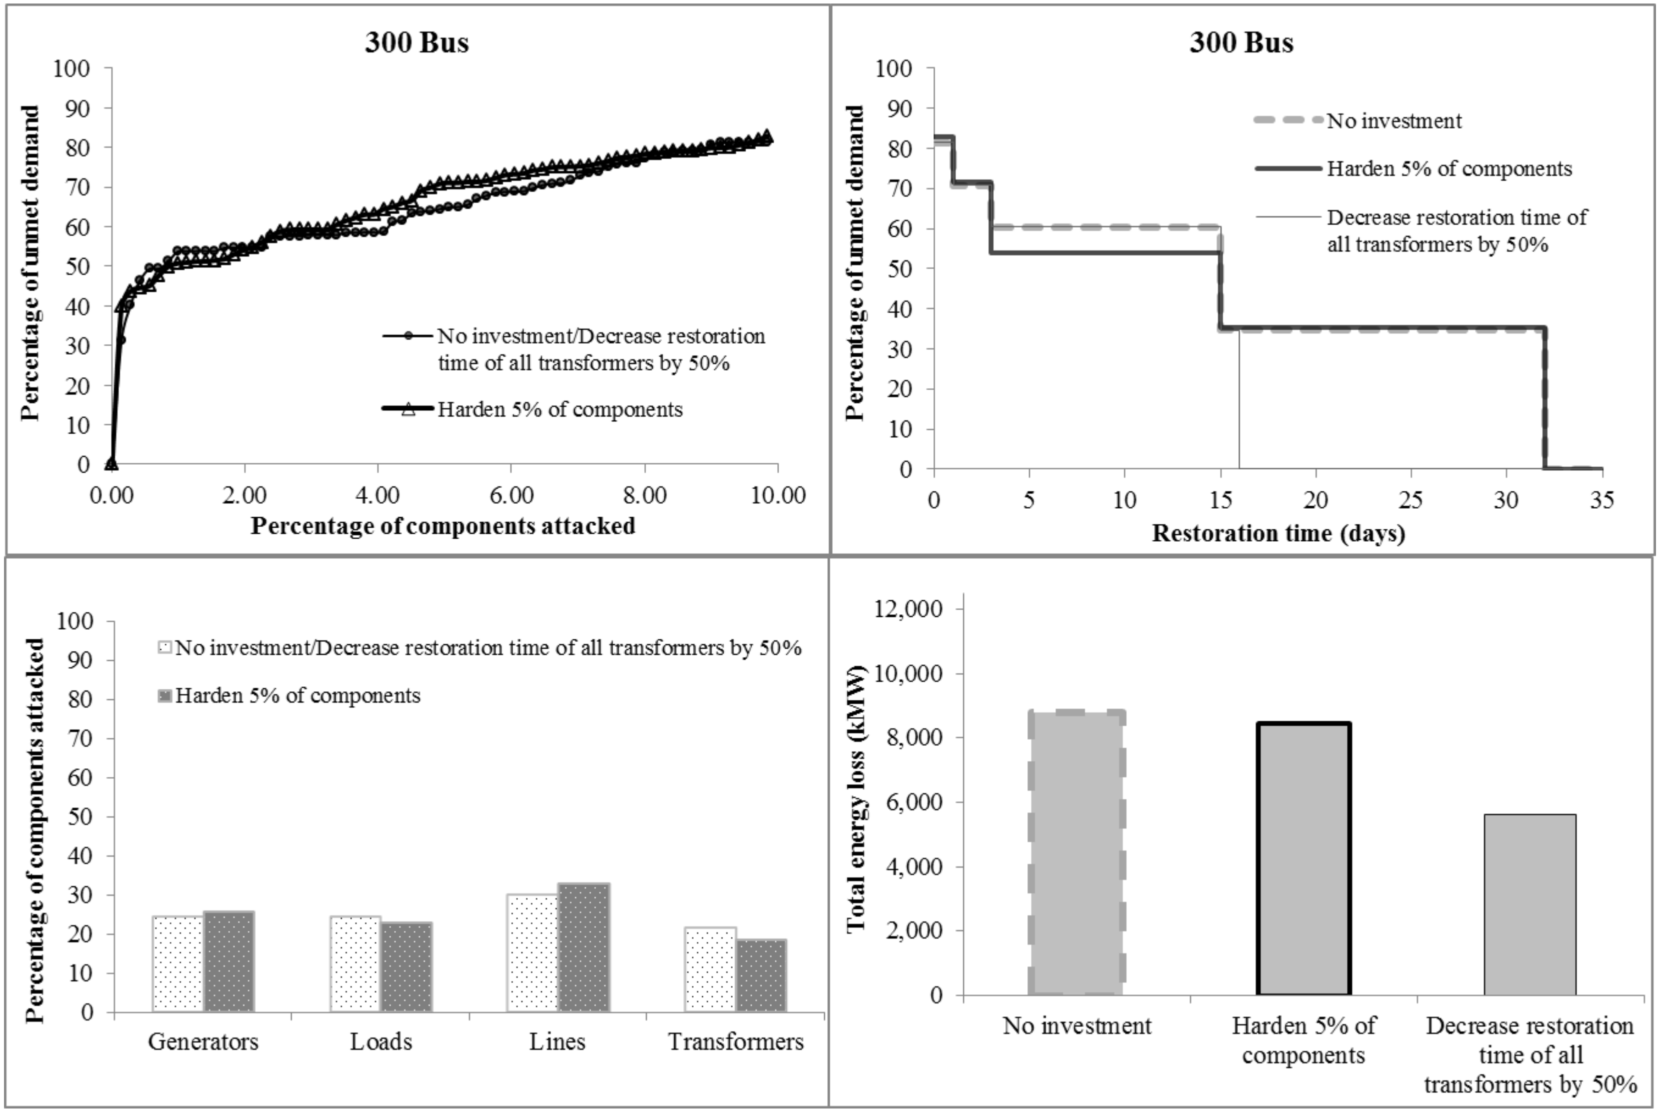

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and not considering restoration times, 24-bus system)

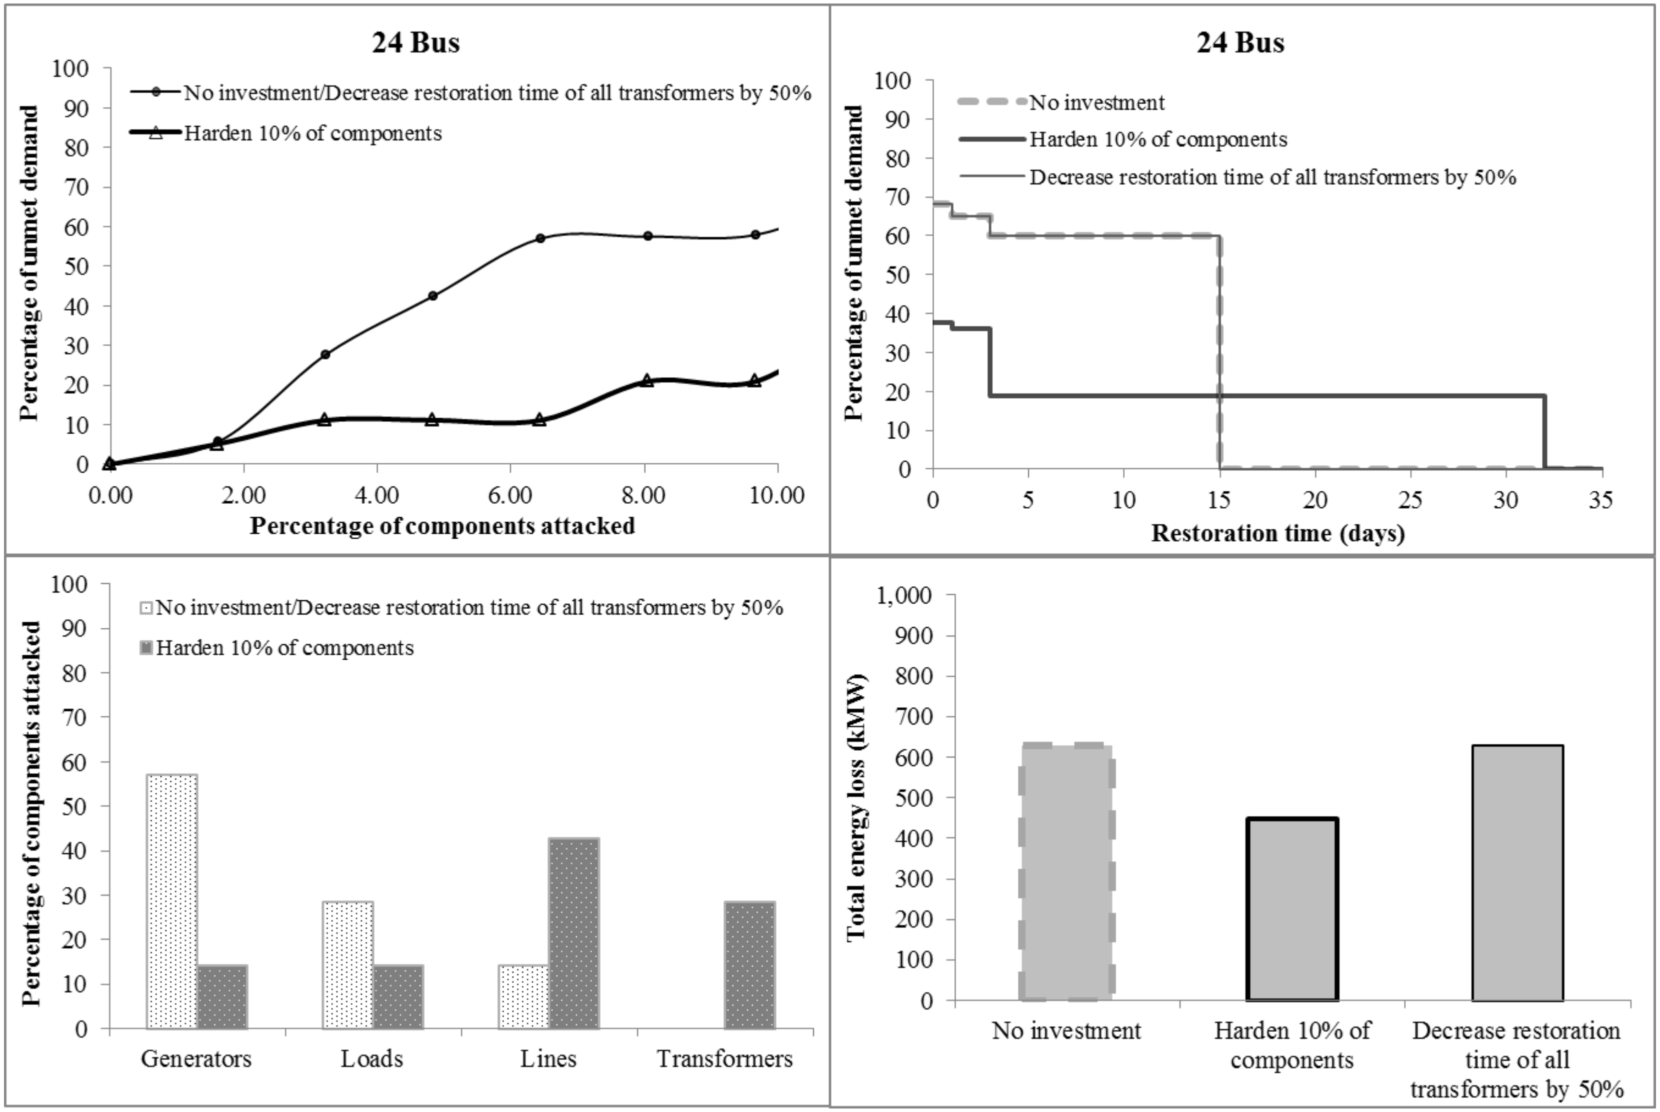

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and not considering restoration times, 48-bus system)

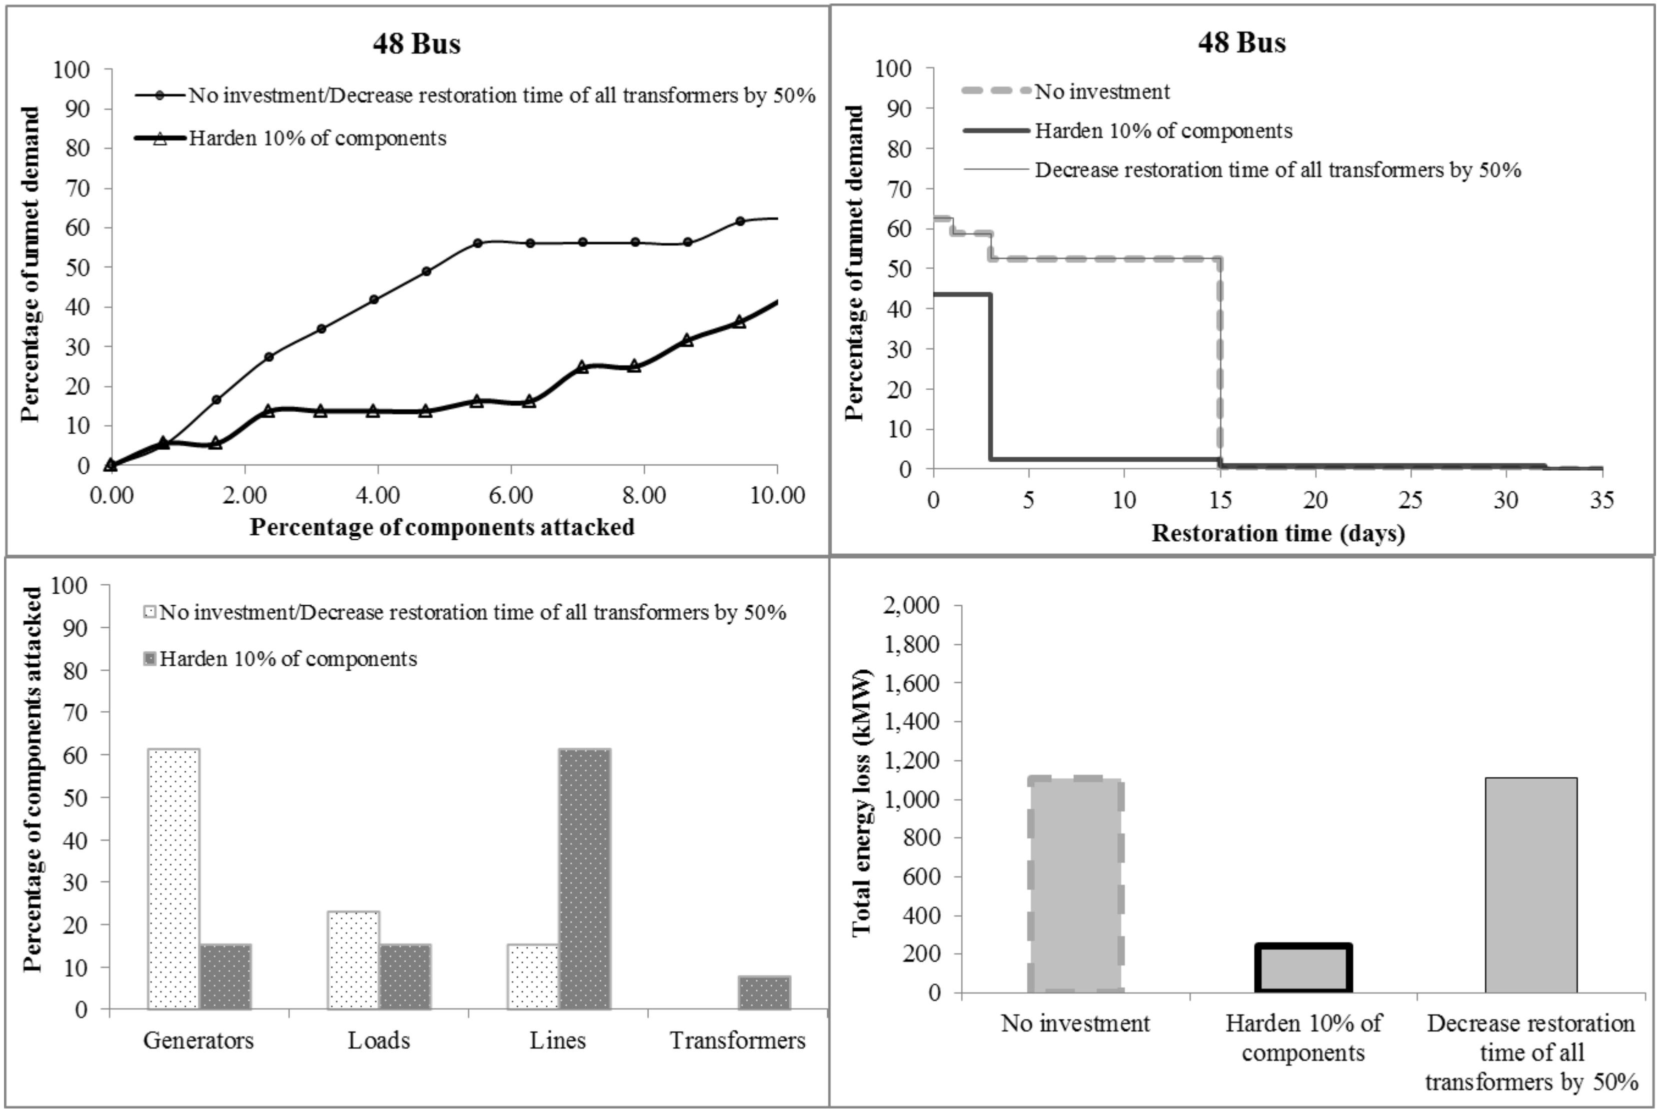

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and not considering restoration times, 118-bus system)

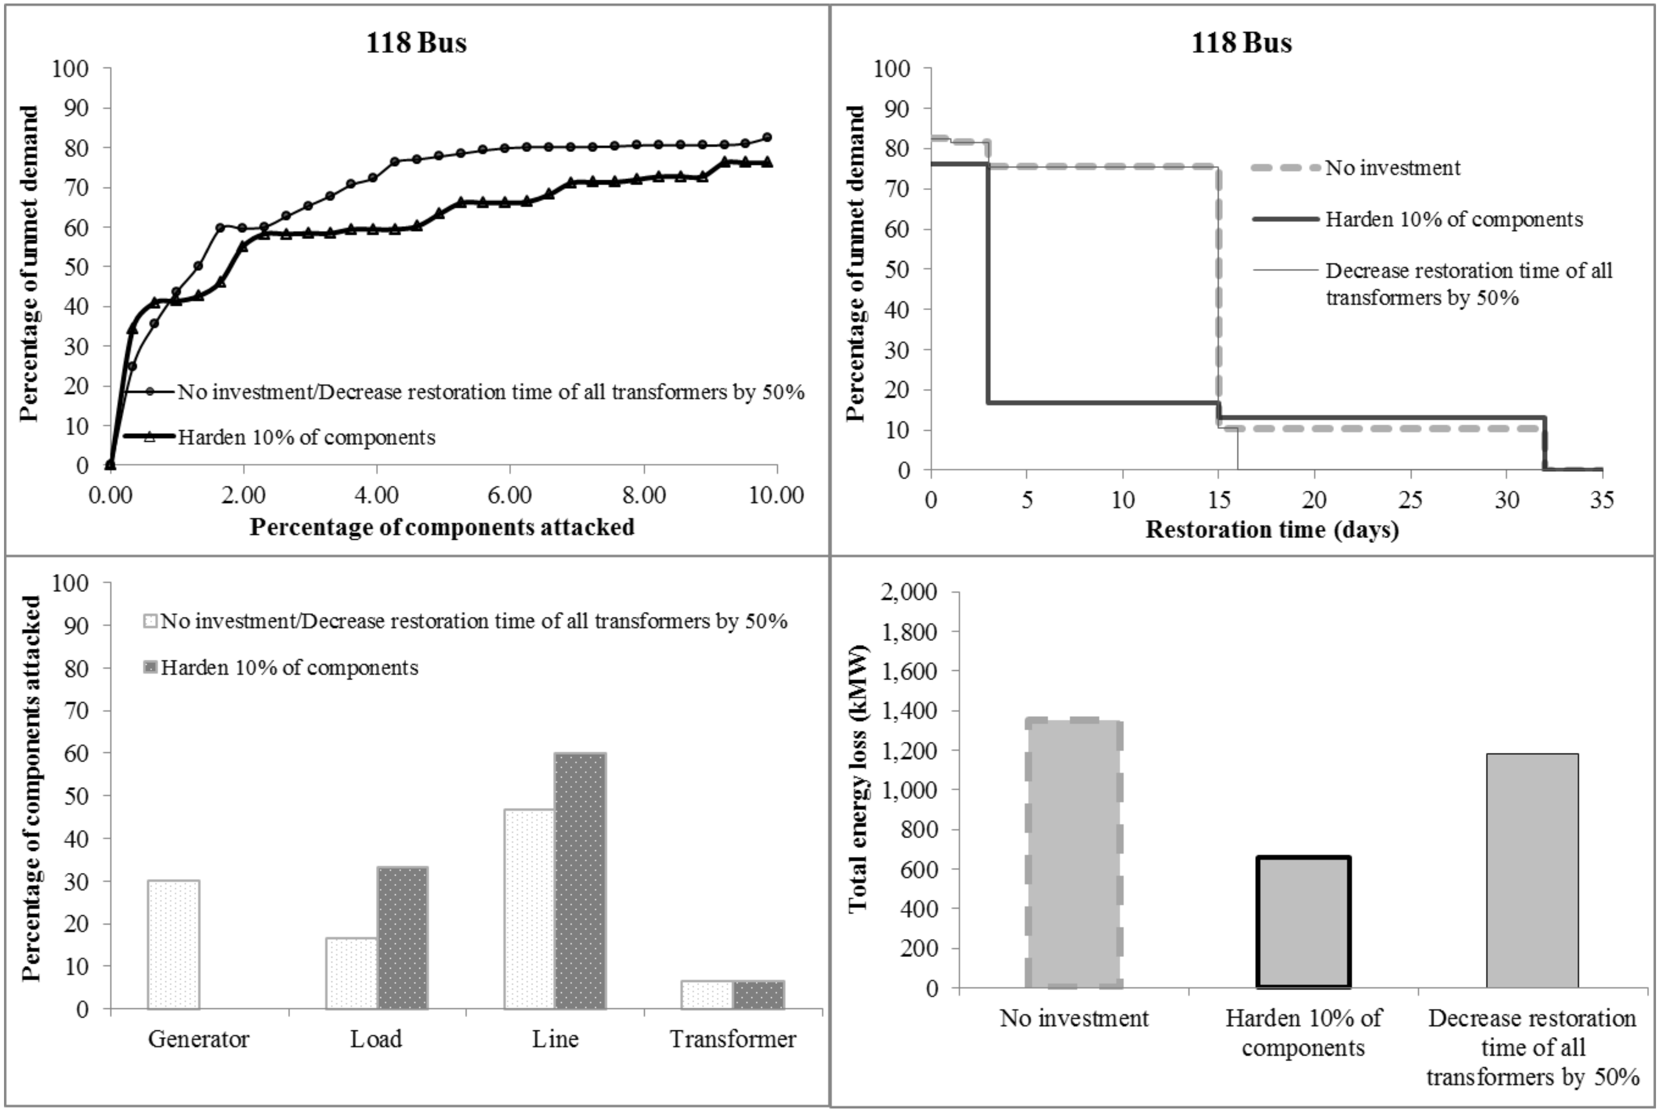

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and not considering restoration times, 300-bus system)

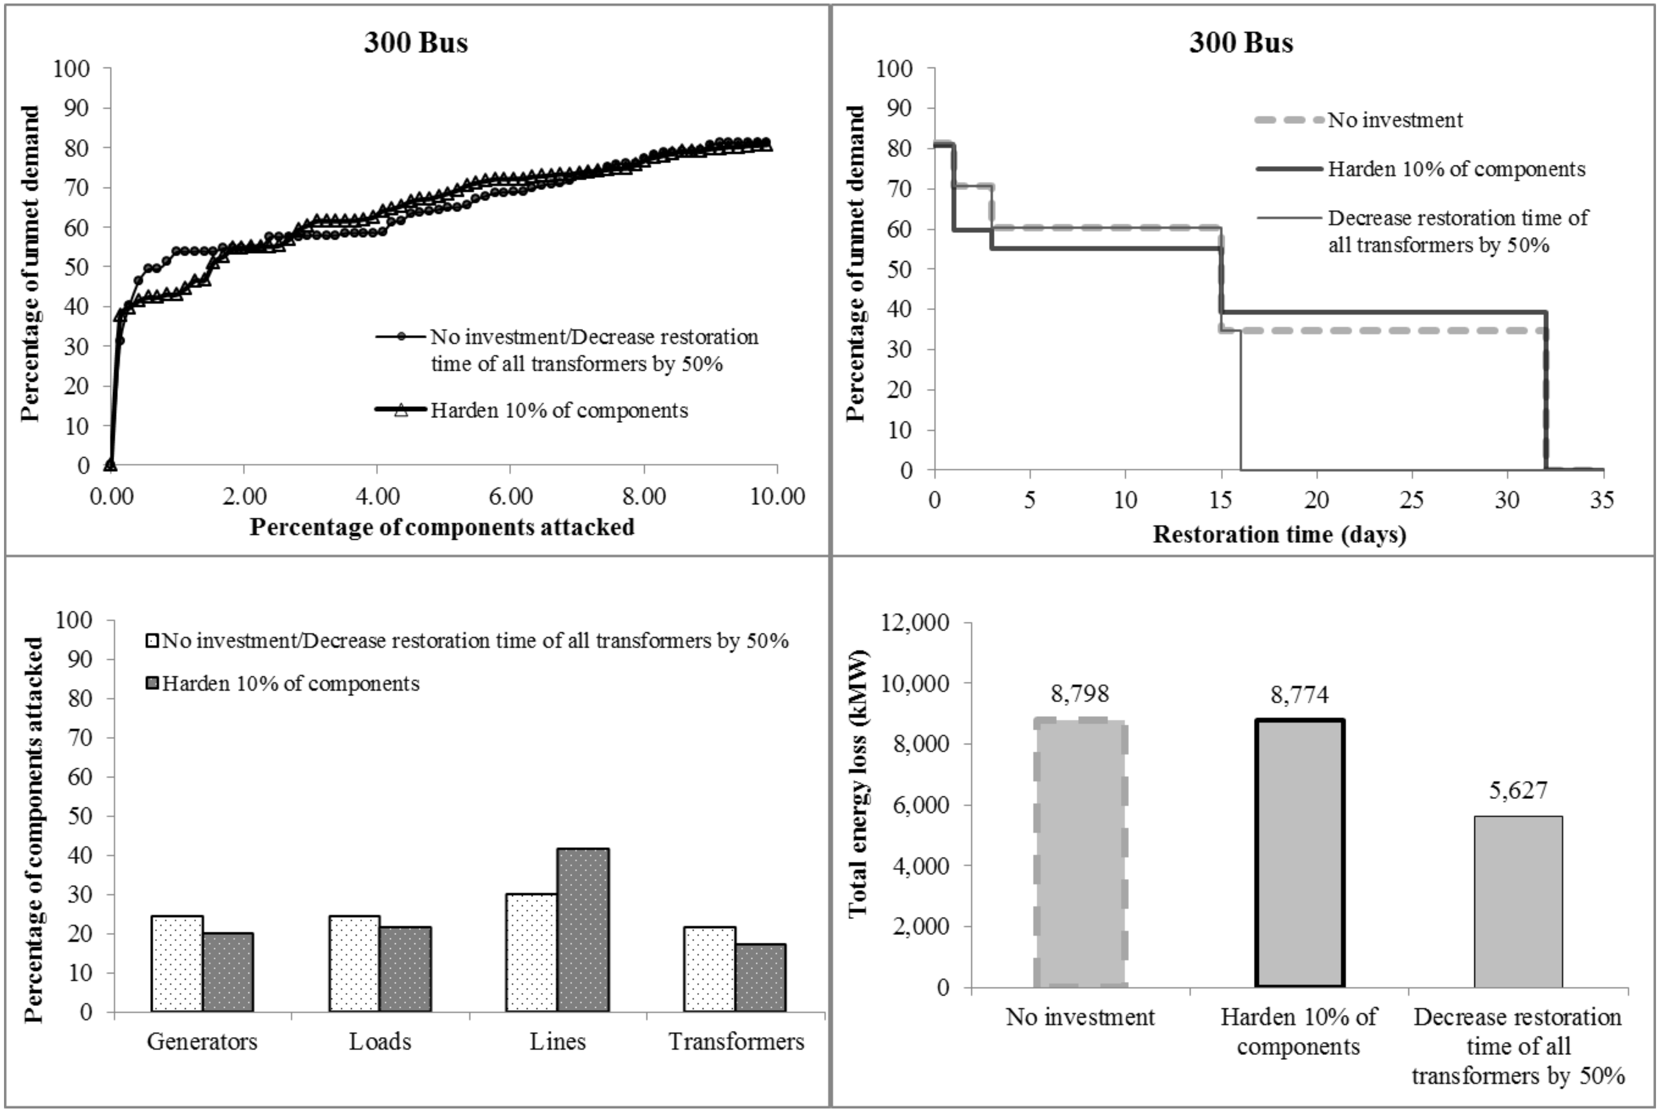

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and not considering restoration times, 24-bus system)

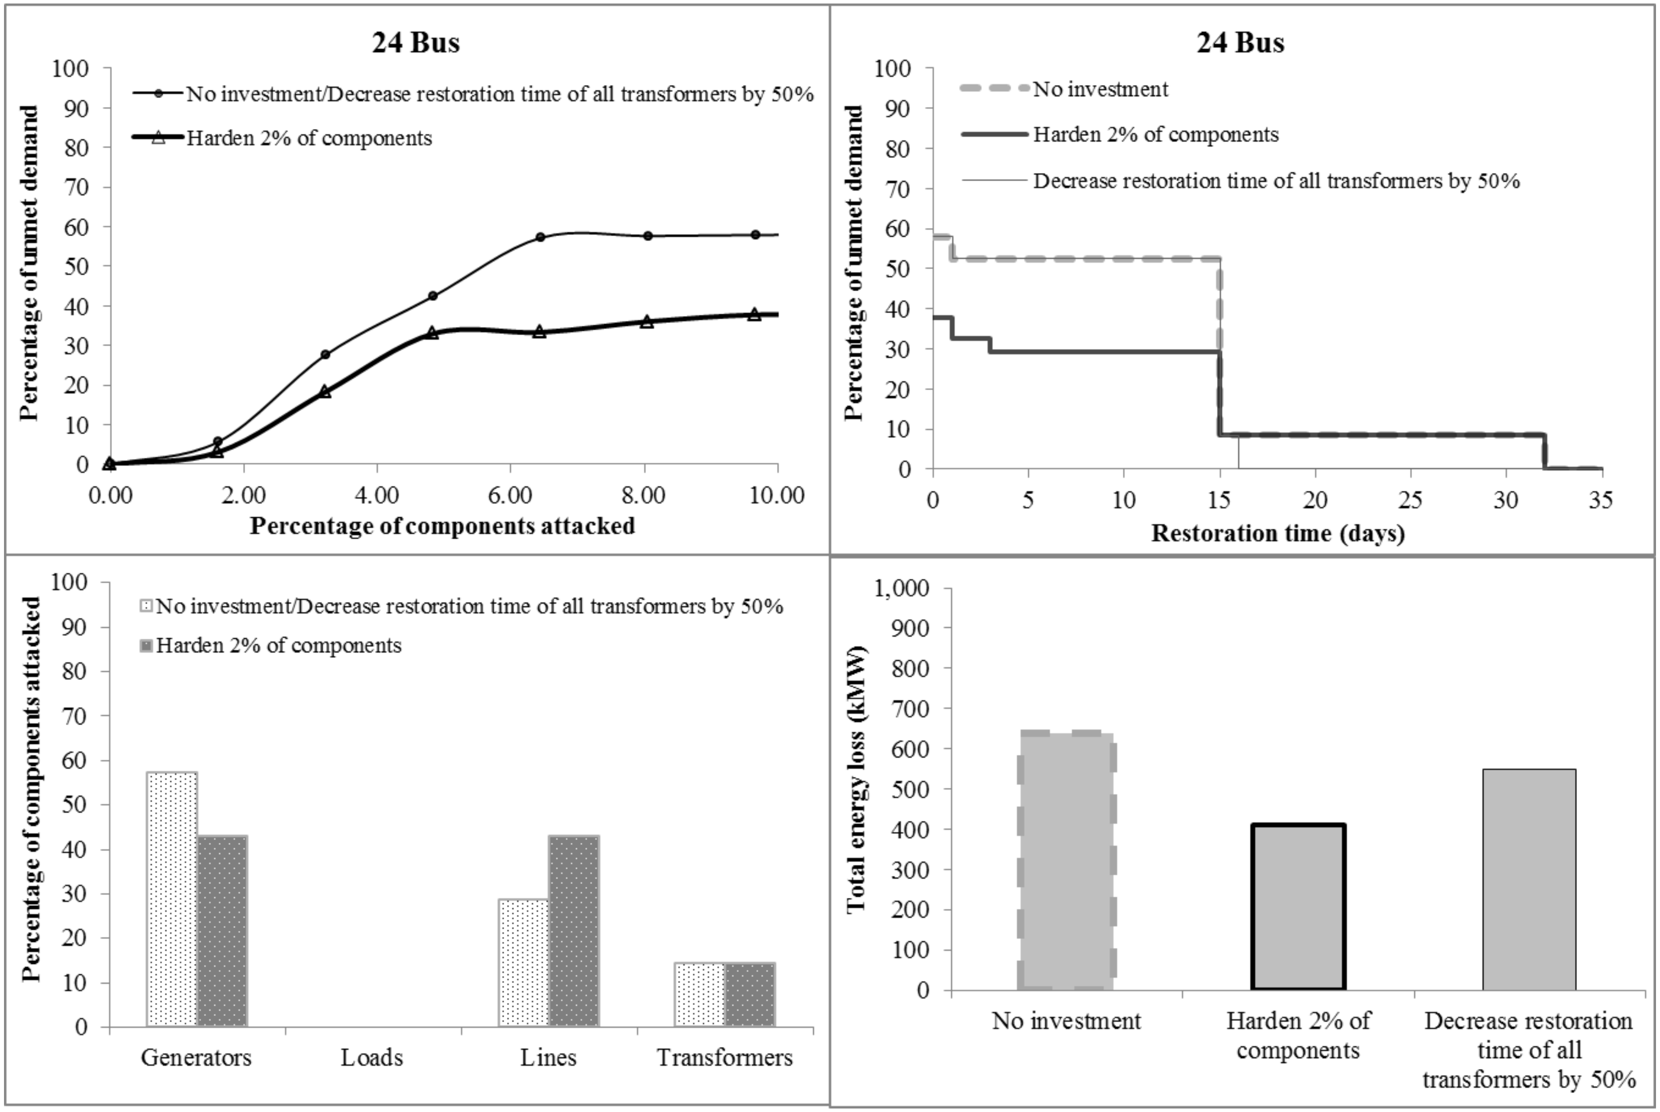

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and not considering restoration times, 48-bus system)

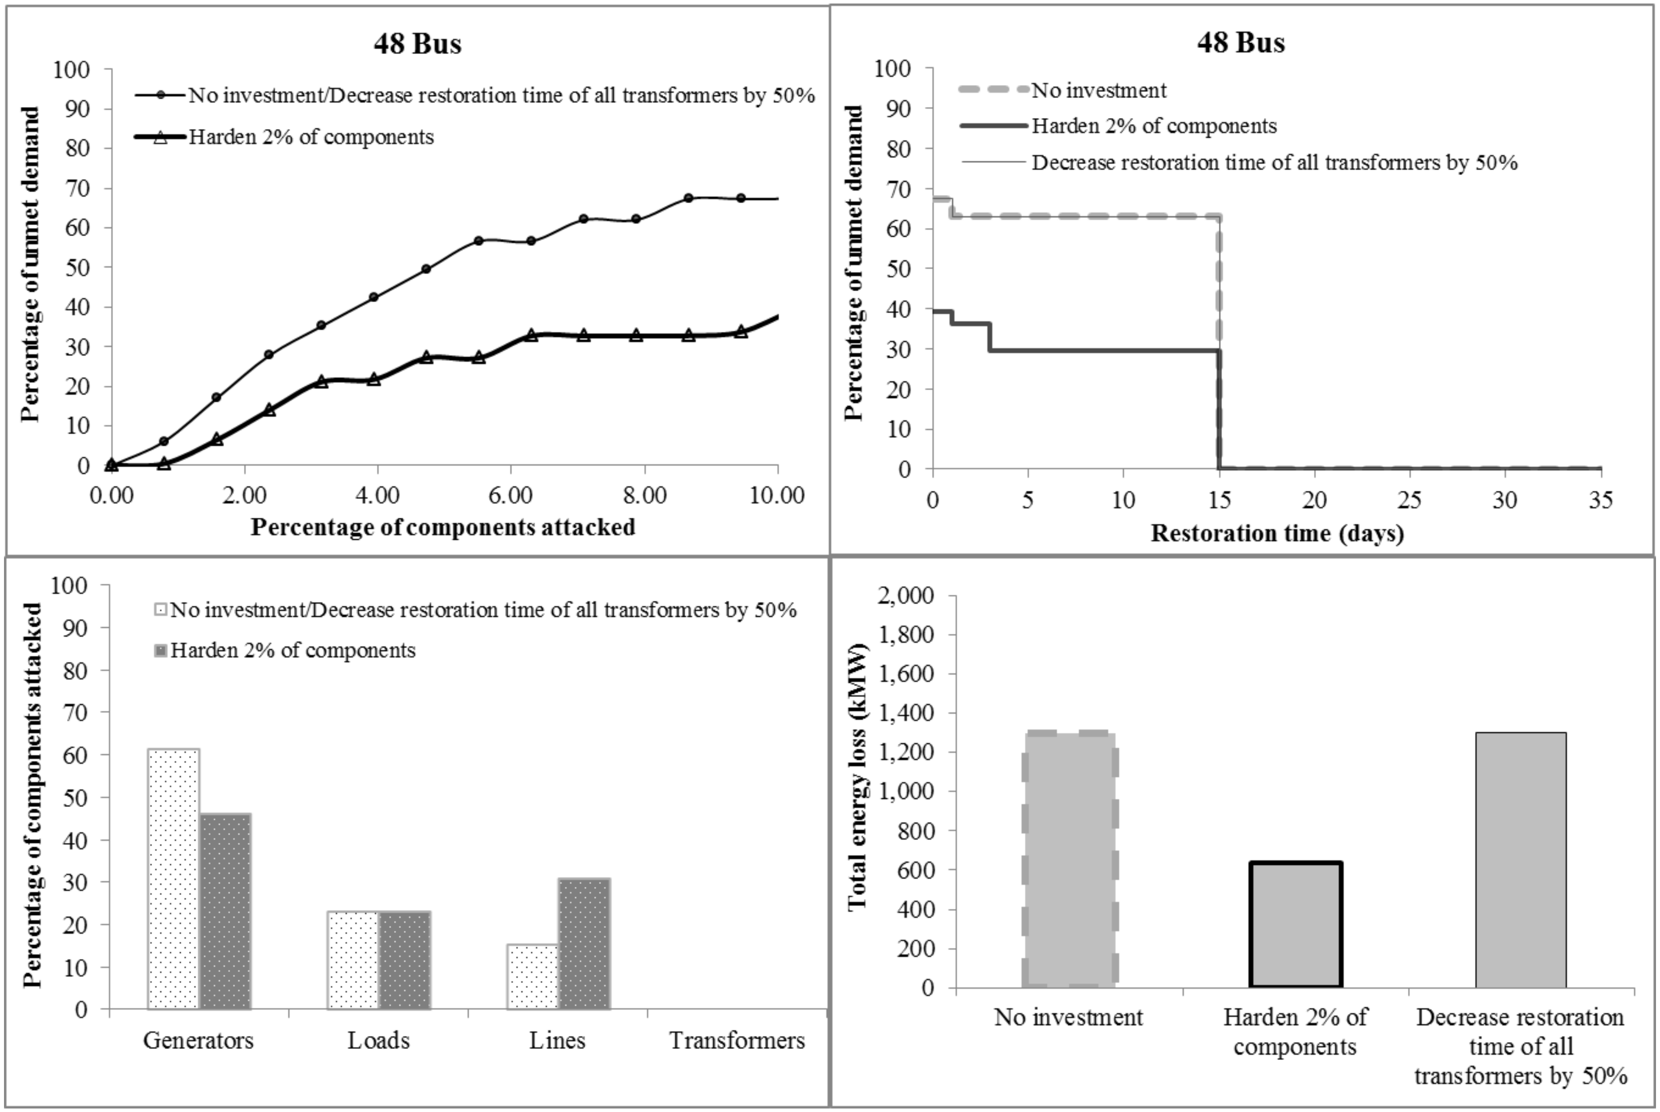

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and not considering restoration times, 118-bus system)

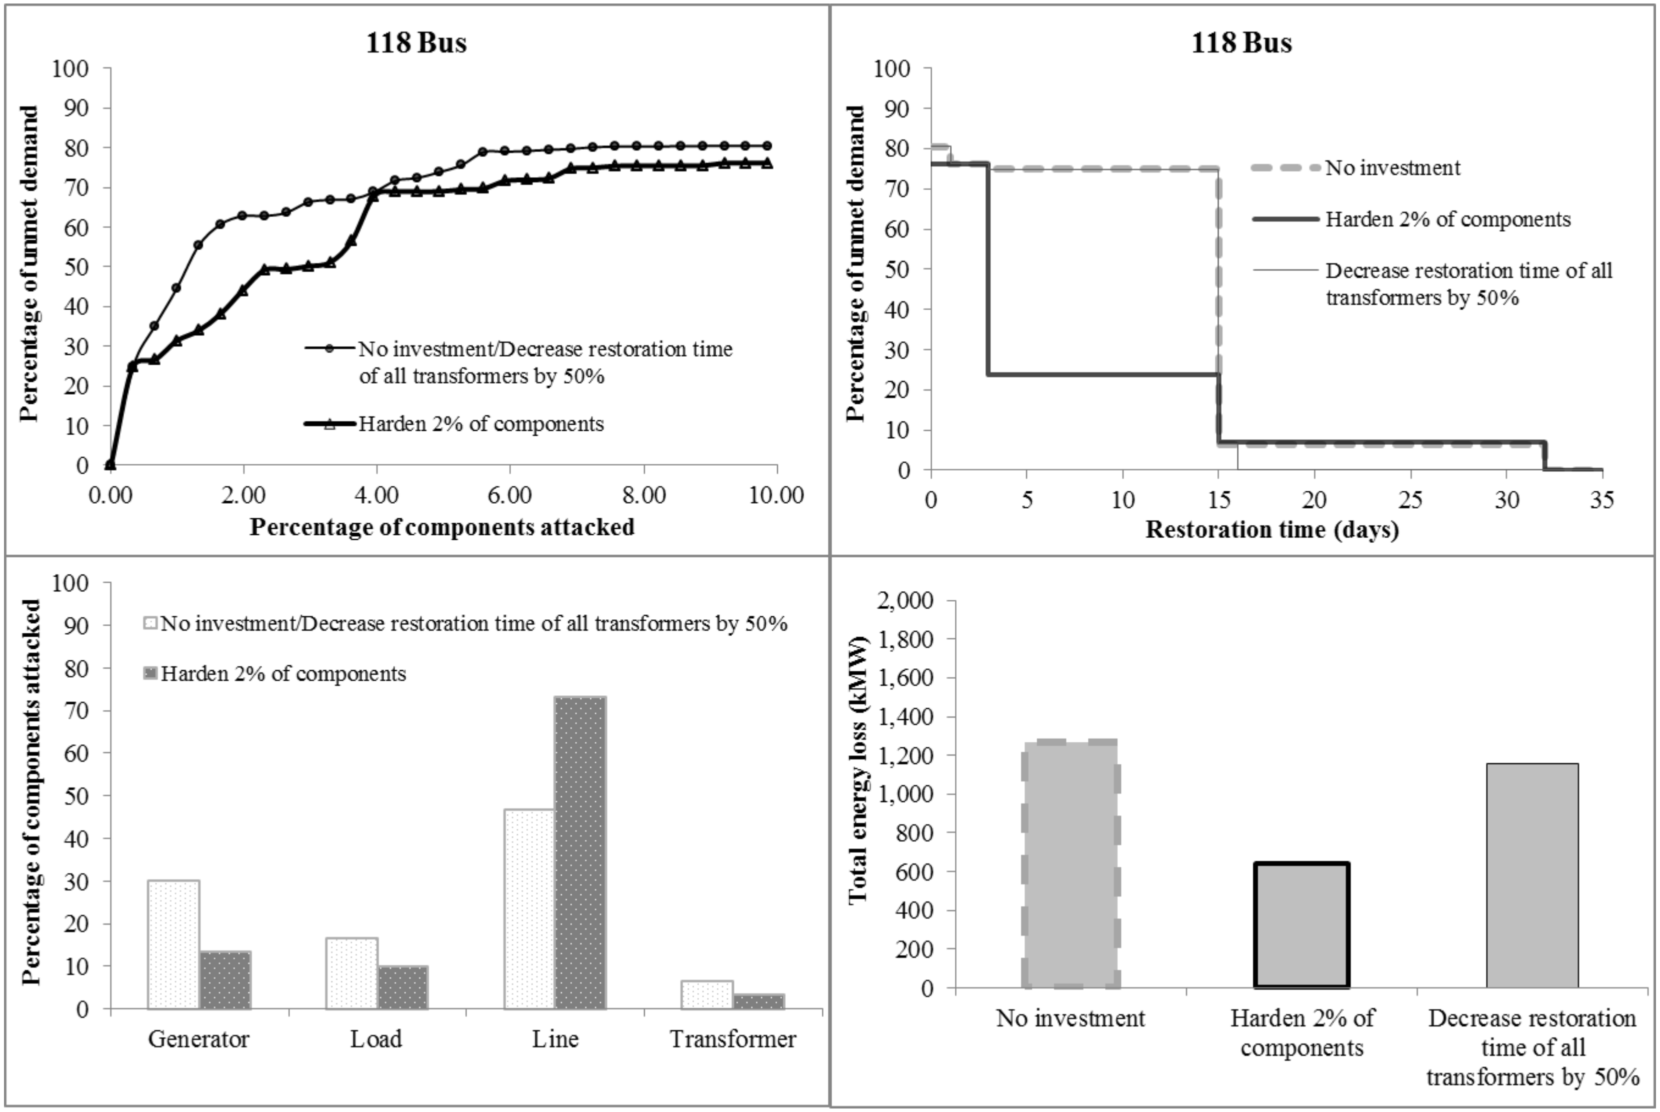

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and not considering restoration times, 300-bus system)

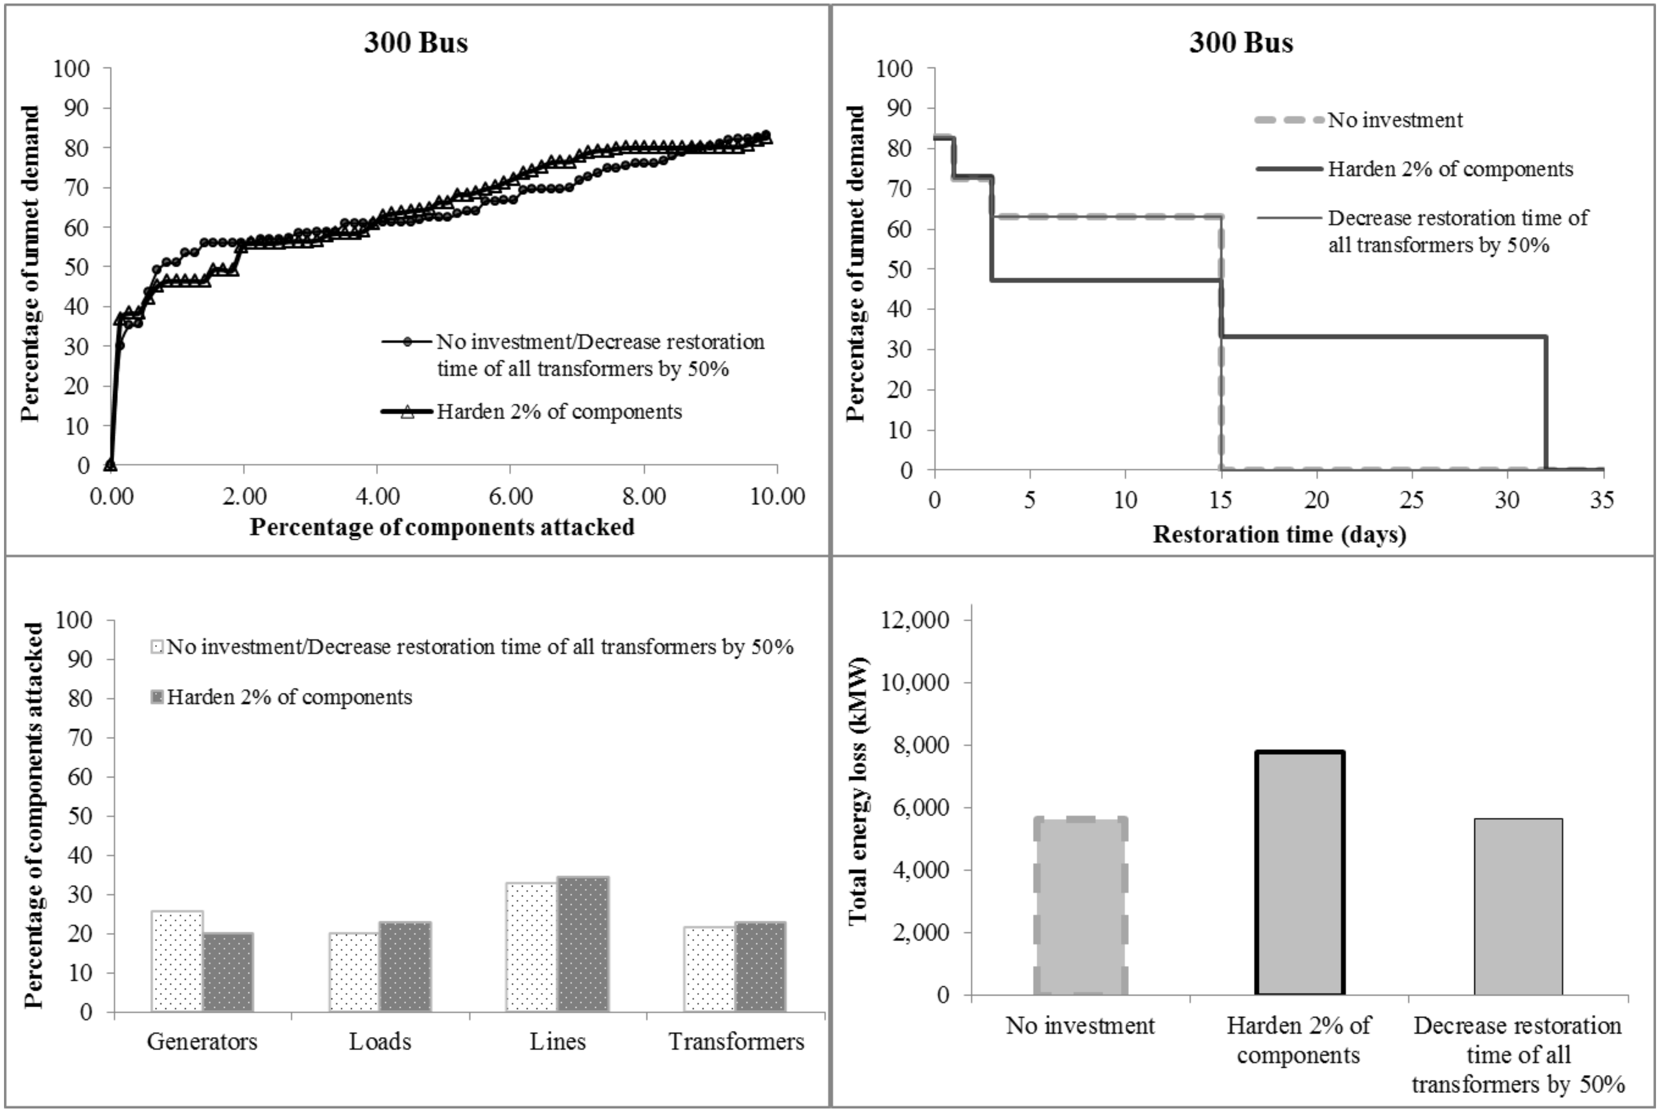

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and not considering restoration times, 24-bus system)

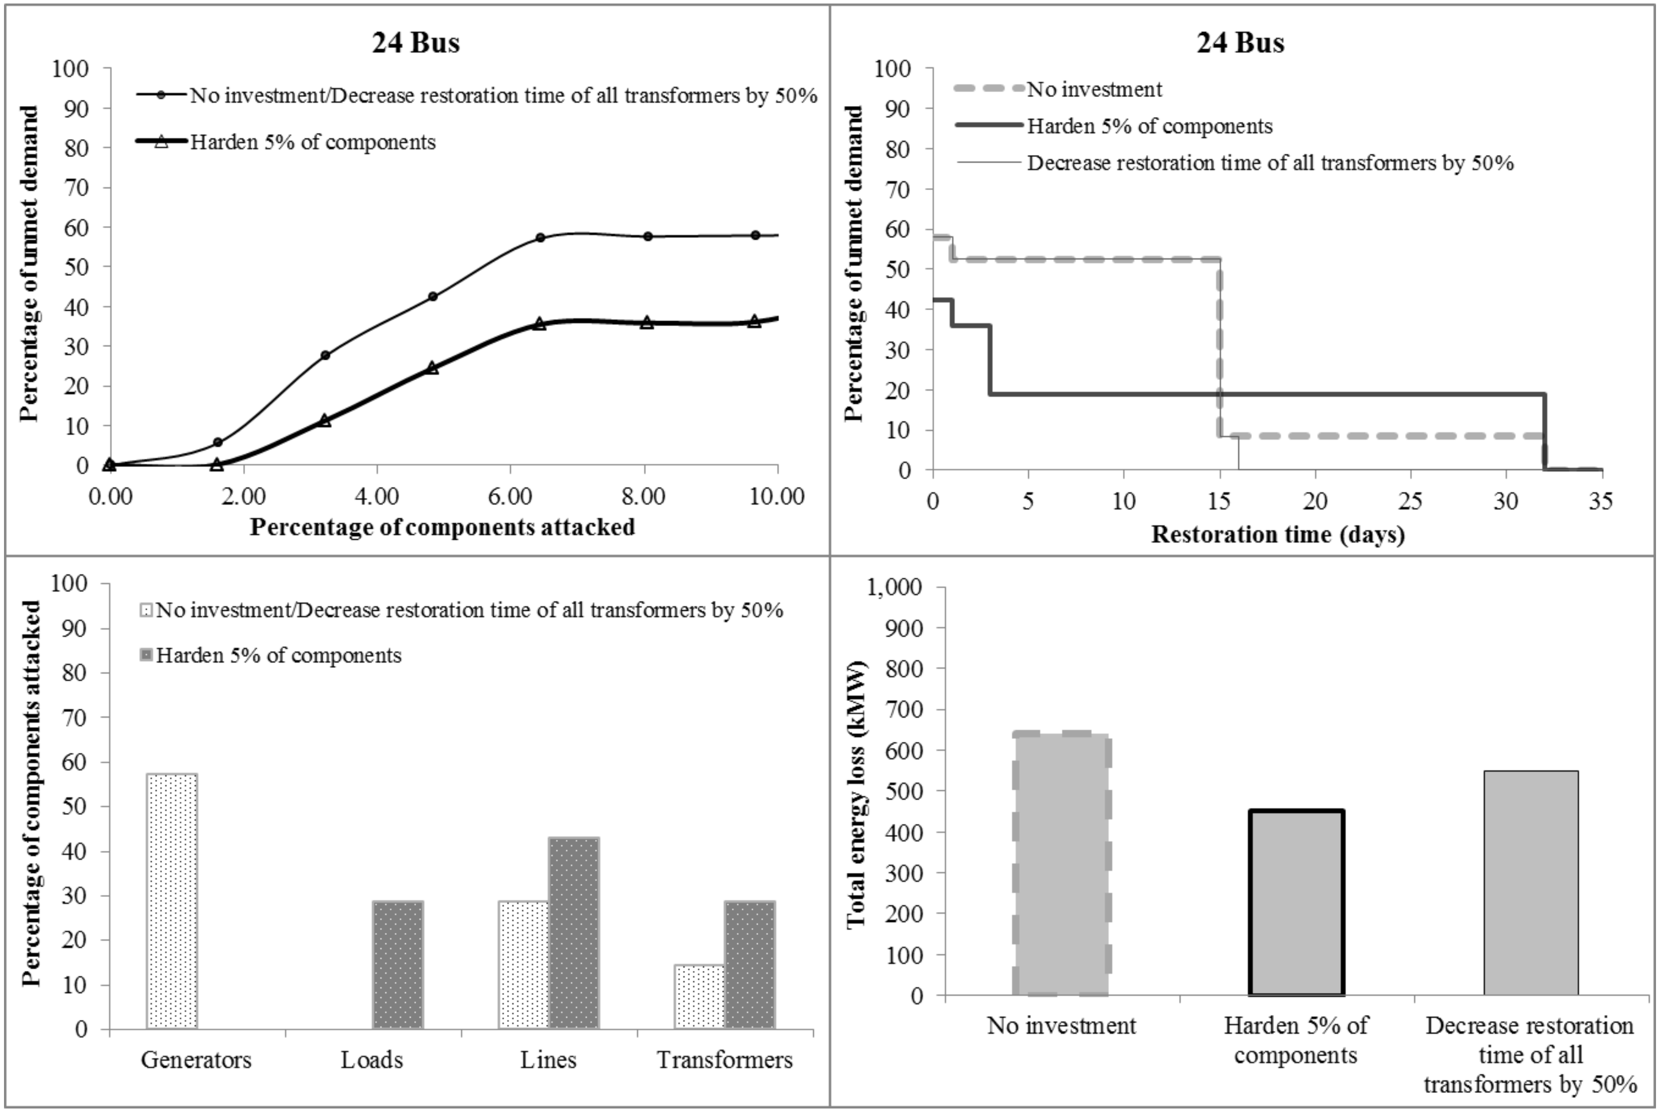

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and not considering restoration times, 48-bus system)

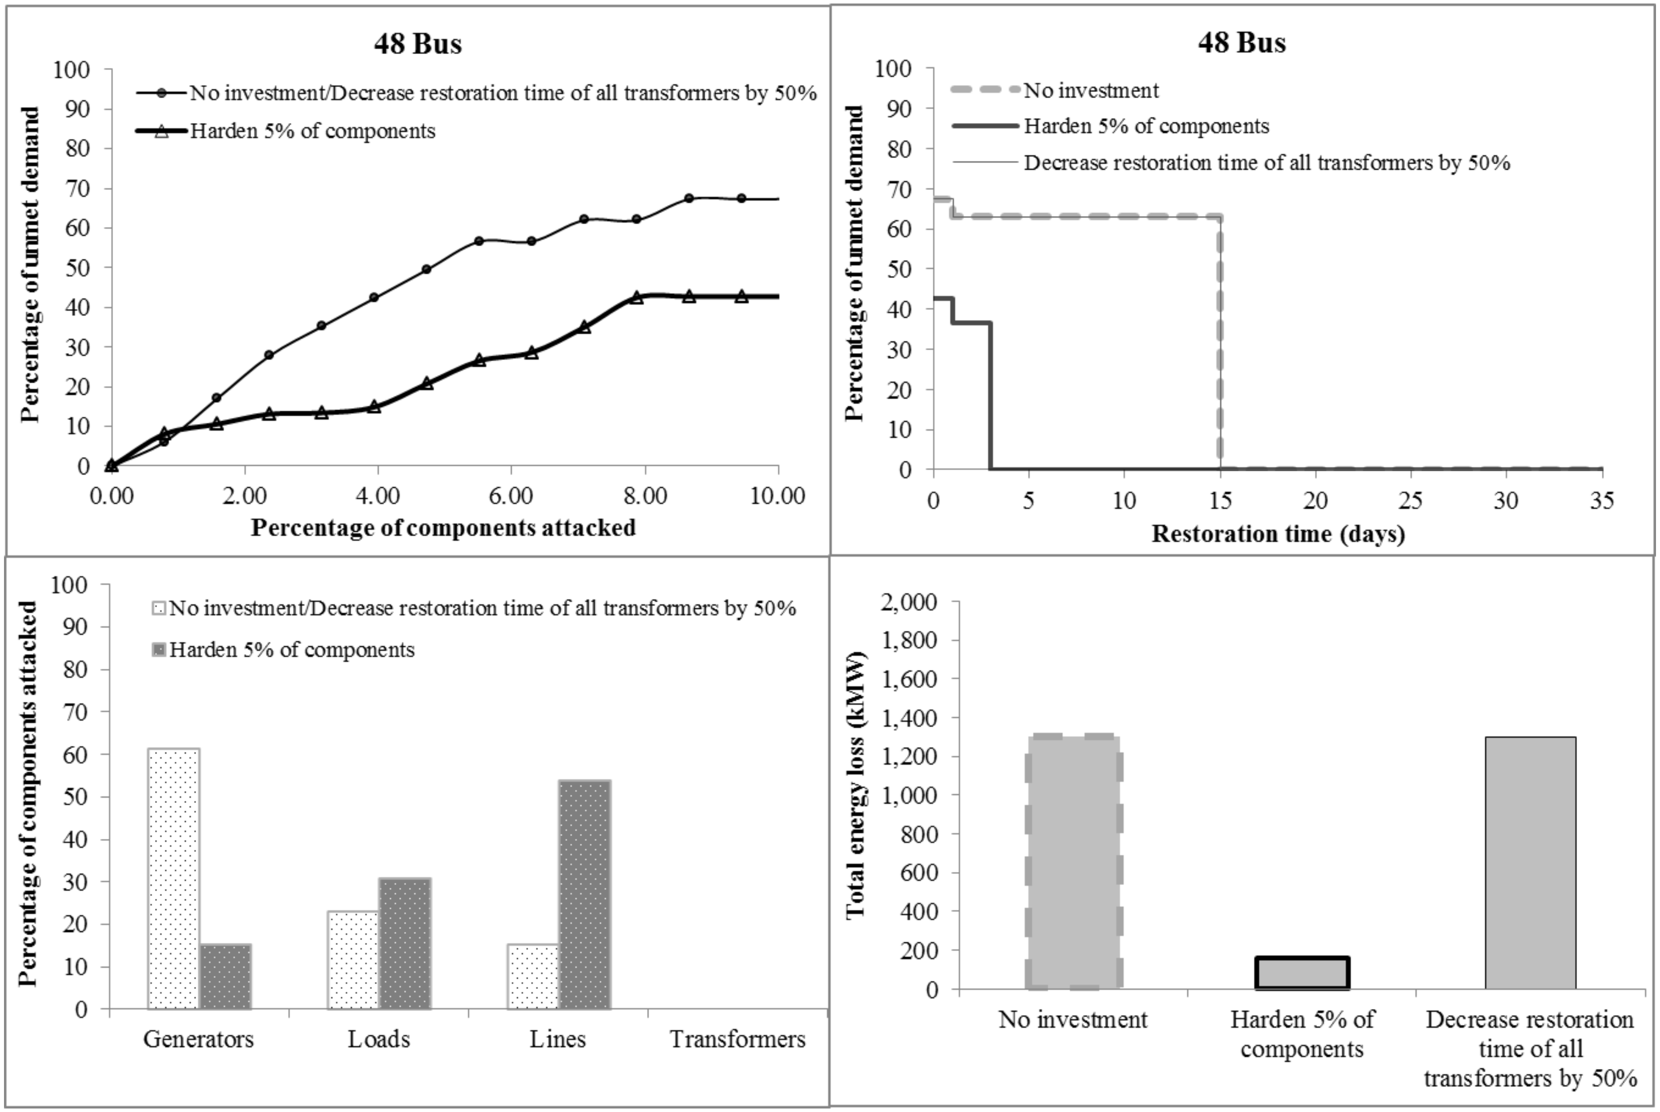

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and not considering restoration times, 118-bus system)

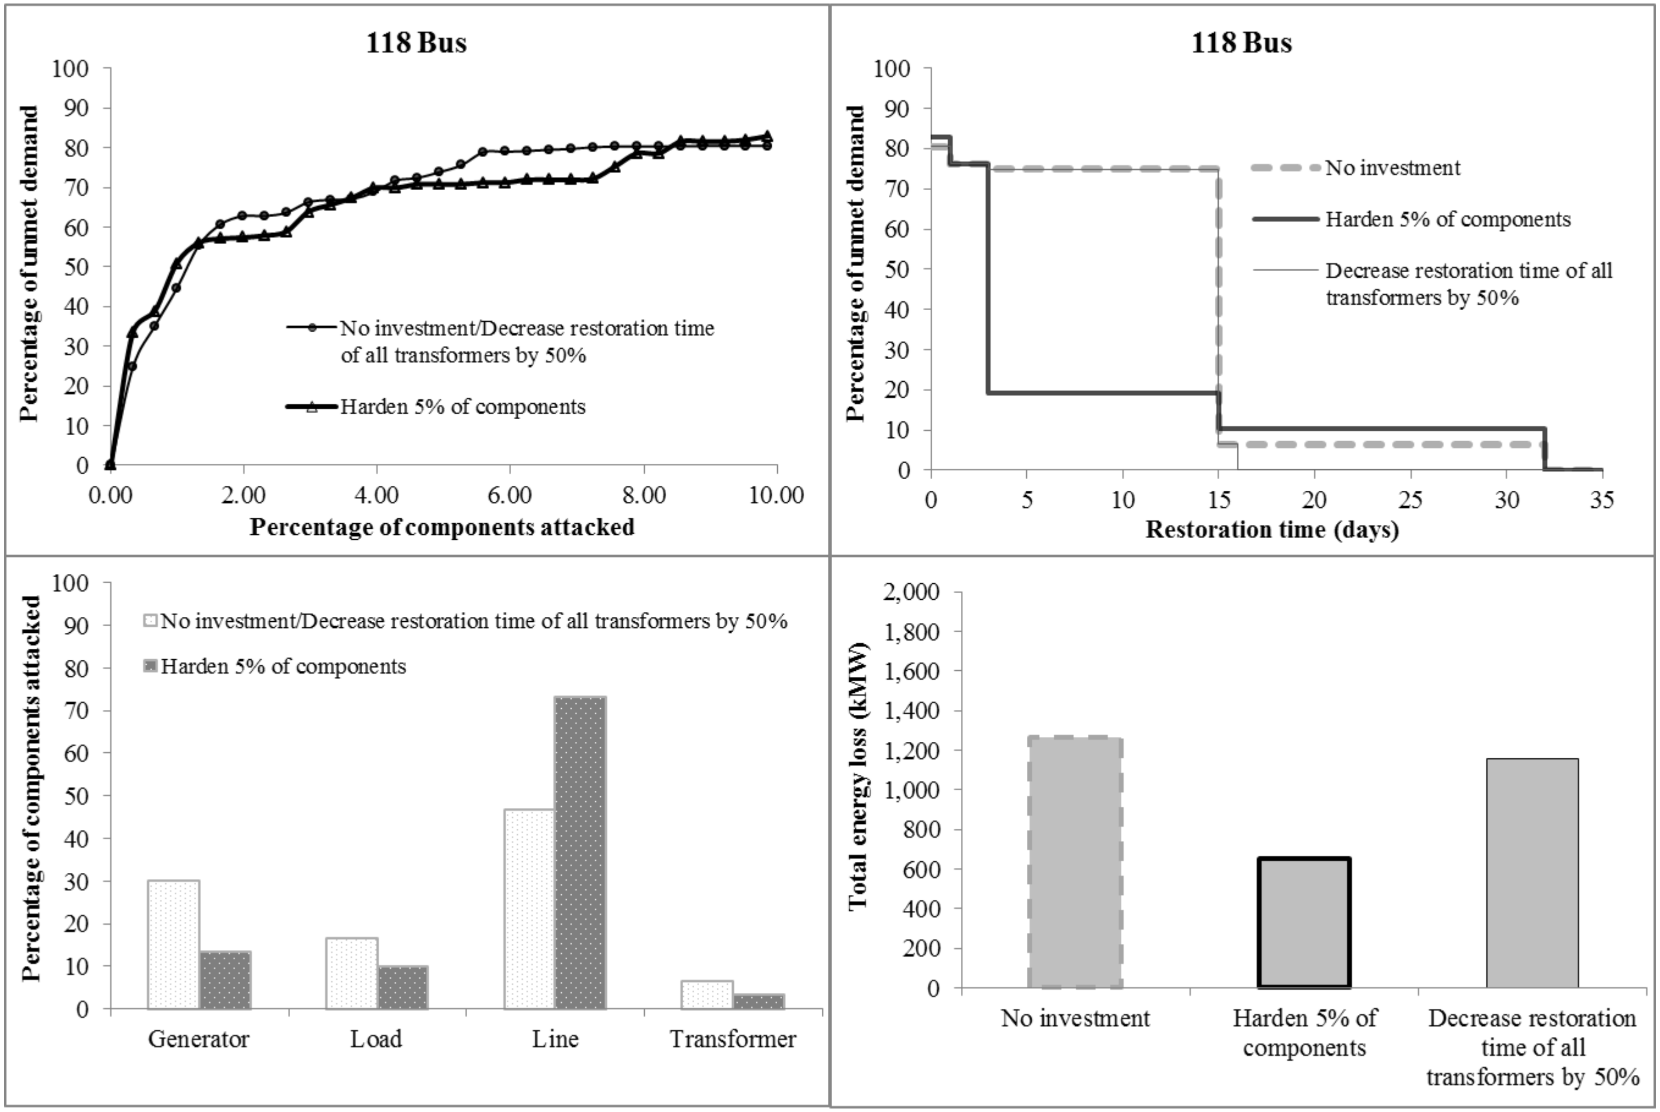

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and not considering restoration times, 300-bus system)

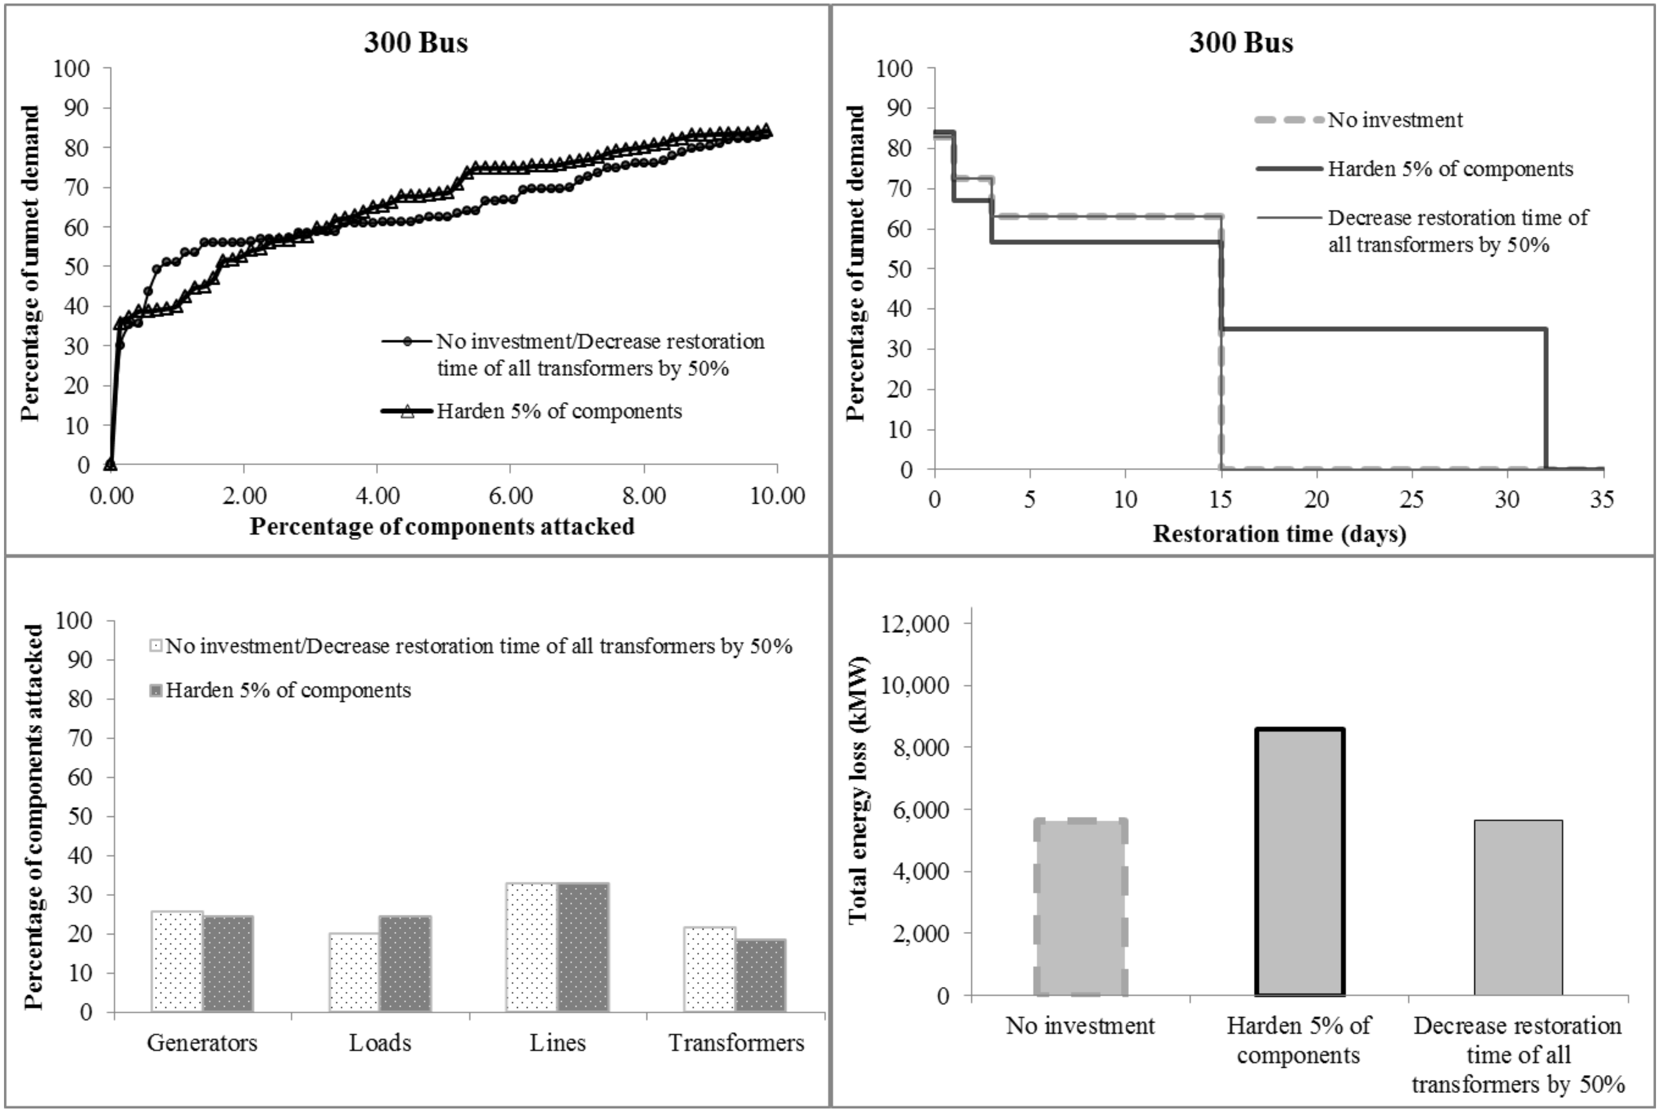

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and not considering restoration times, 24-bus system)

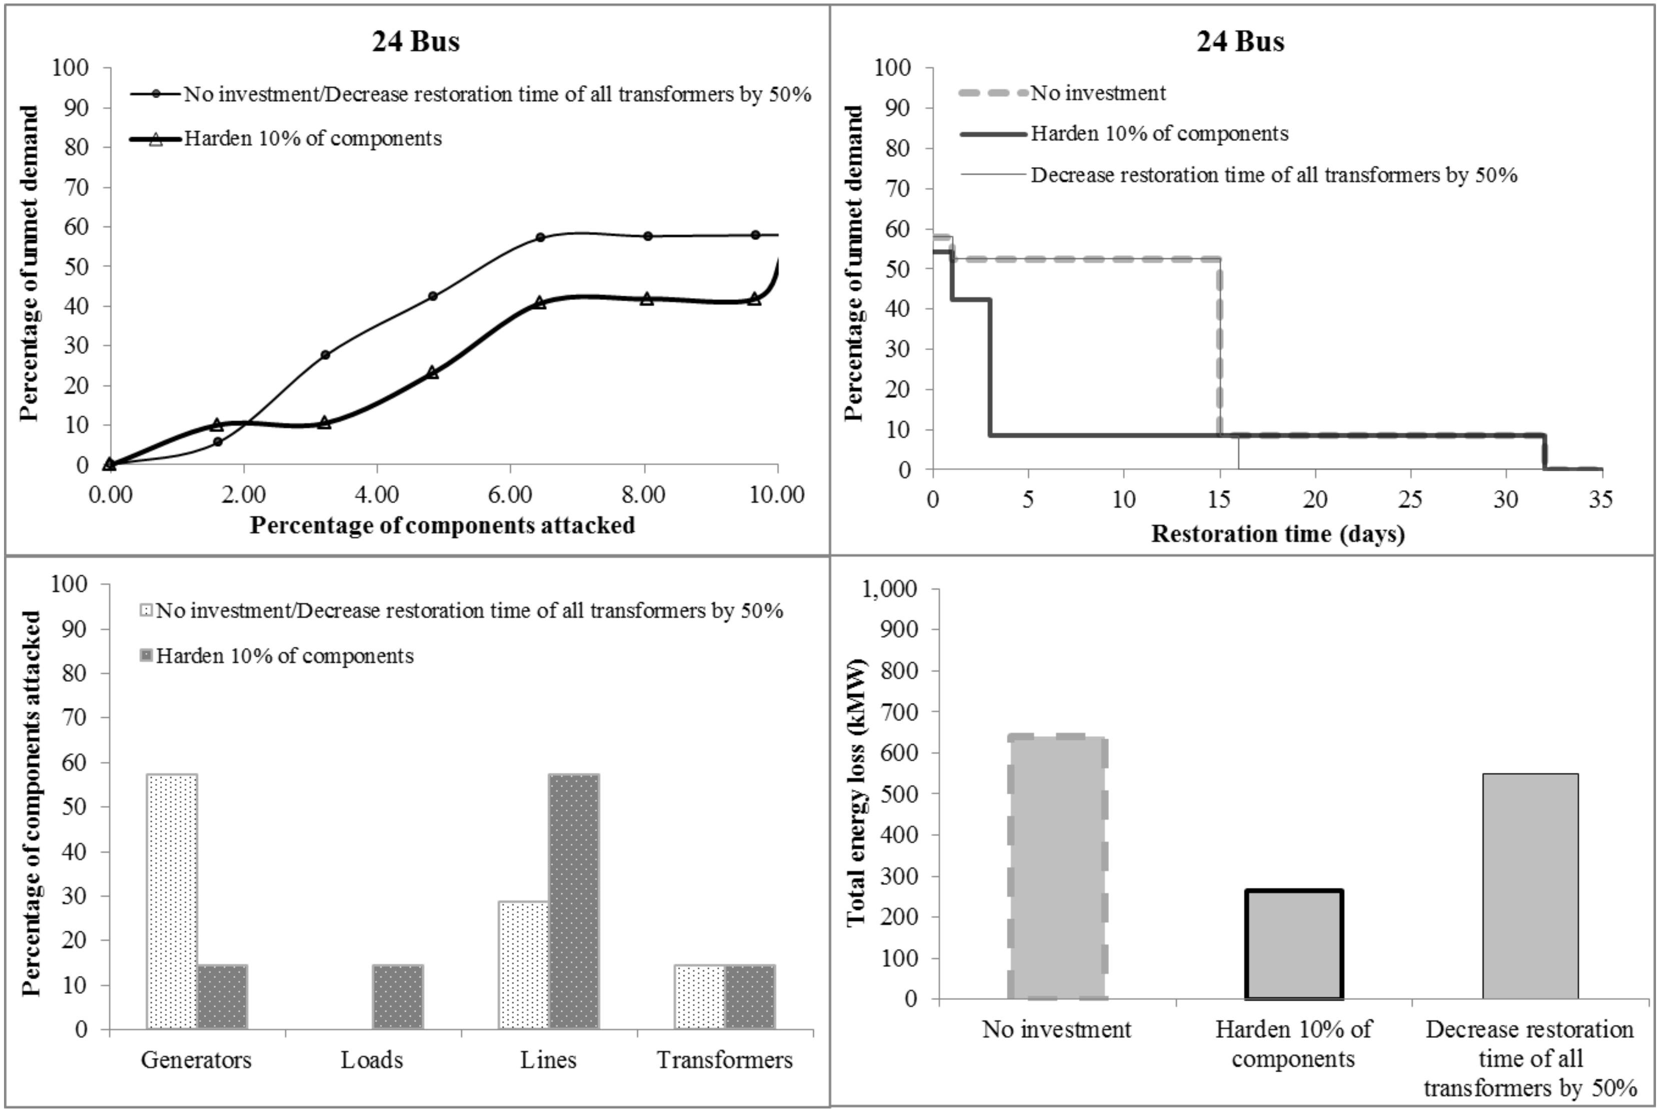

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and not considering restoration times, 48-bus system)

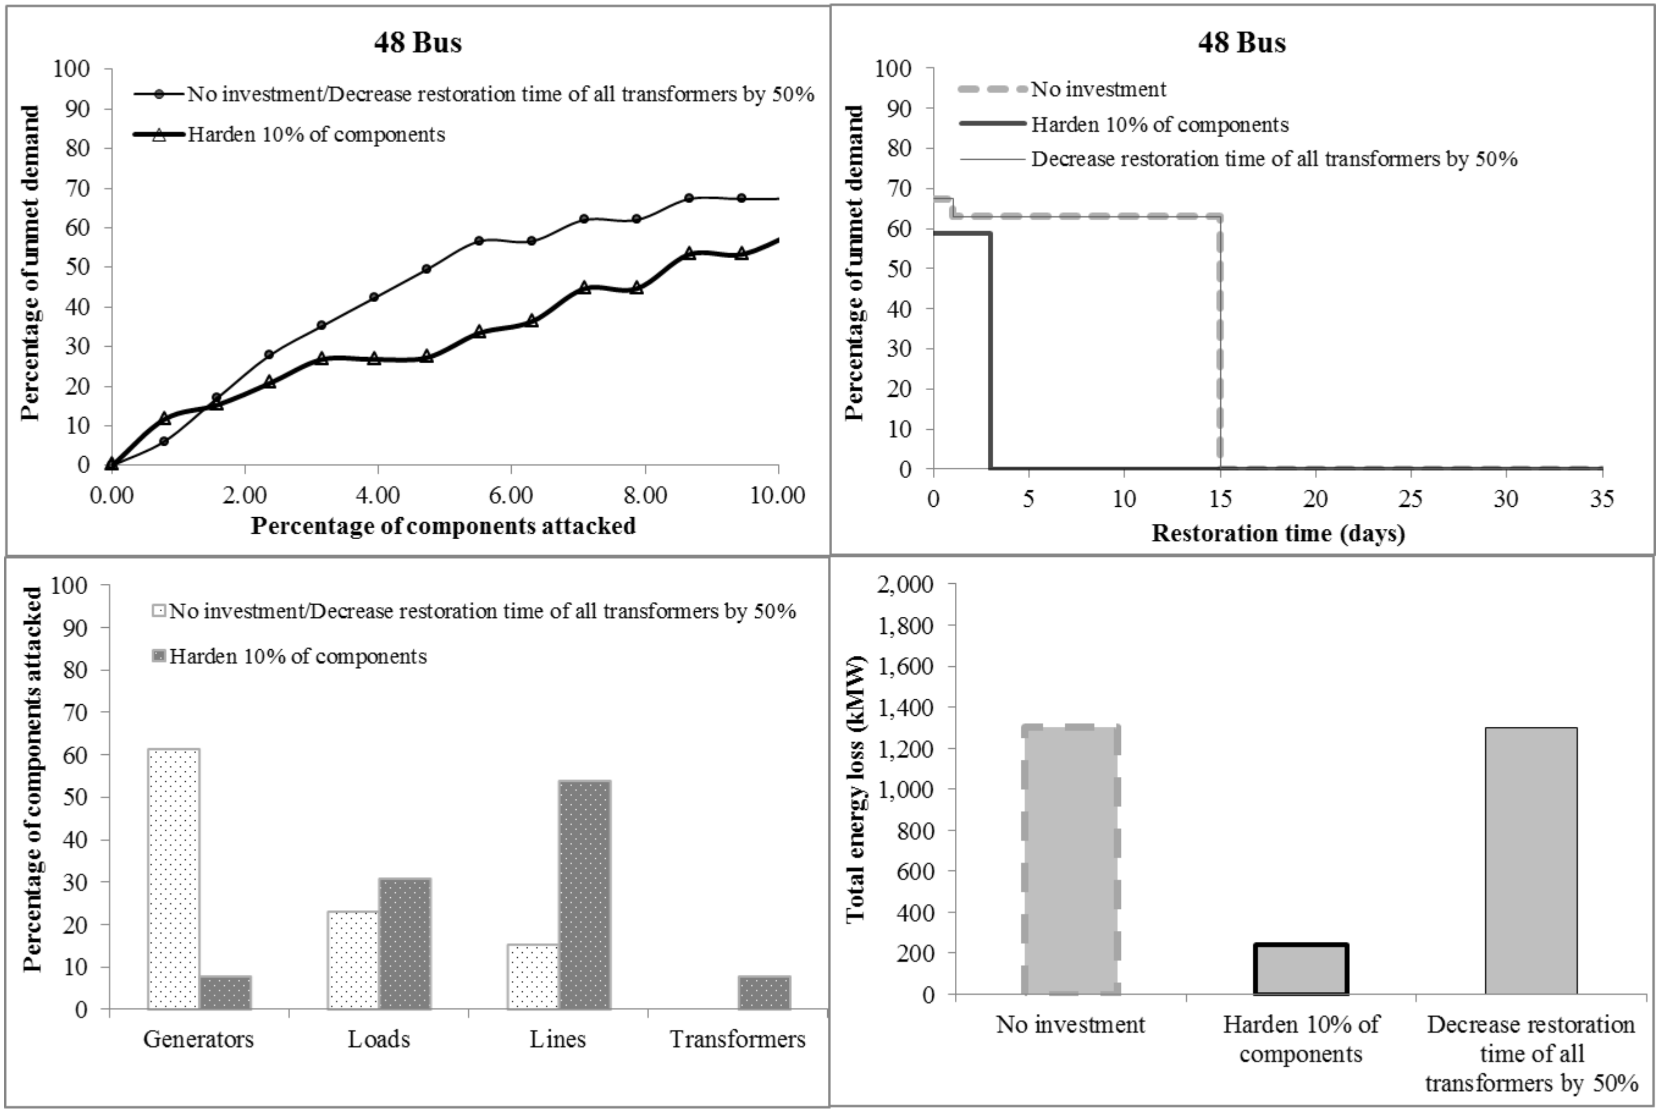

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and not considering restoration times, 118-bus system)

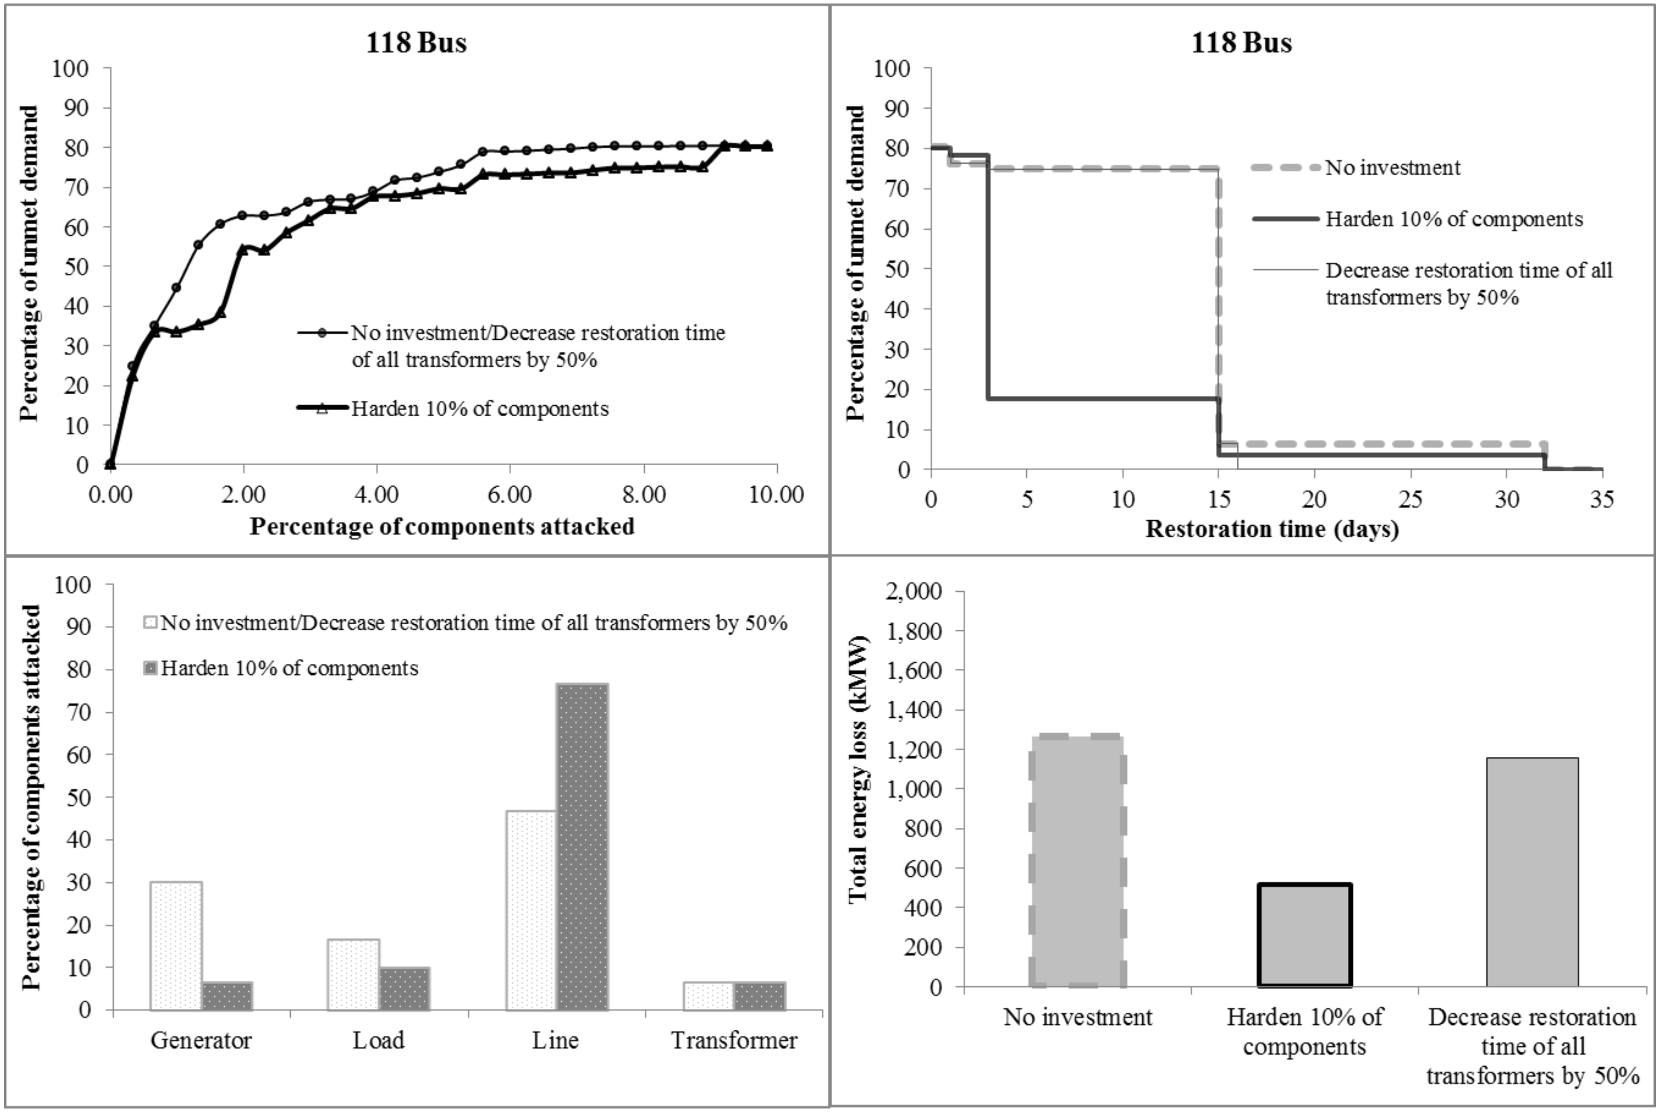

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and not considering restoration times, 300-bus system)

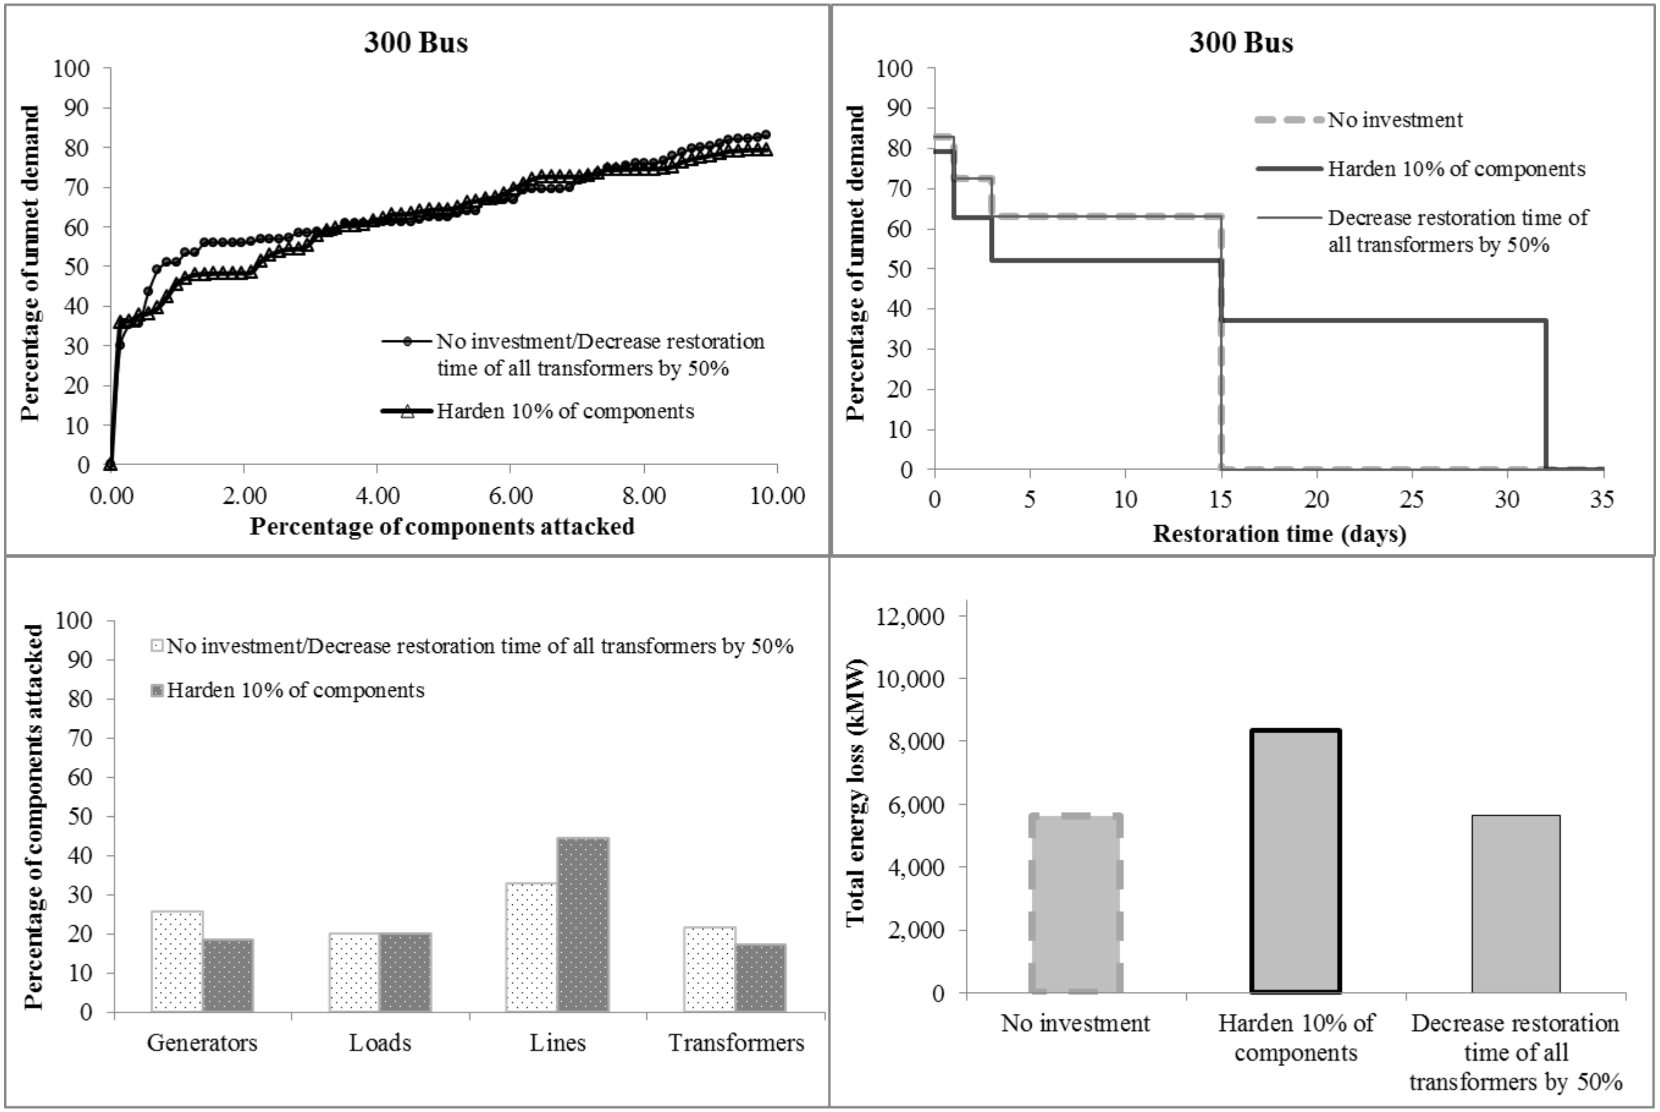

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and considering restoration times, 24-bus system)

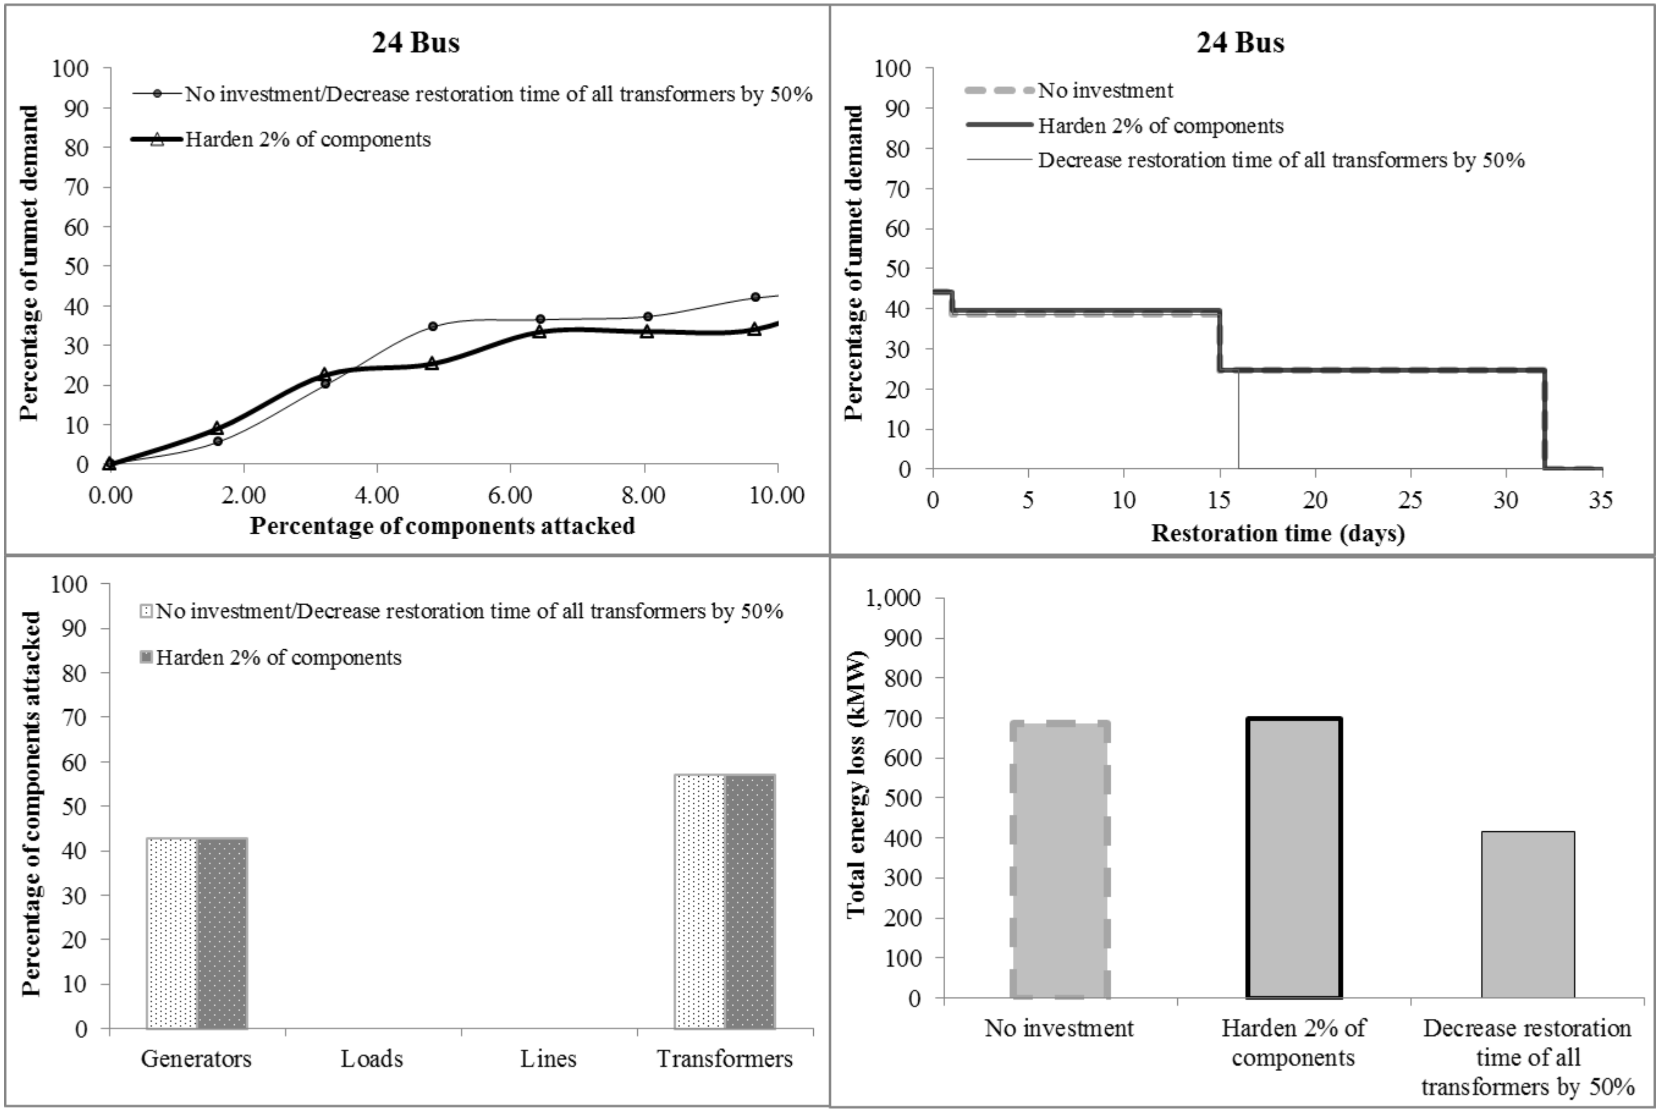

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and considering restoration times, 48-bus system)

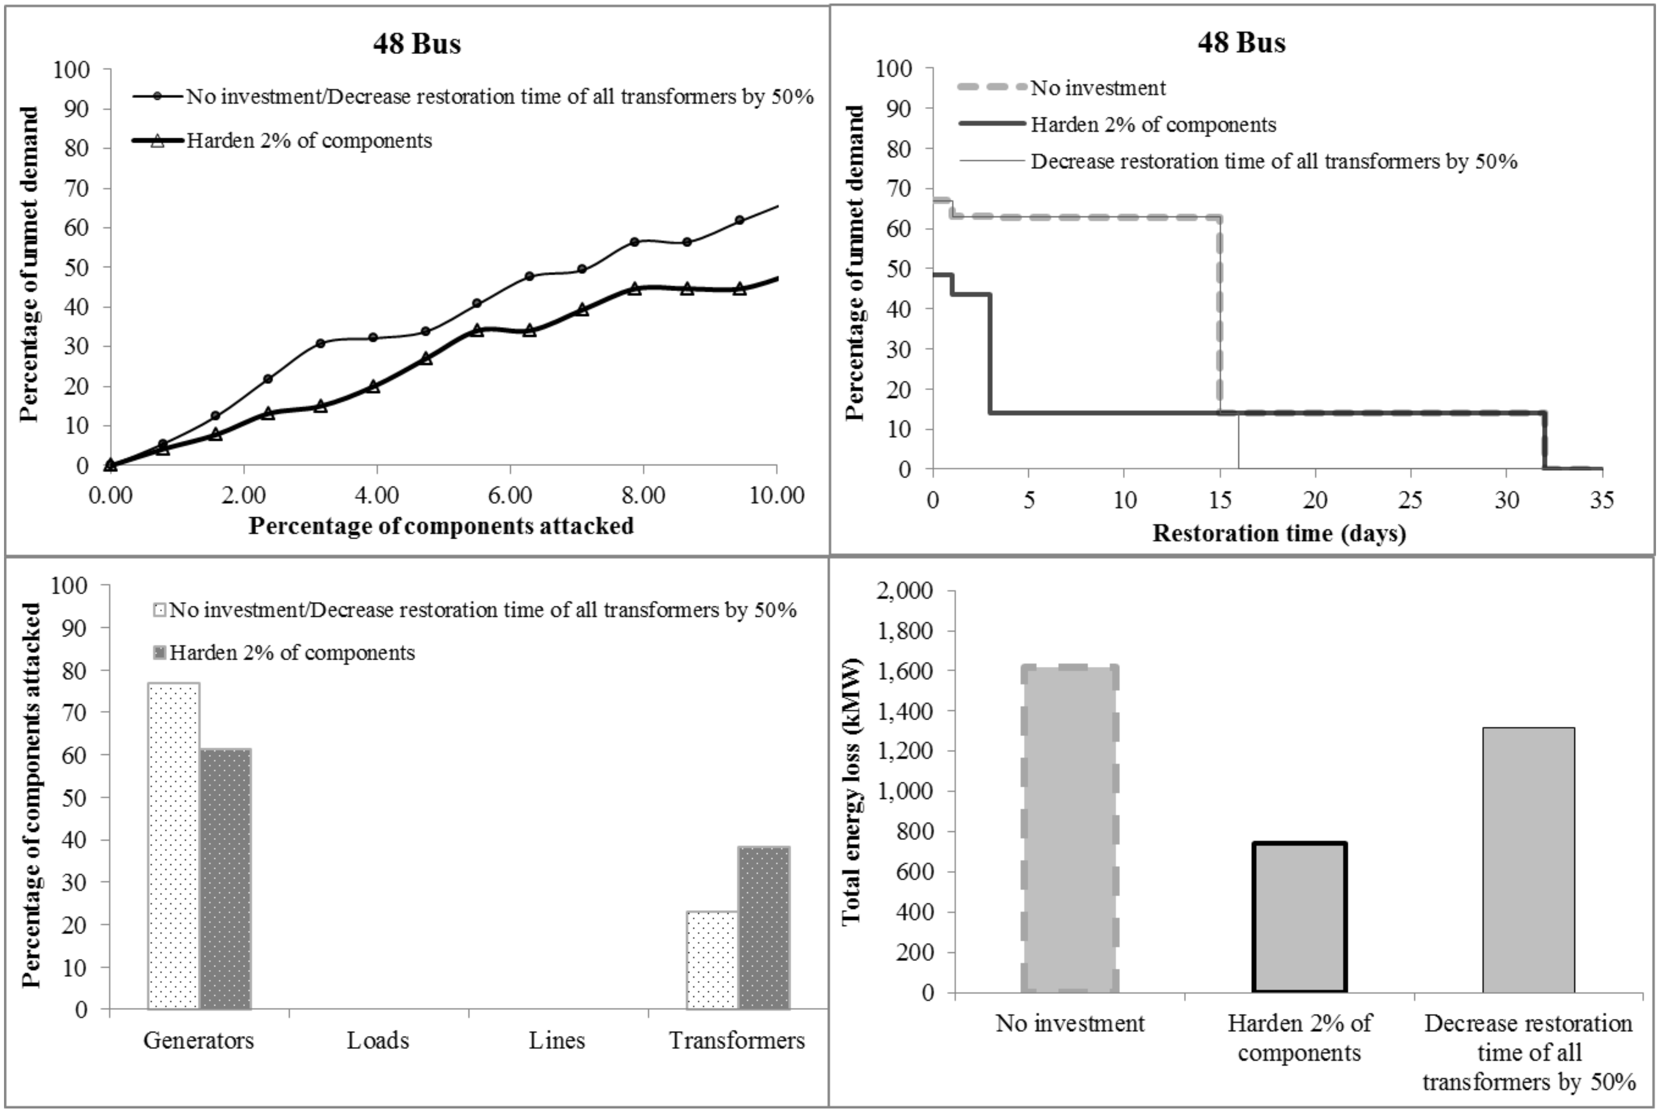

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and considering restoration times, 118-bus system)

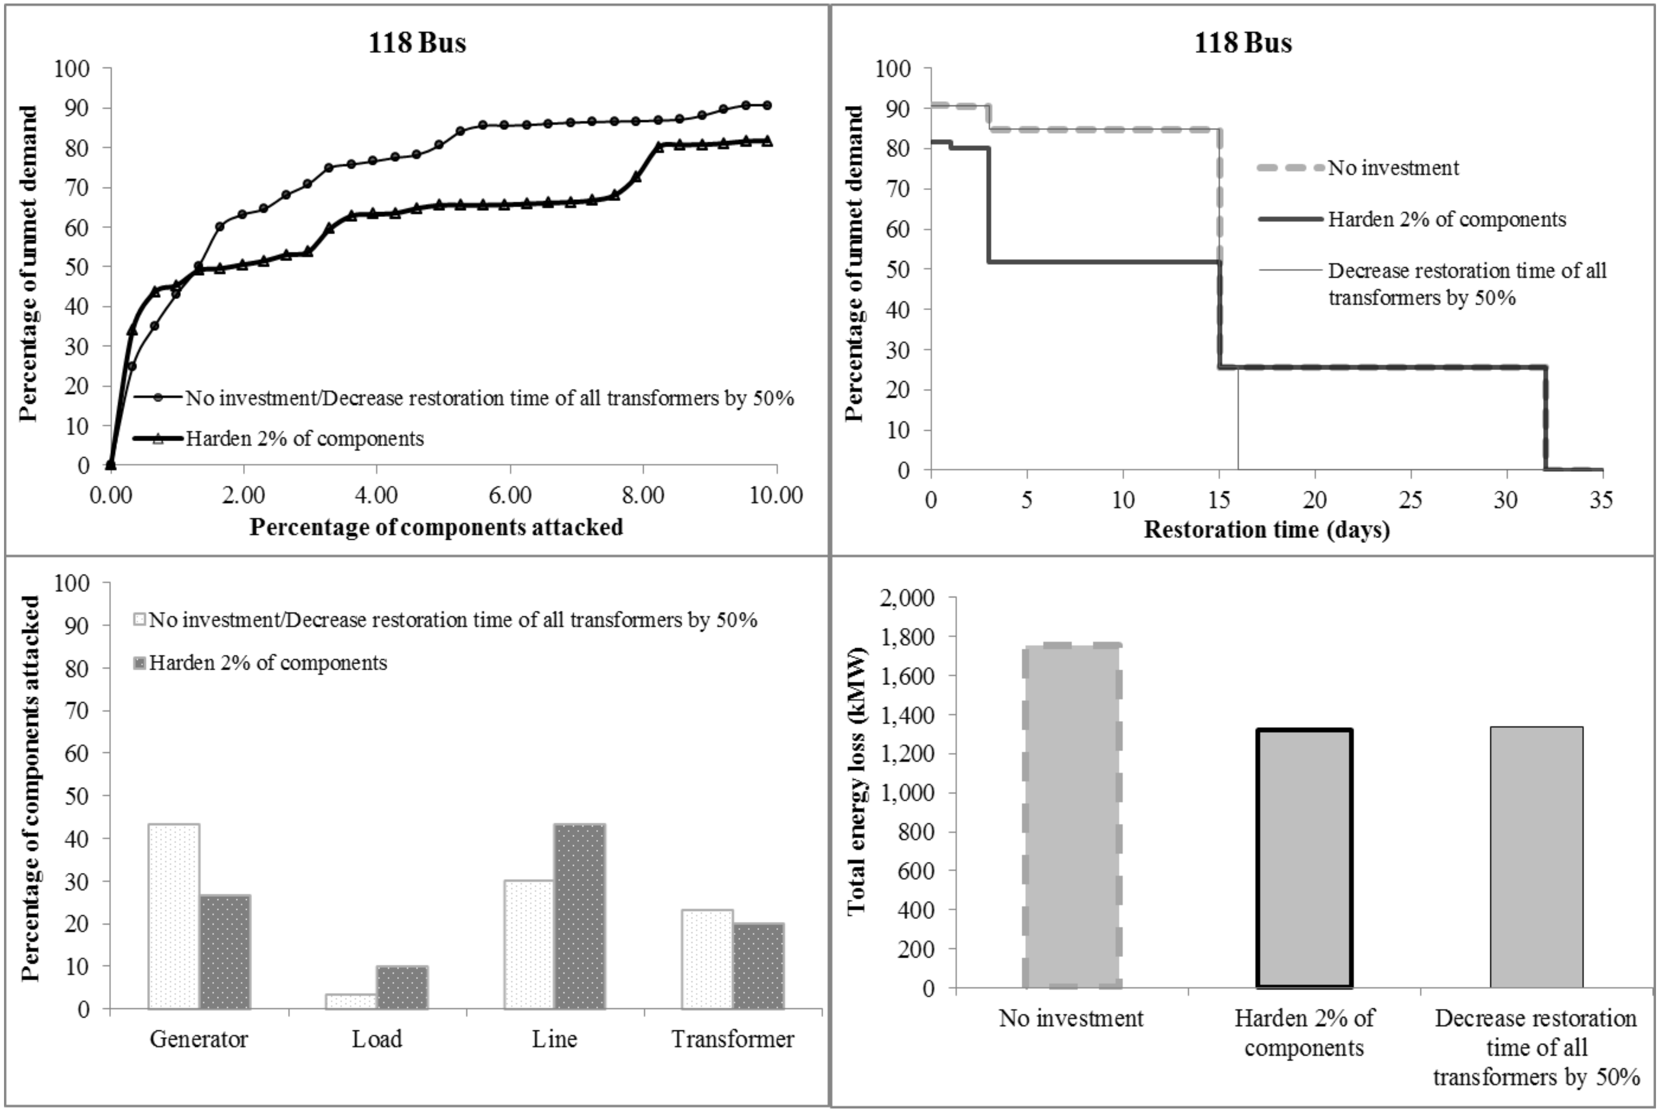

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with no cascading knowledge and considering restoration times, 300-bus system)

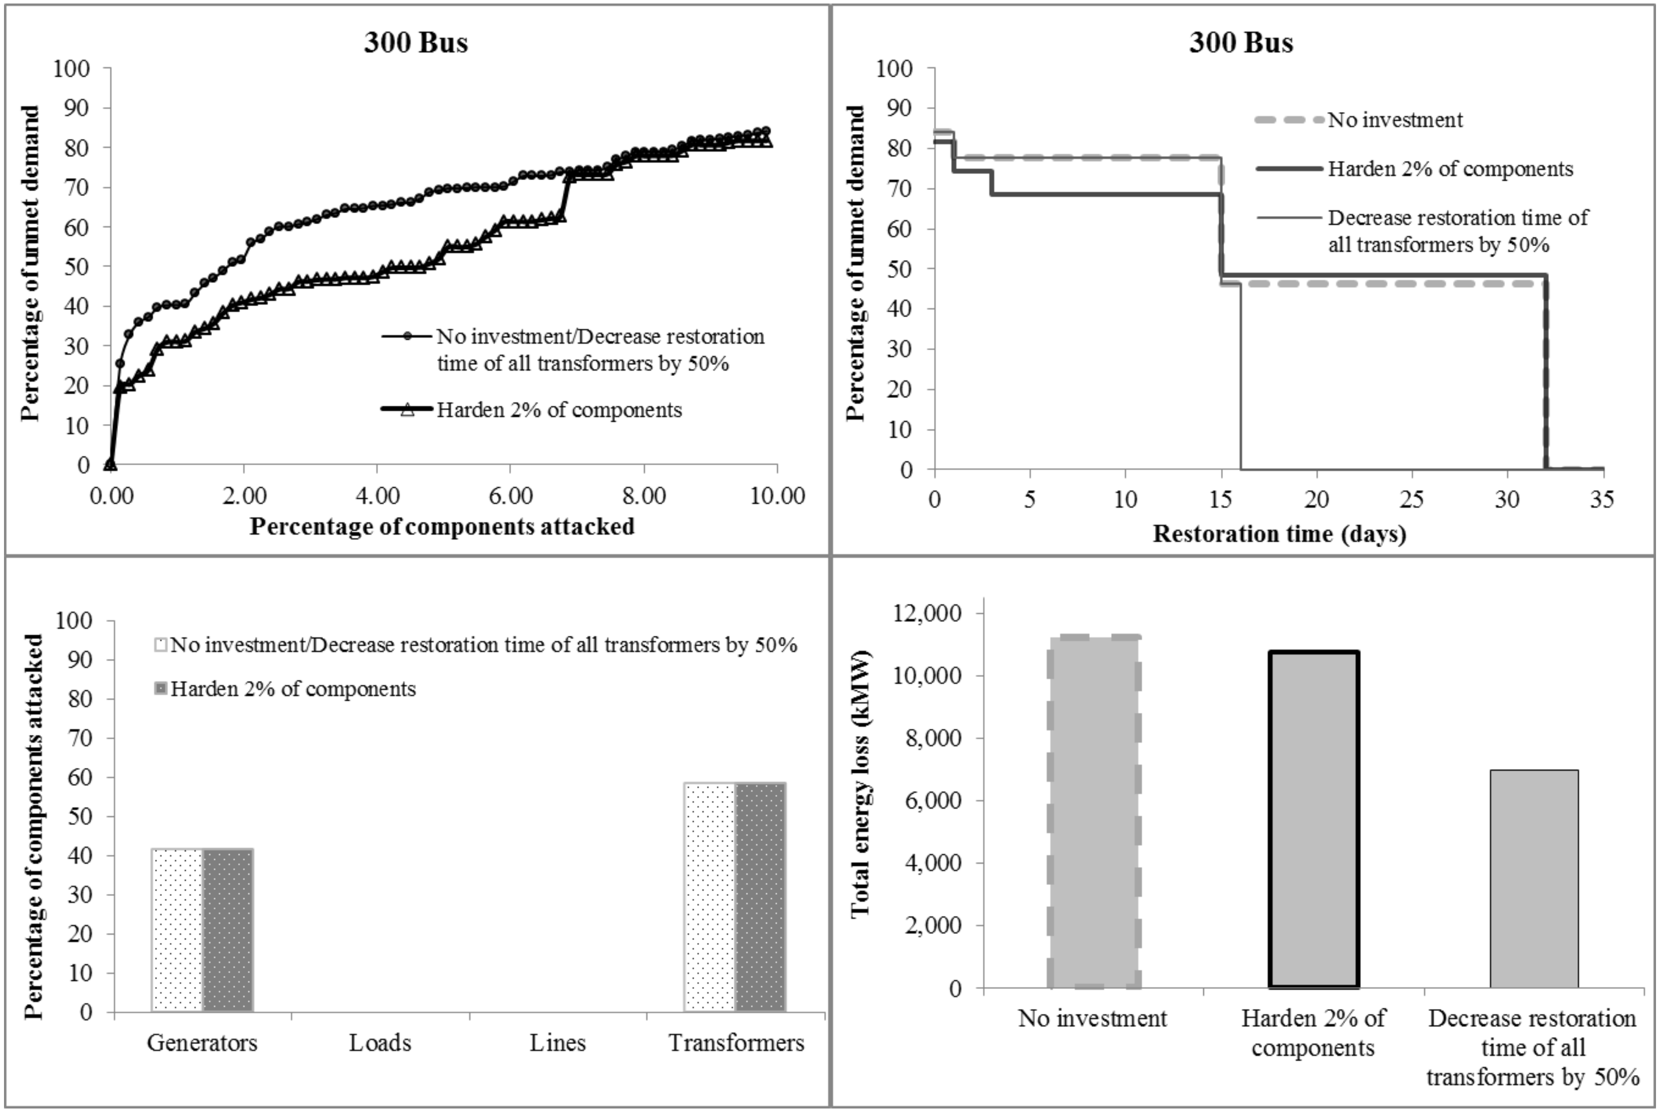

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and considering restoration times, 24-bus system)

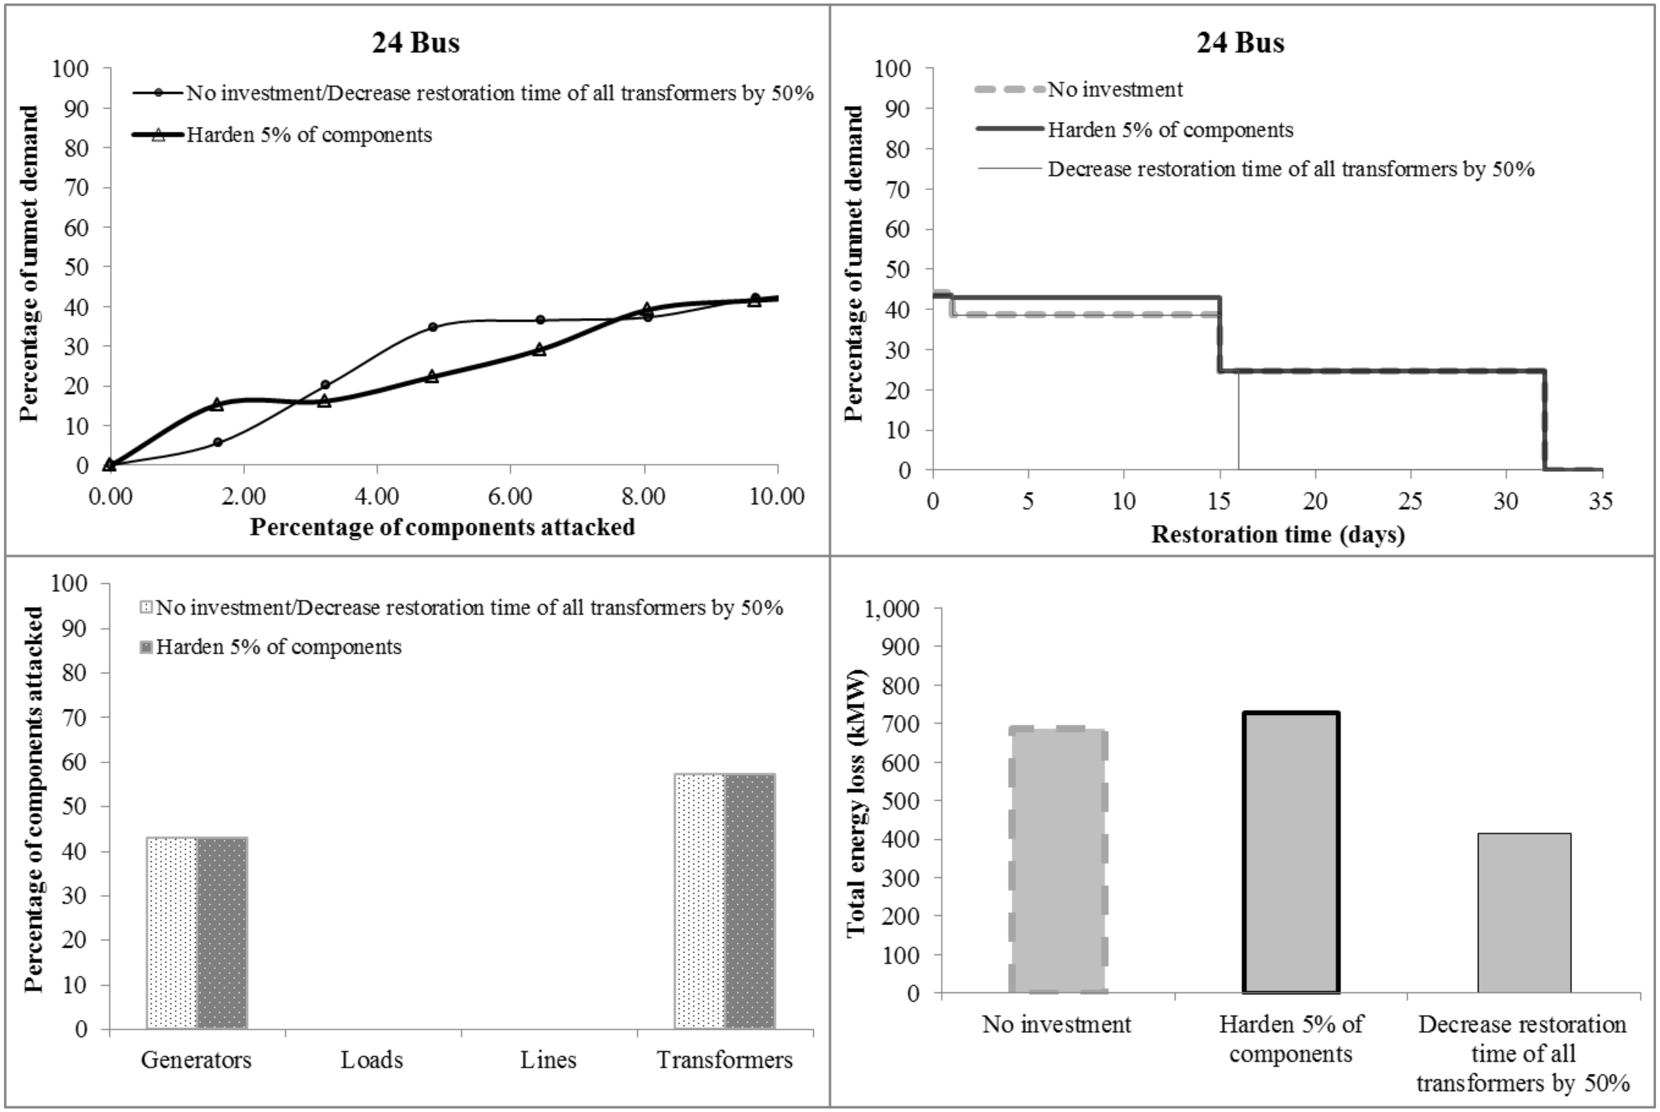

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and considering restoration times, 48-bus system)

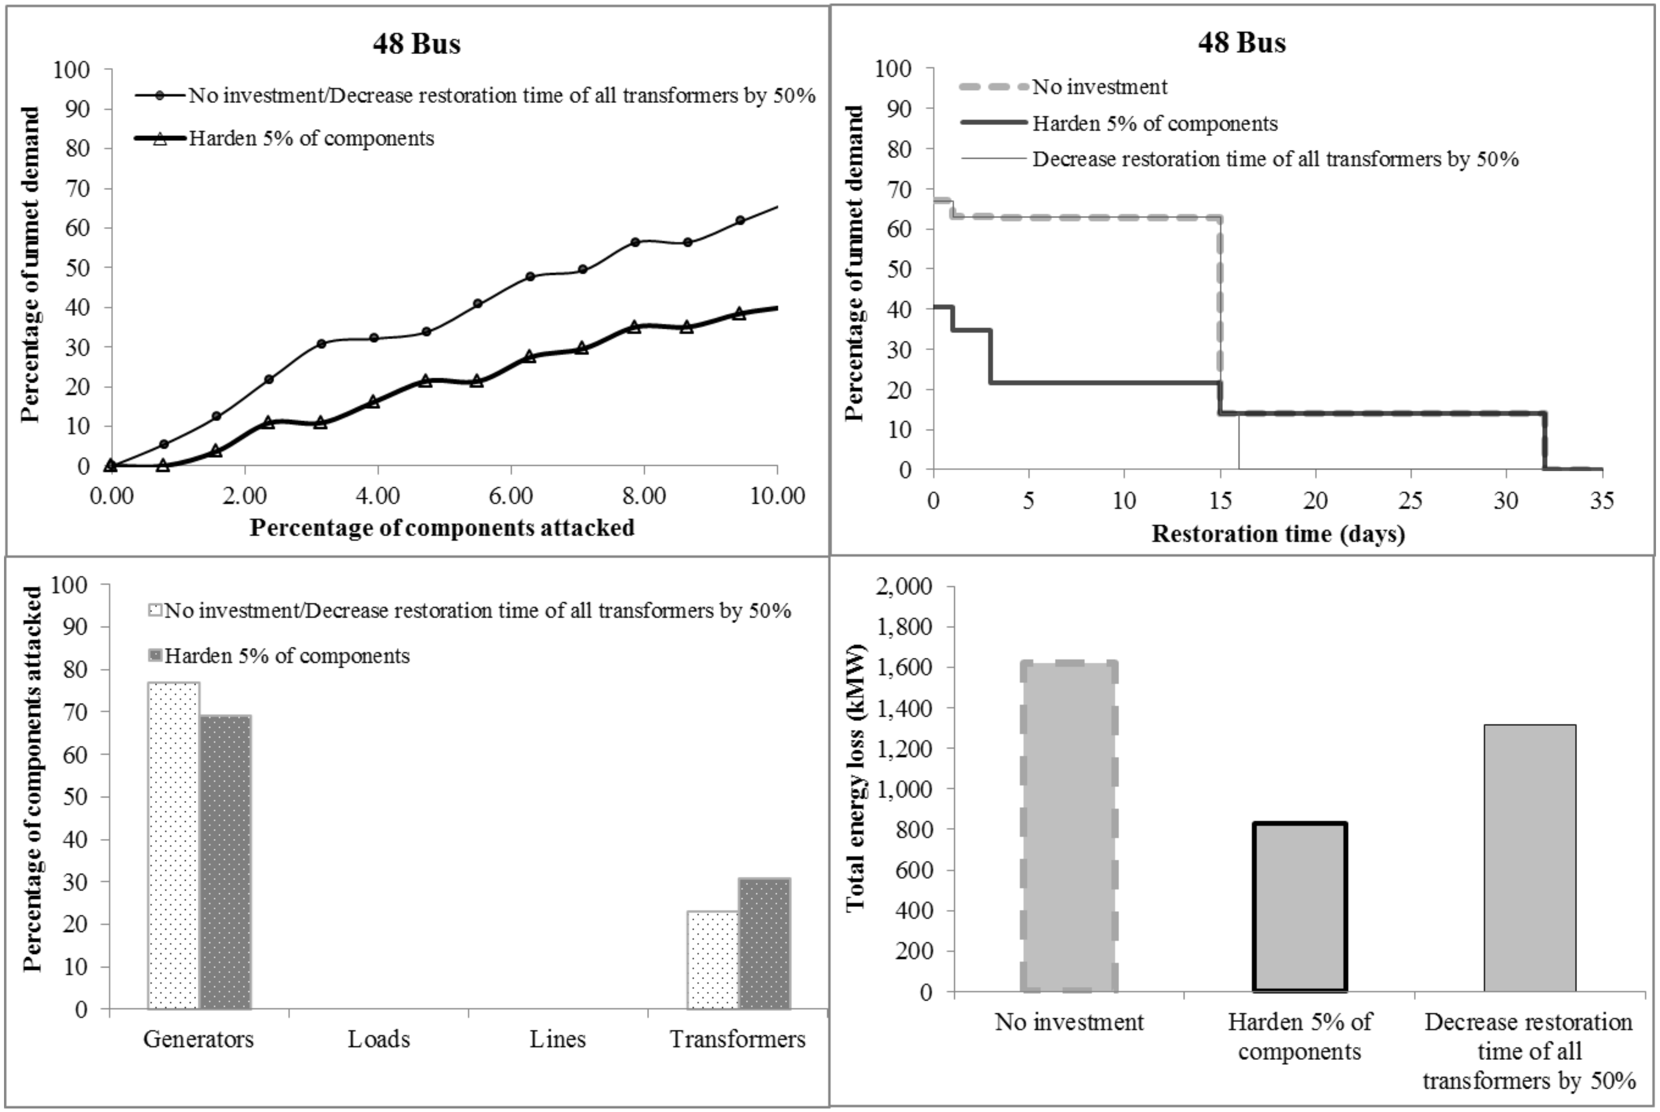

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and considering restoration times, 118-bus system)

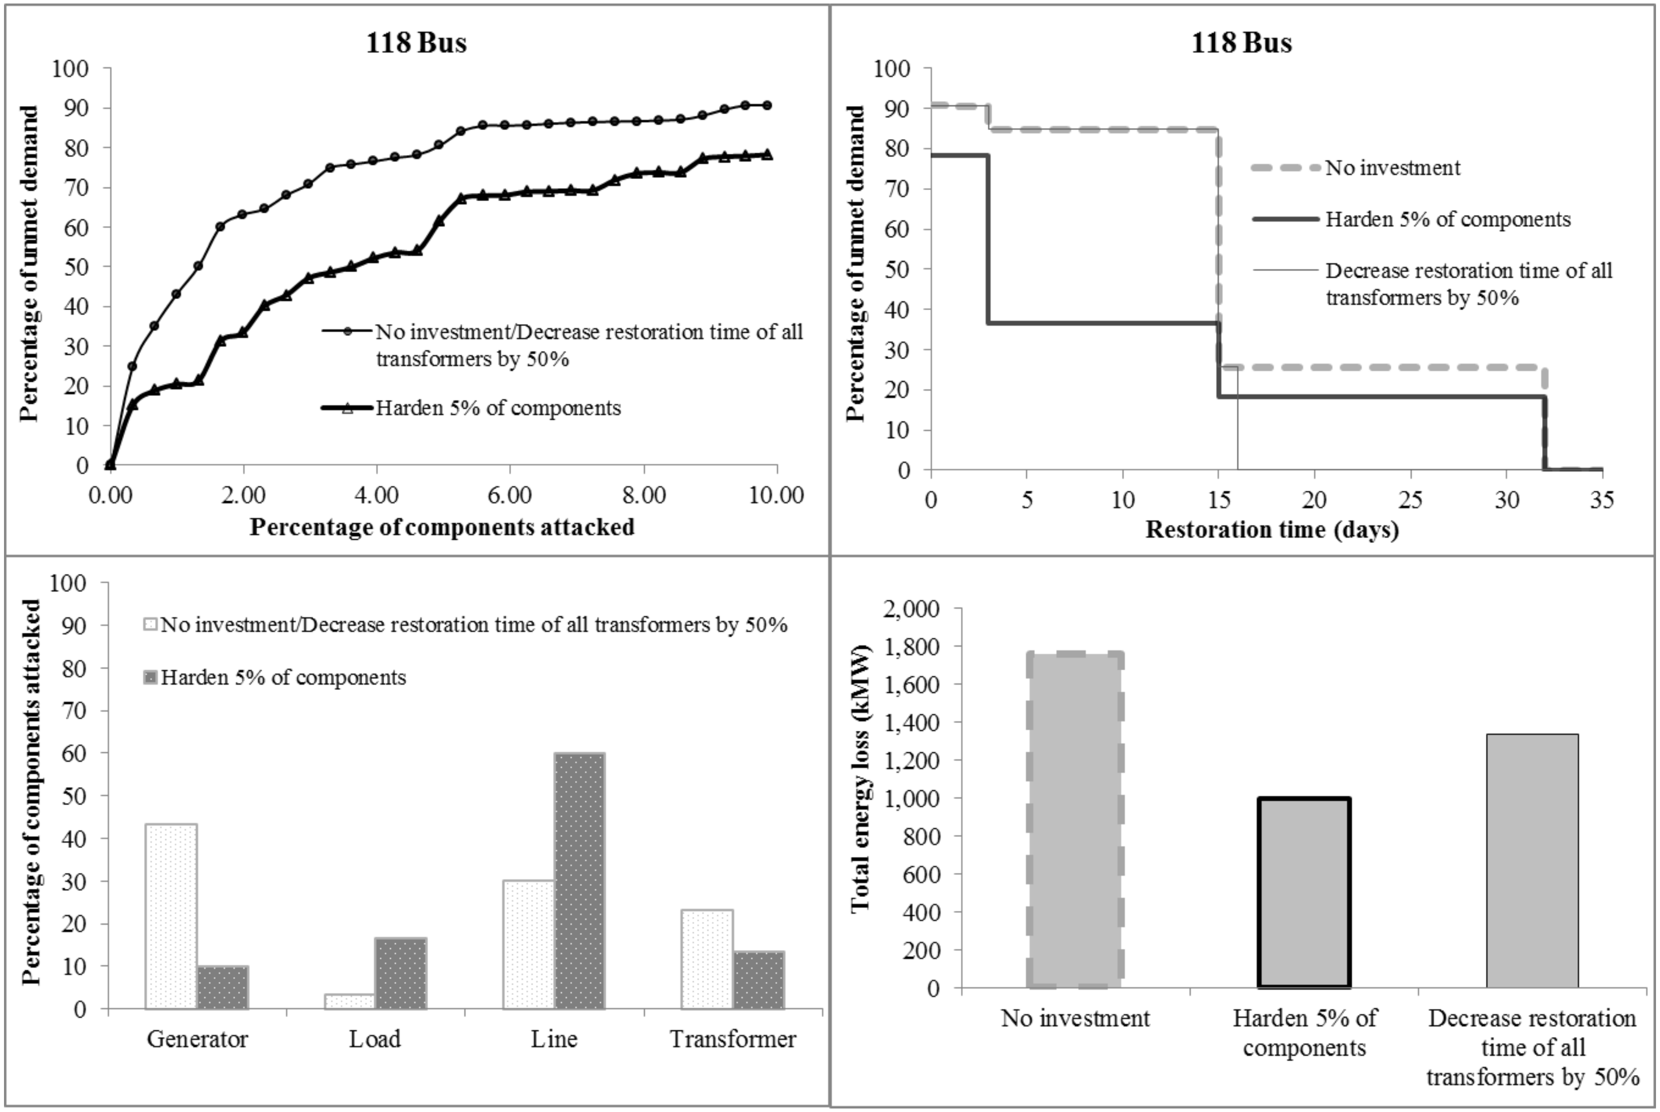

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with no cascading knowledge and considering restoration times, 300-bus system)

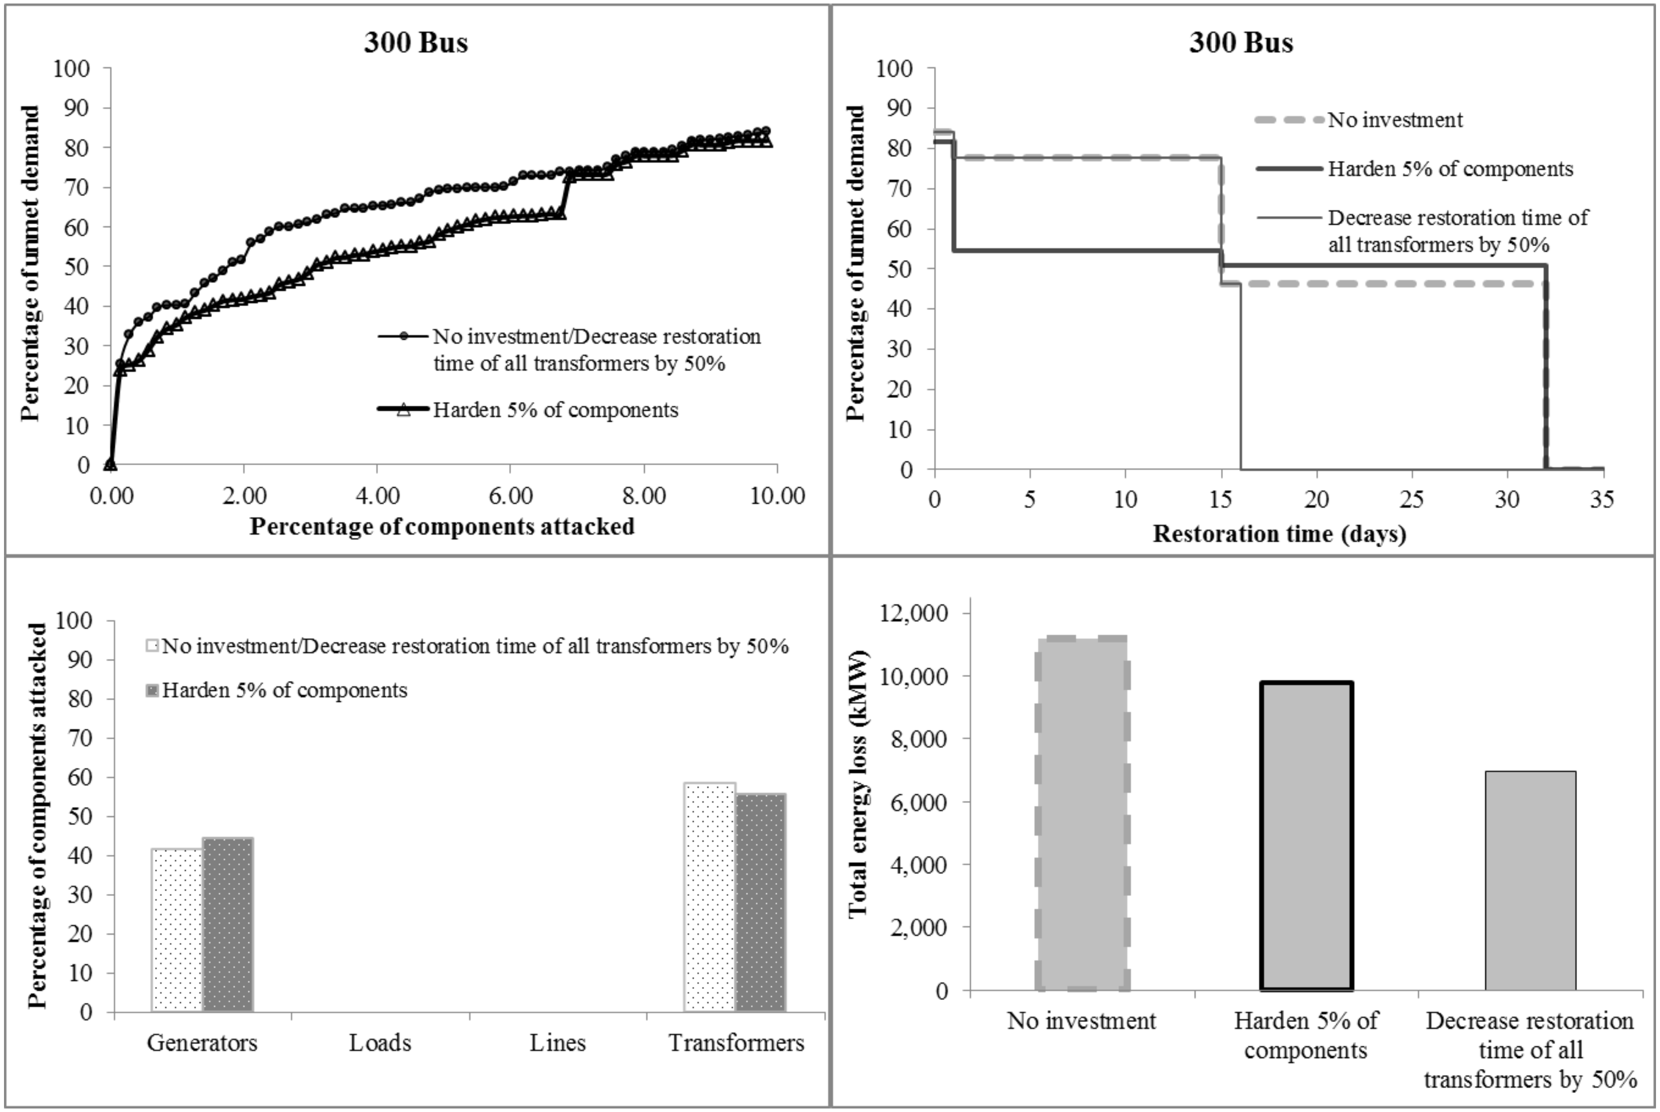

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and considering restoration times, 24-bus system)

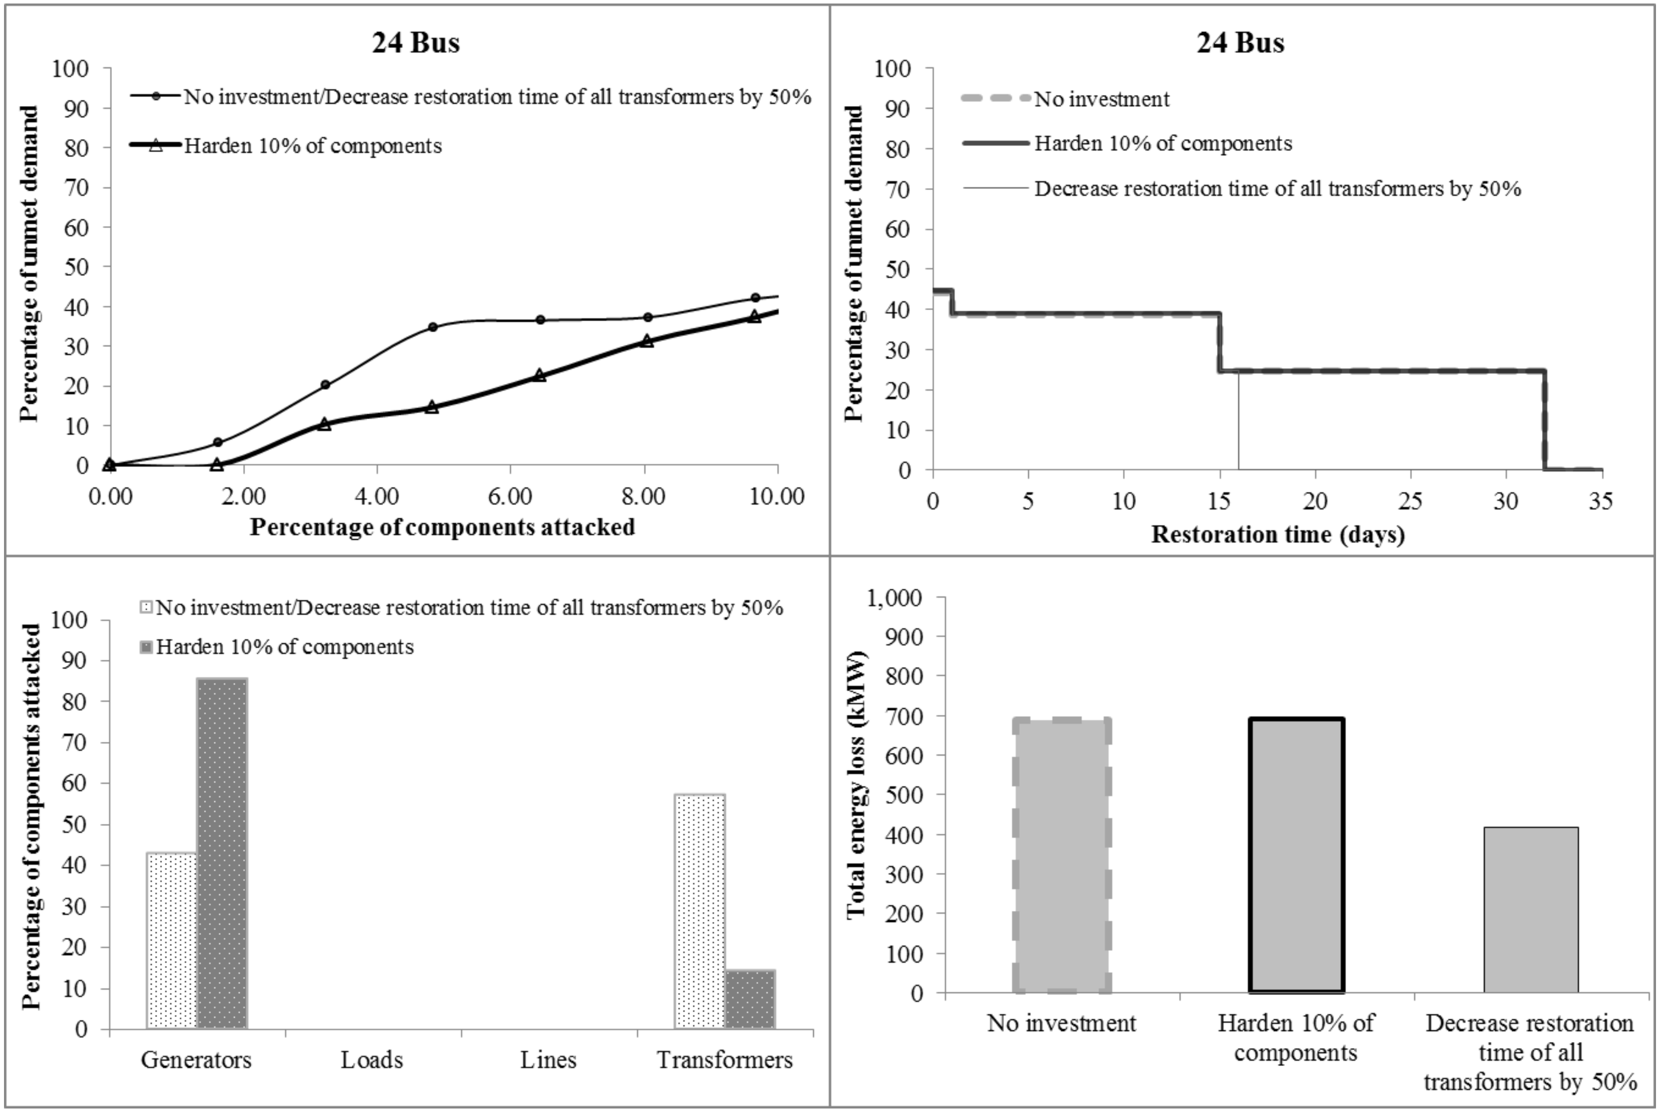

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and considering restoration times, 48-bus system)

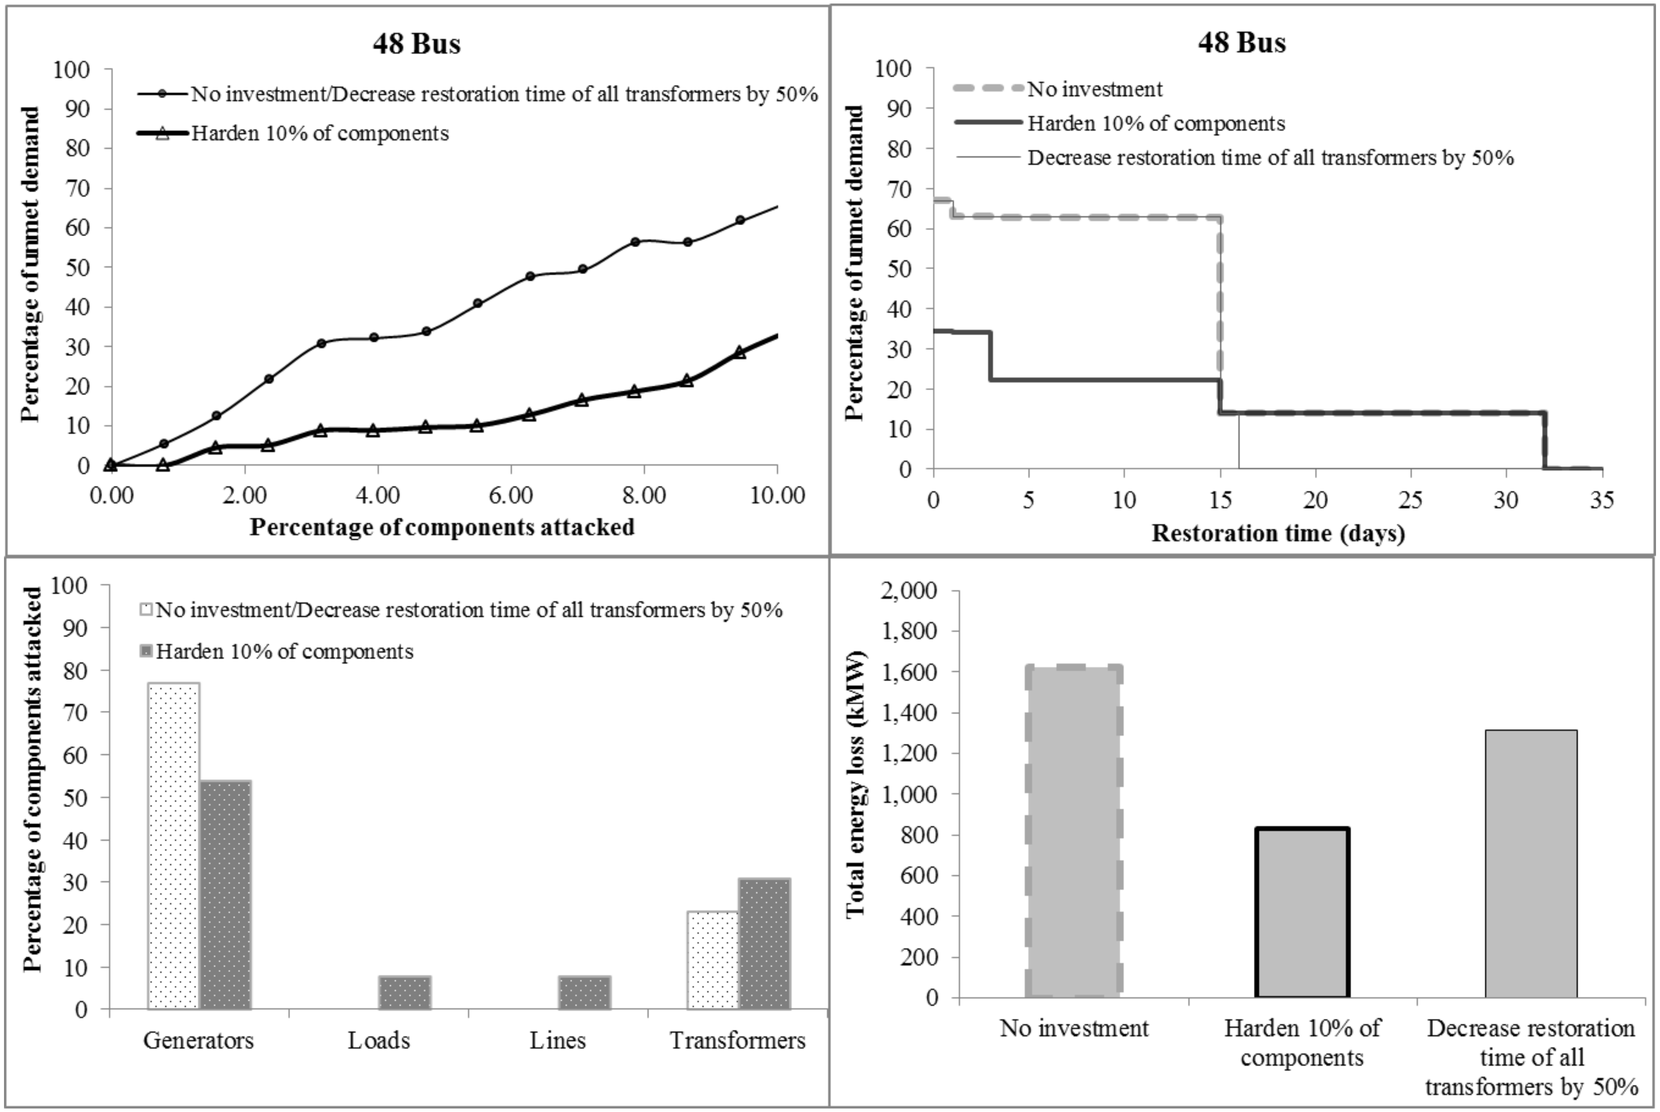

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and considering restoration times, 118-bus system)

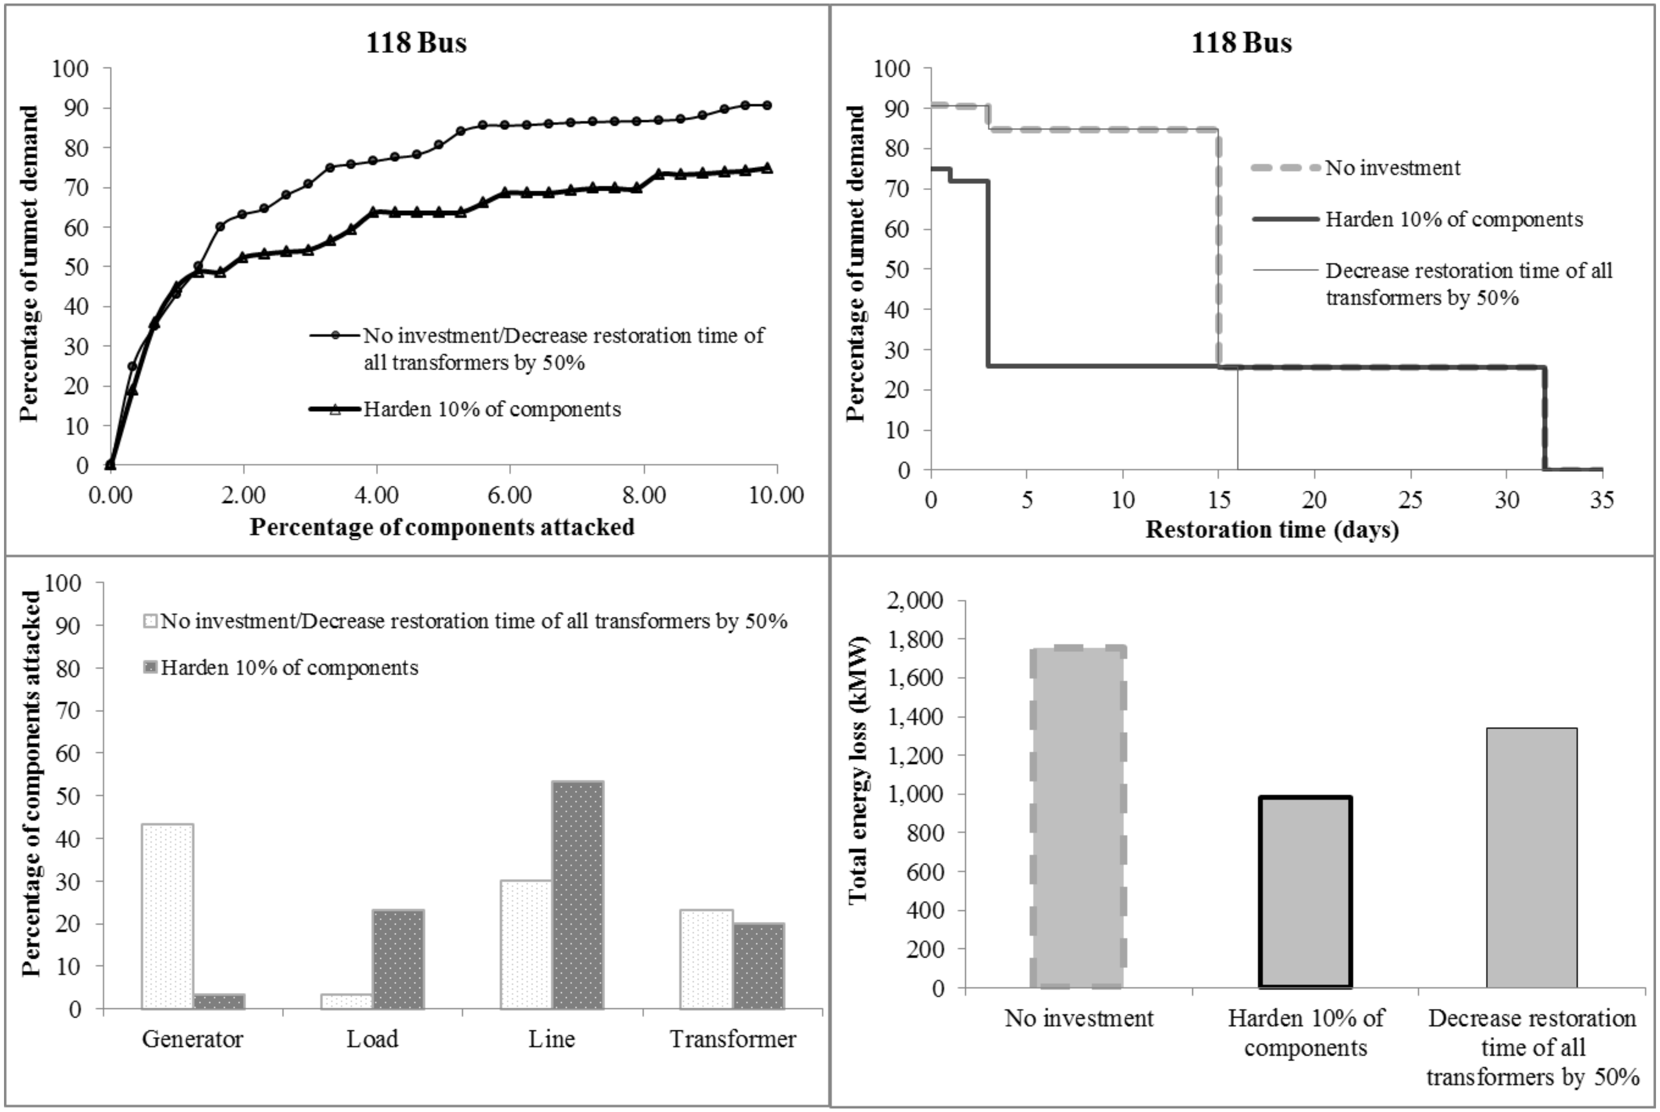

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with no cascading knowledge and considering restoration times, 300-bus system)

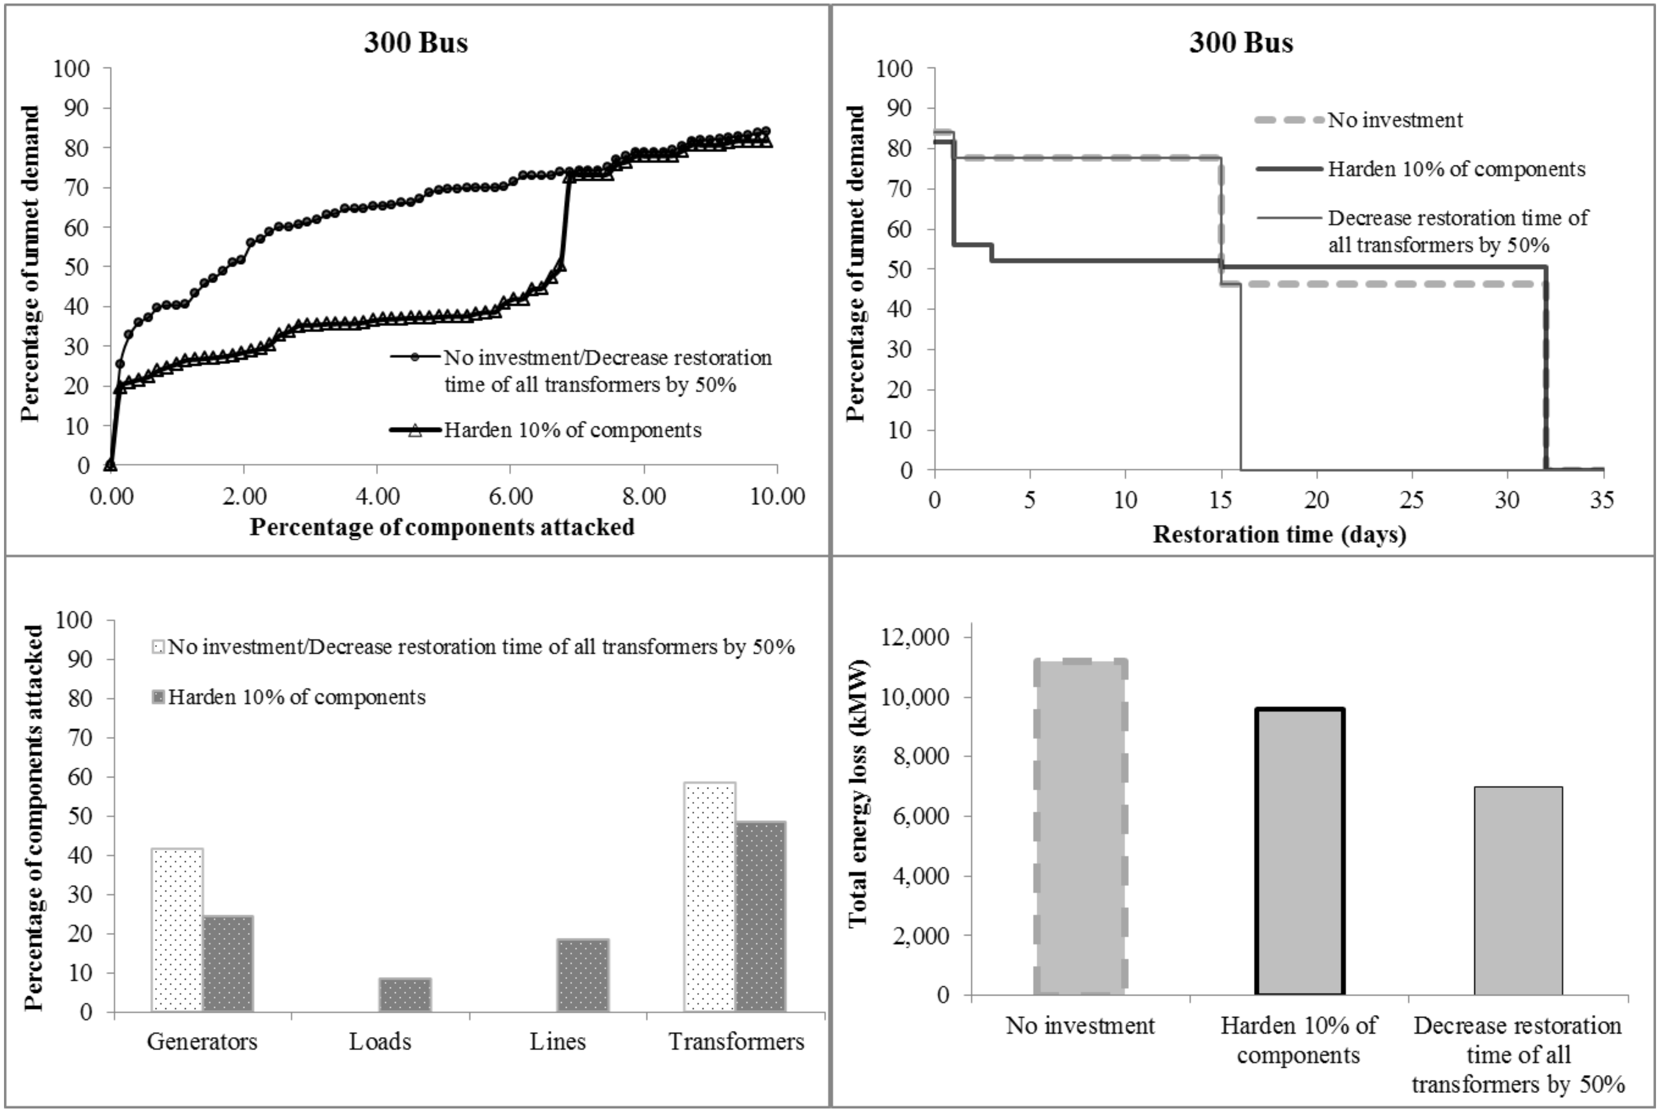

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and considering restoration times, 24-bus system)

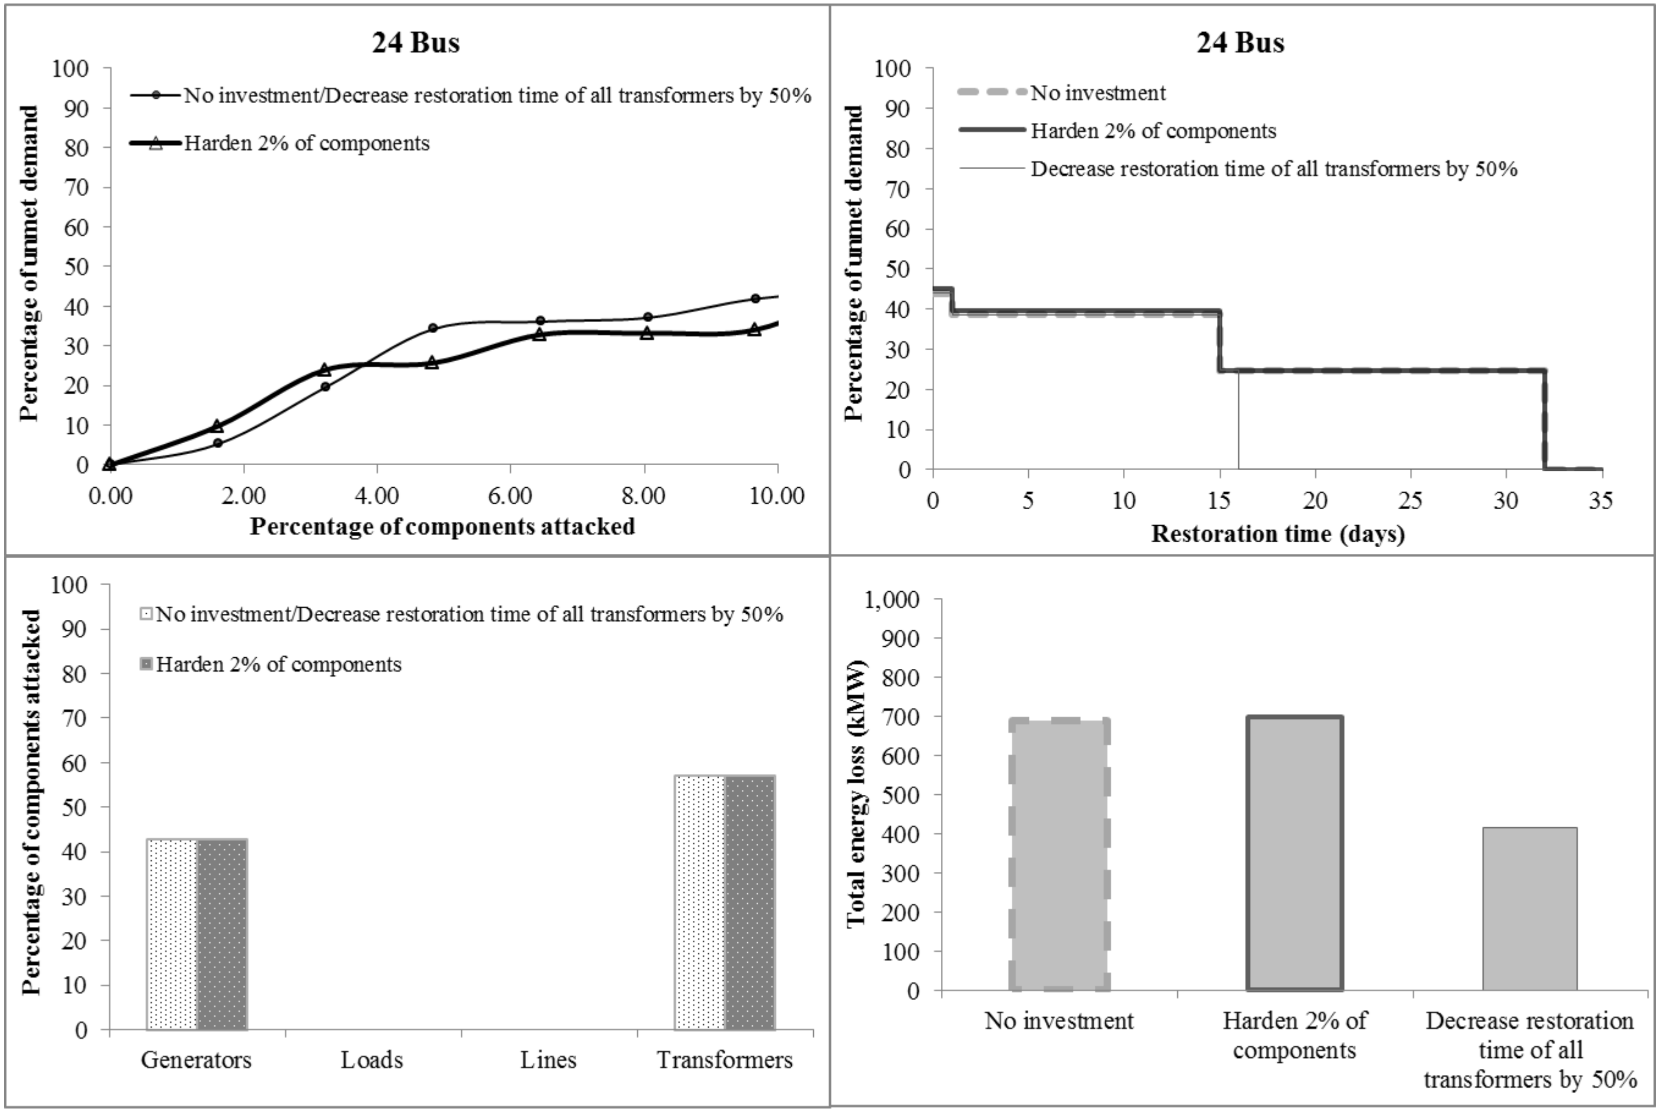

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and considering restoration times, 48-bus system)

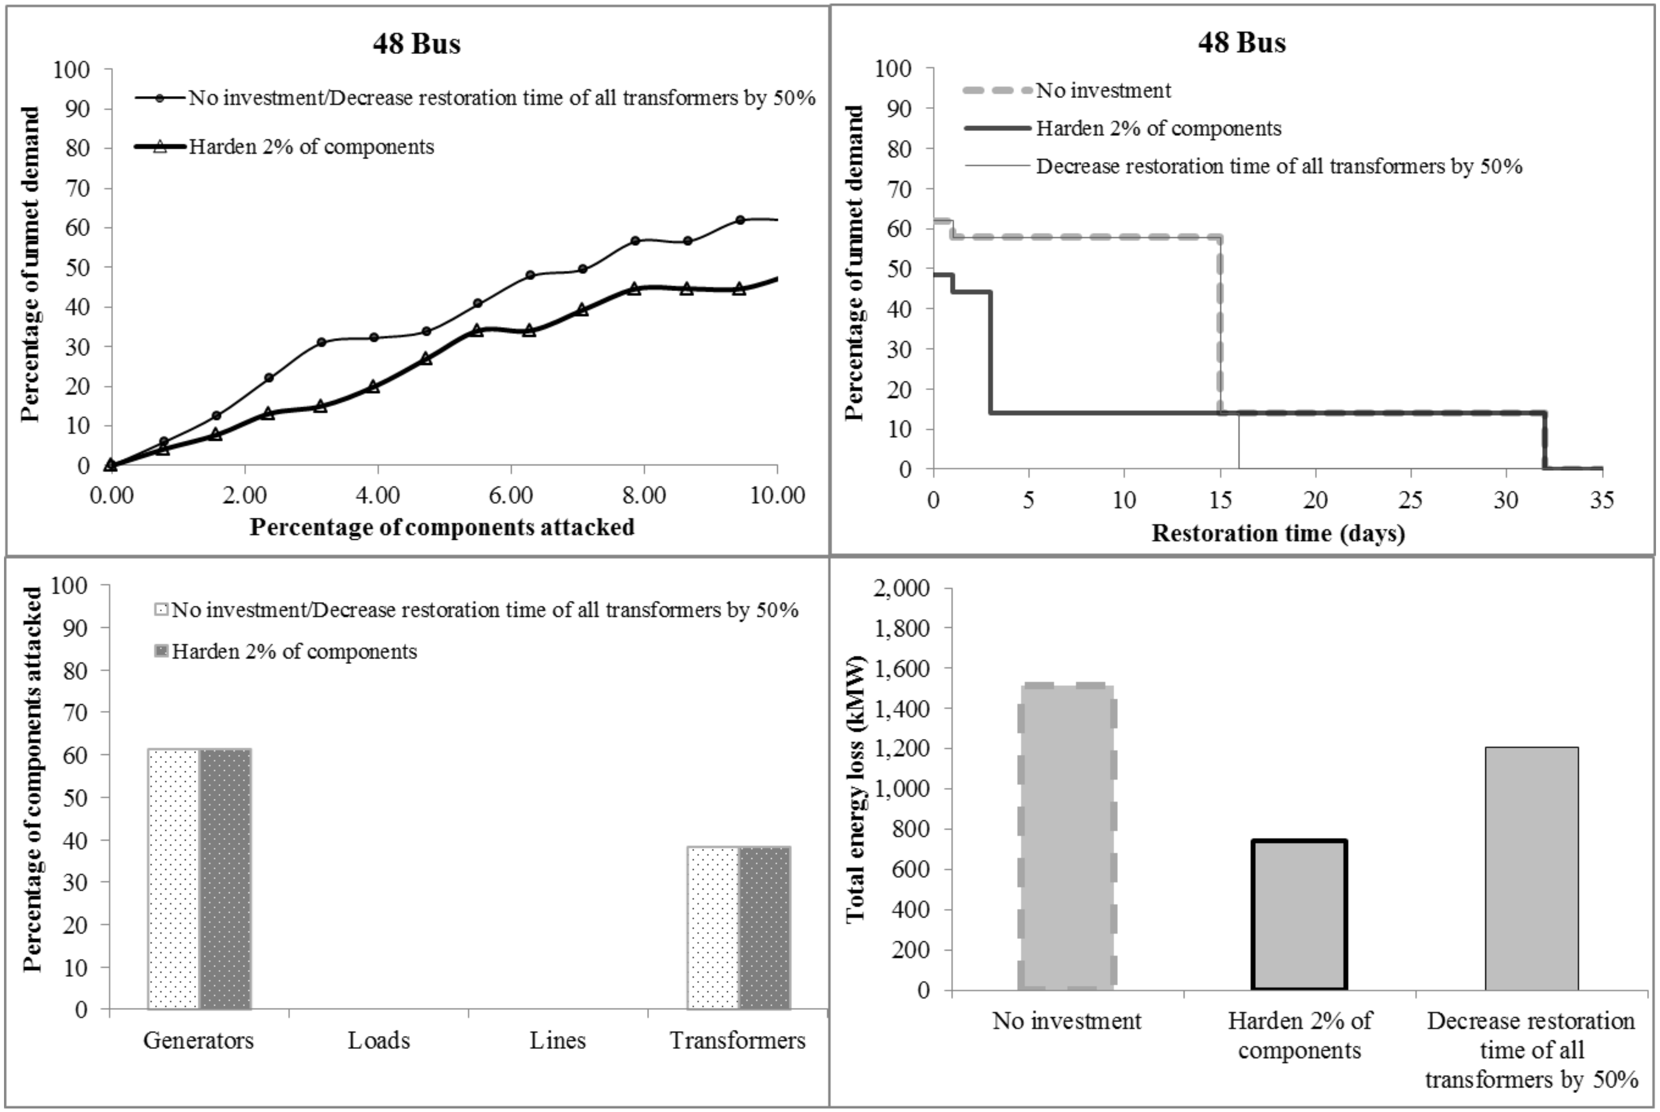

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and considering restoration times, 118-bus system)

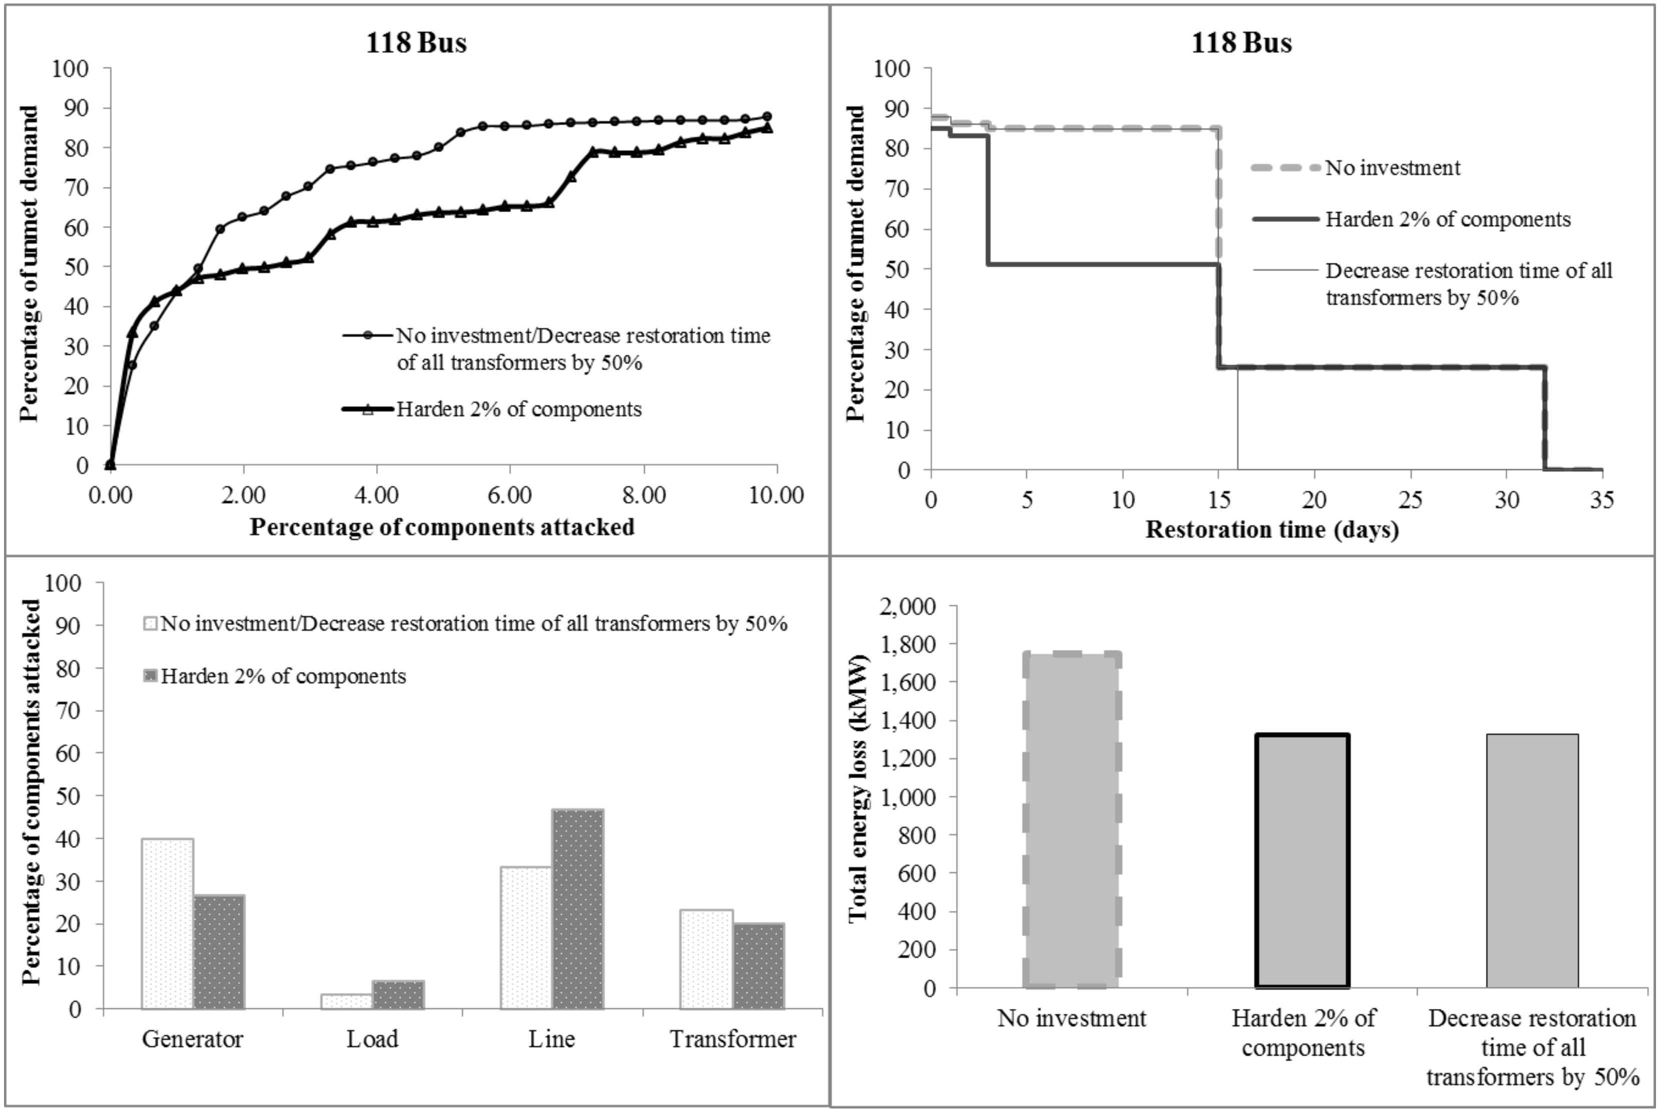

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 2% of the components against a static attacker with cascading knowledge and considering restoration times, 300-bus system)

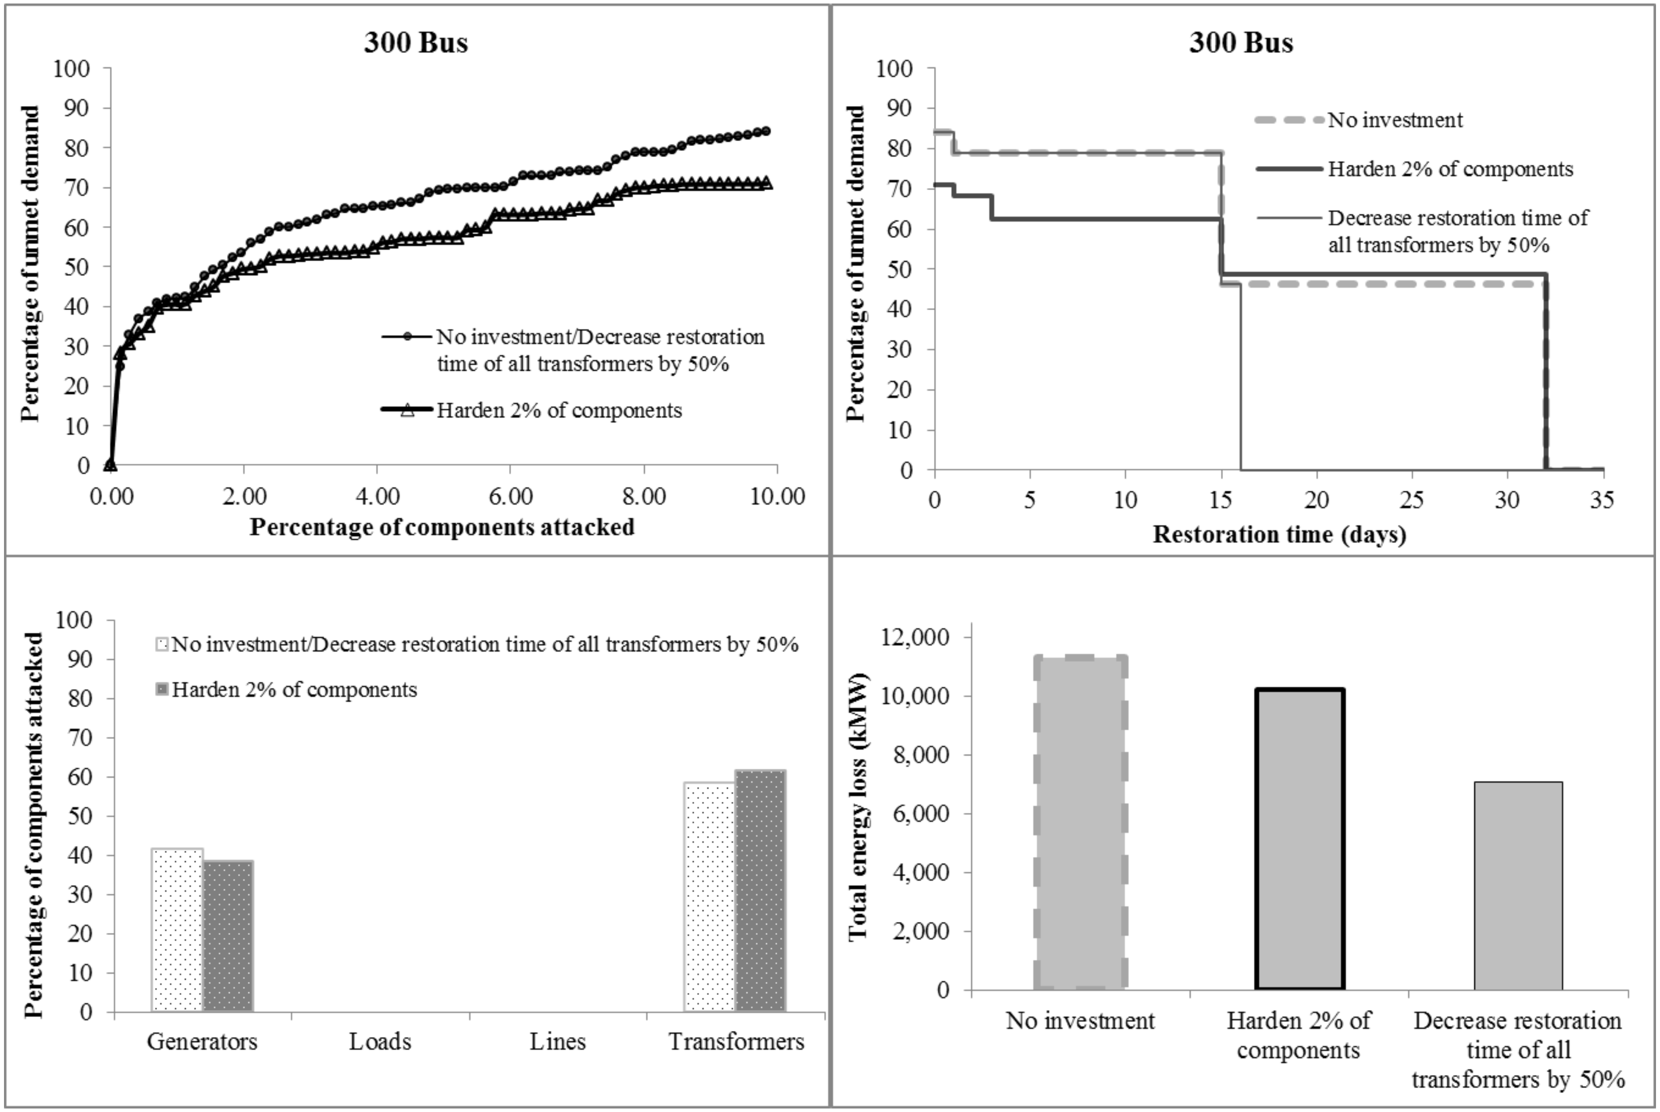

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and considering restoration times, 24-bus system)

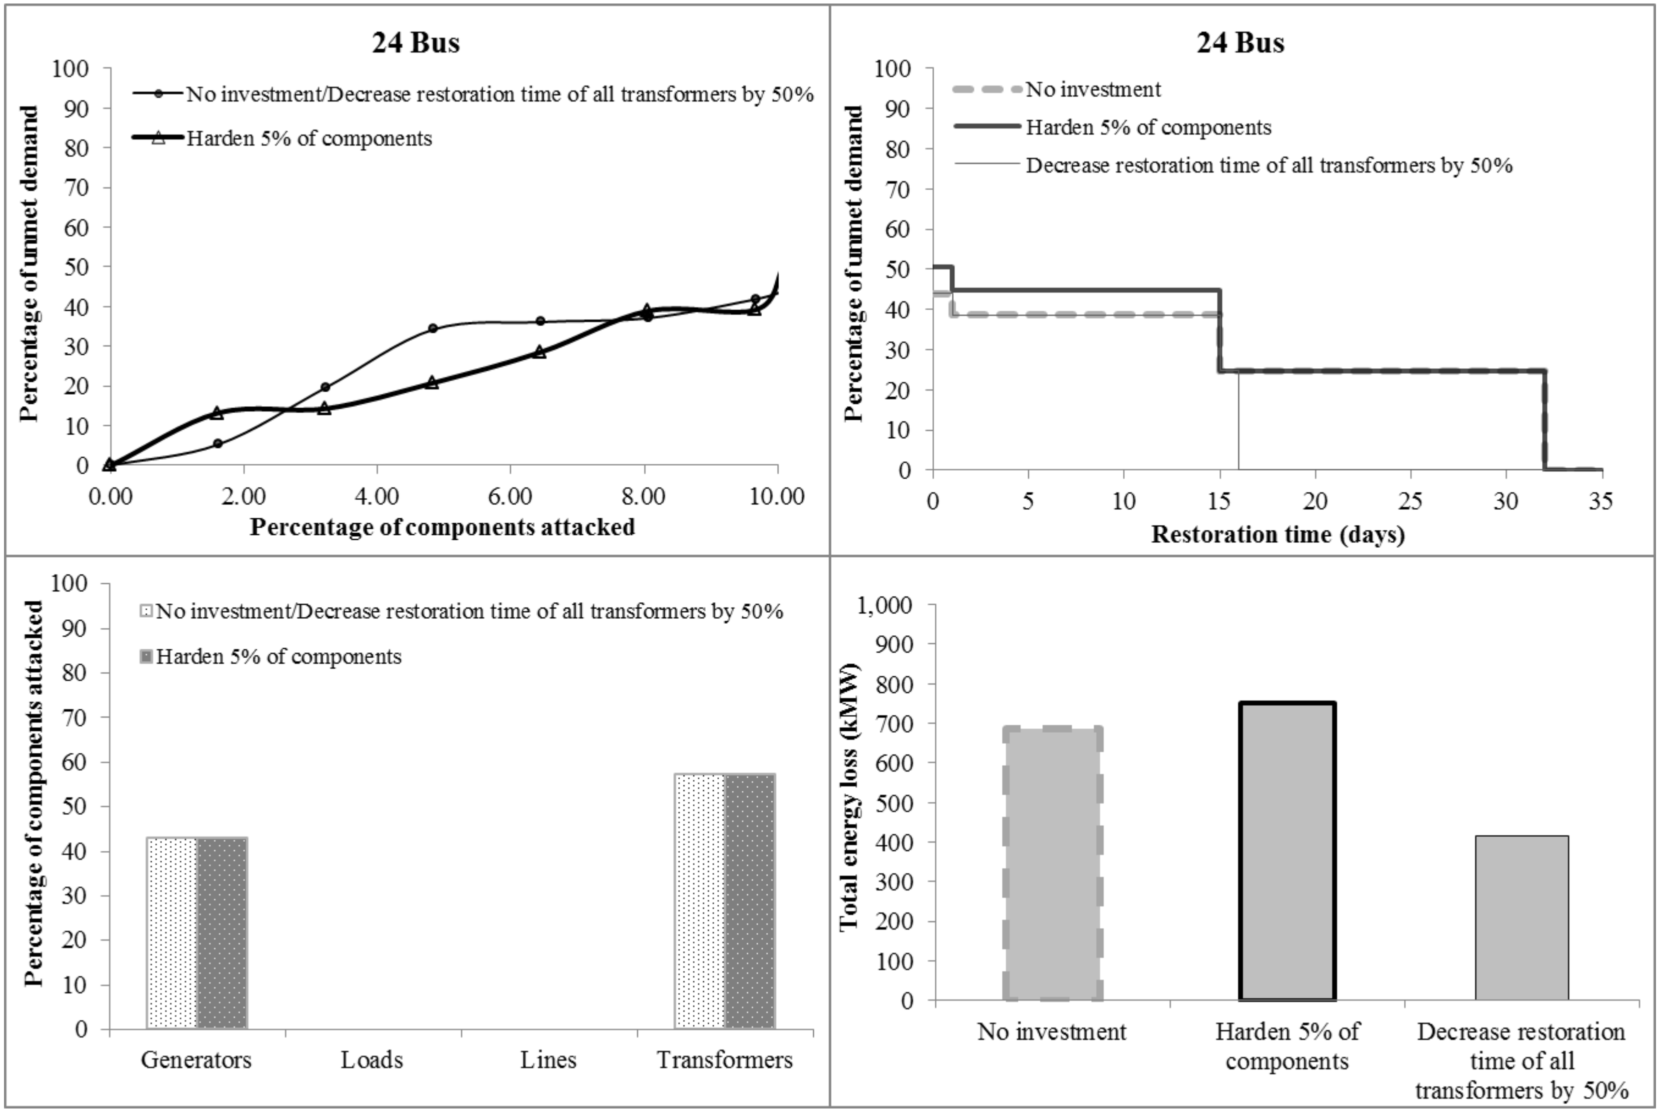

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and considering restoration times, 48-bus system)

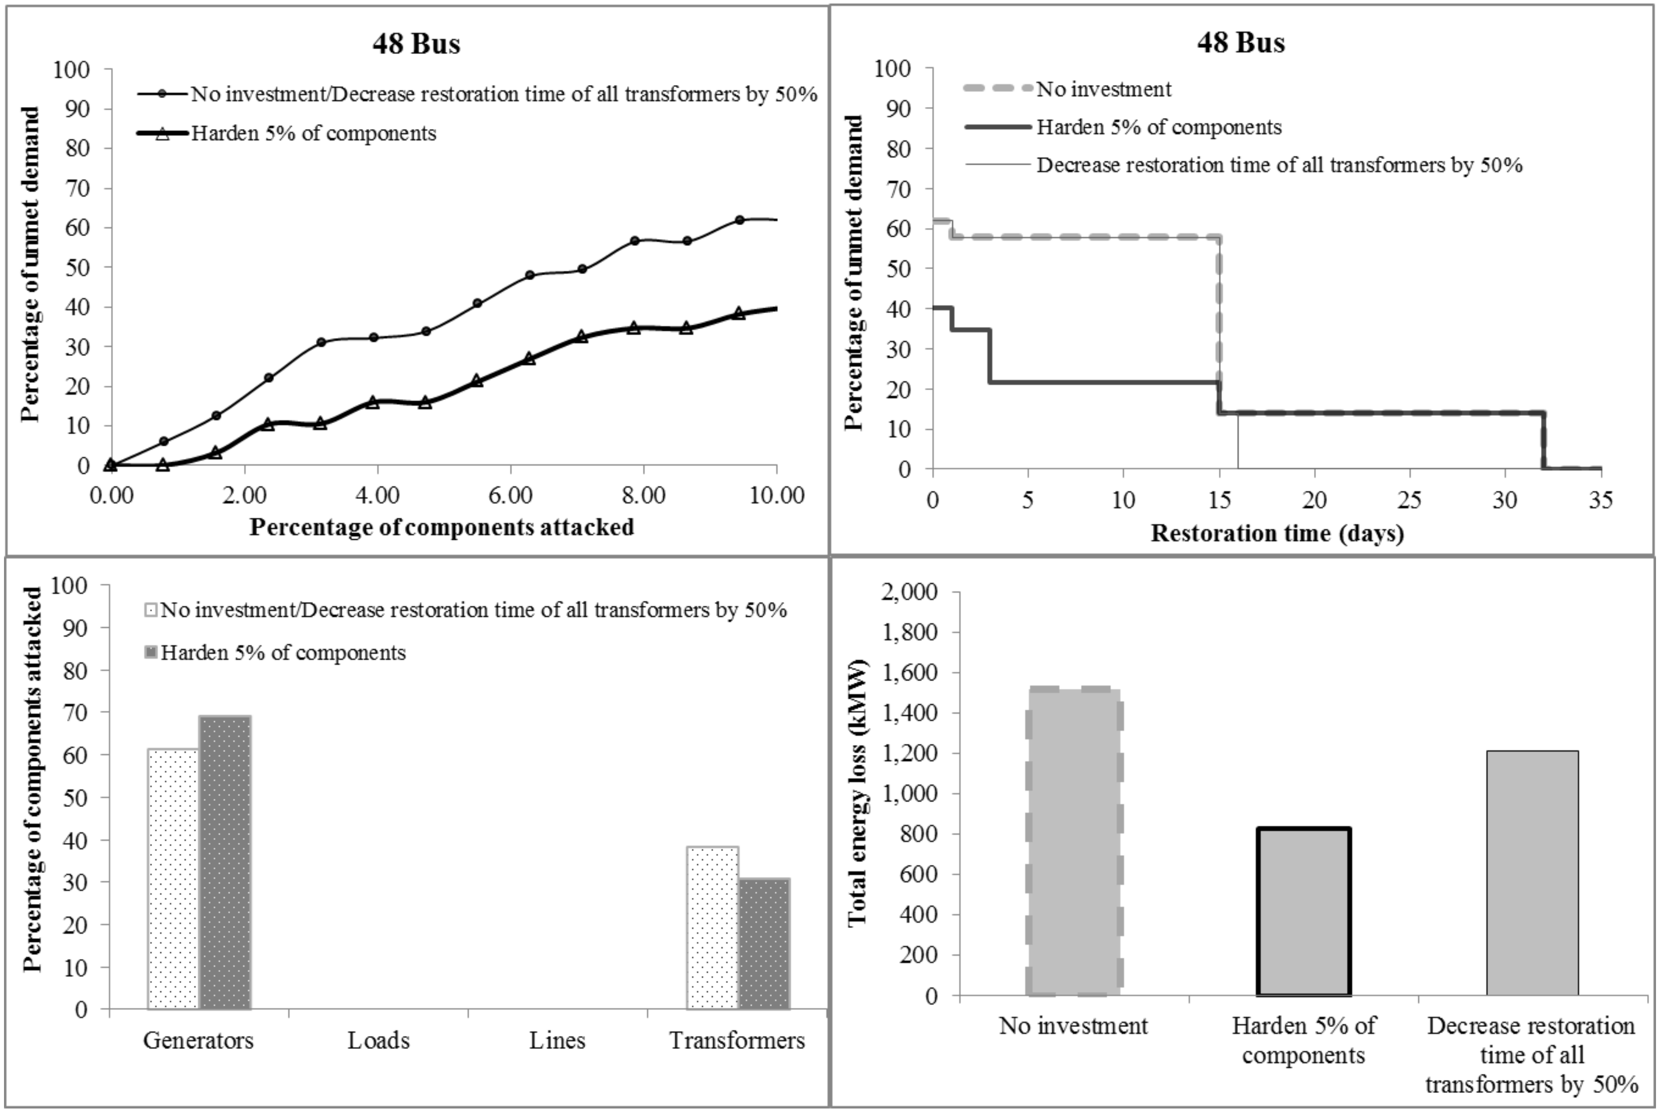

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and considering restoration times, 118-bus system)

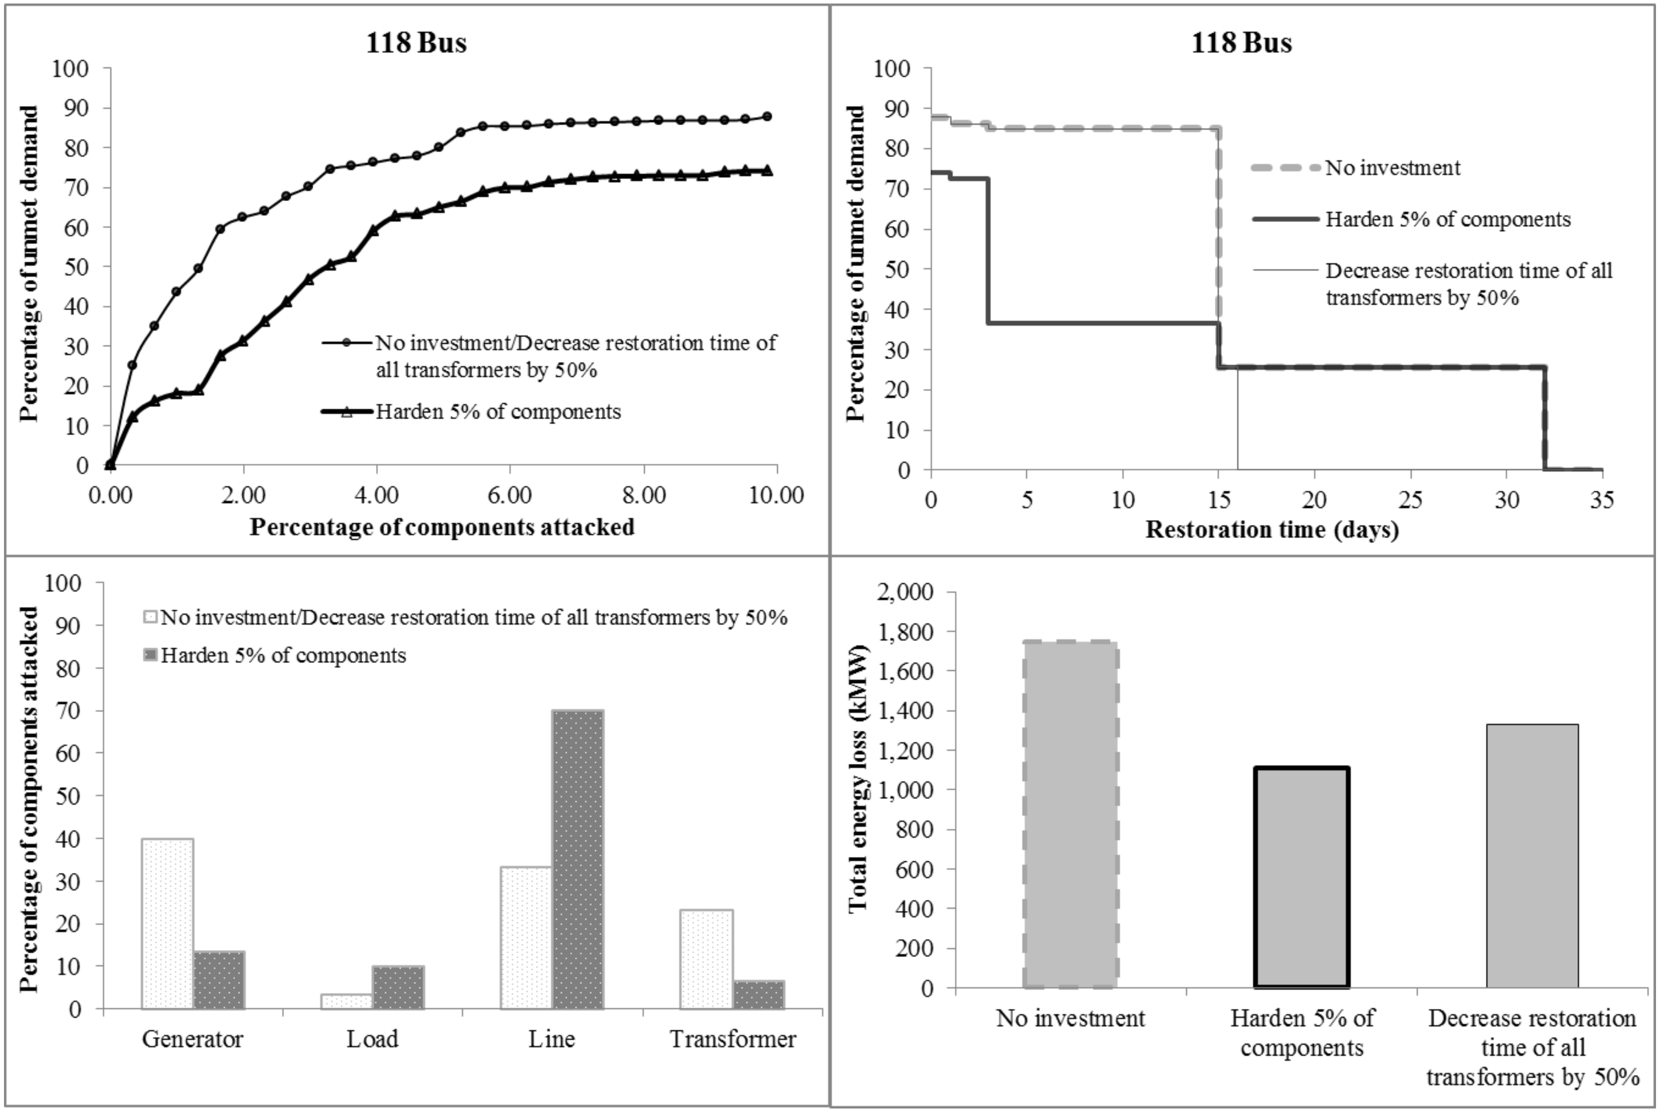

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 5% of the components against a static attacker with cascading knowledge and considering restoration times, 300-bus system)

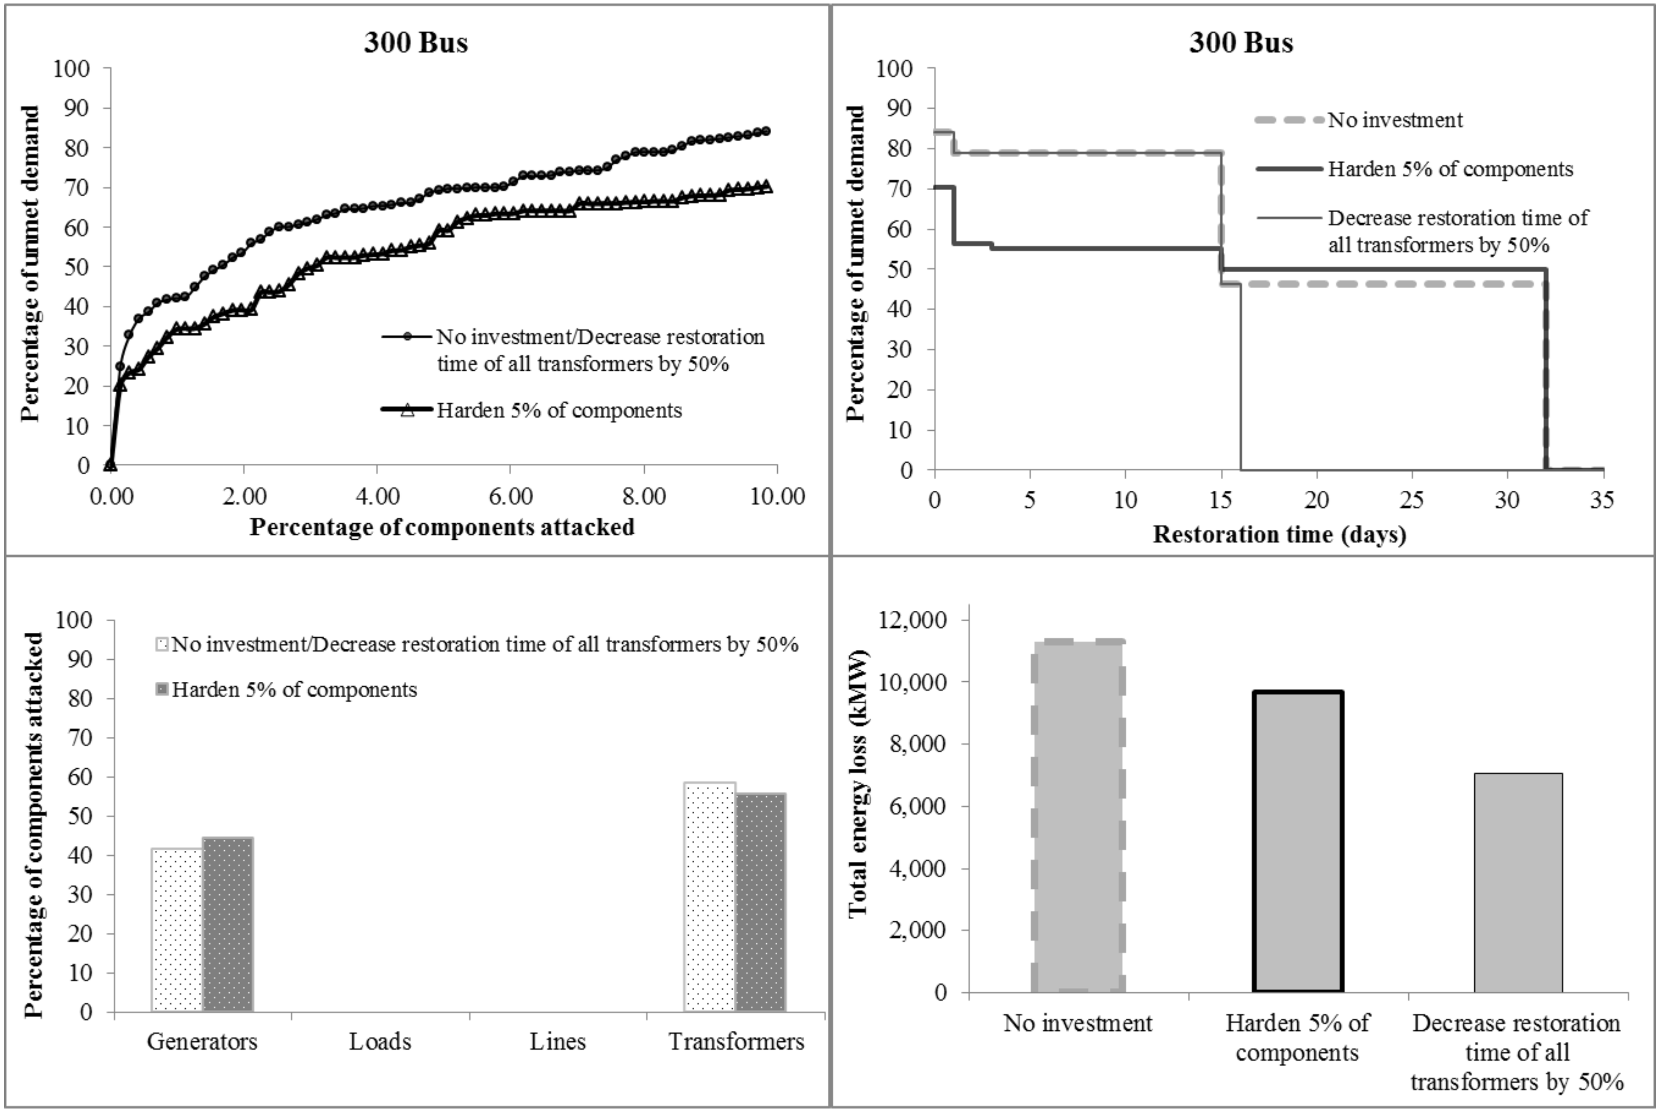

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and considering restoration times, 24-bus system)

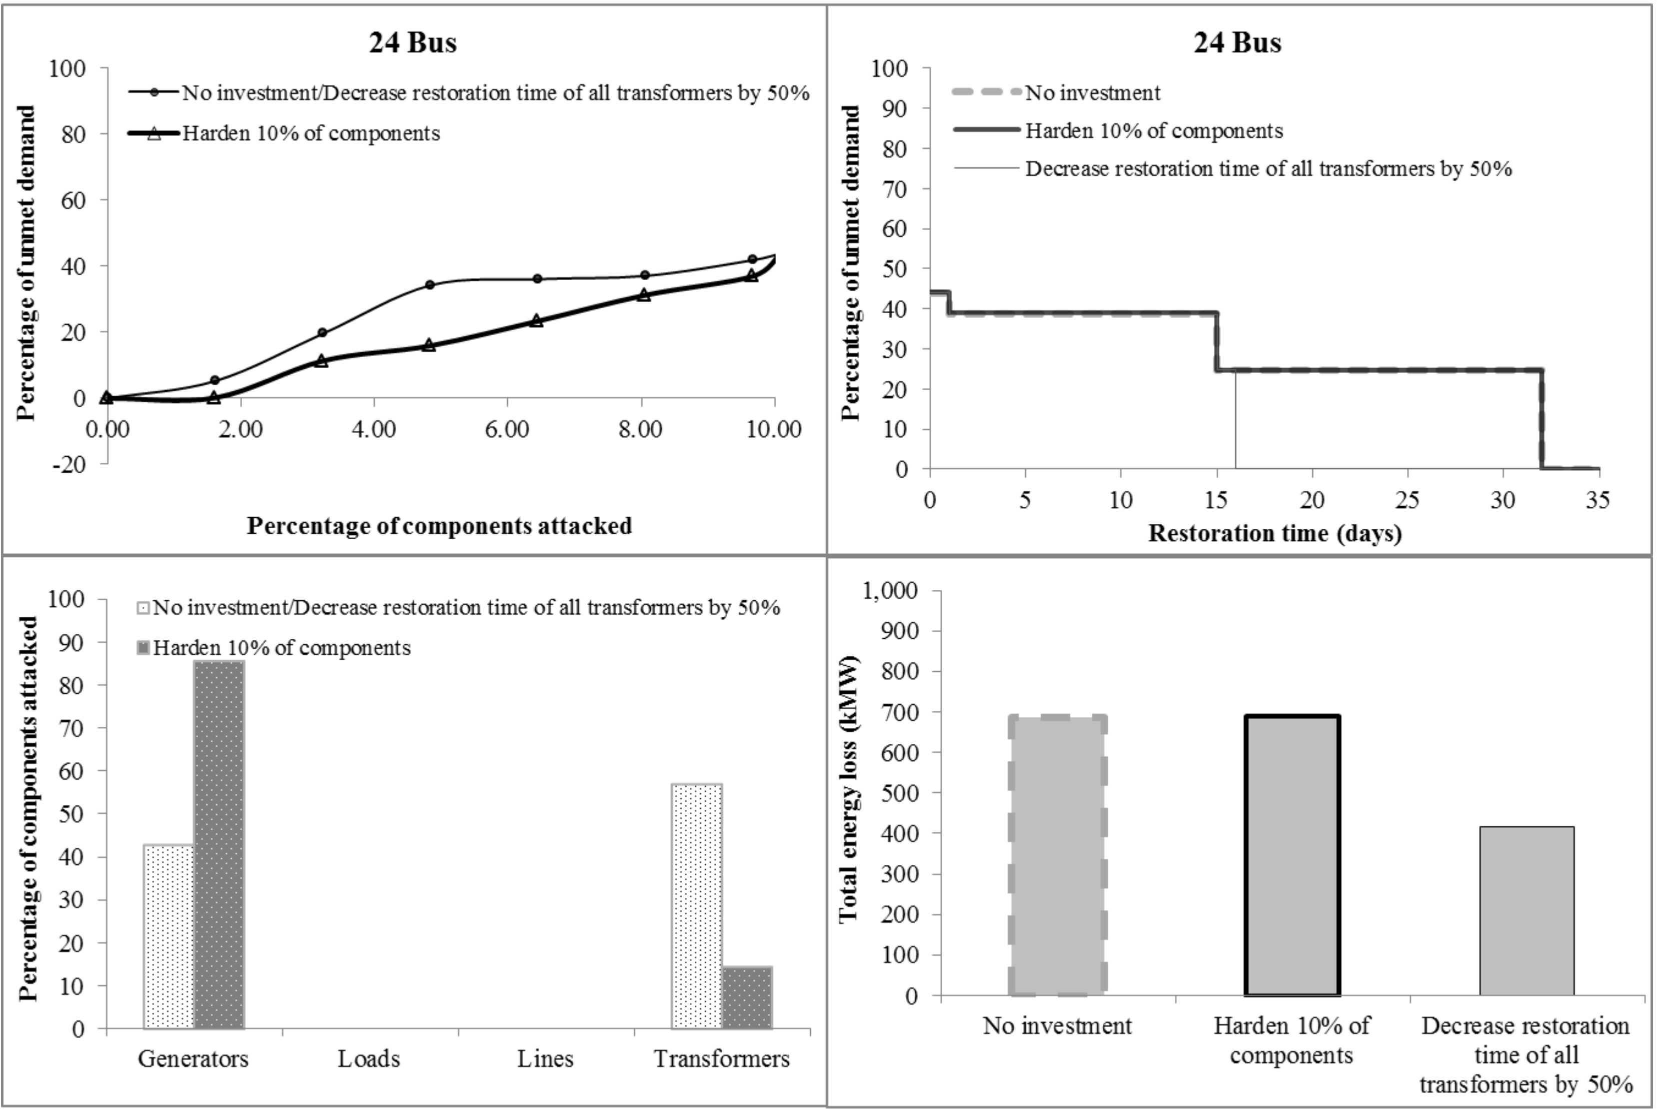

*

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and considering restoration times, 48-bus system)

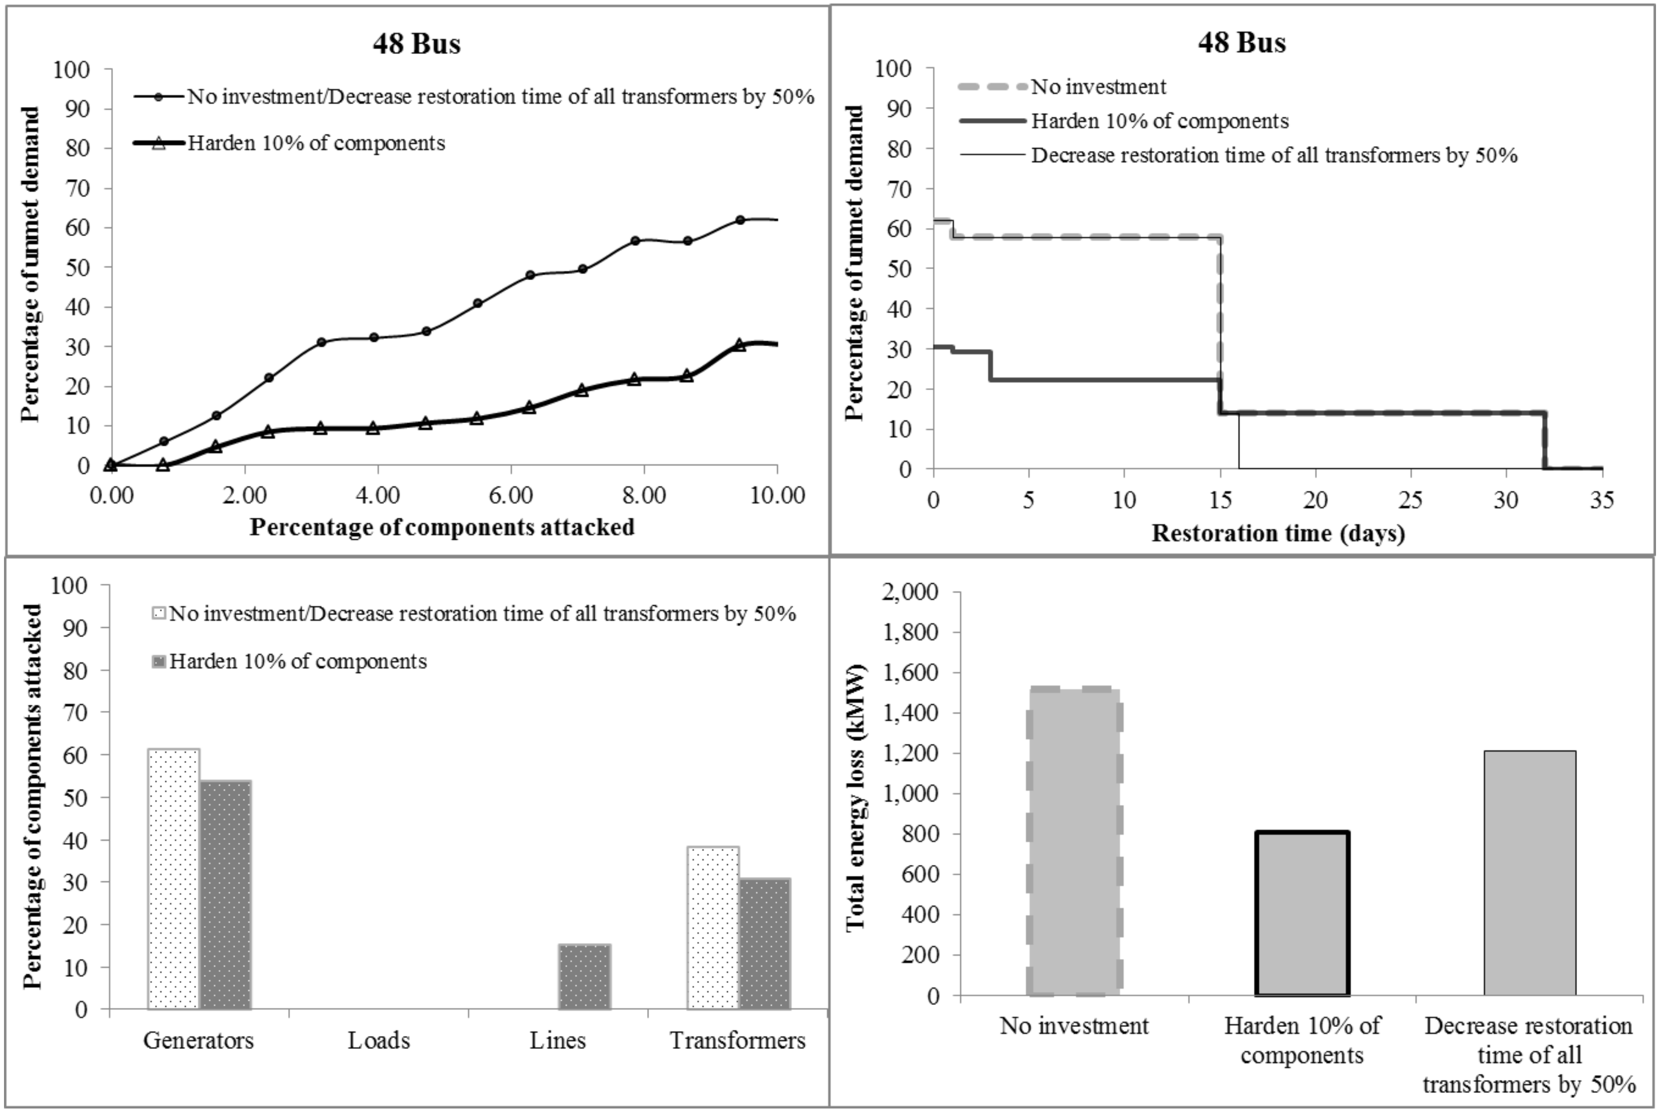

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and considering restoration times, 118-bus system)

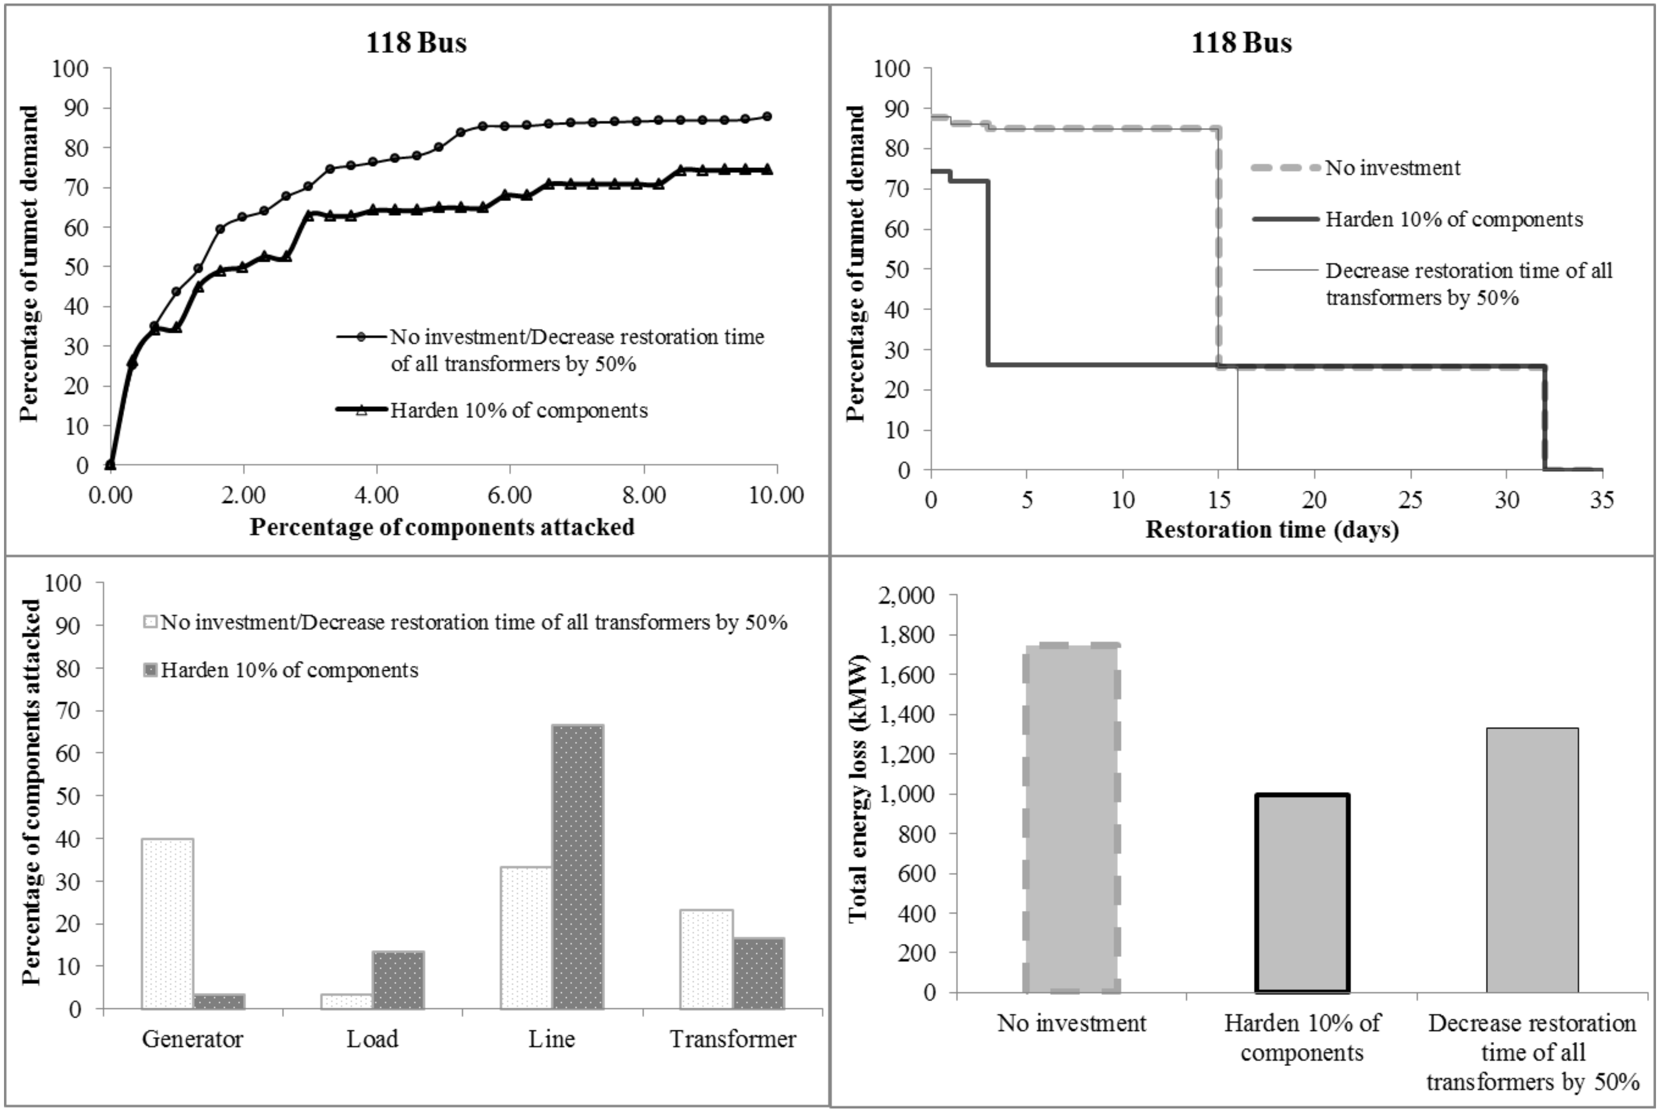

Appendix M. Hardening versus decreasing restoration times of transformers (hardening 10% of the components against a static attacker with cascading knowledge and considering restoration times, 300-bus system)

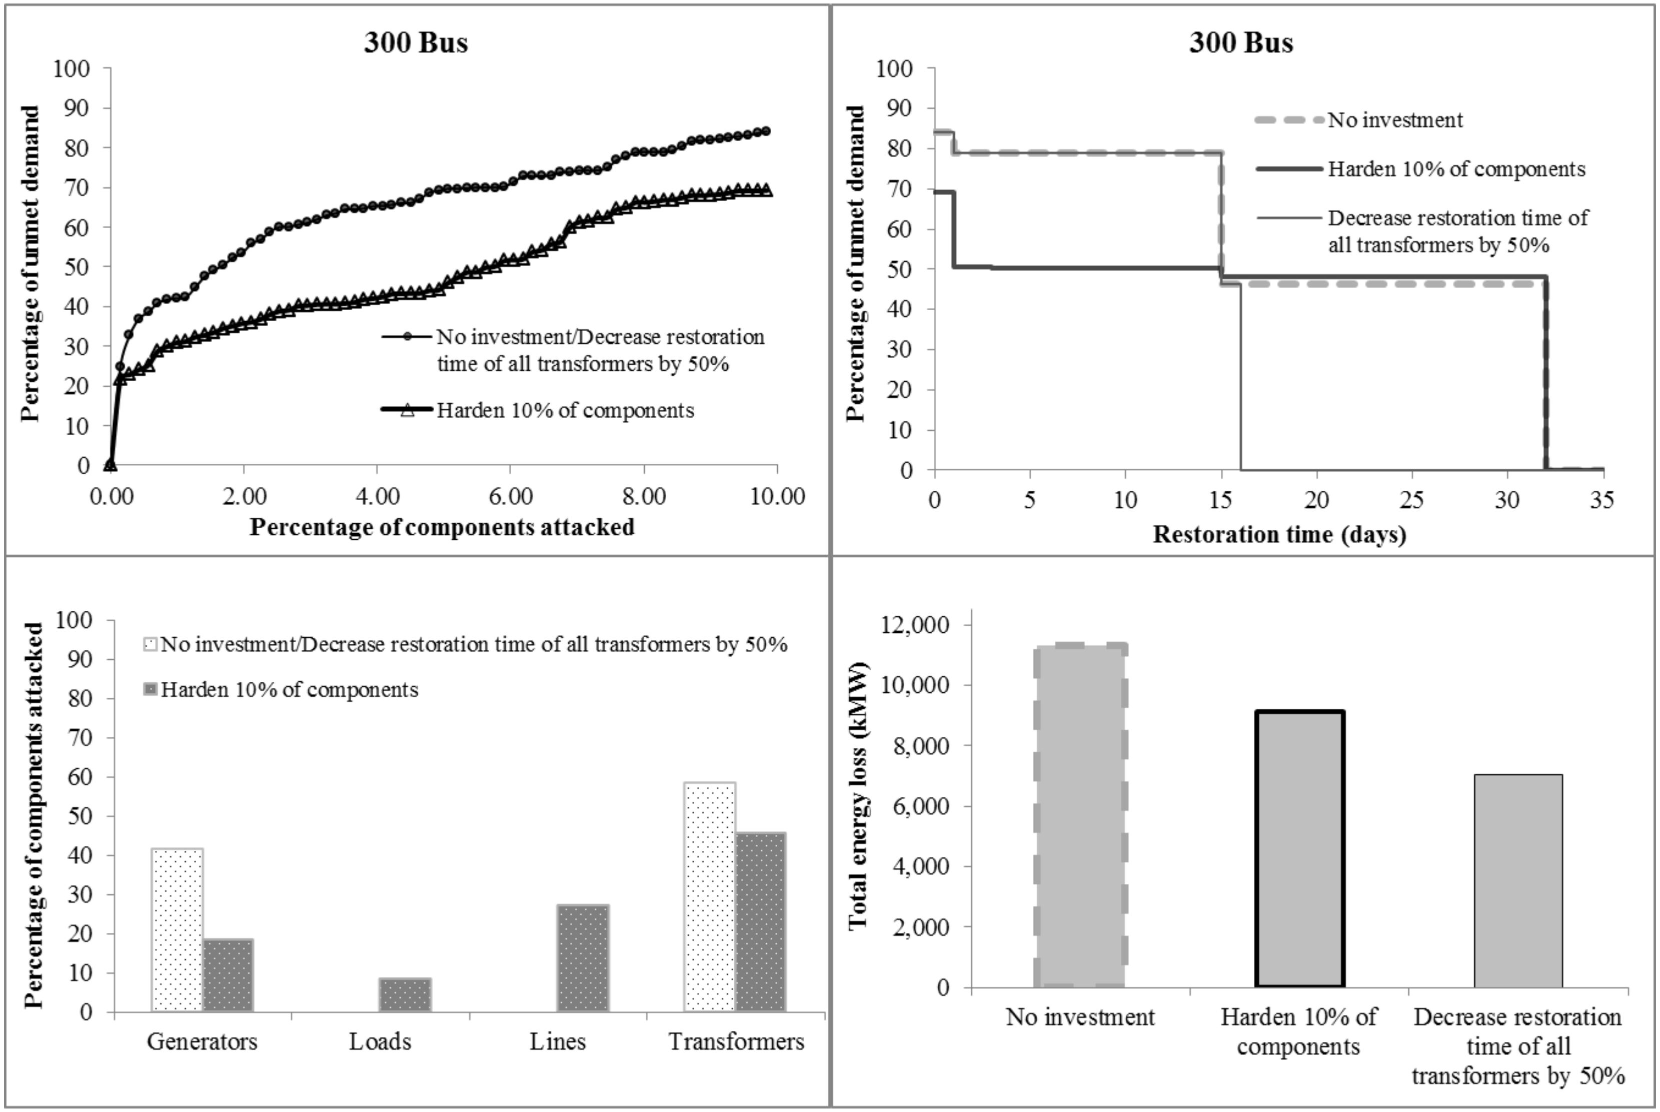

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a dynamic attacker, 24-bus system)

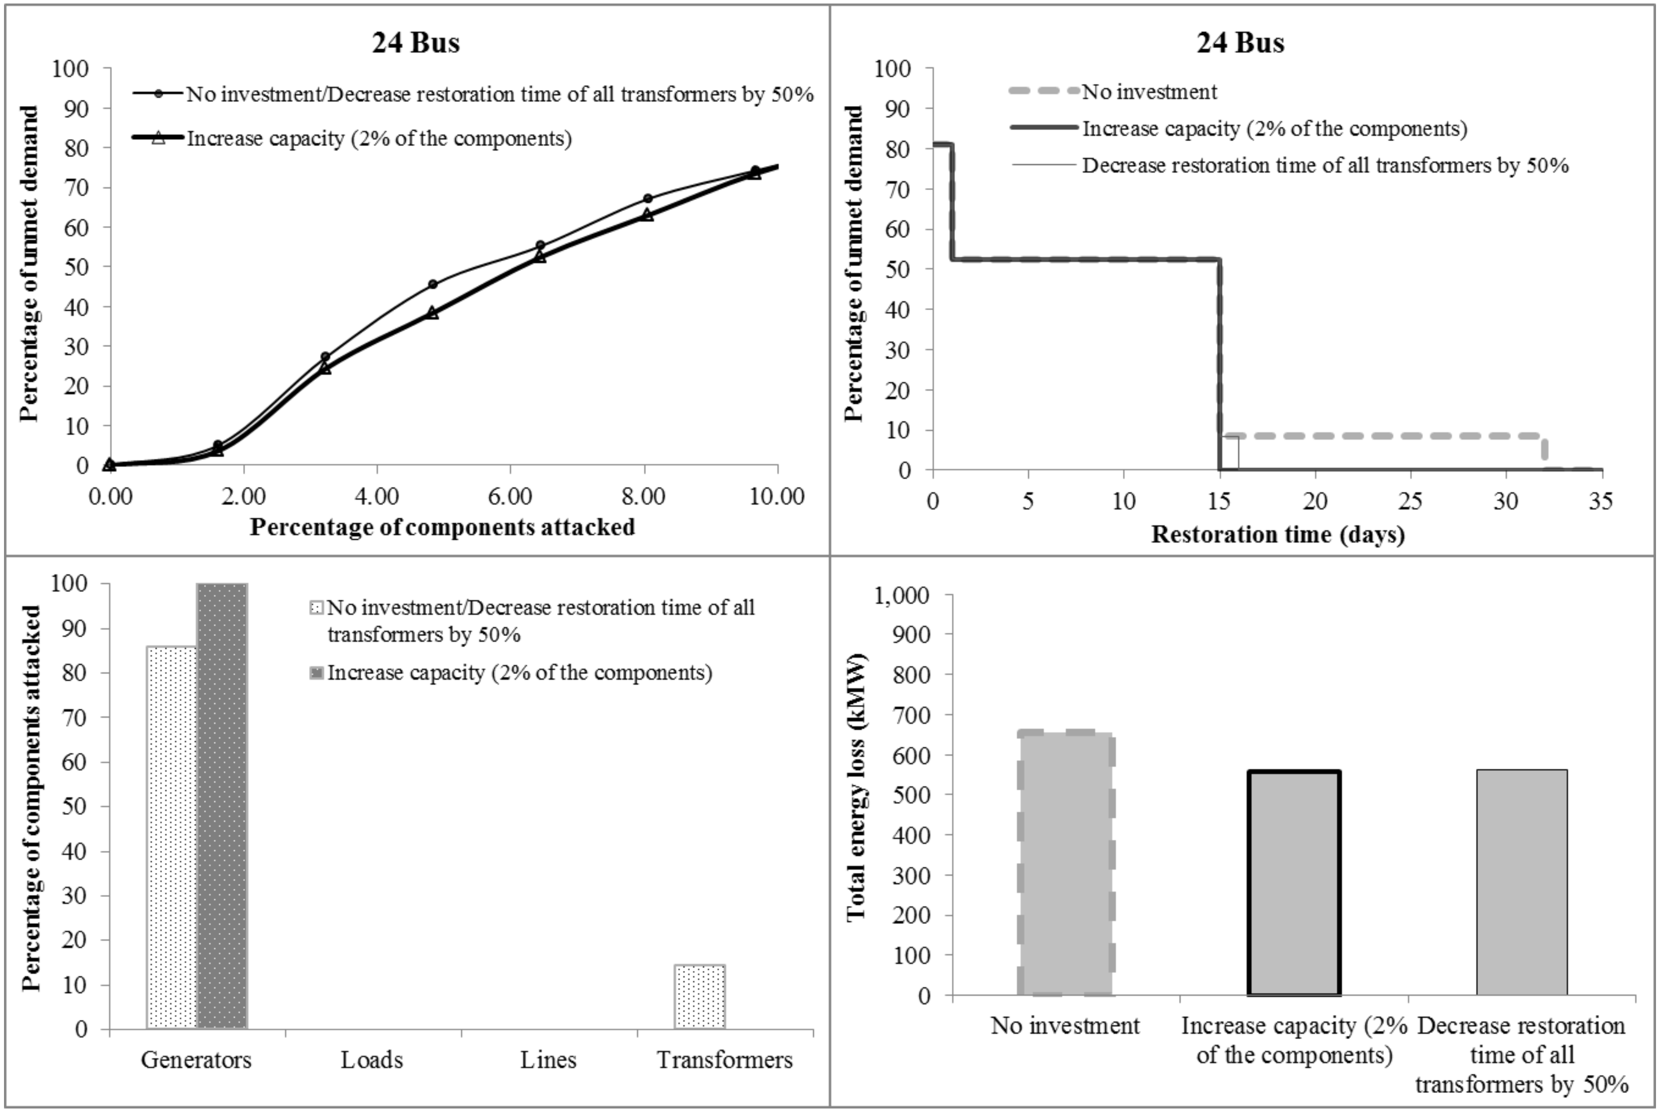

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a dynamic attacker, 48-bus system)

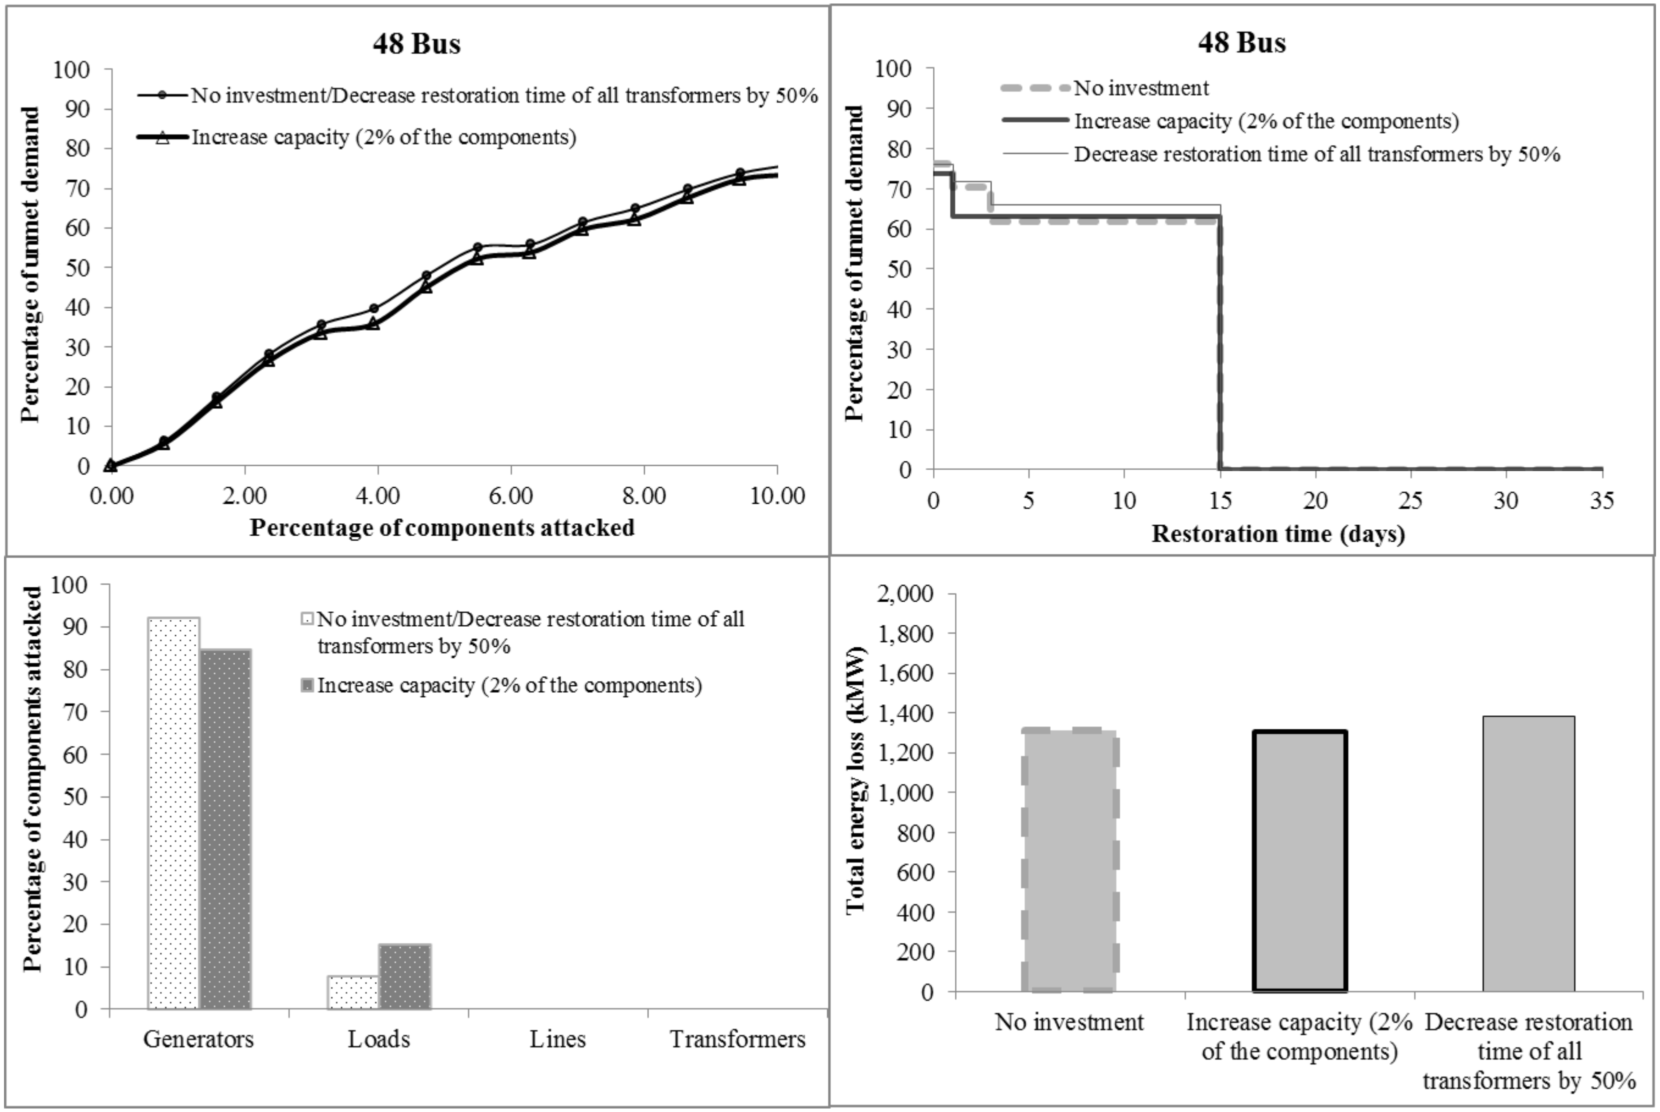

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a dynamic attacker, 118-bus system)

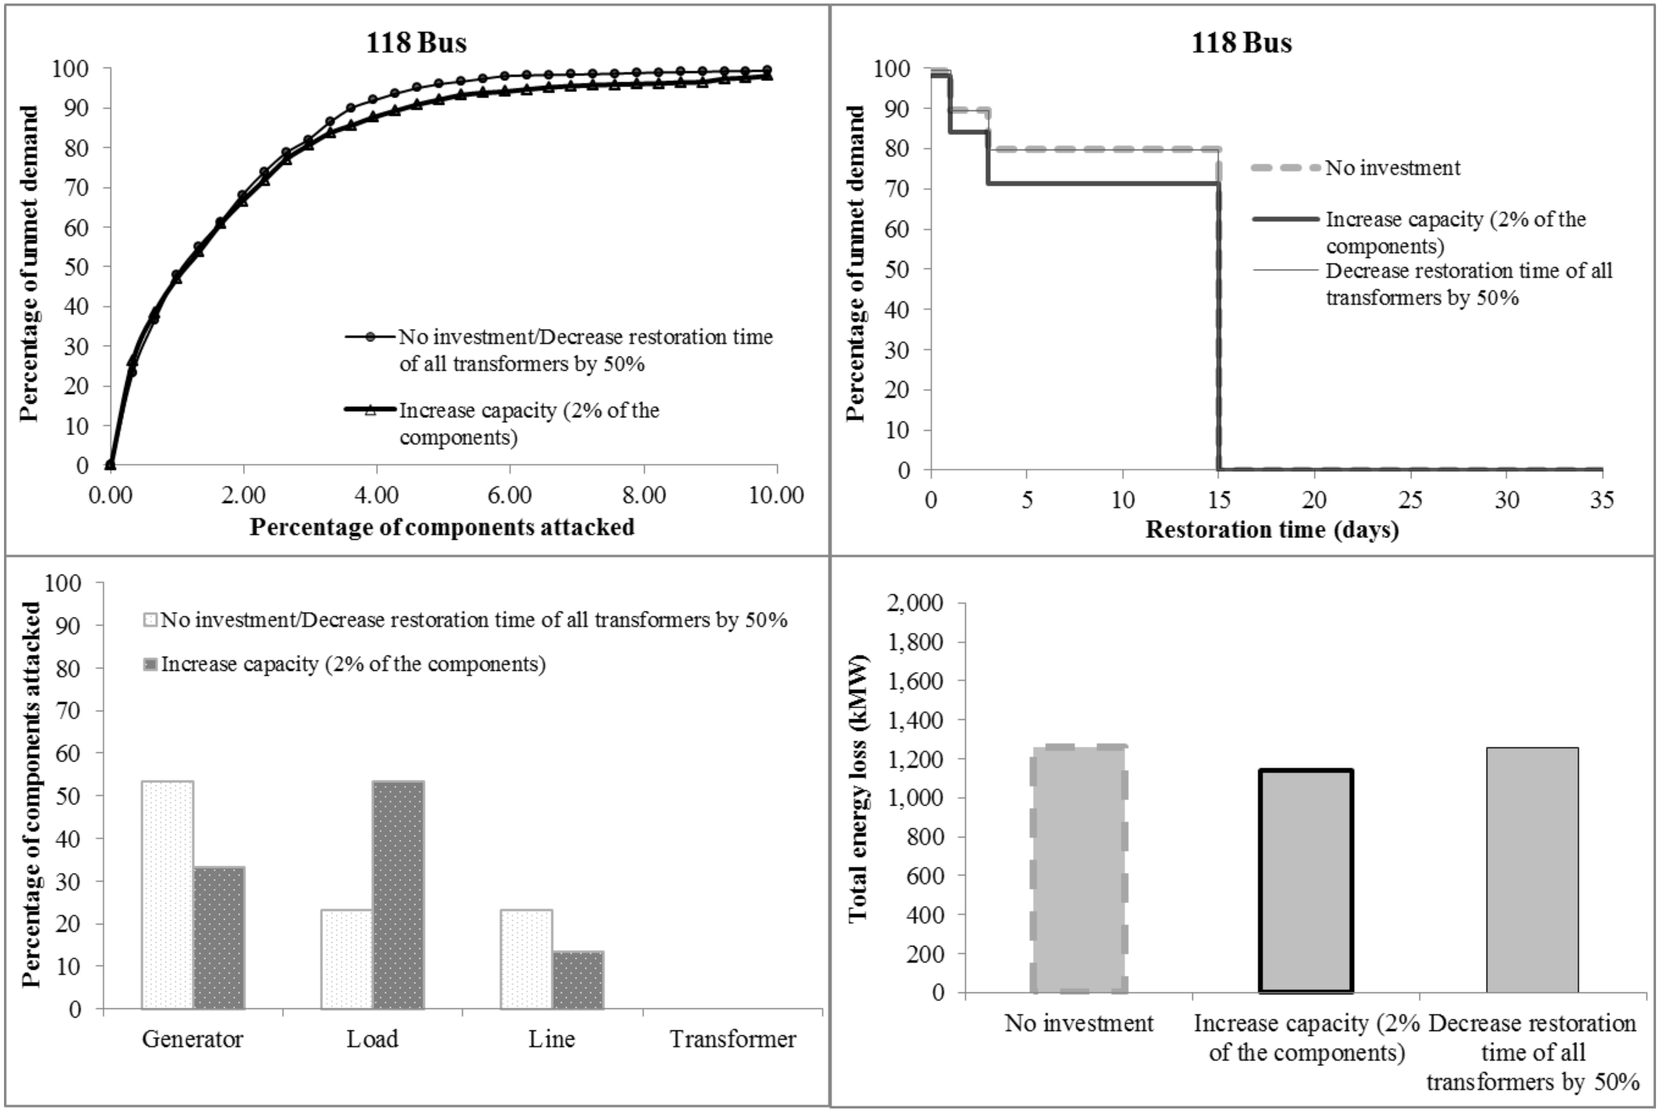

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a dynamic attacker, 300-bus system)

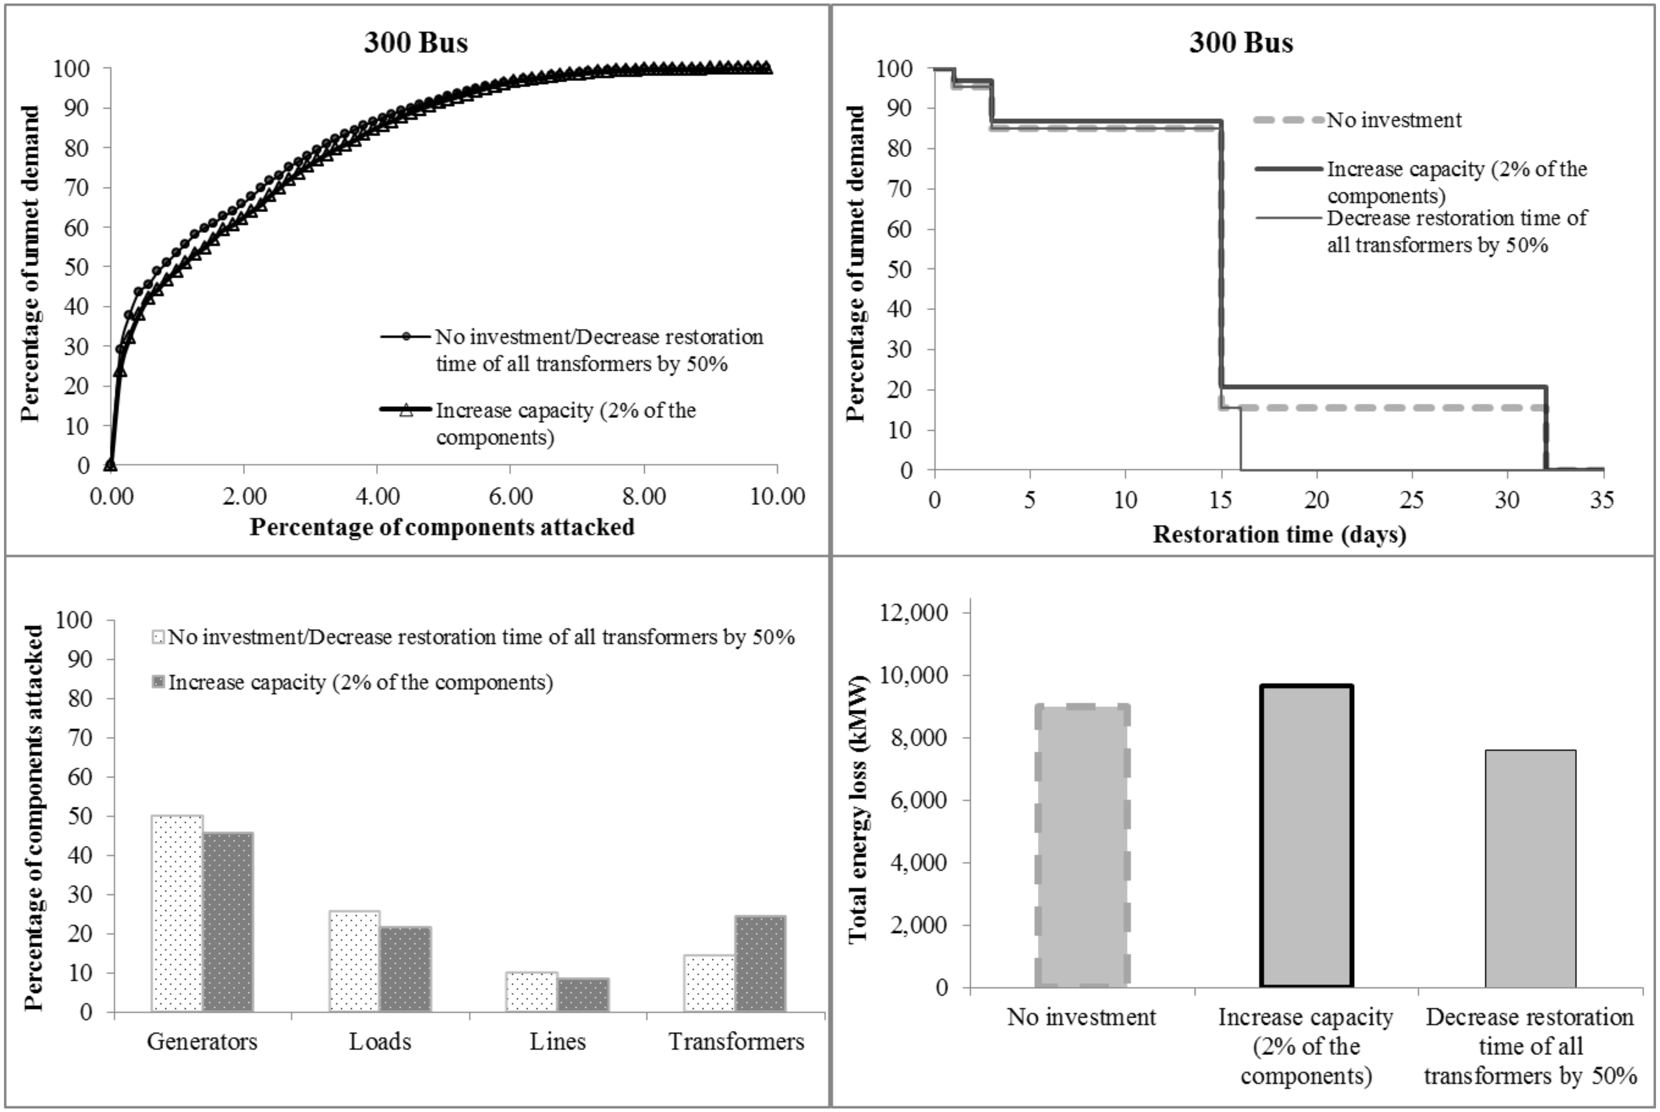

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a dynamic attacker, 24-bus system)

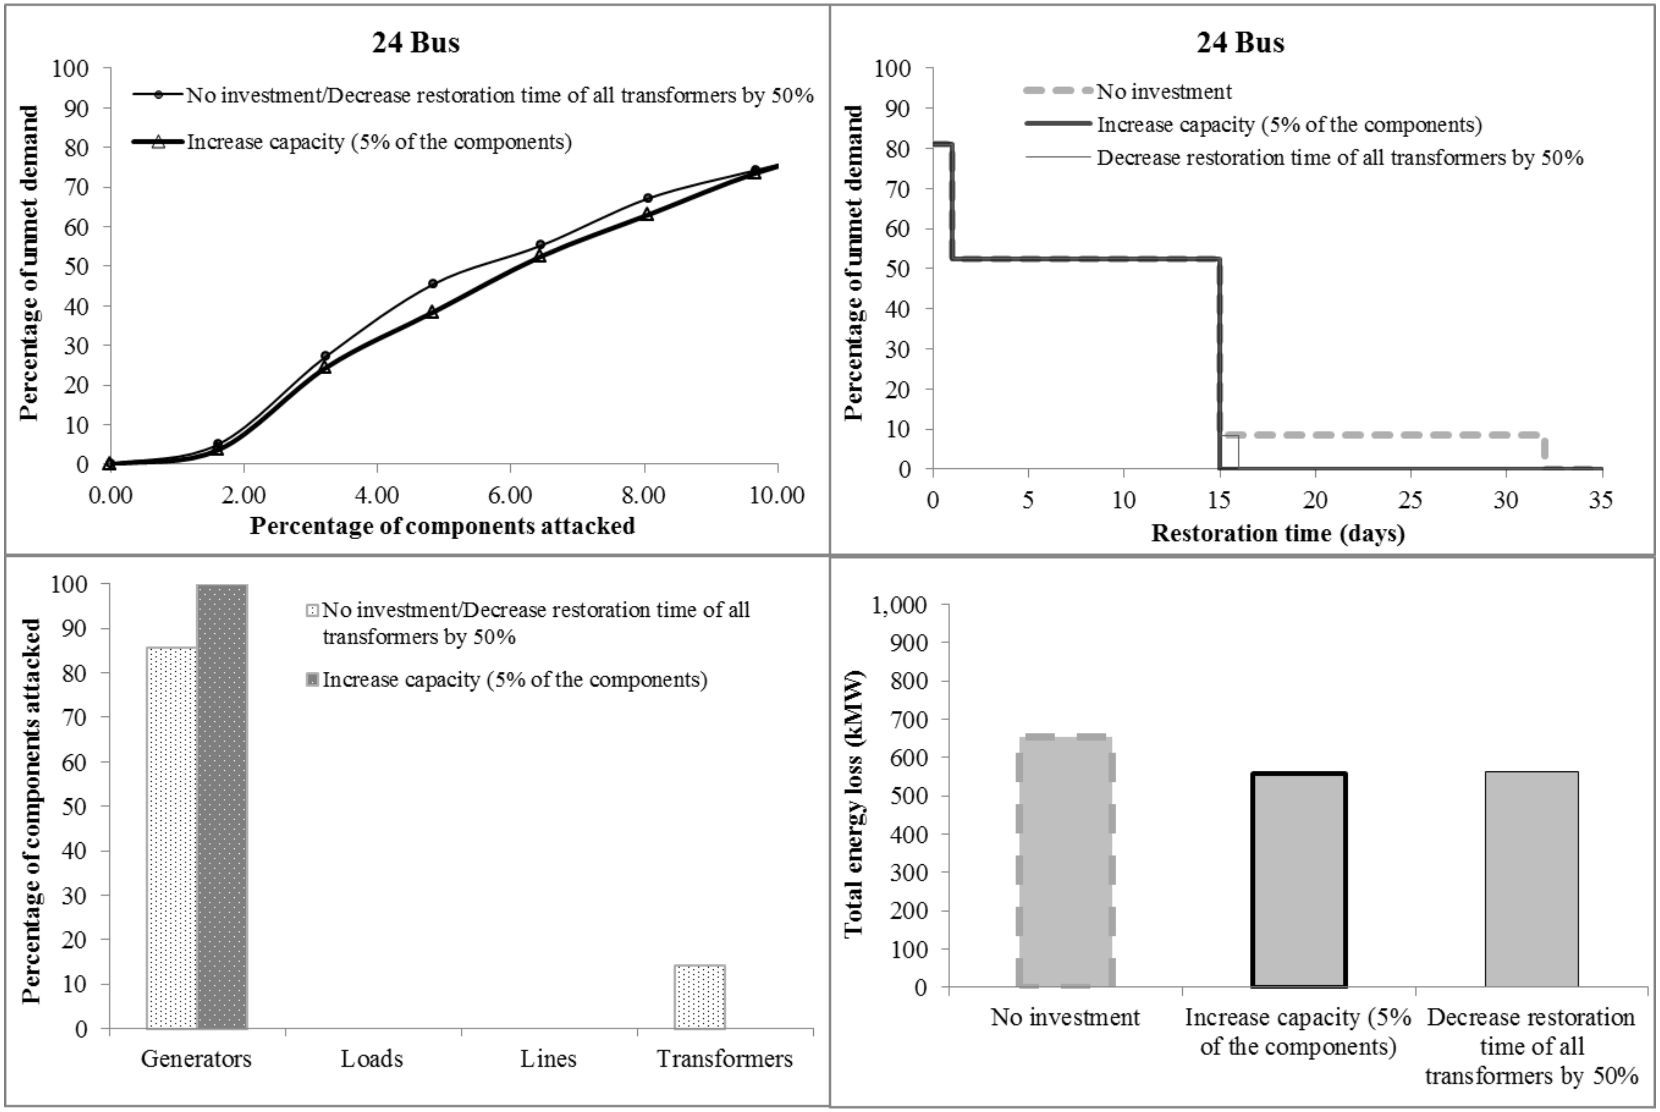

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a dynamic attacker, 48-bus system)

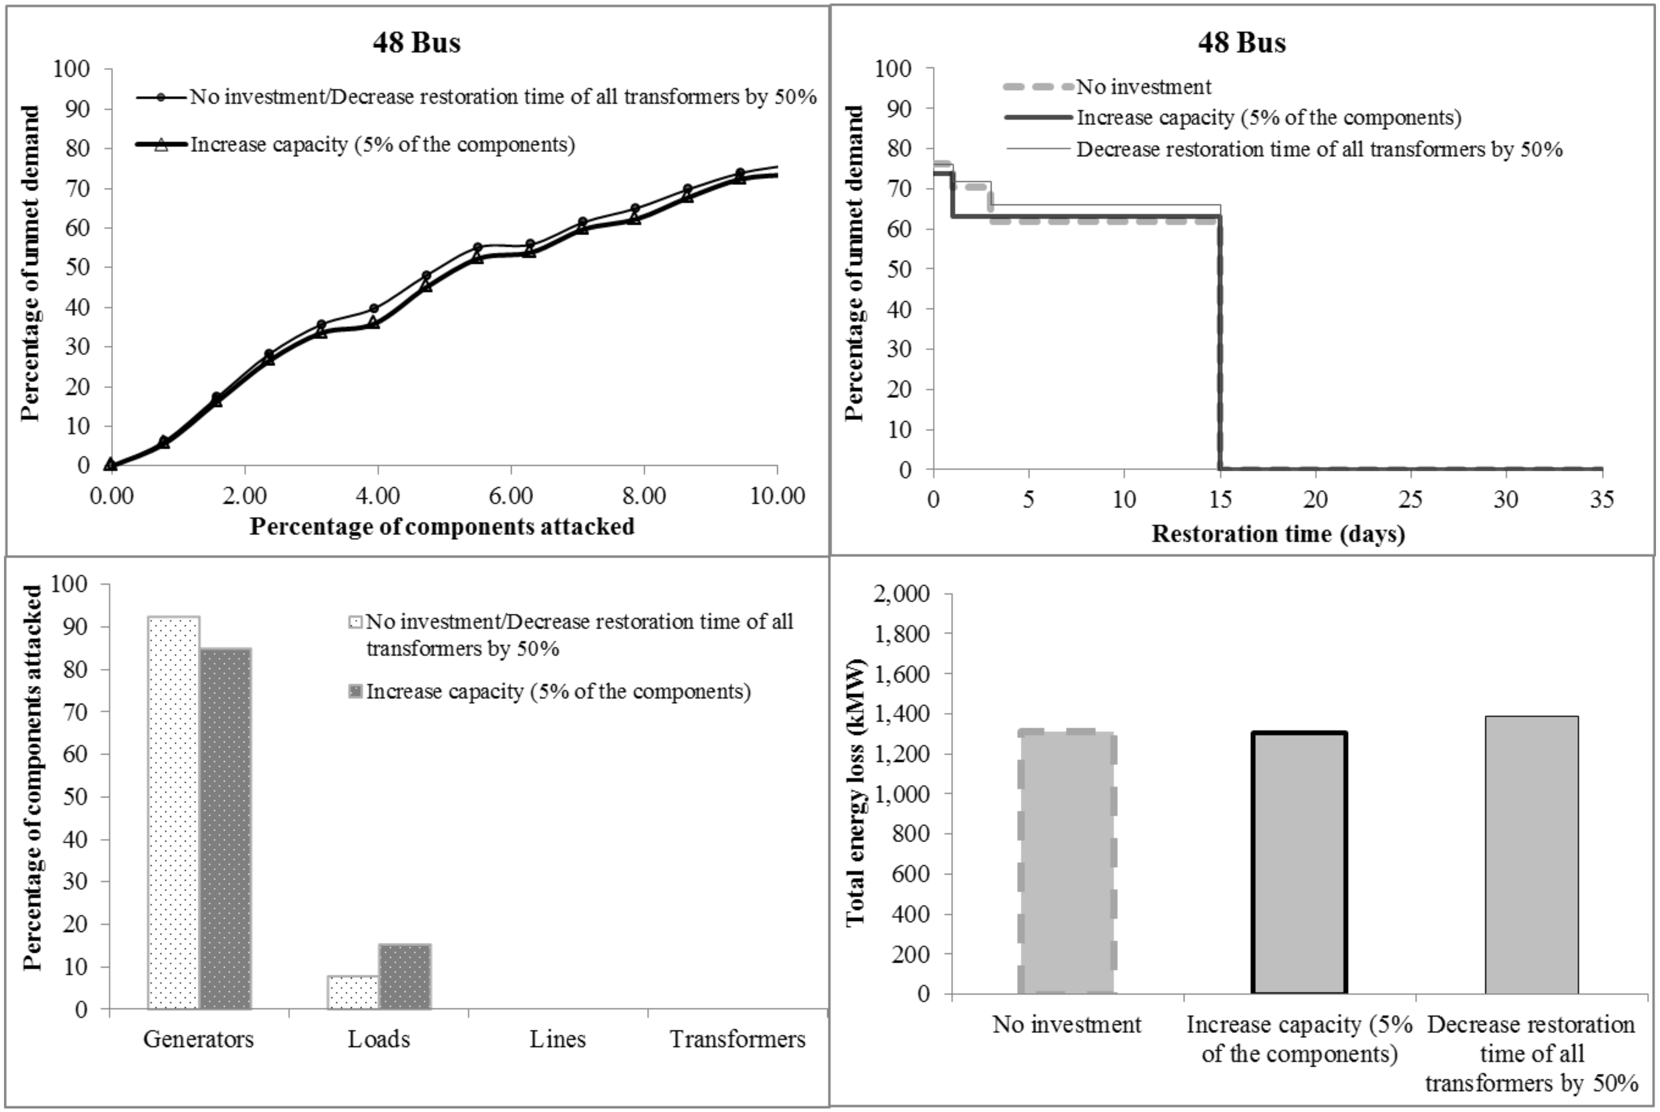

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a dynamic attacker, 118-bus system)

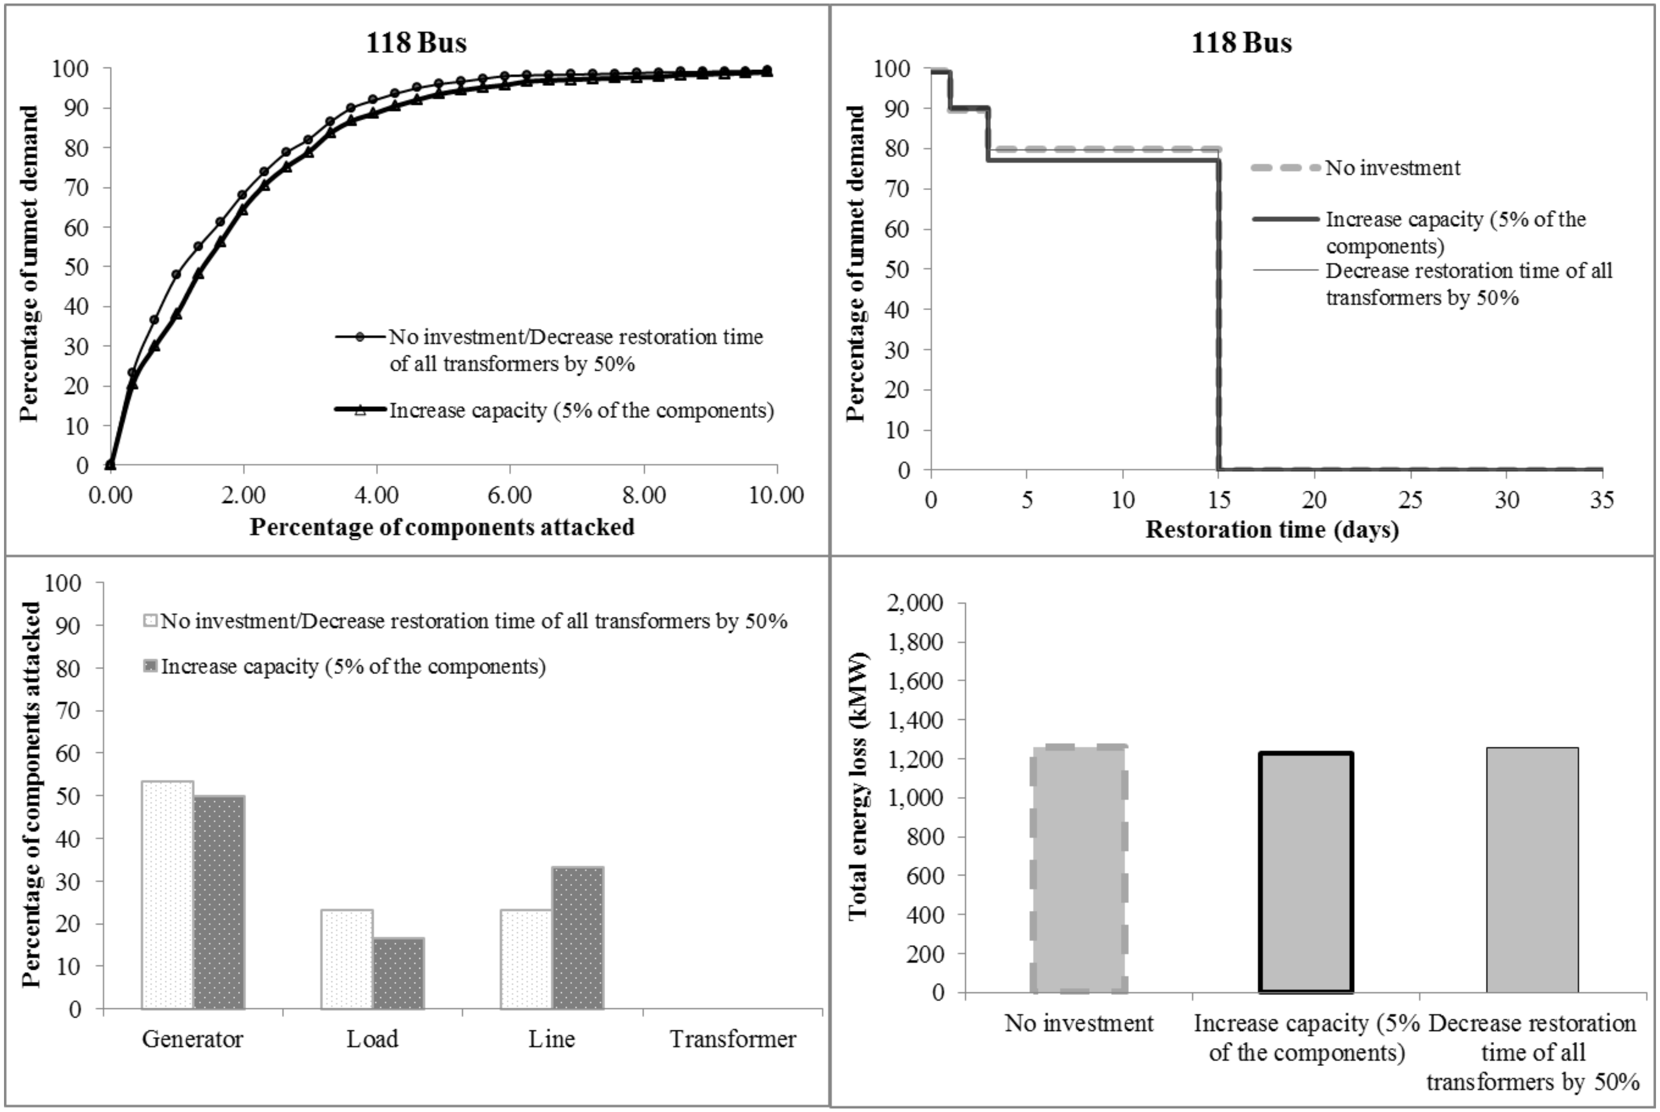

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a dynamic attacker, 300-bus system)

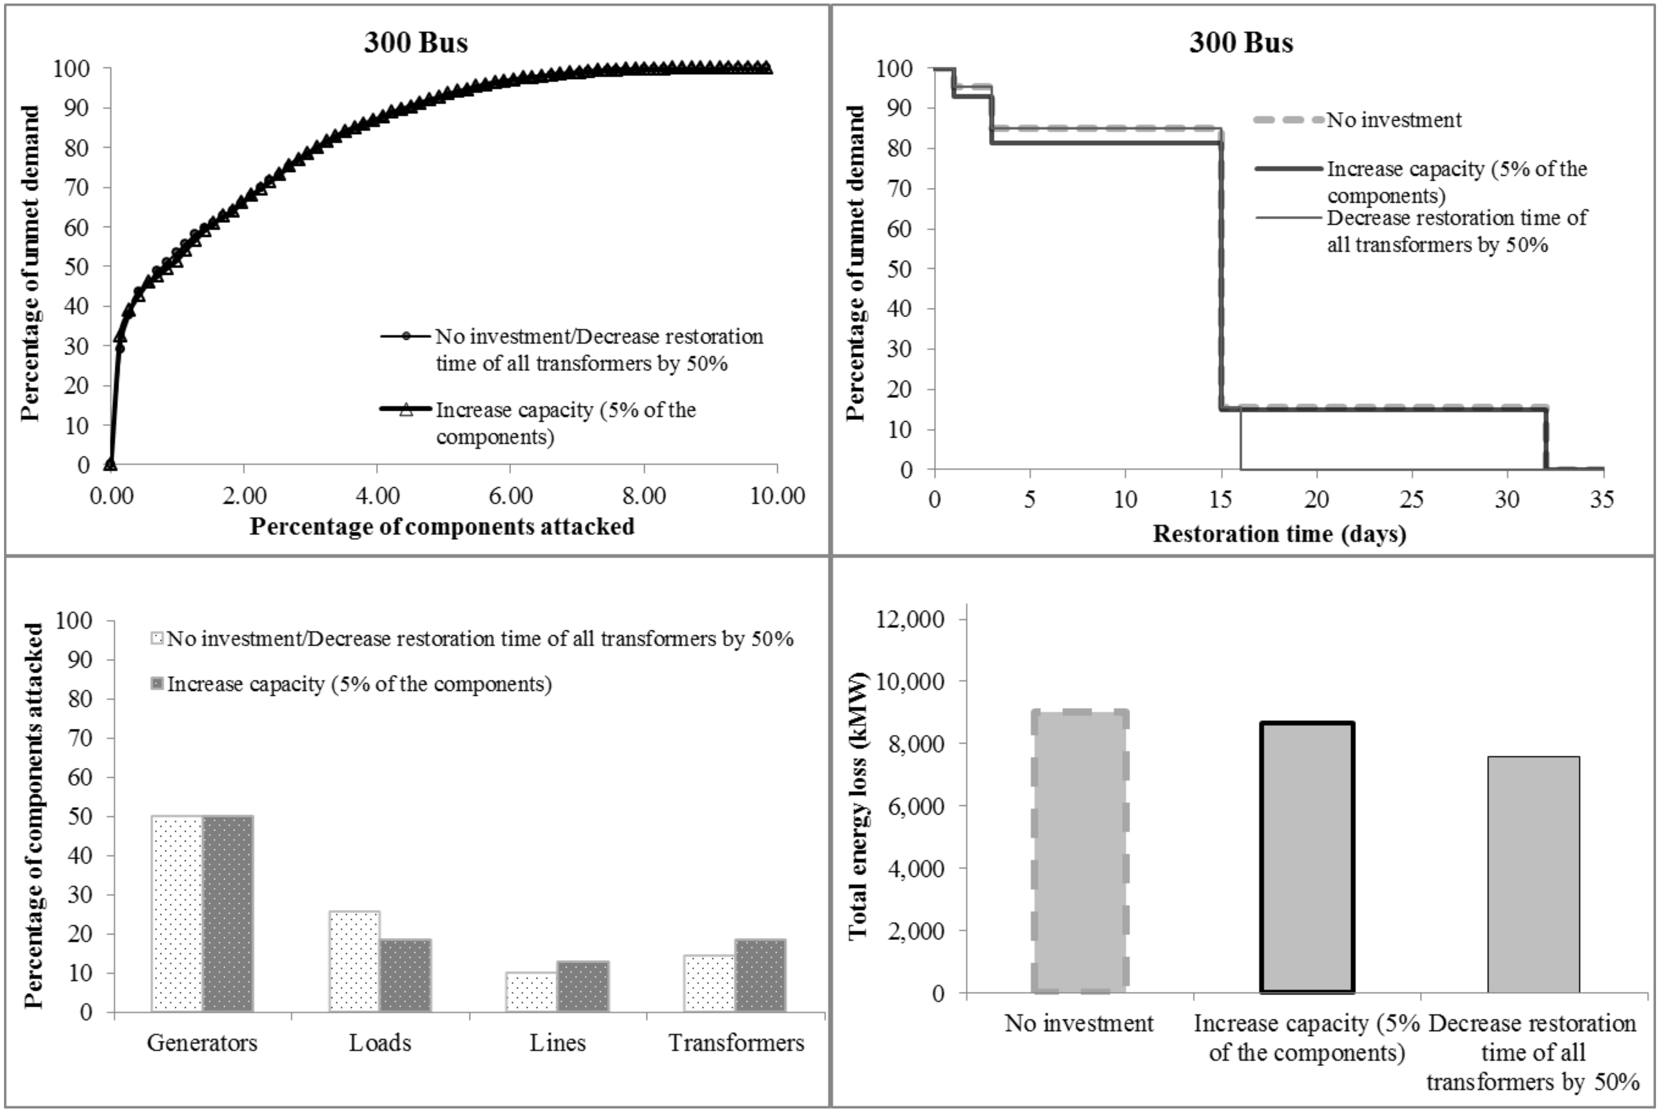

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a dynamic attacker, 24-bus system)

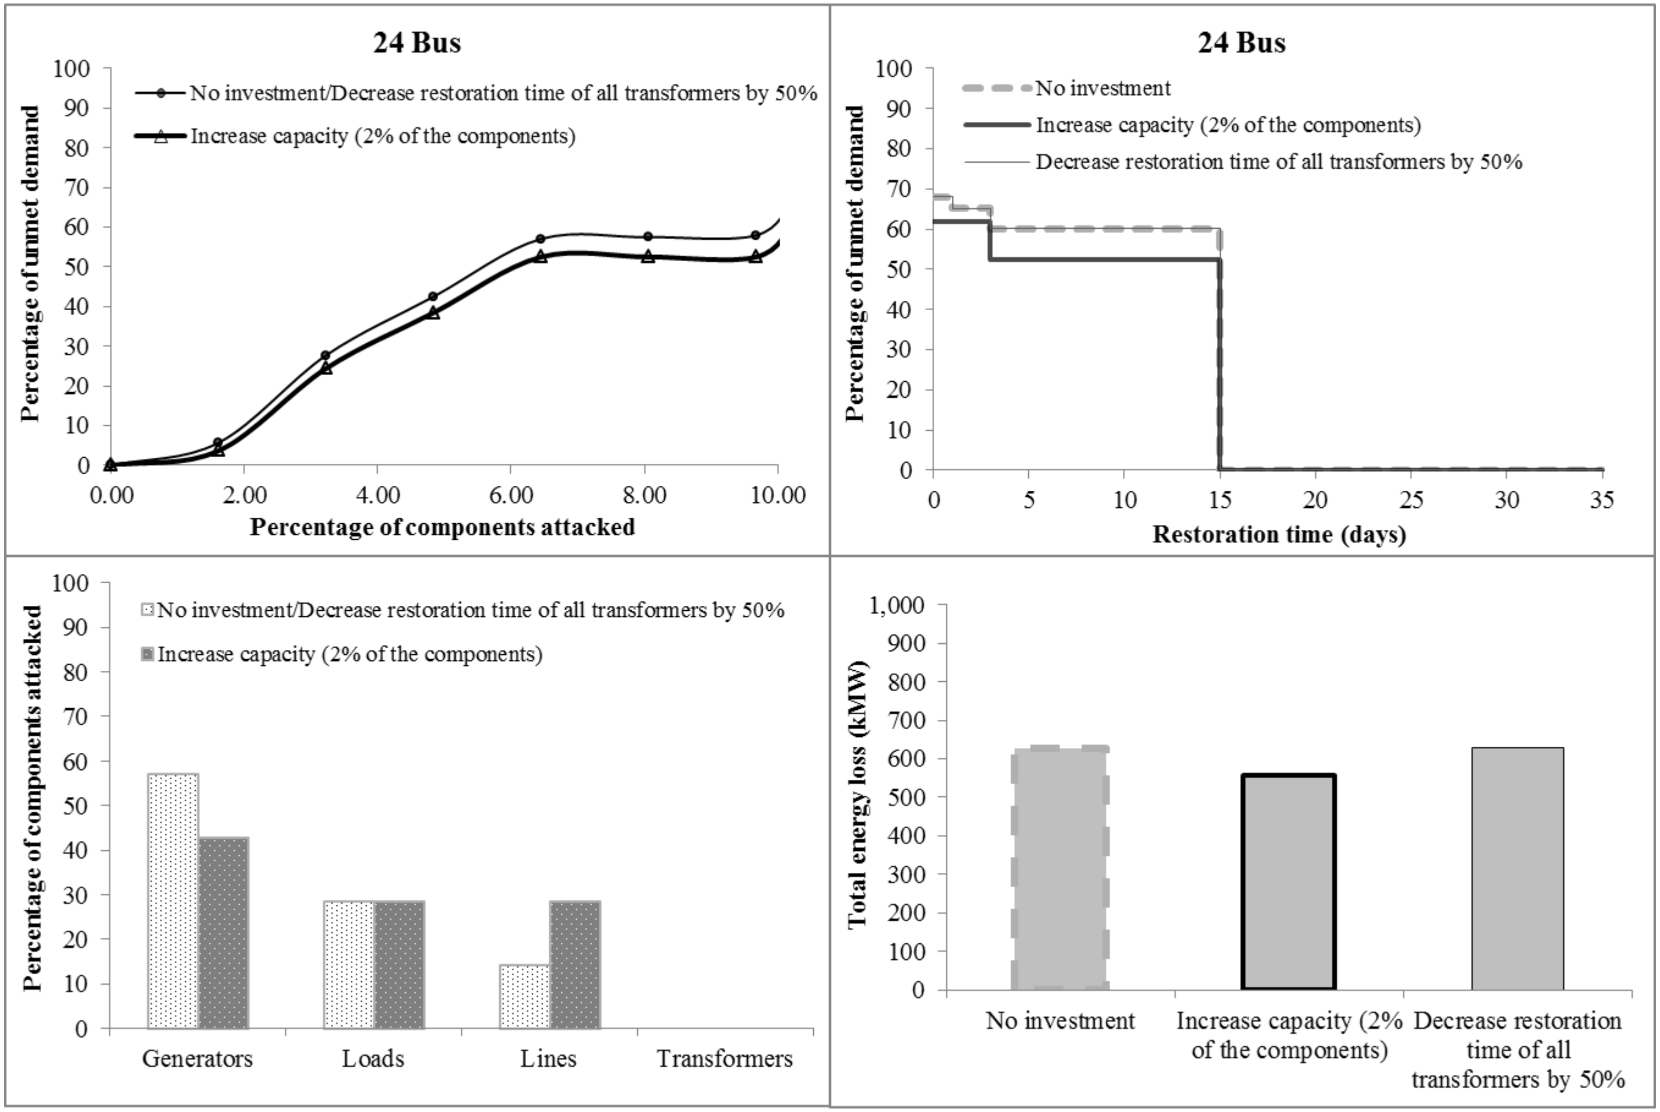

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a dynamic attacker, 48-bus system)

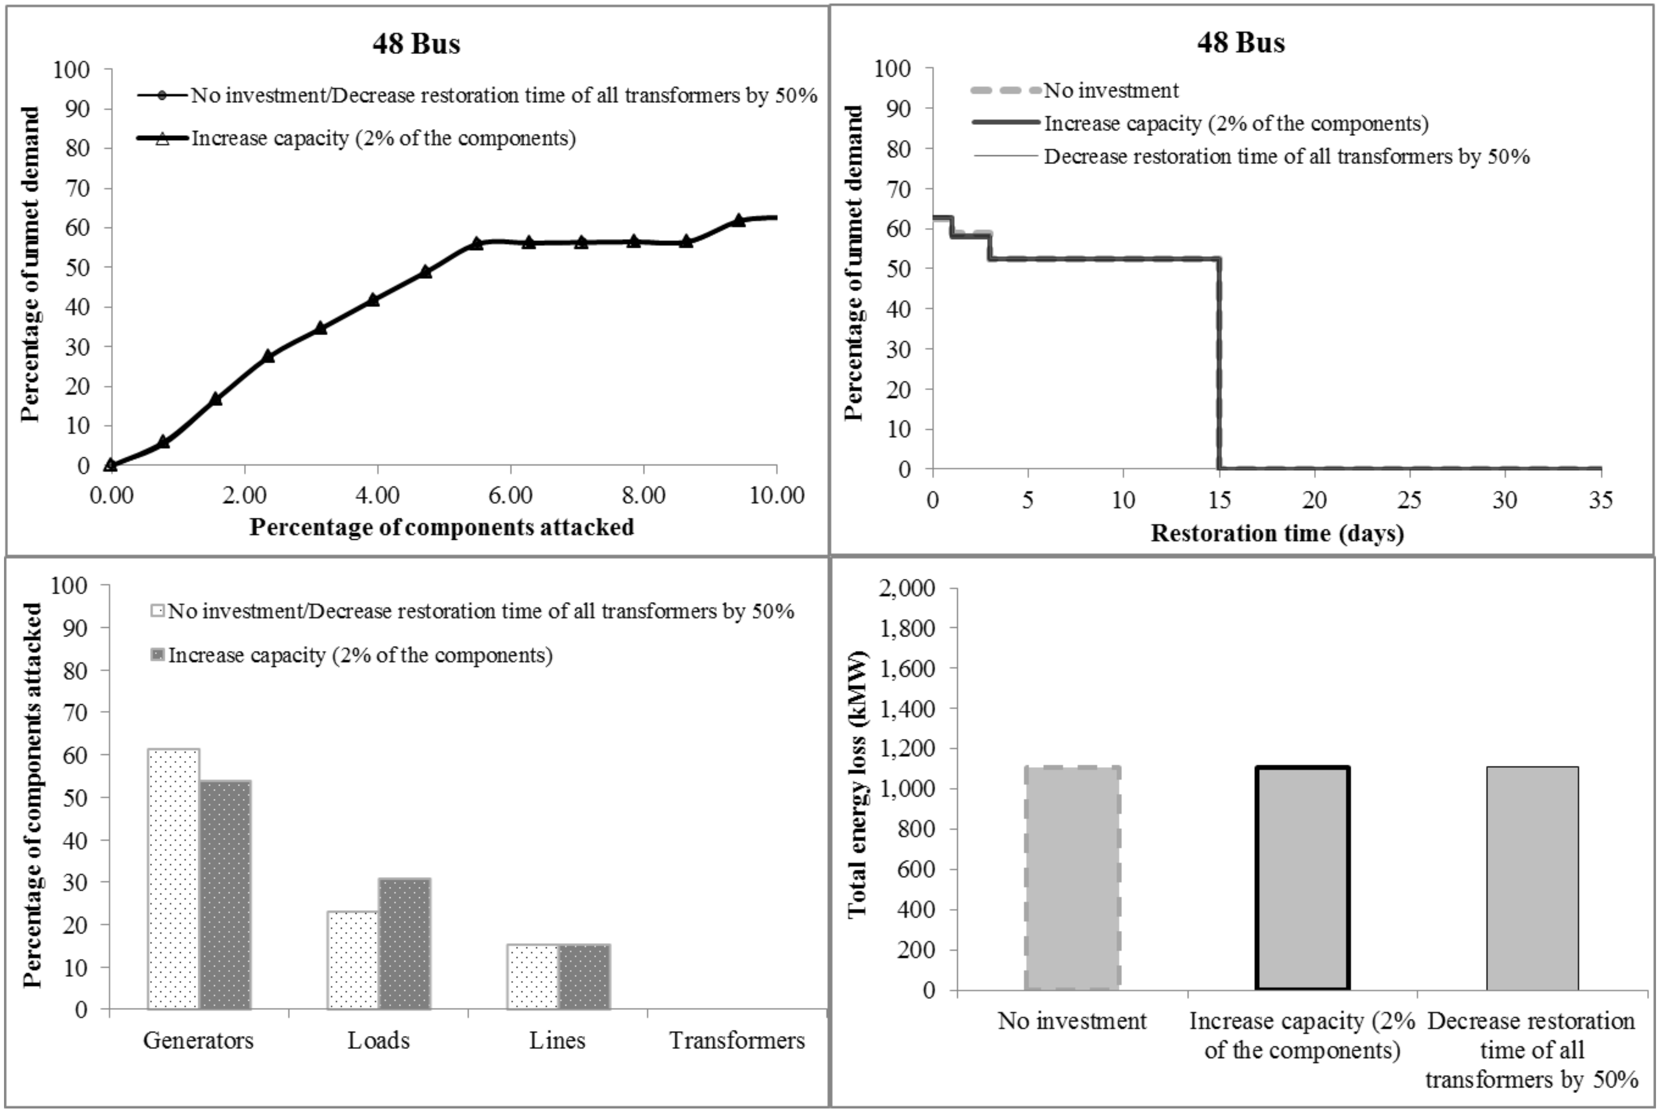

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a dynamic attacker, 118-bus system)

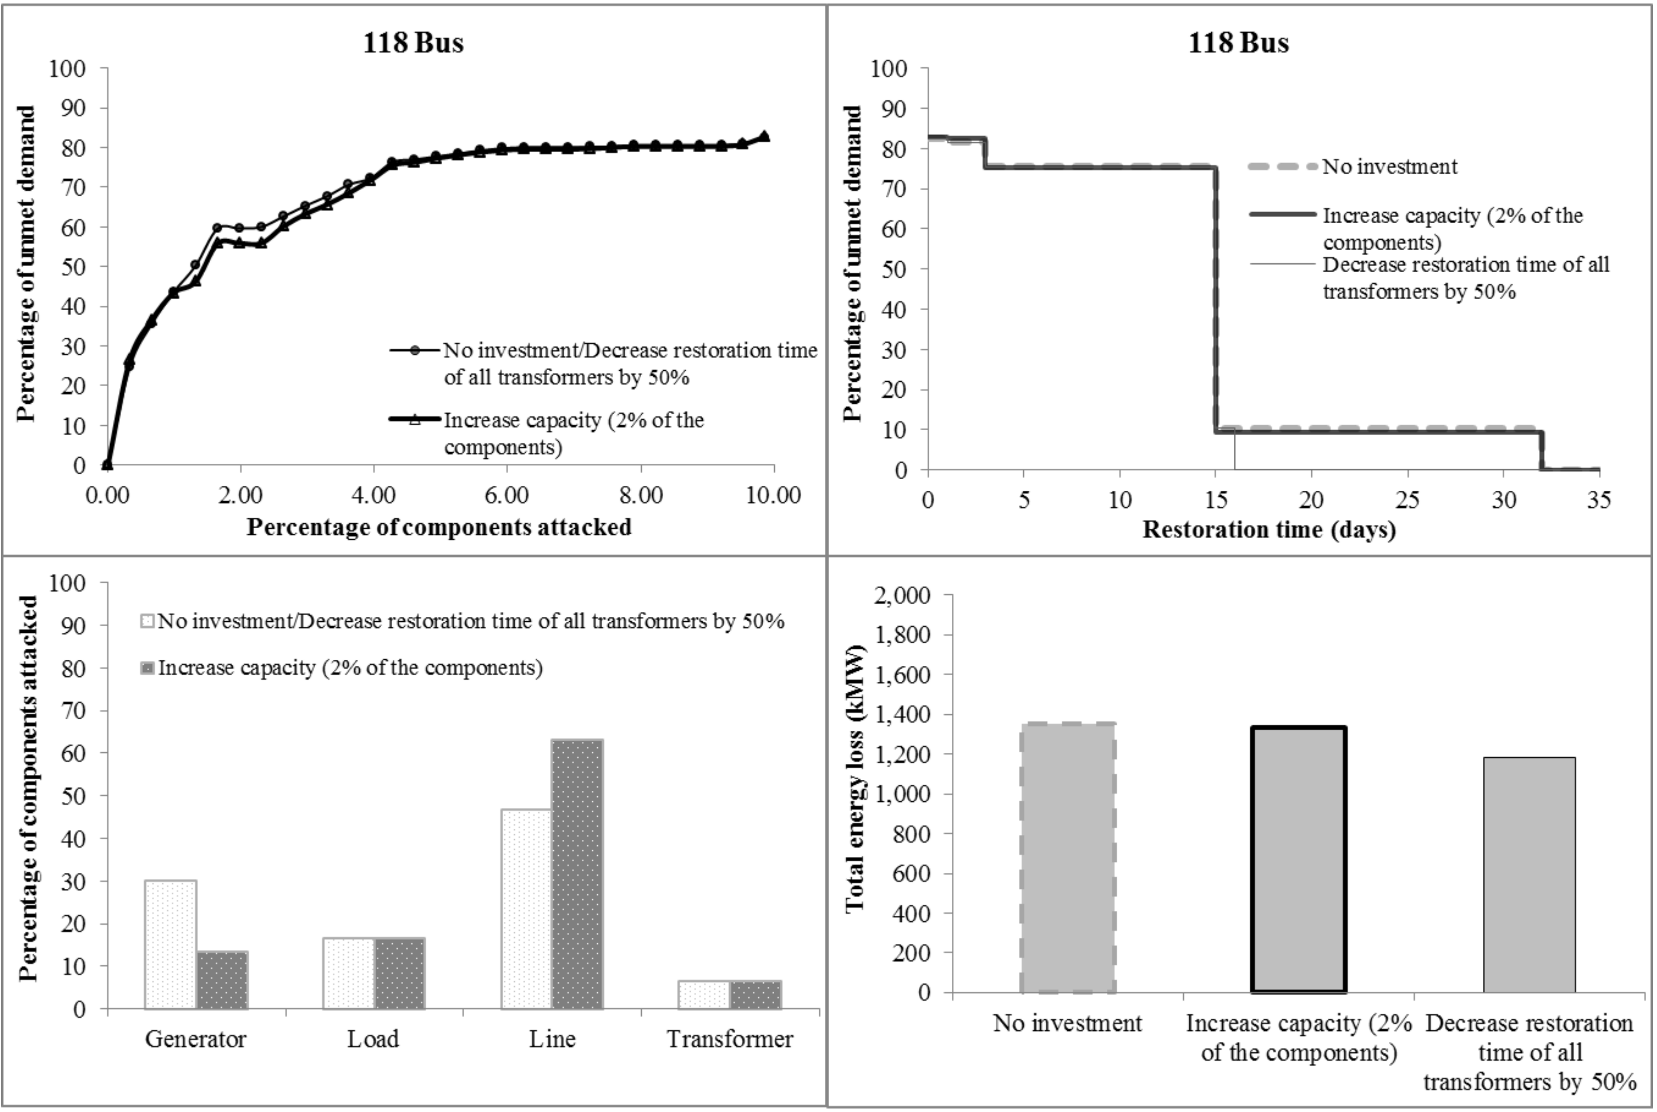

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a dynamic attacker, 300-bus system)

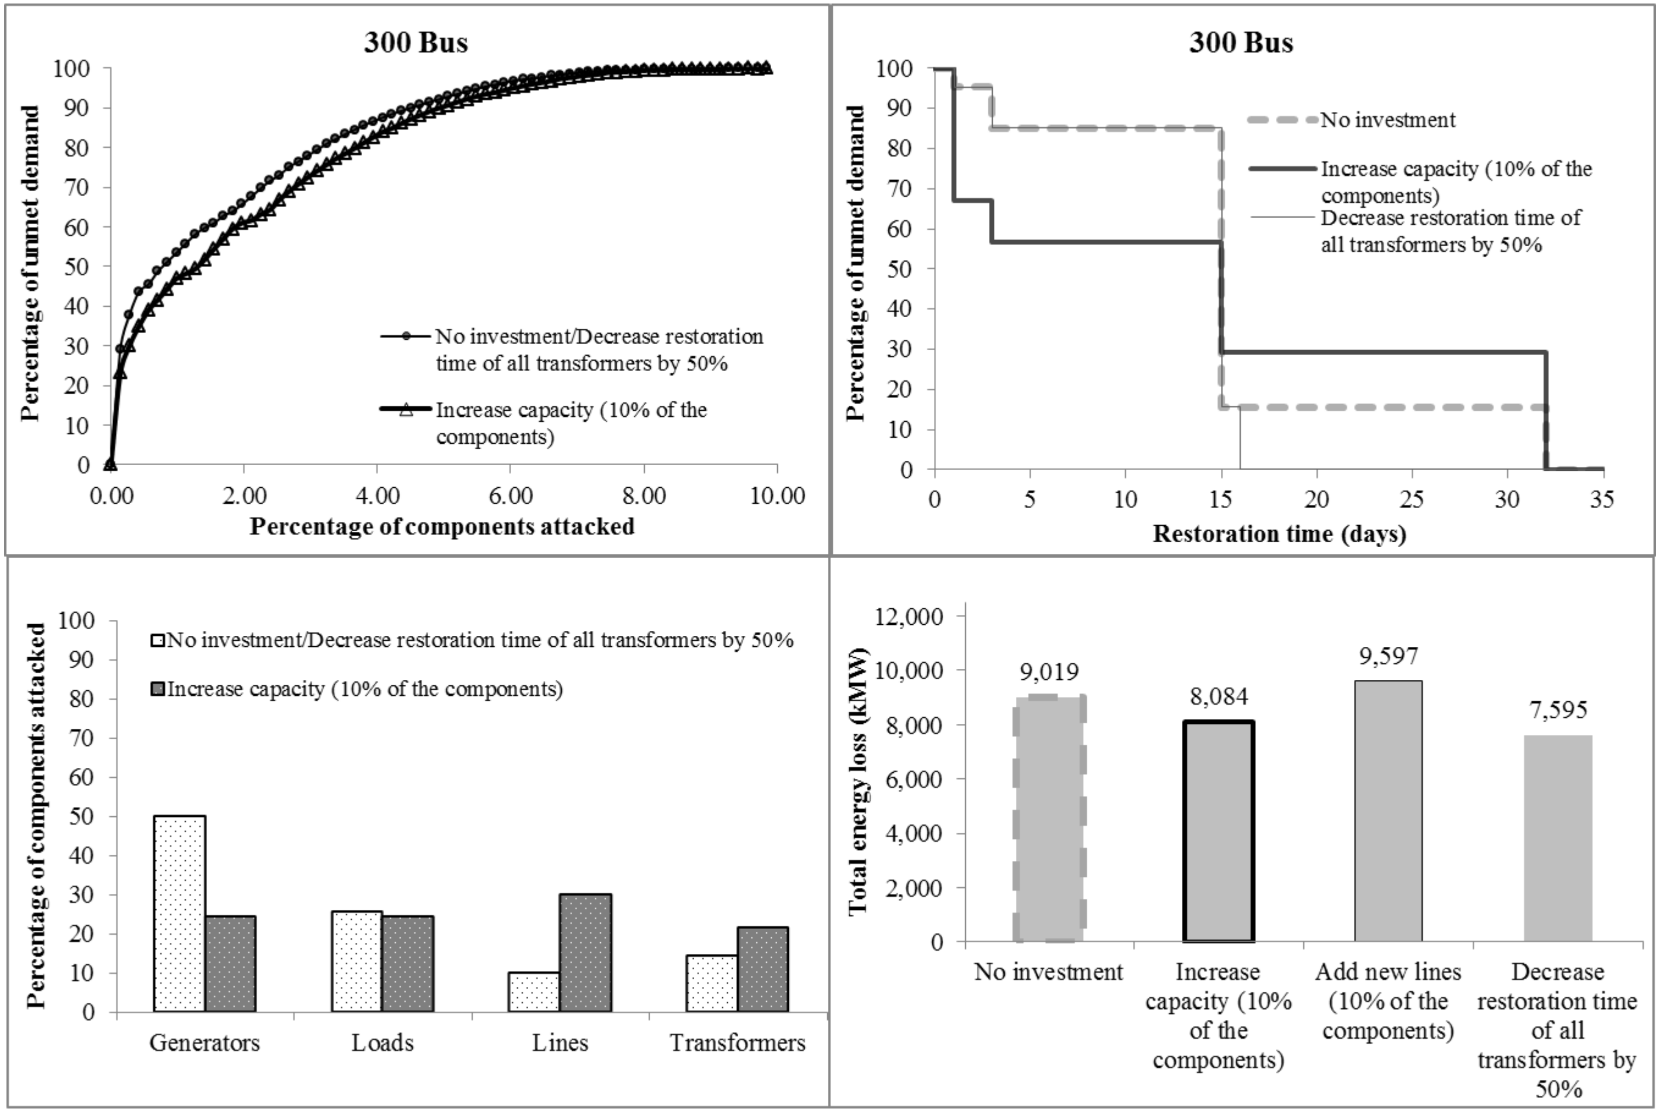

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 24-bus system)

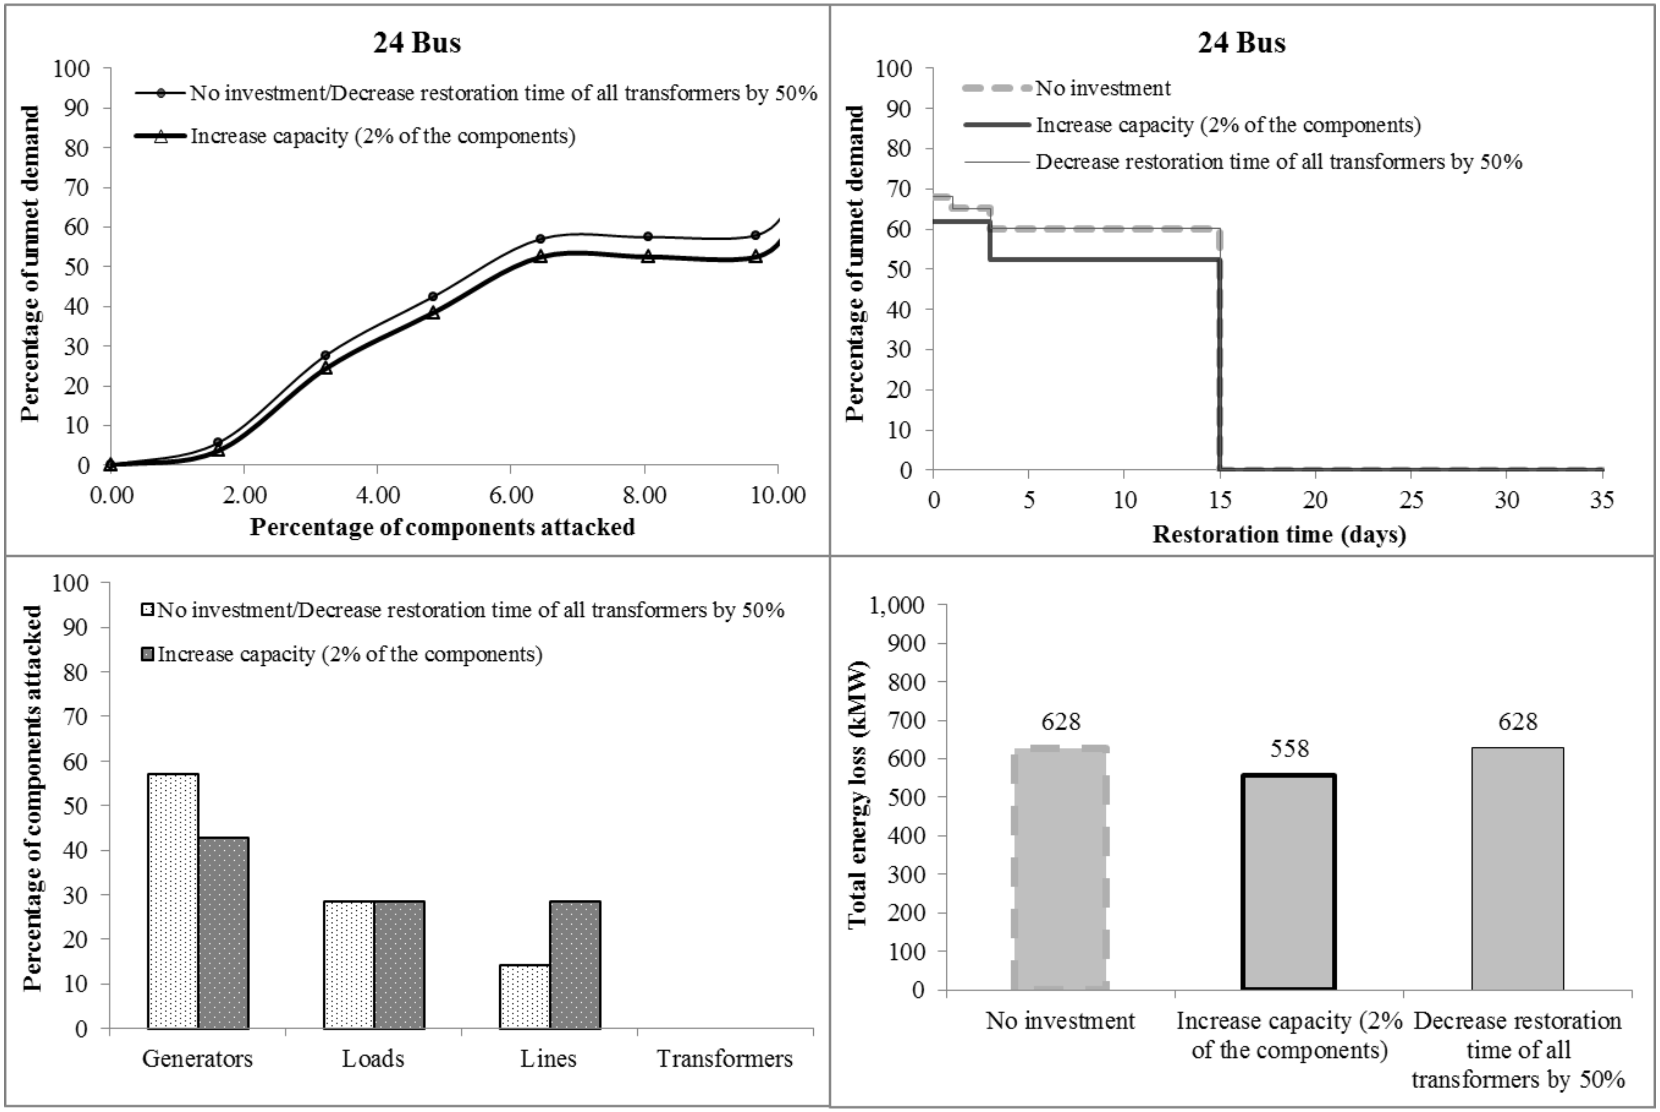

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 48-bus system)

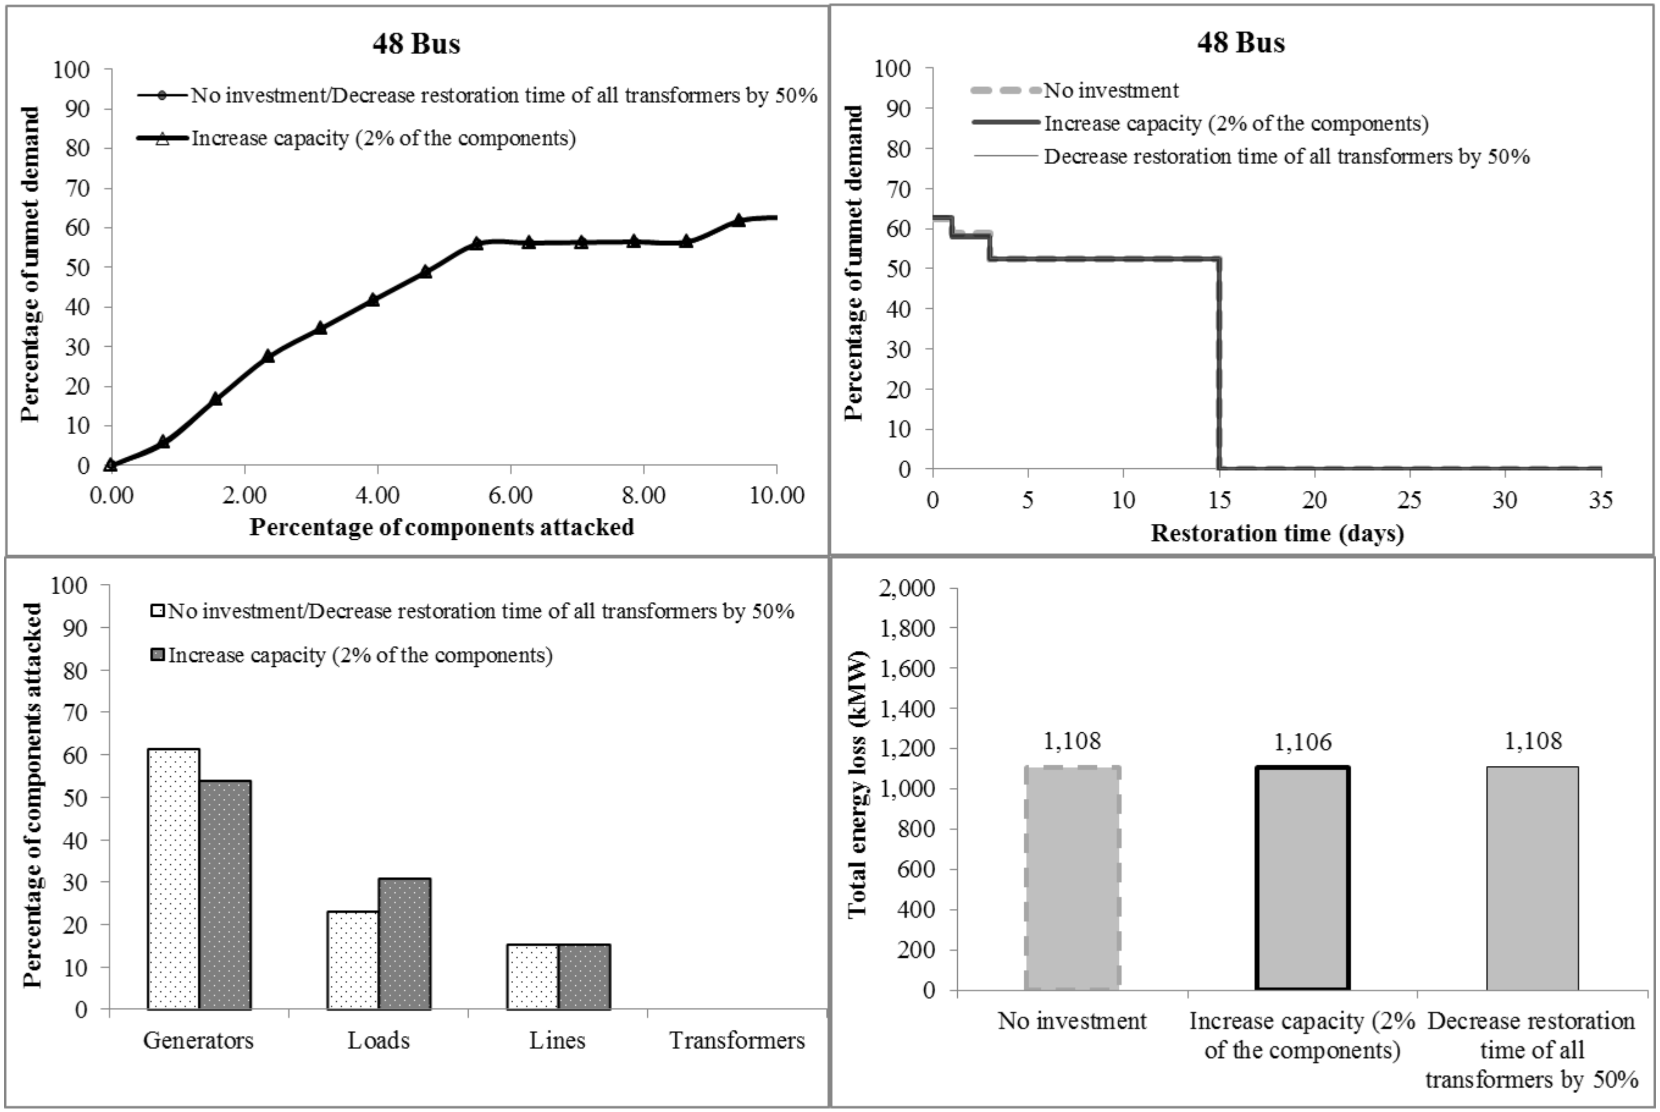

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 118-bus system)

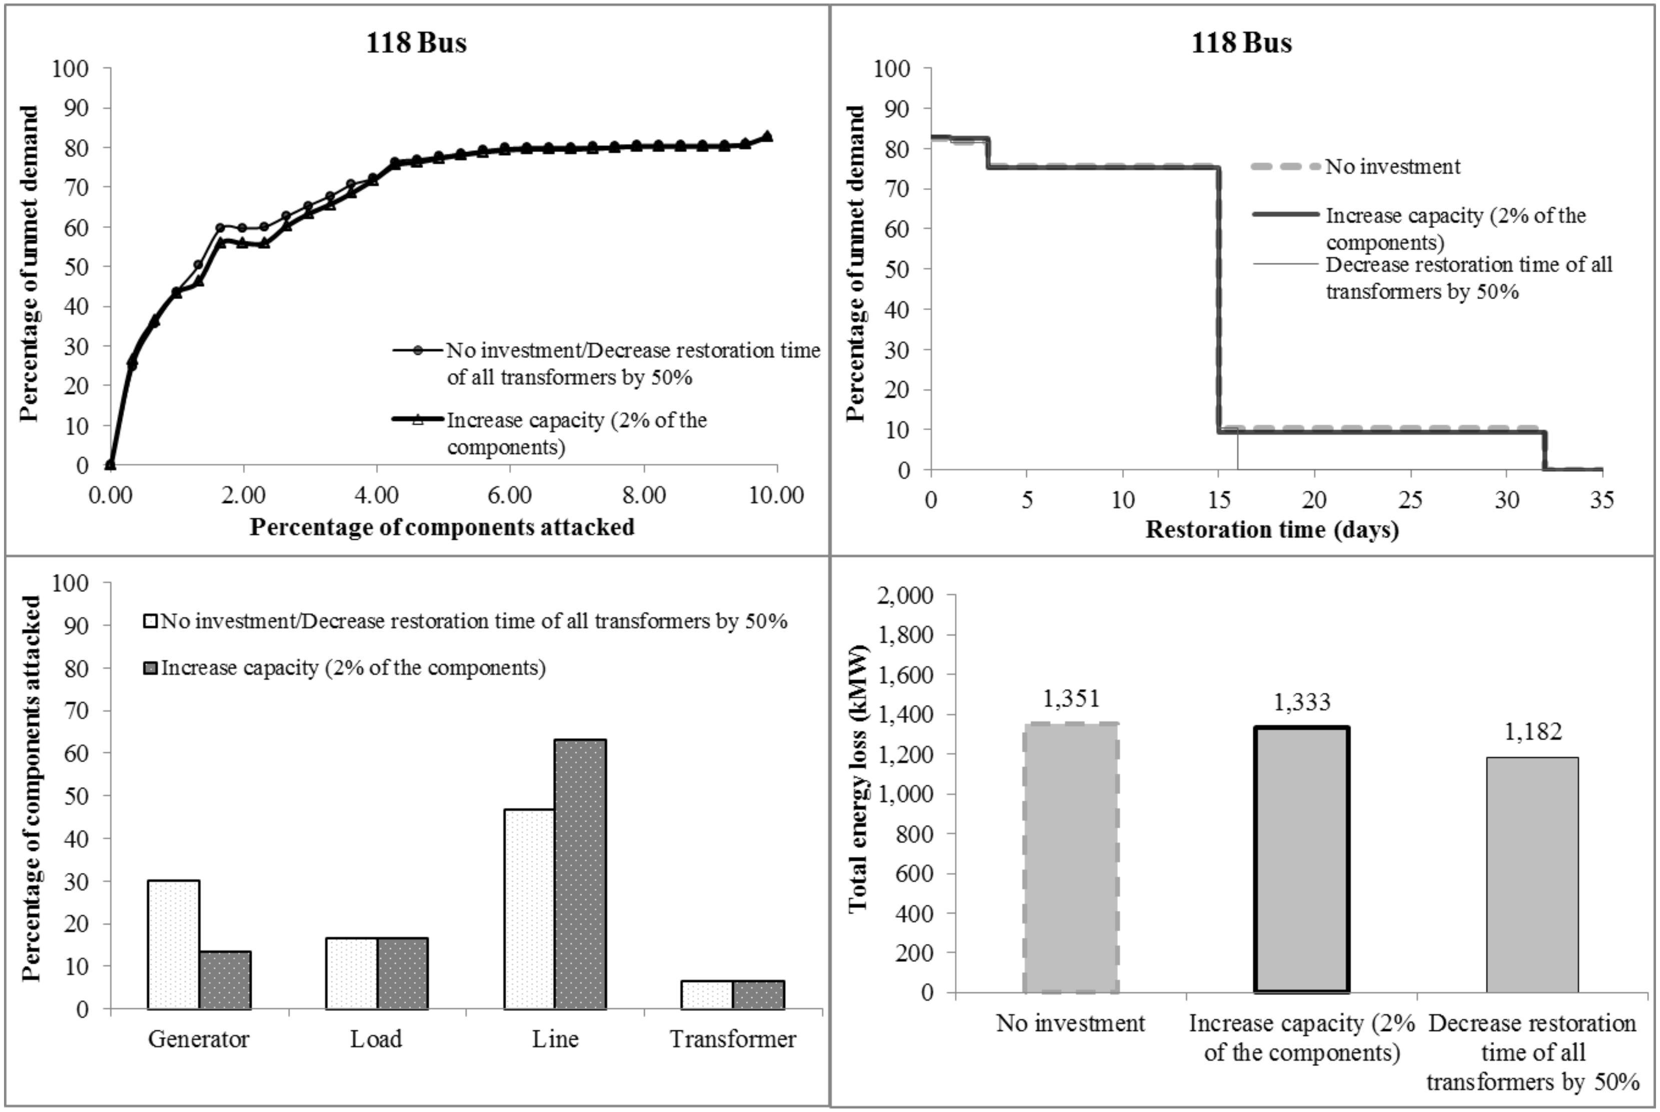

Appendix N. Increasing capacity versus decreasing restoration times of transformers
(increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 300-bus system)

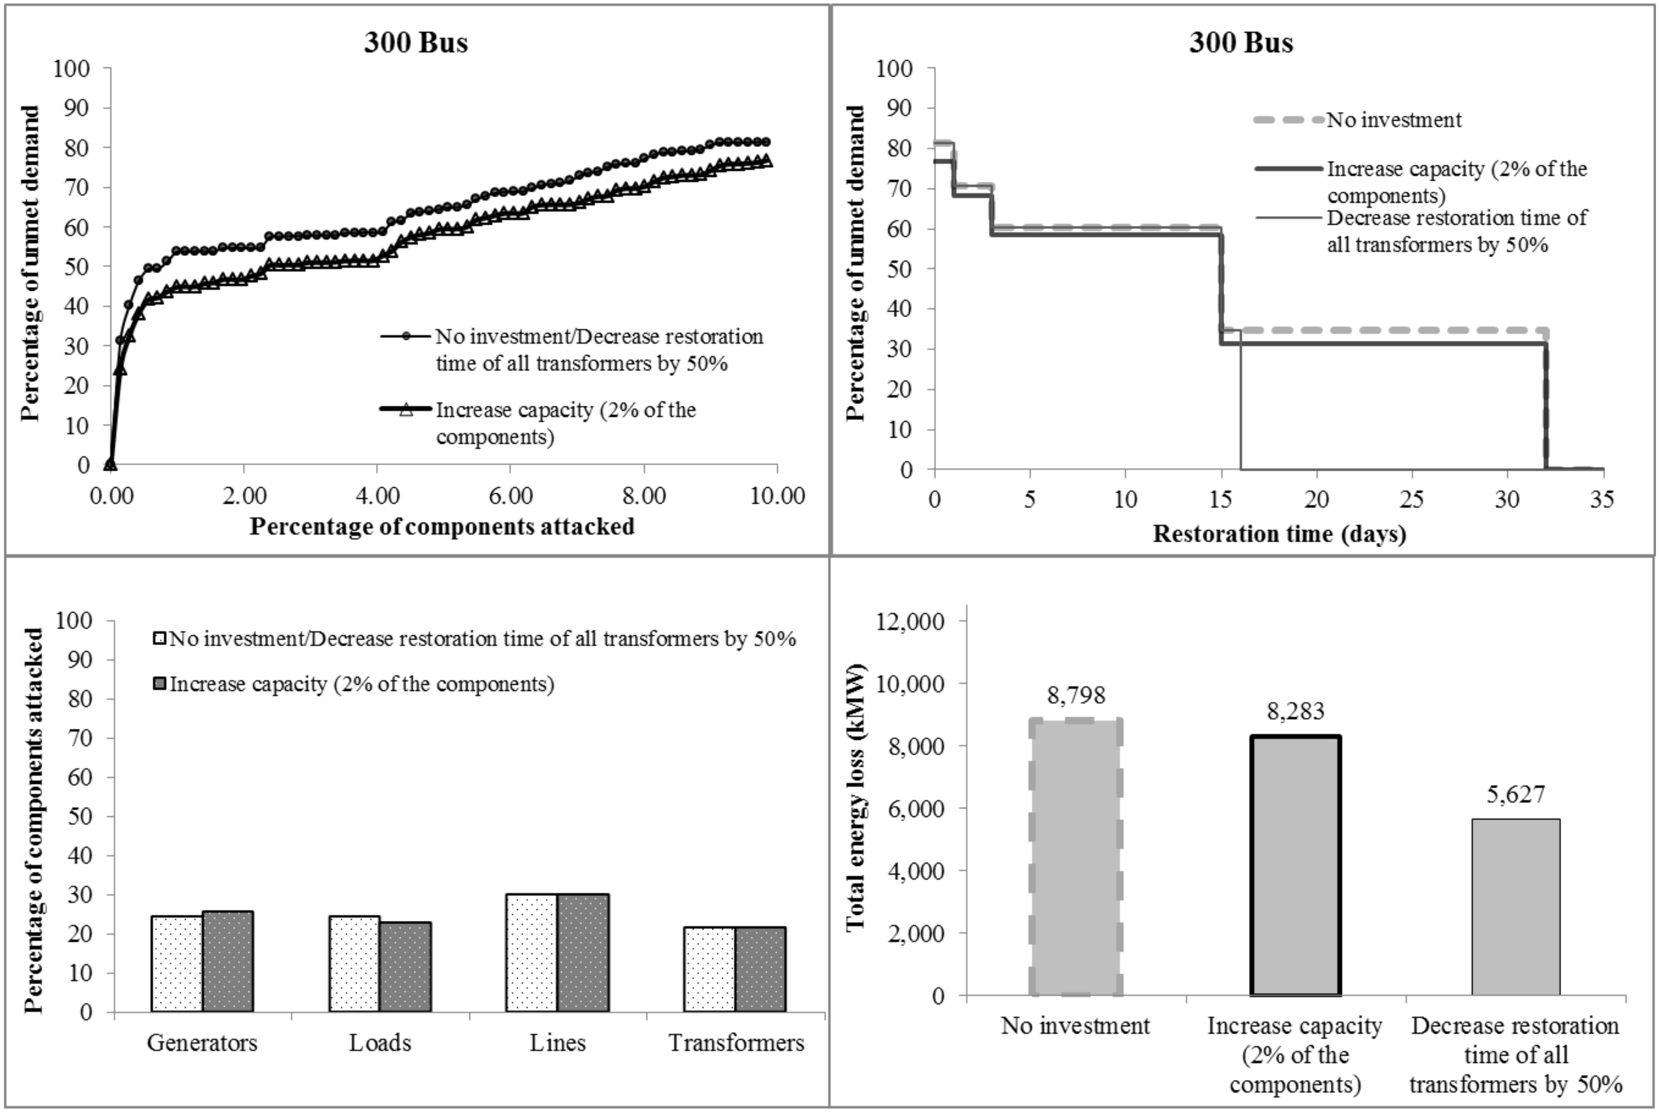

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 24-bus system)

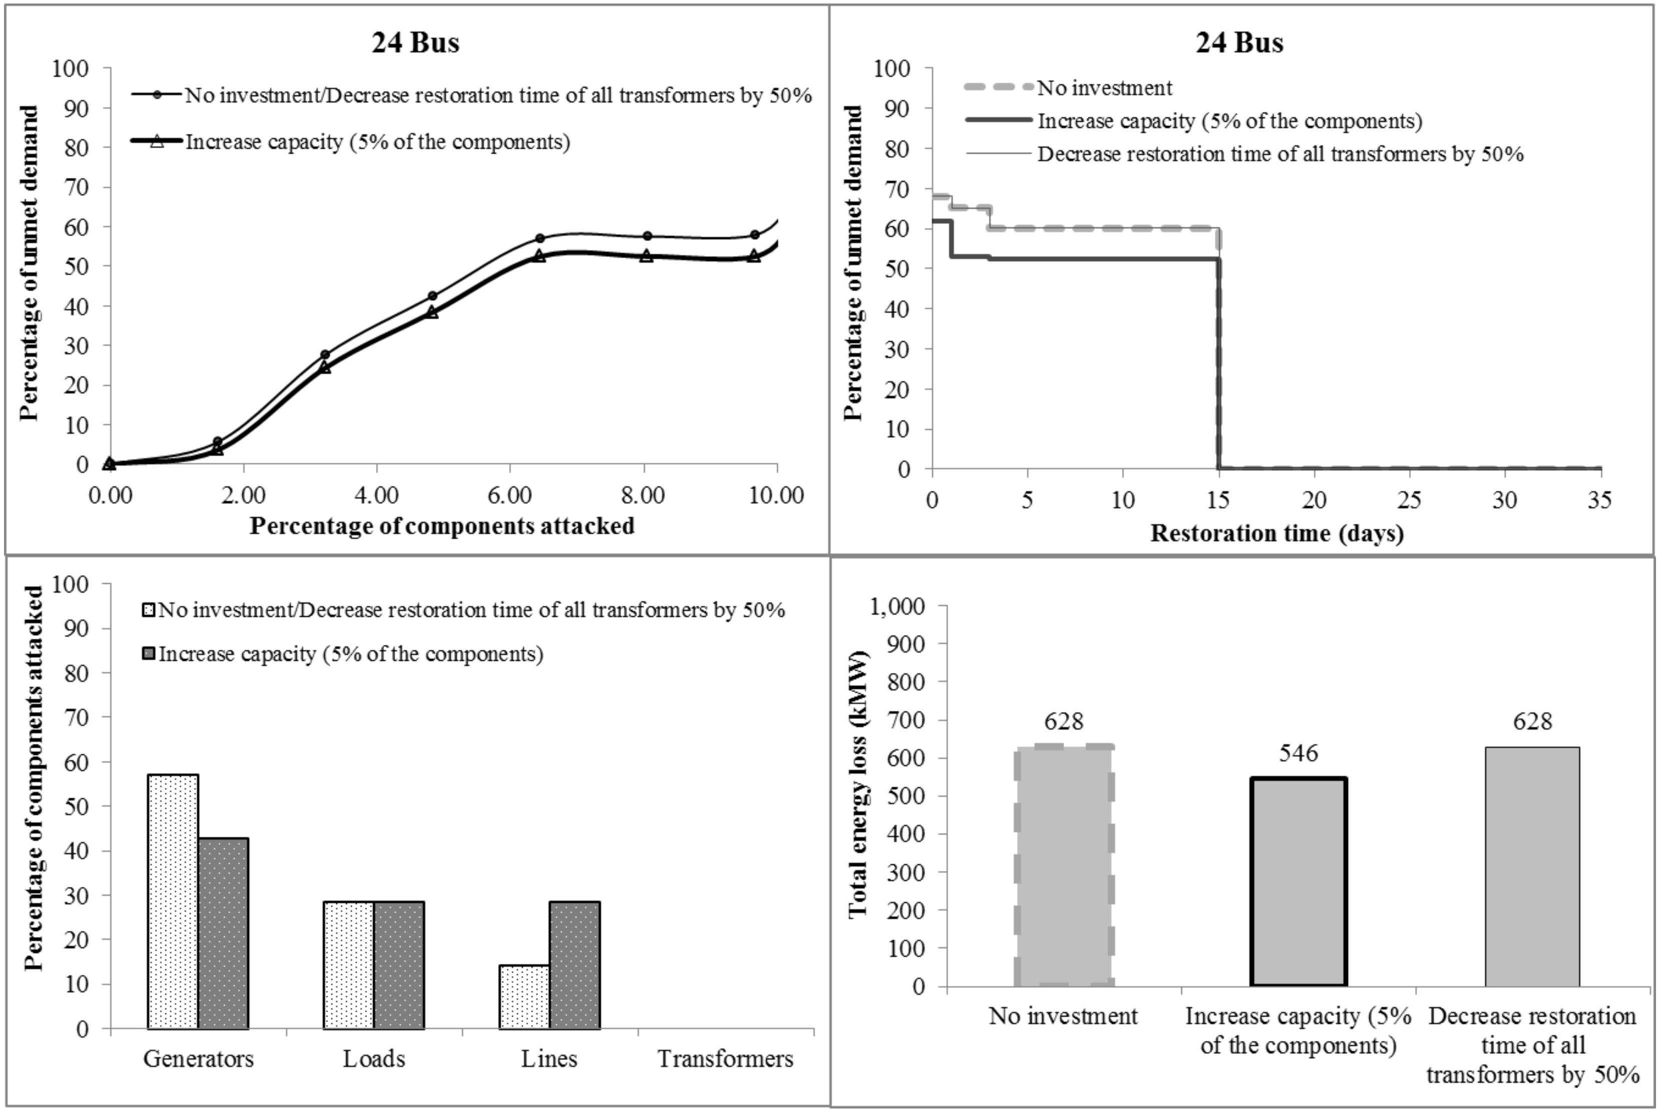

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 48-bus system)

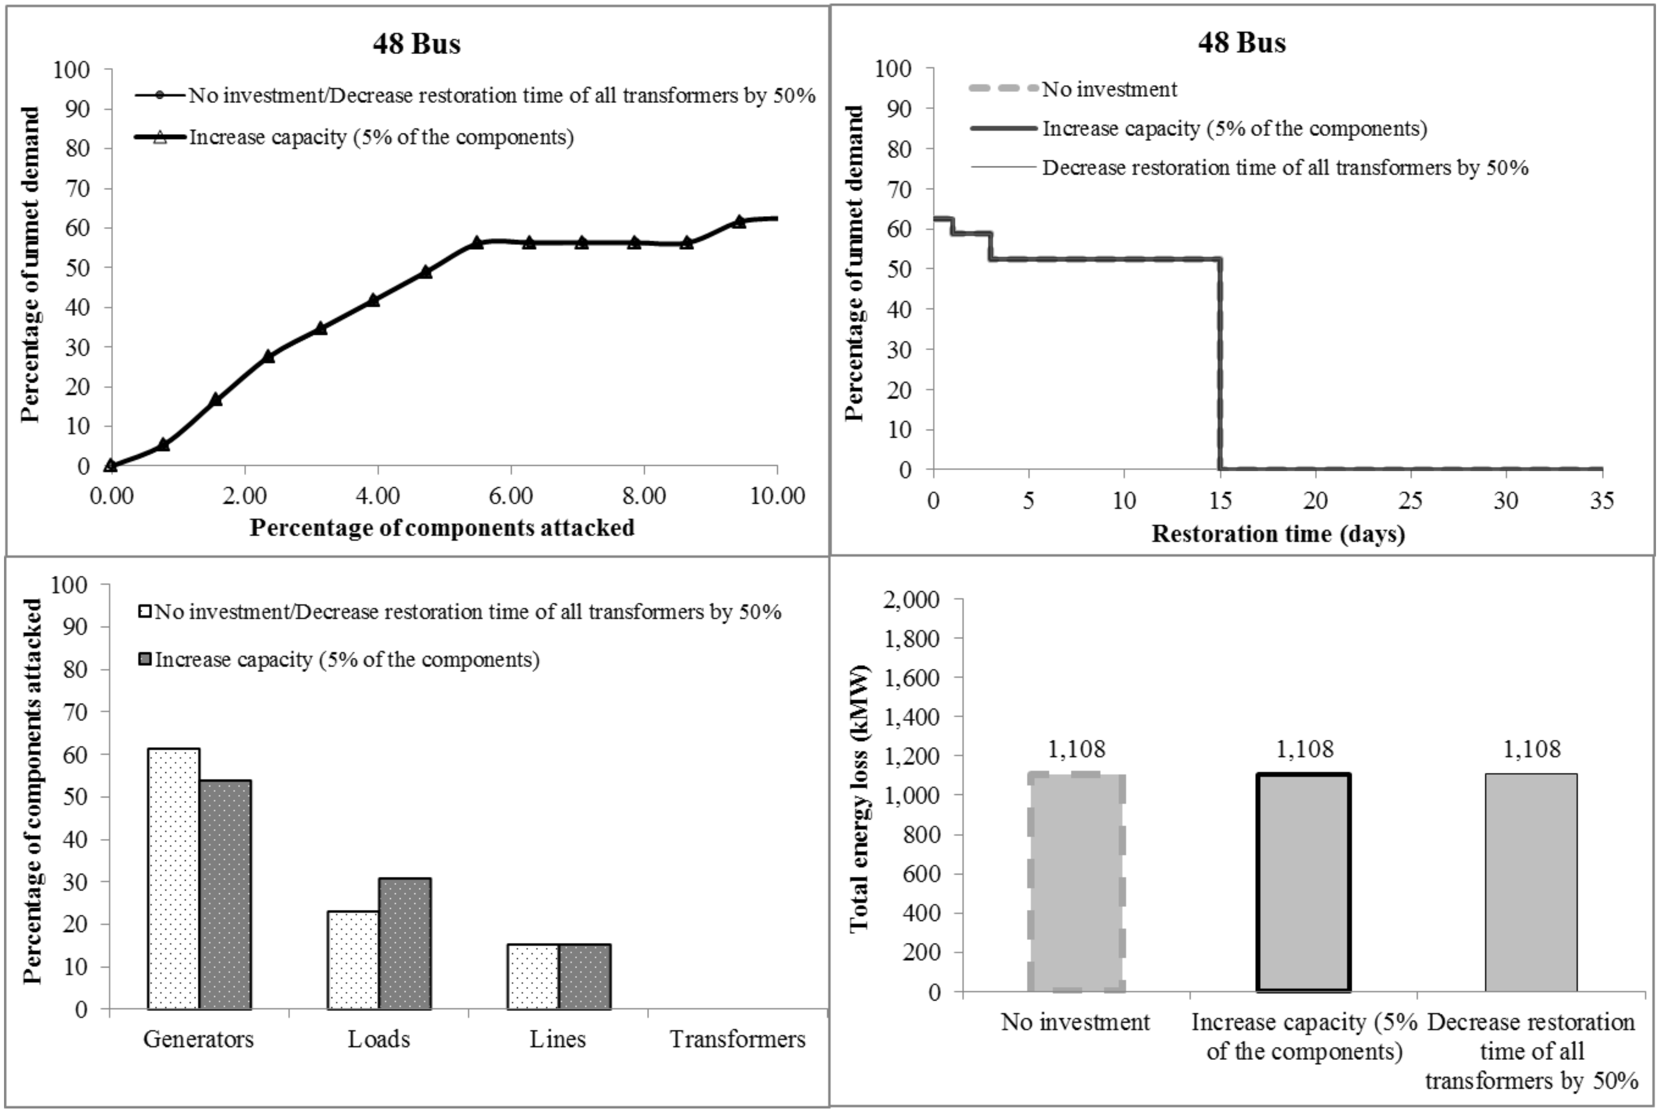

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 118-bus system)

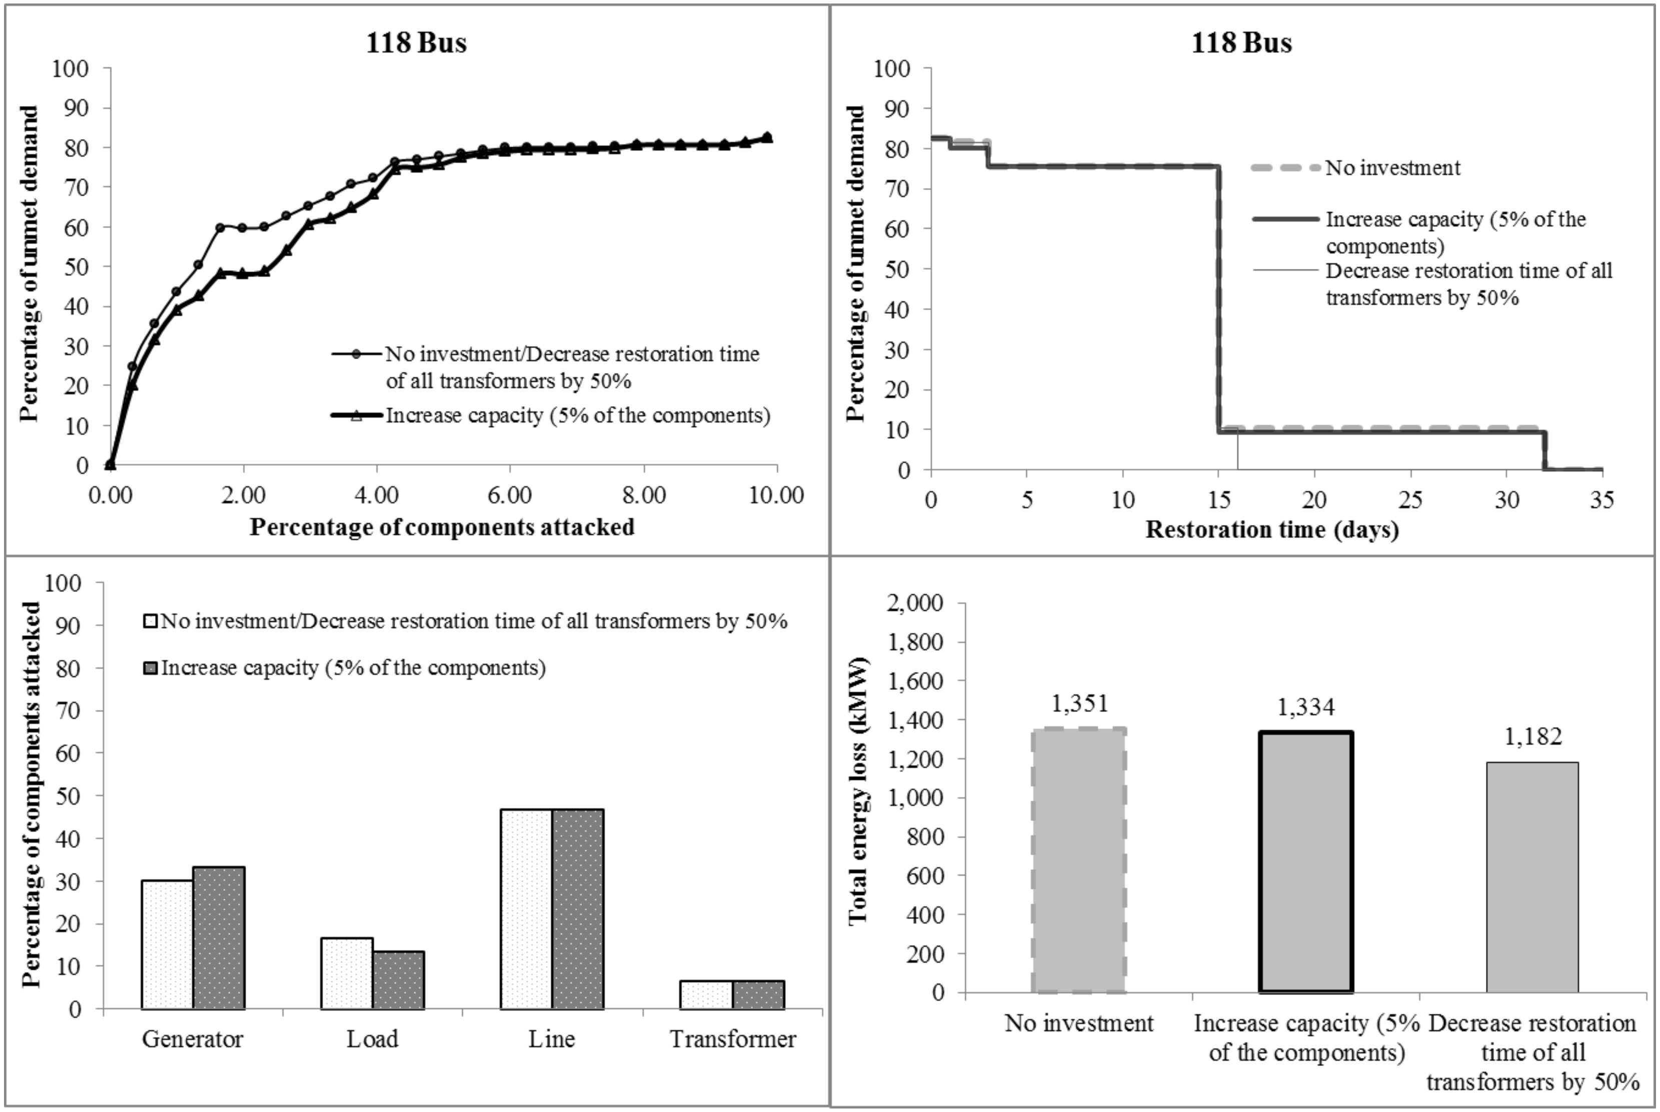

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 300-bus system)

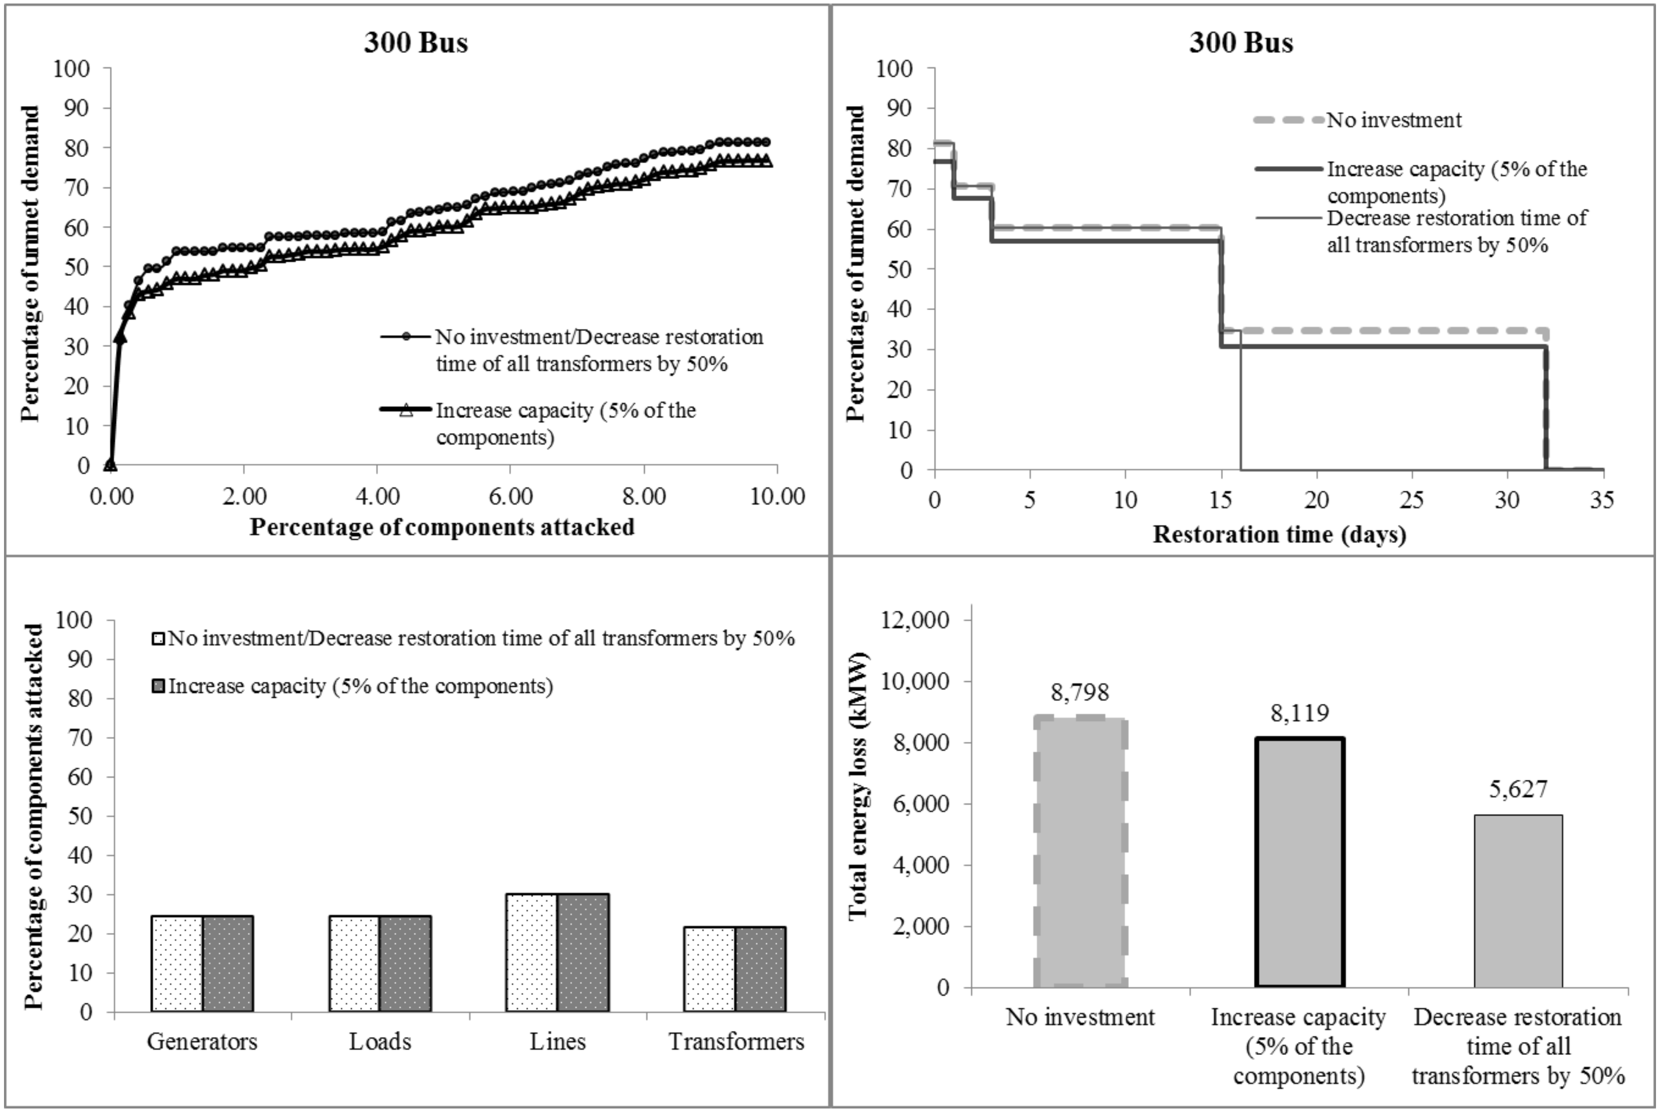

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 24-bus system)

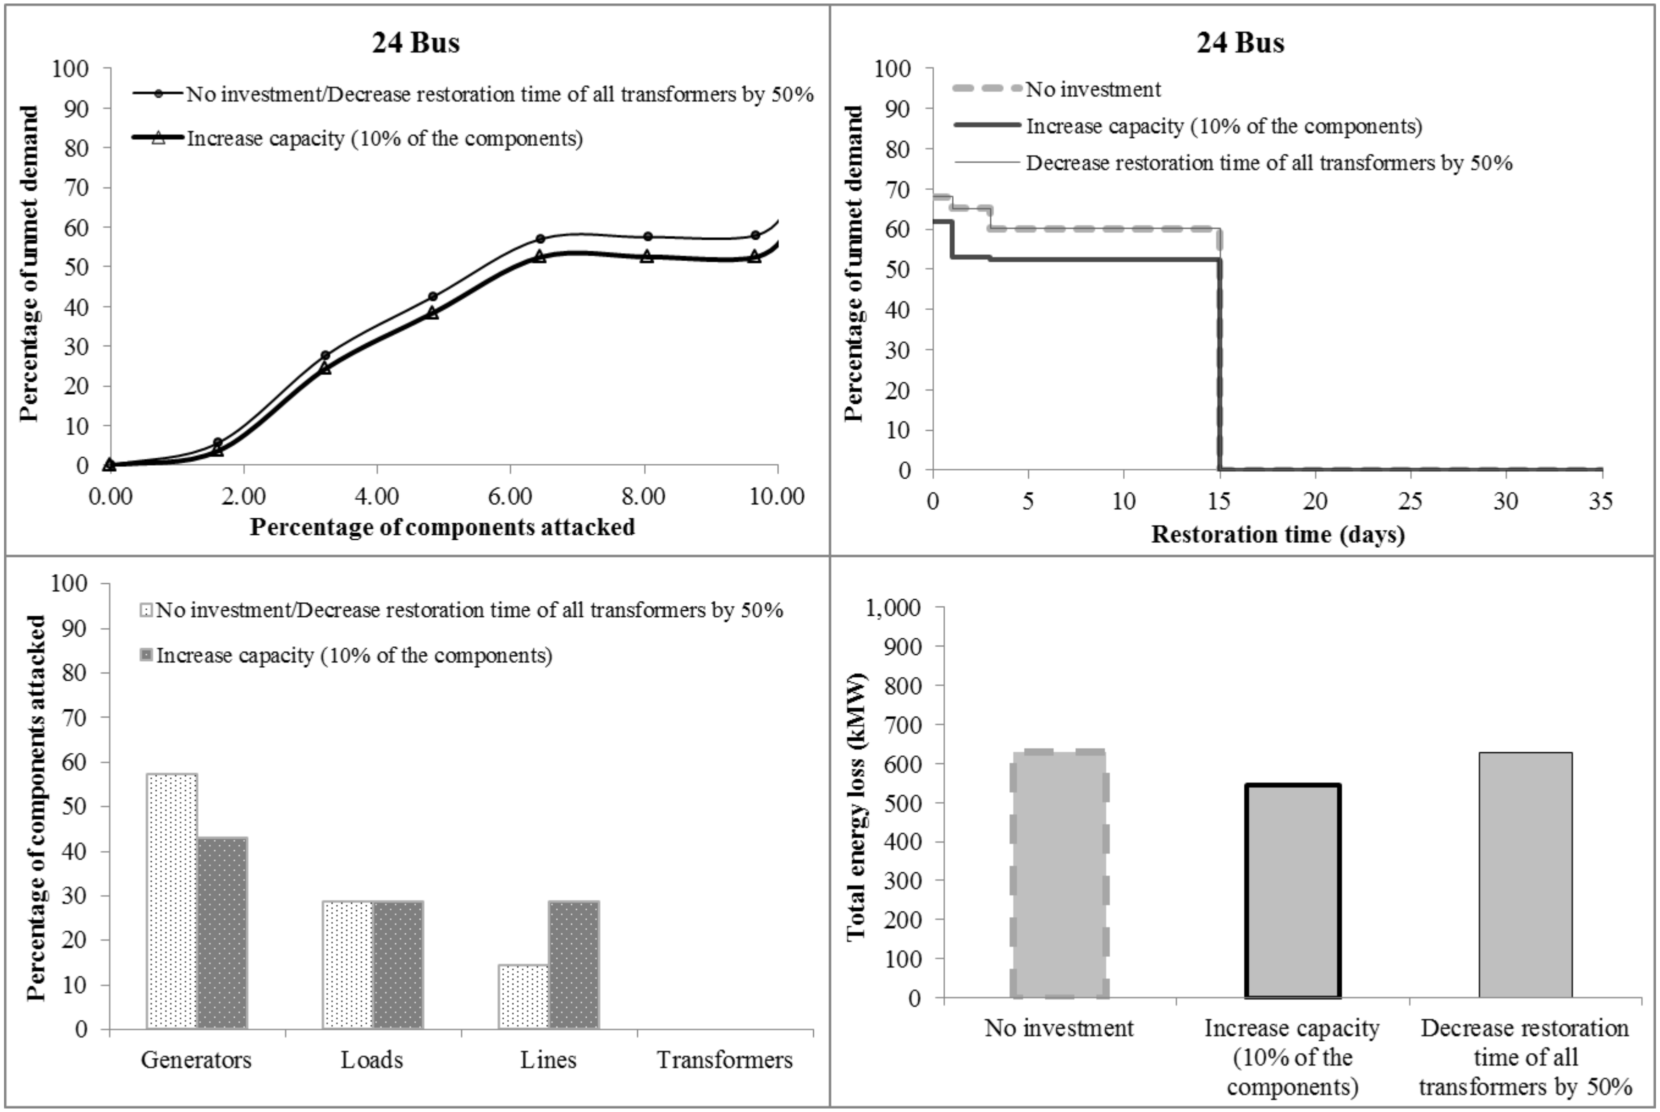

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 48-bus system)

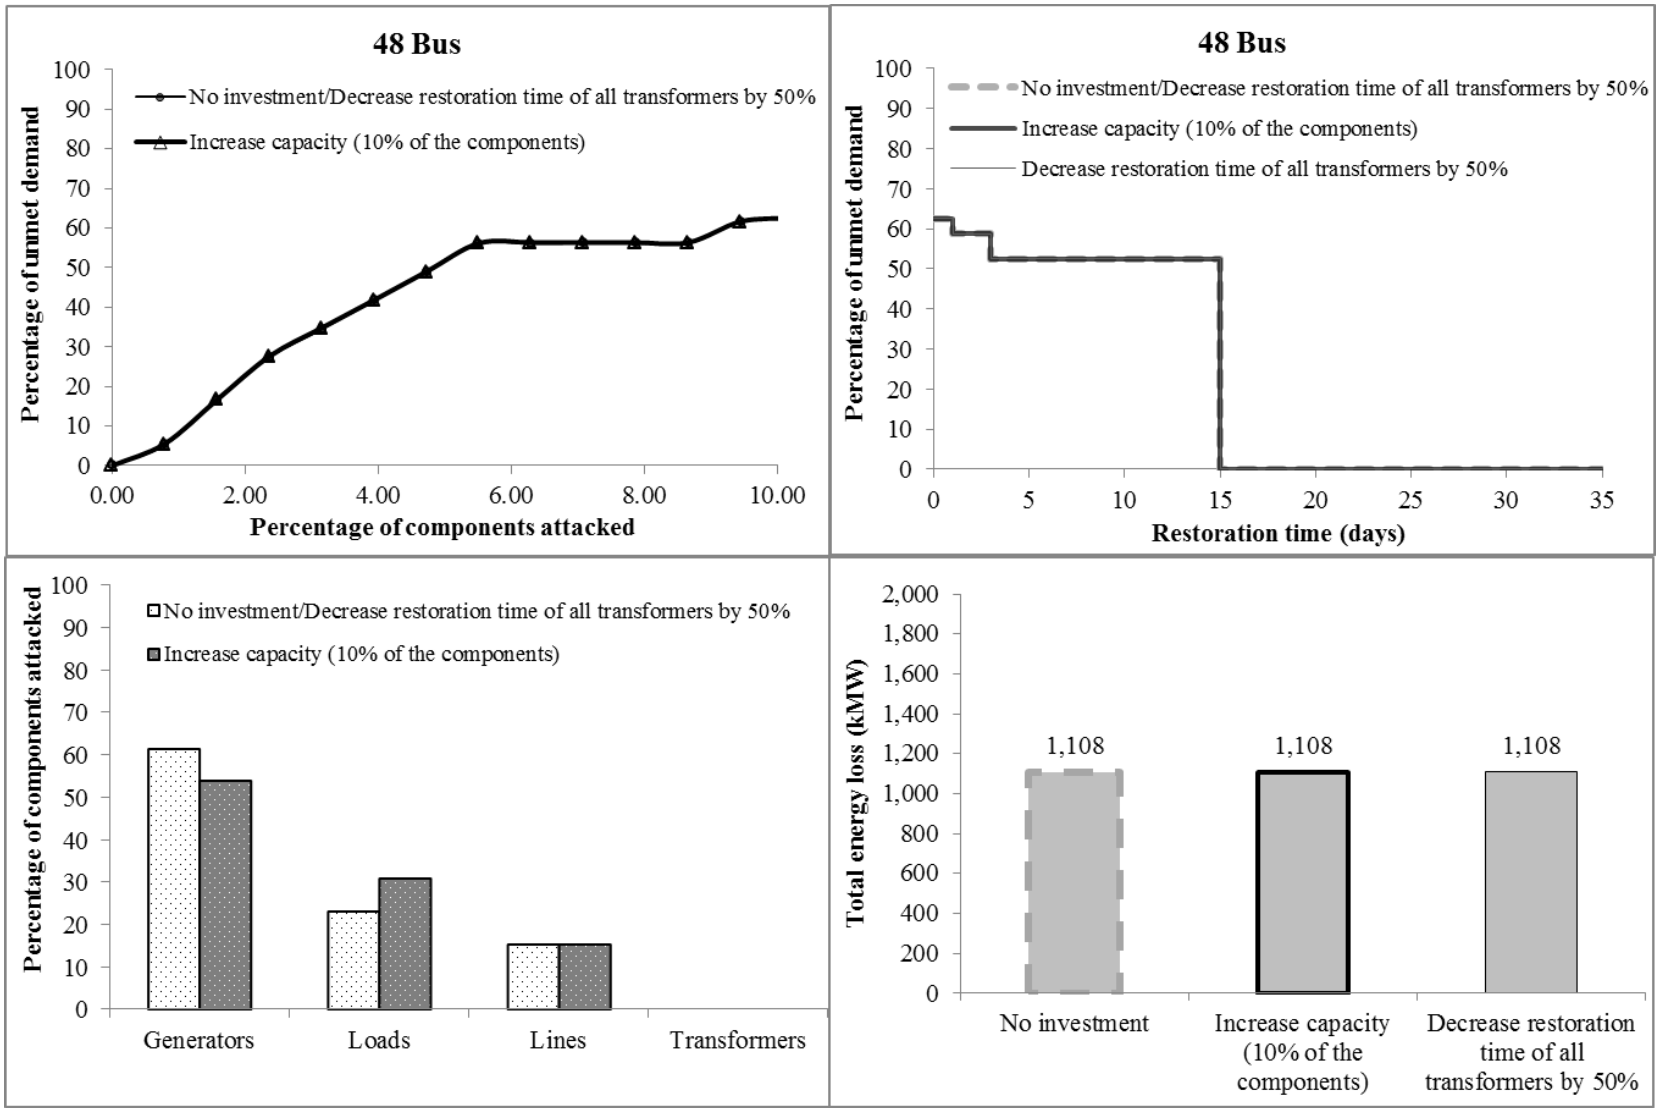

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 118-bus system)

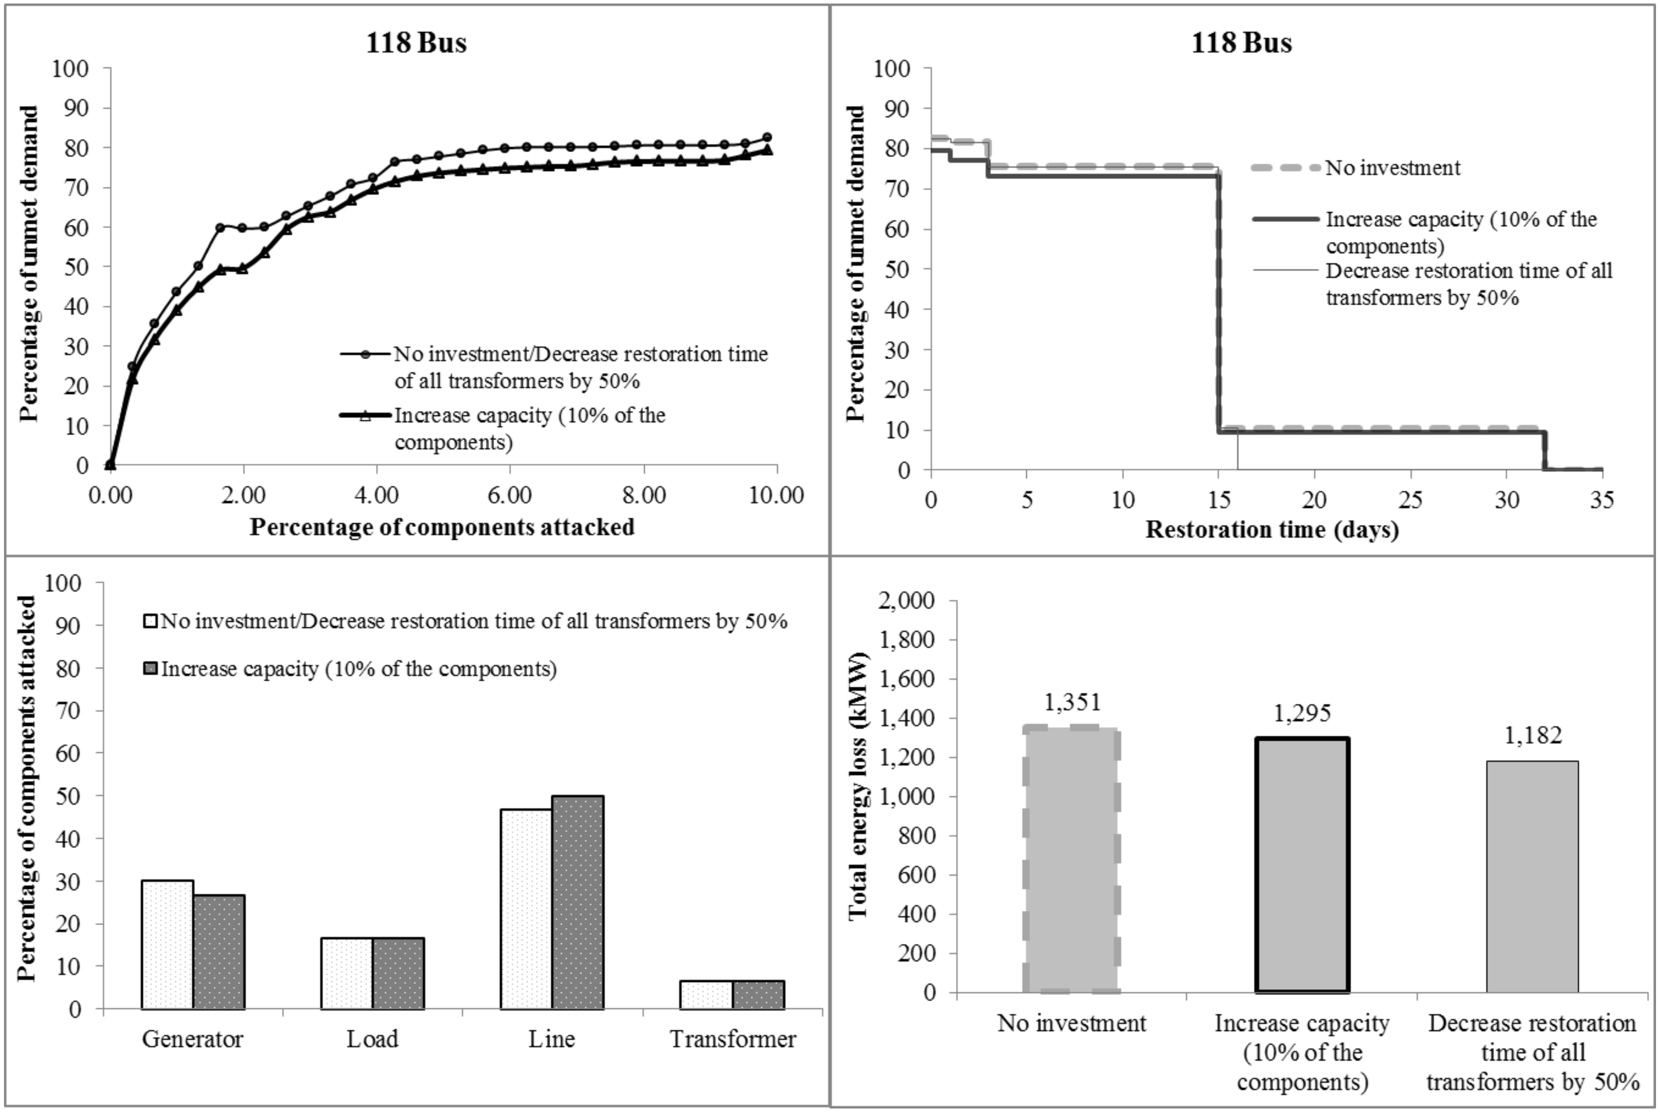

Appendix N. Increasing capacity versus decreasing restoration times of transformers
(increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 300-bus system)

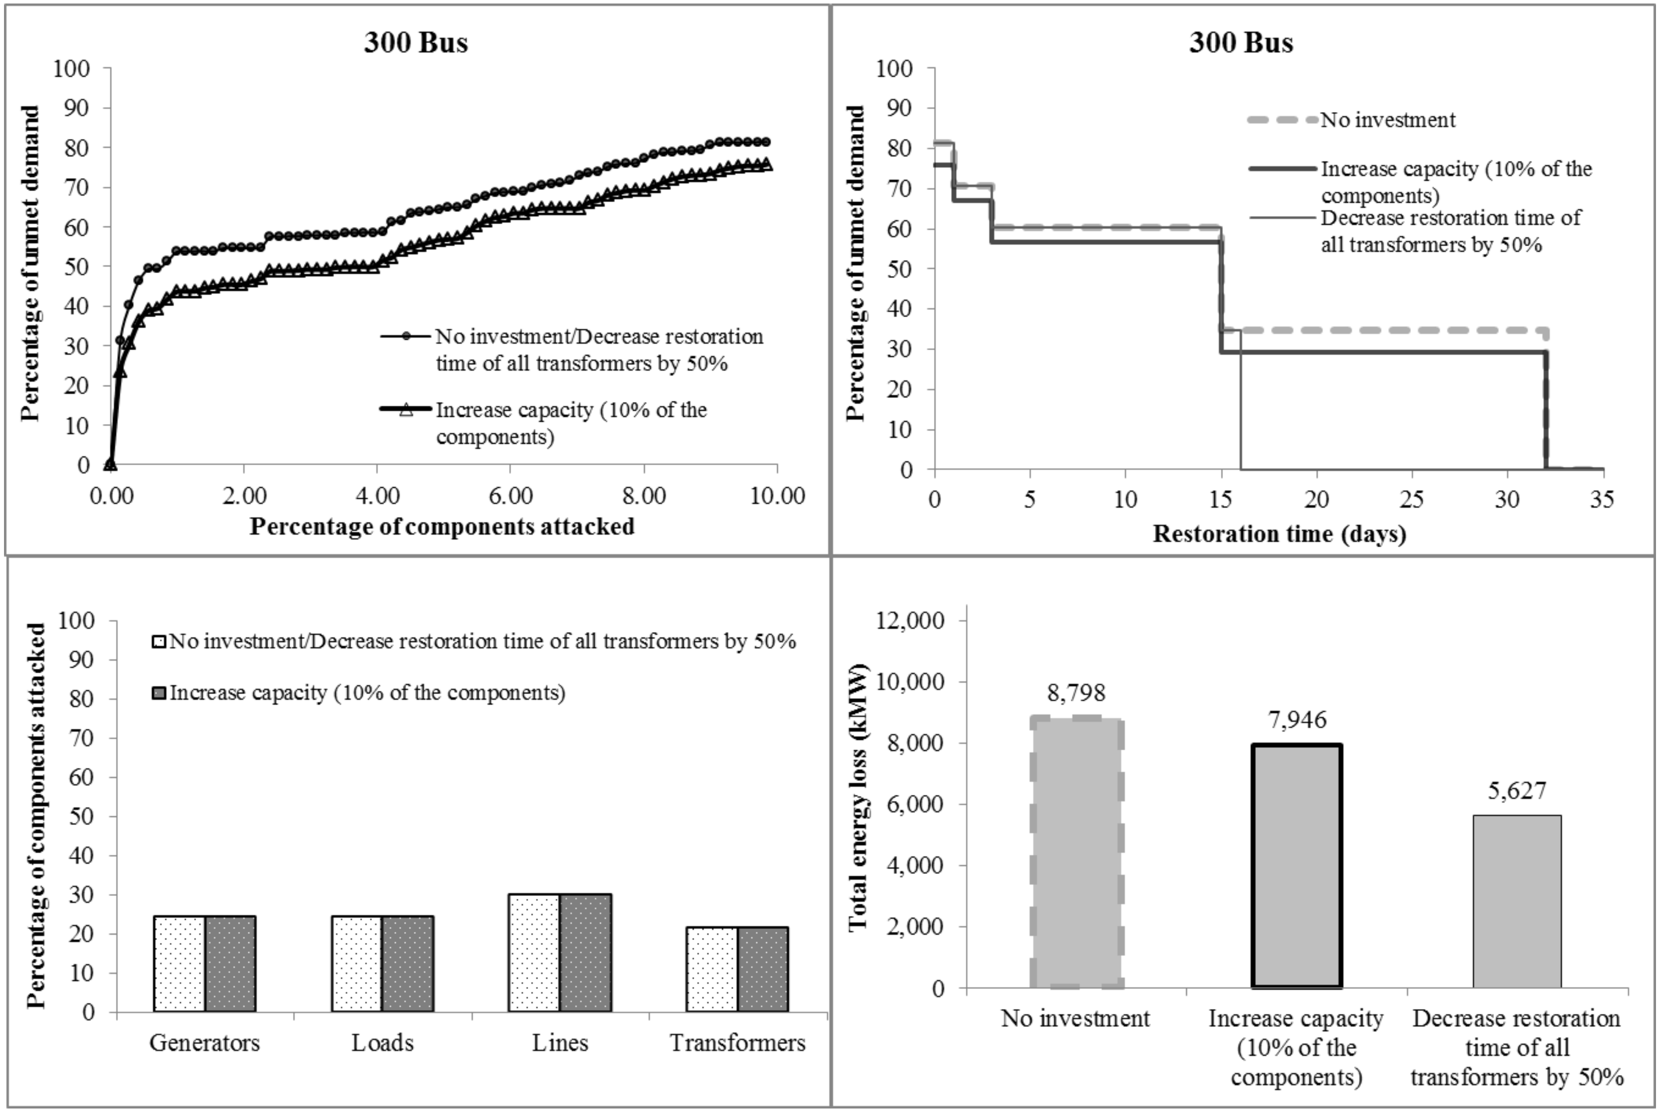

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 24-bus system)

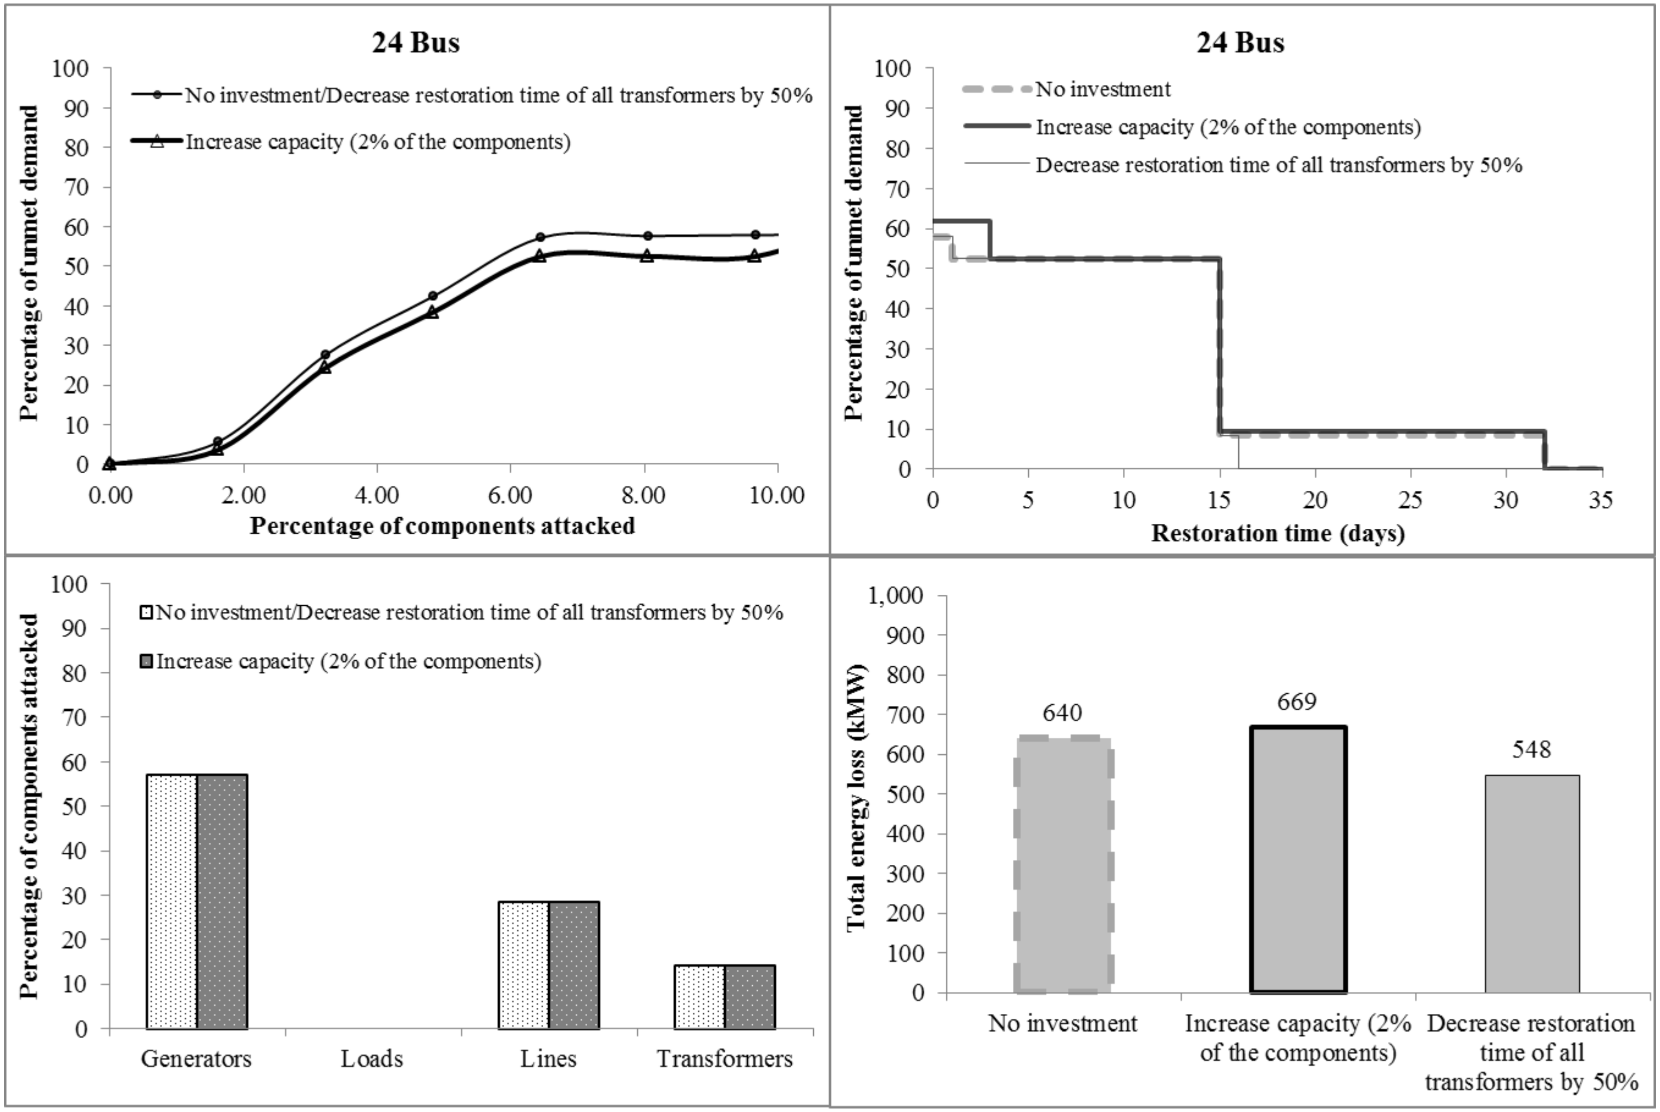

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 48-bus system)

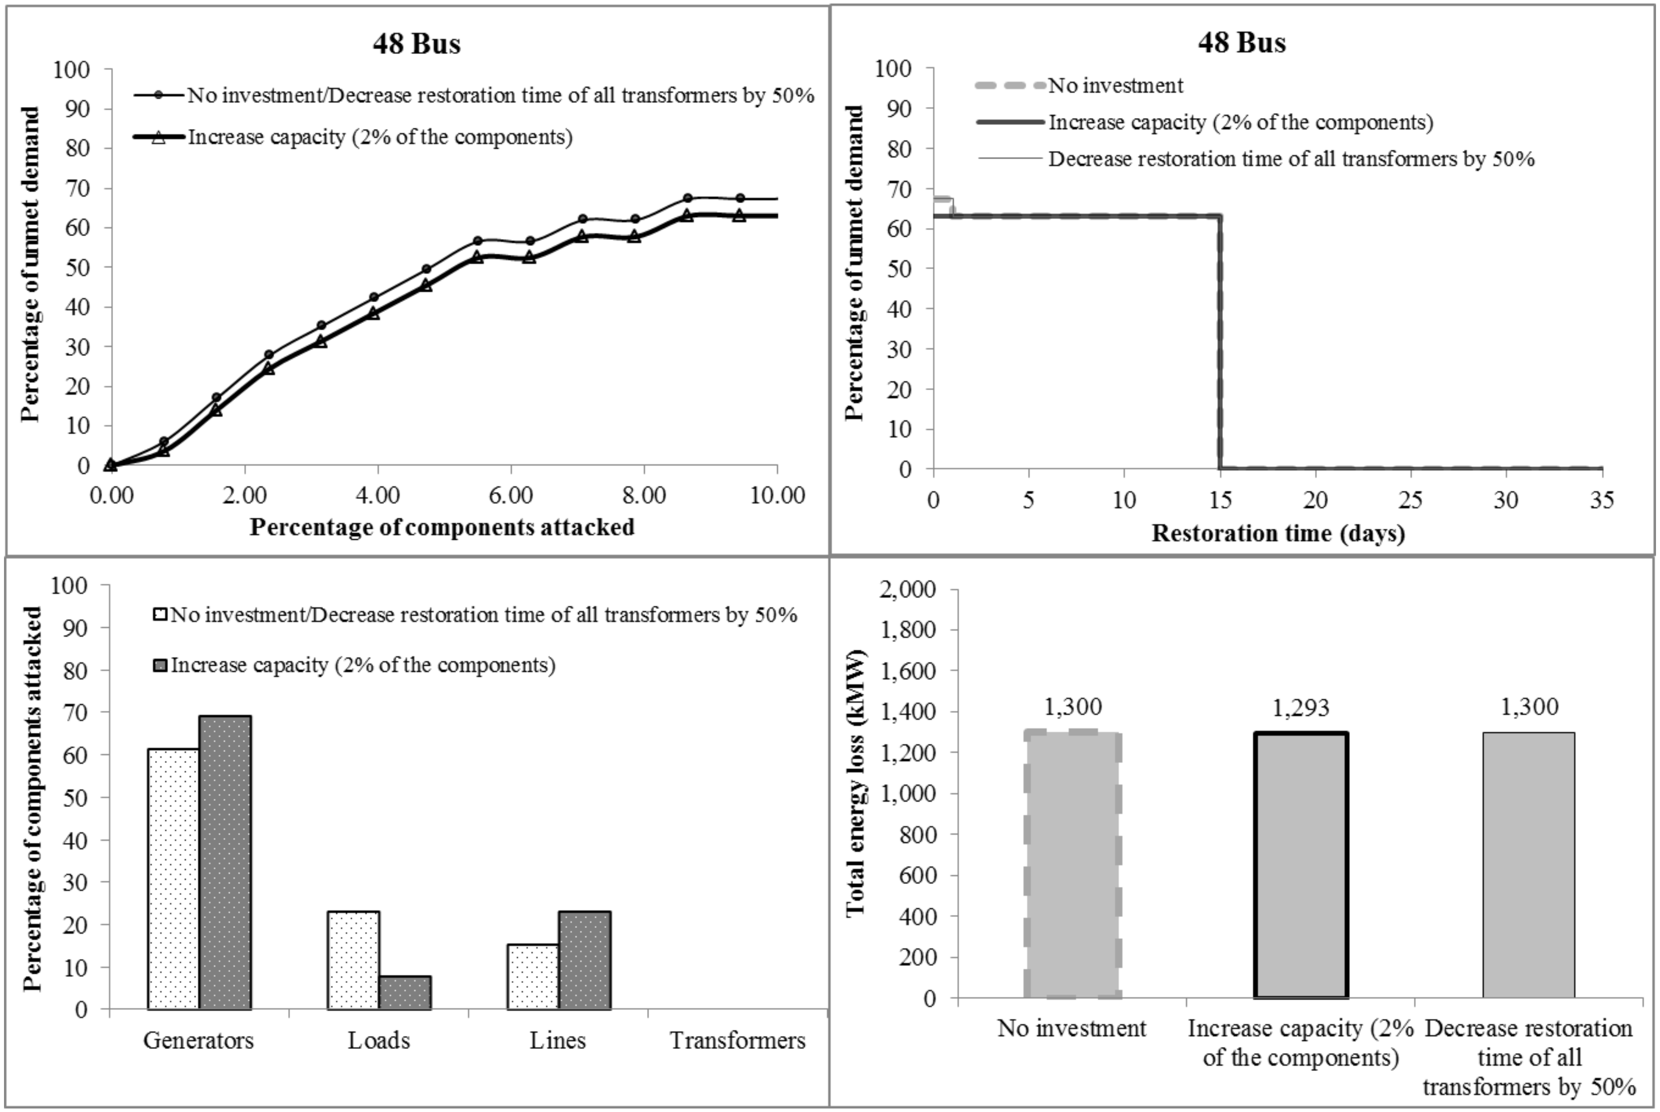

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 118-bus system)

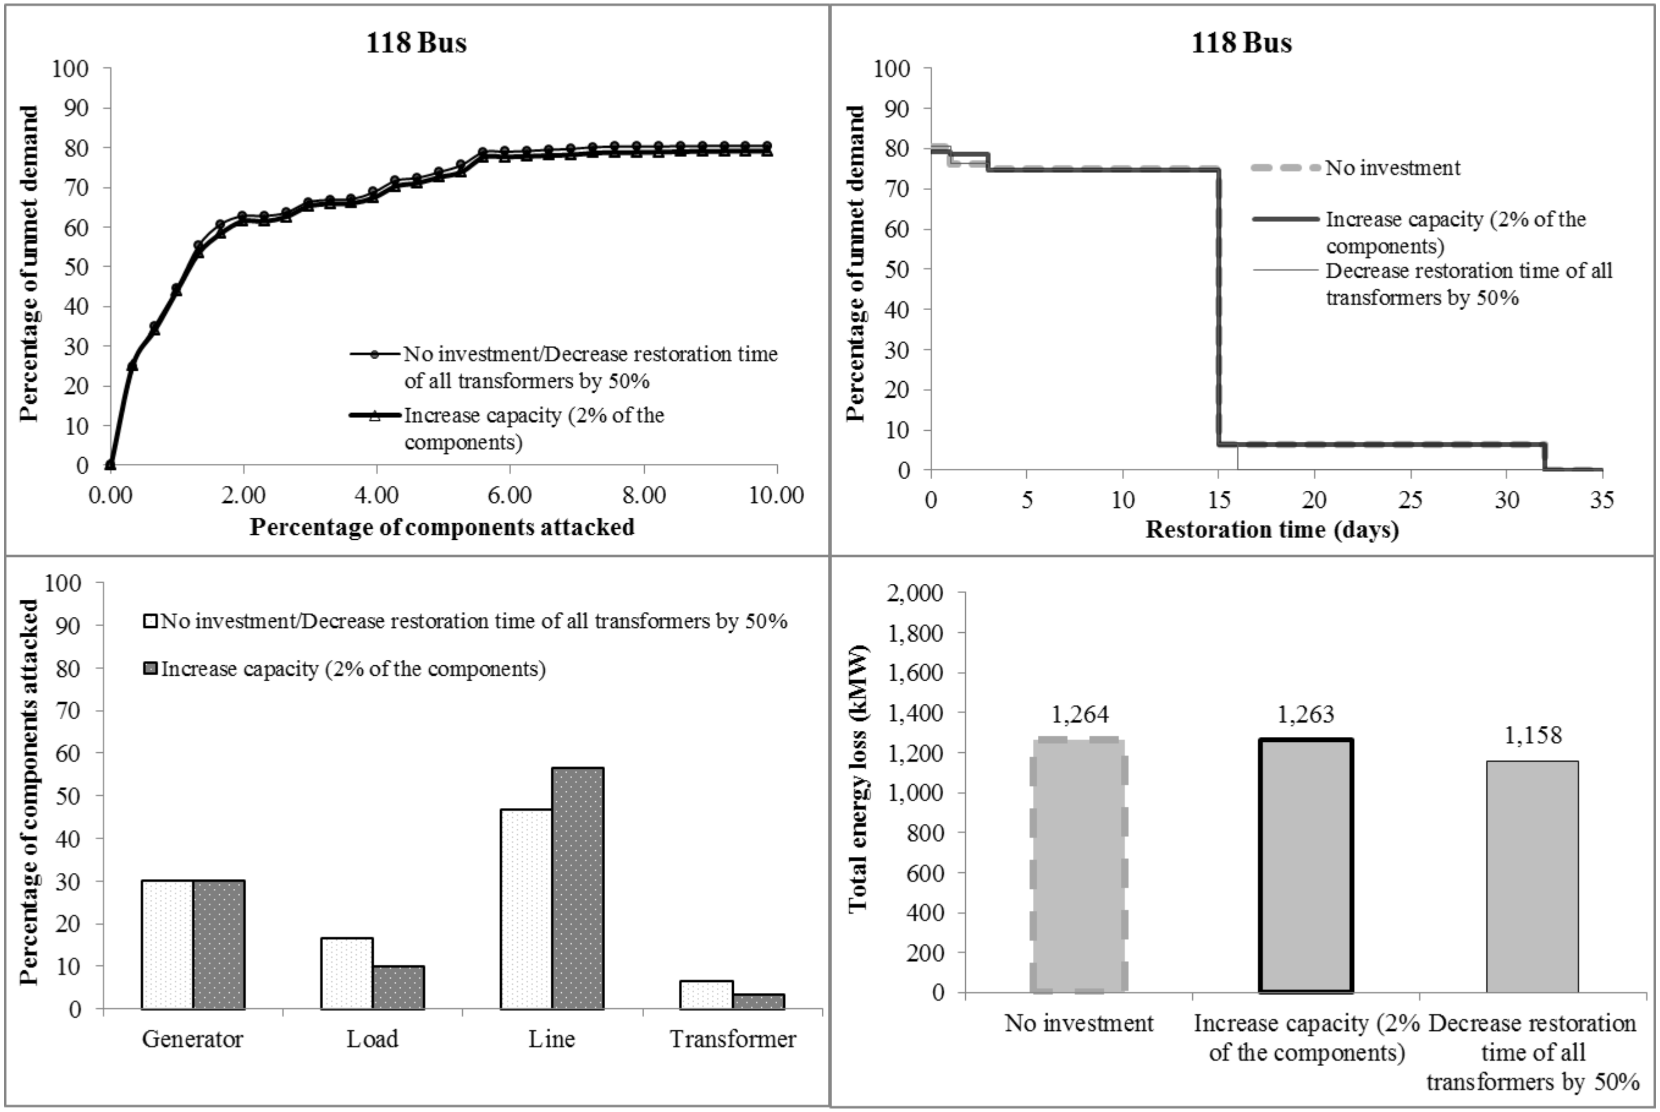

Appendix N. Increasing capacity versus decreasing restoration times of transformers
(increasing capacity of the 2% of the components against a static attacker with no cascading knowledge, 300-bus system)

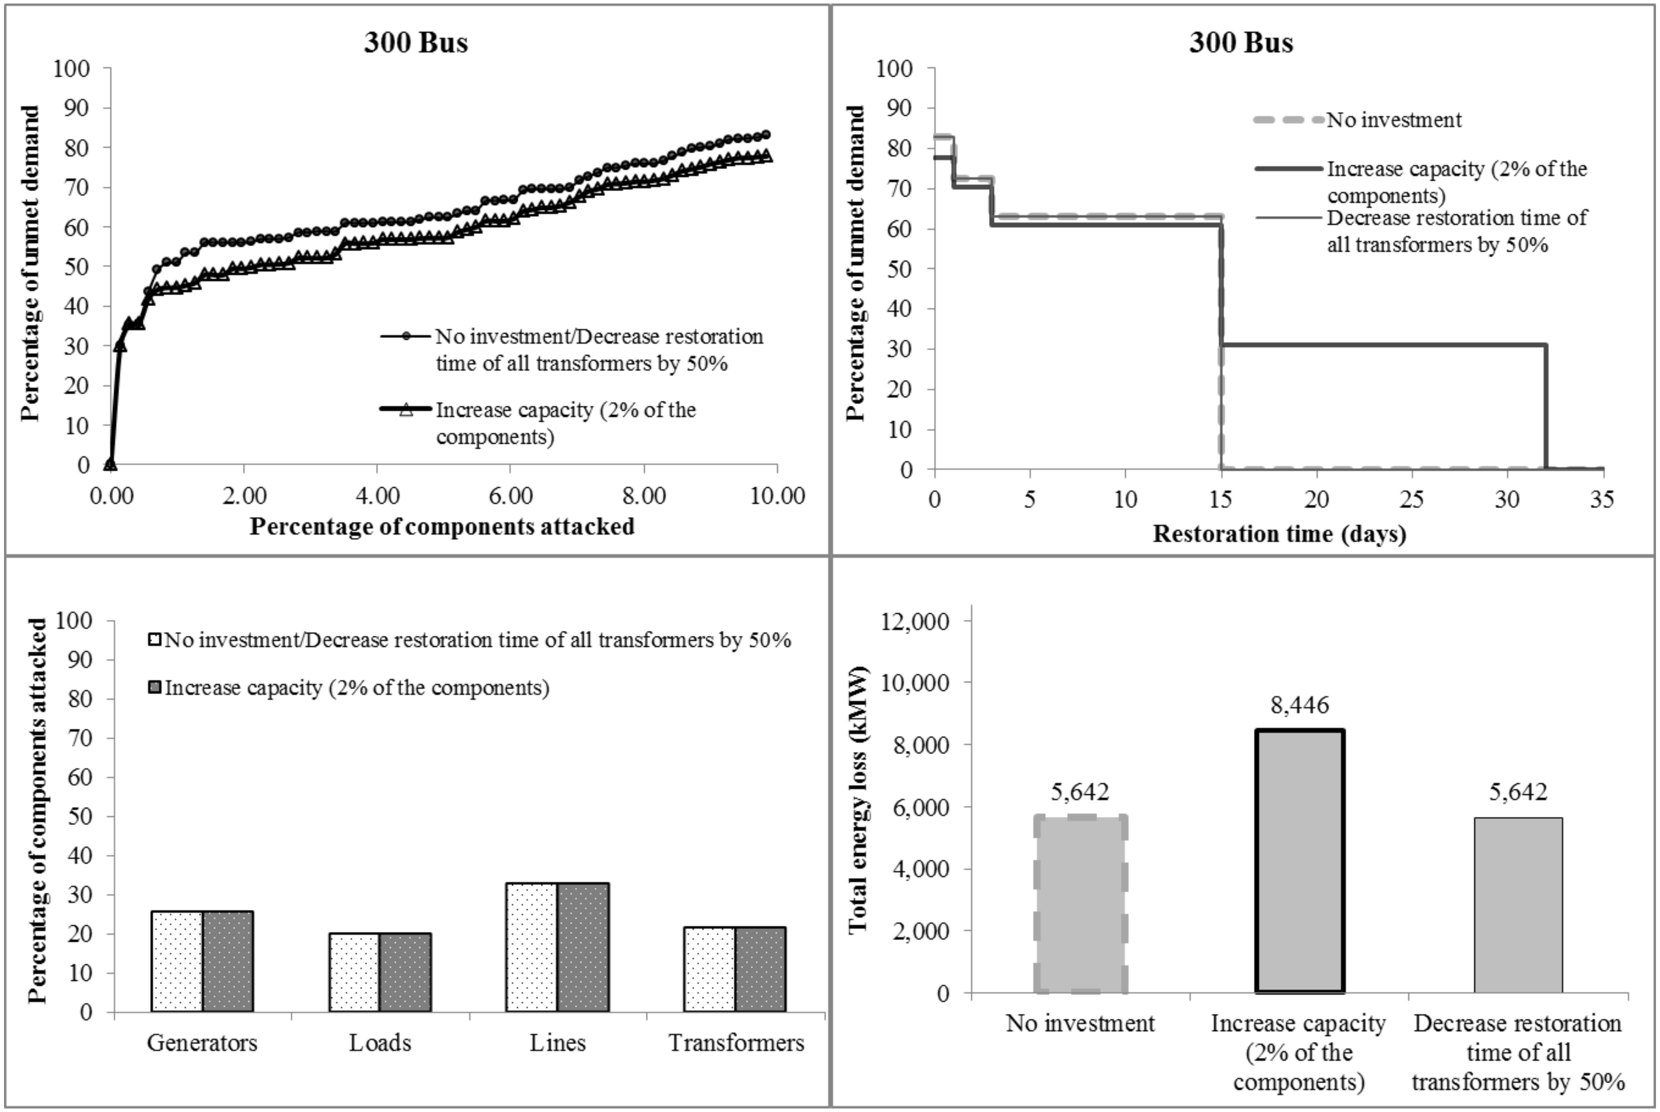

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 24-bus system)

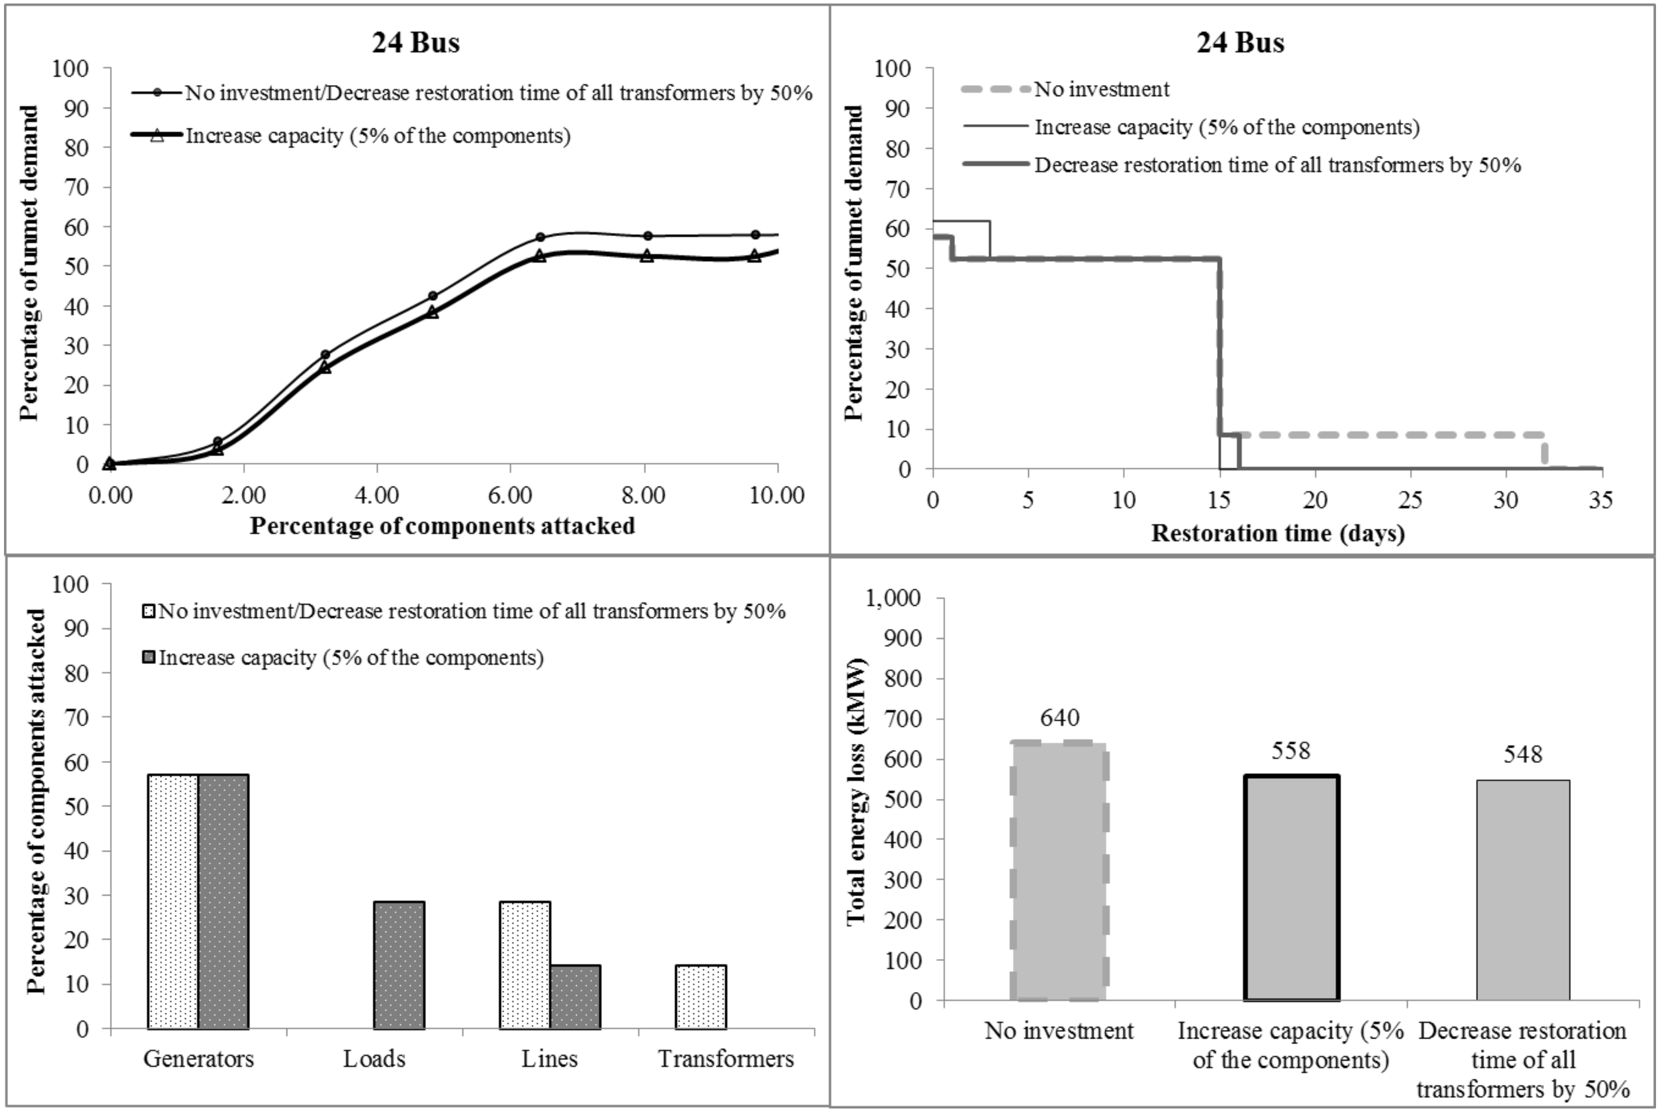

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 48-bus system)

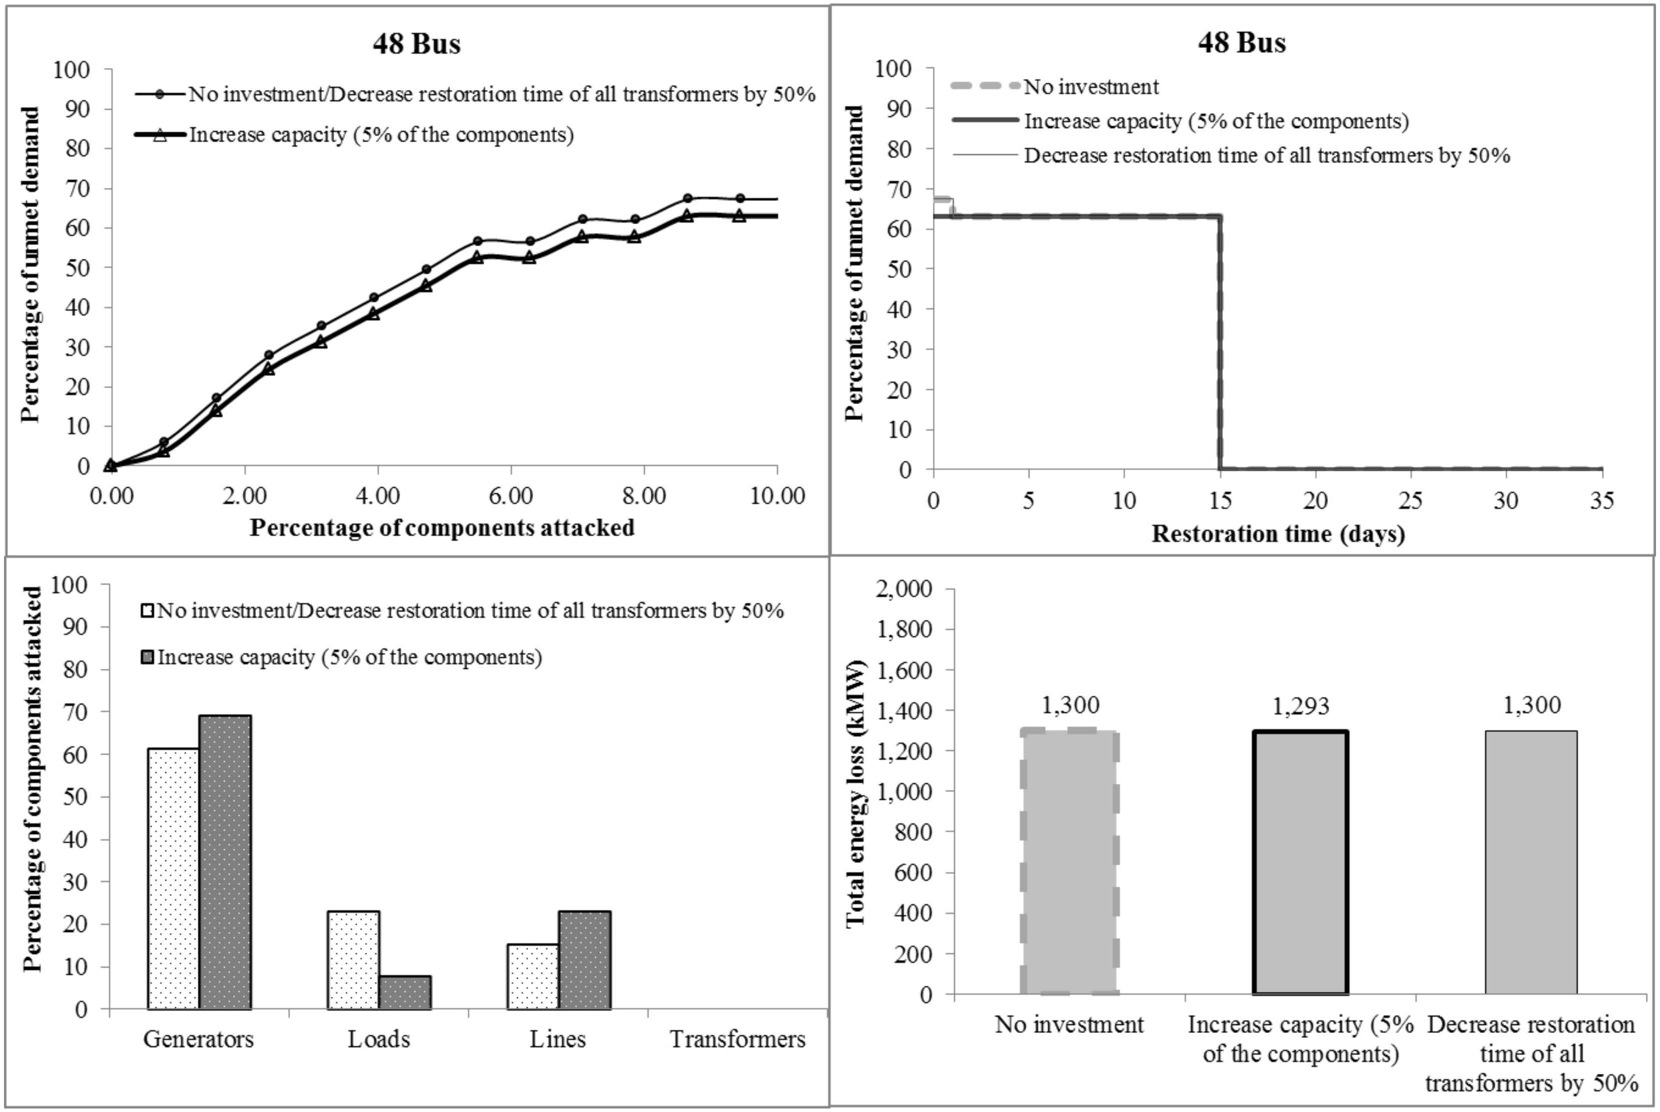

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 118-bus system)

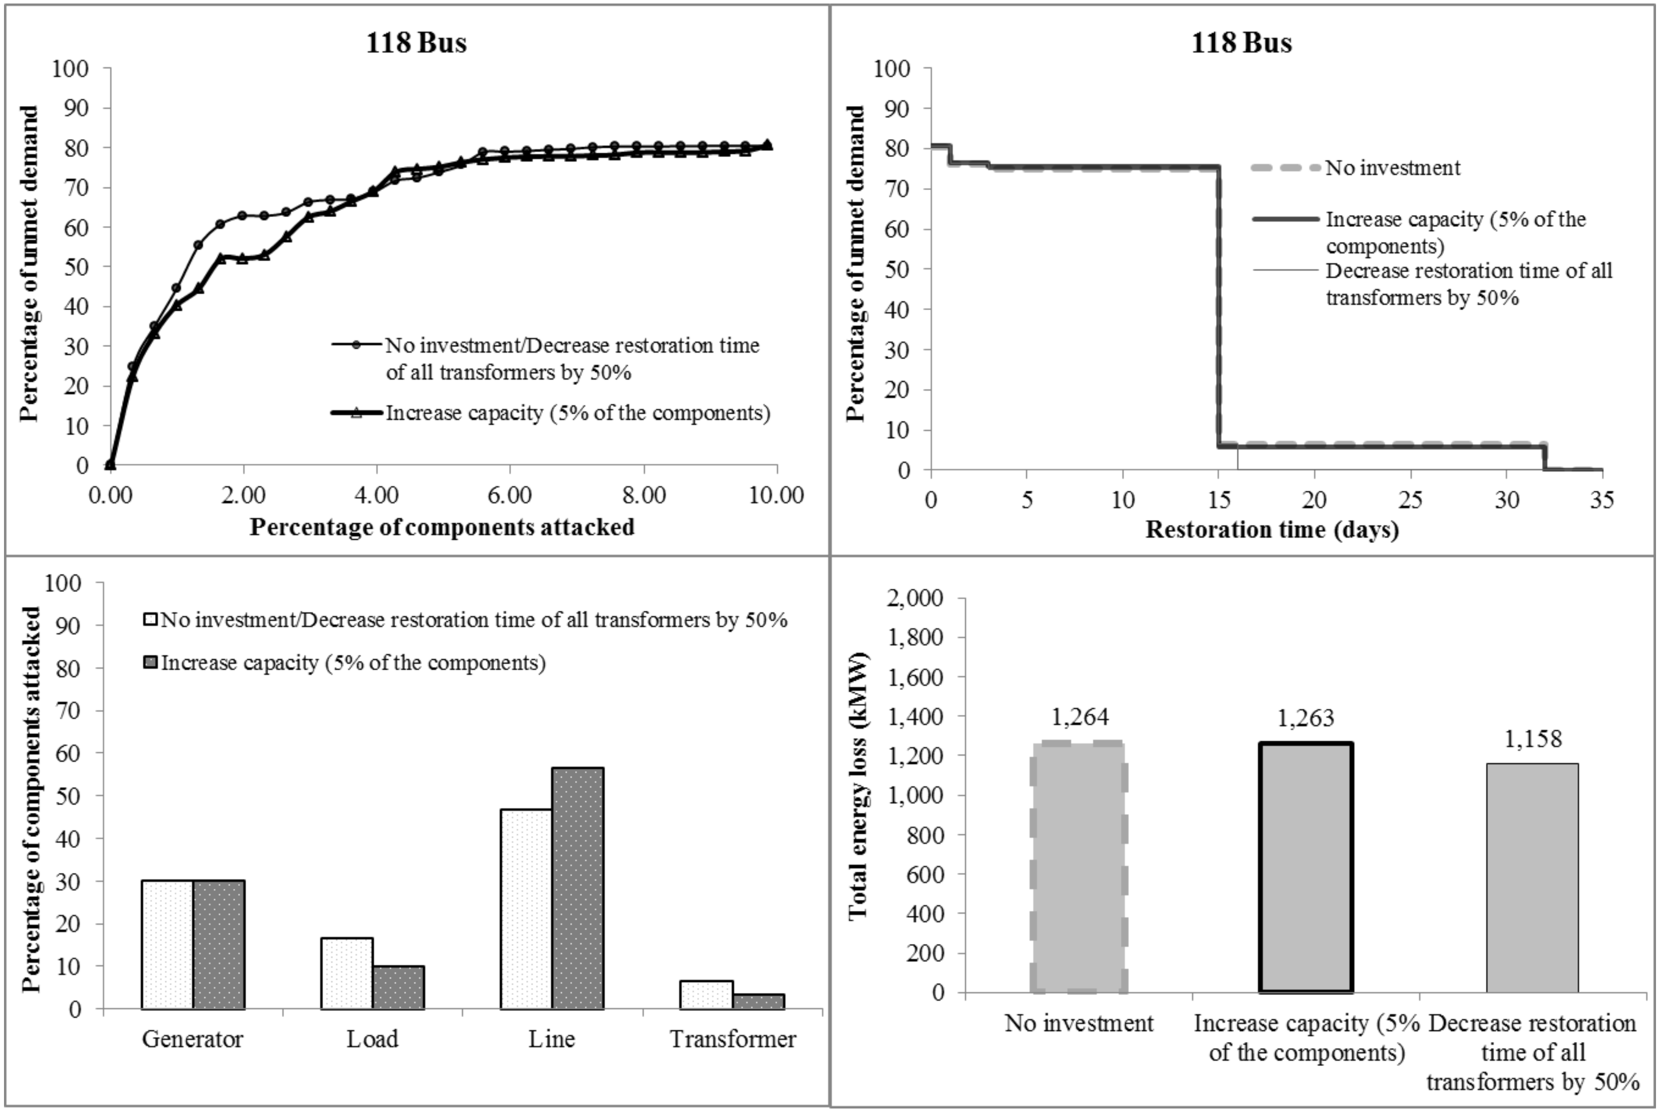

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 5% of the components against a static attacker with no cascading knowledge, 300-bus system)

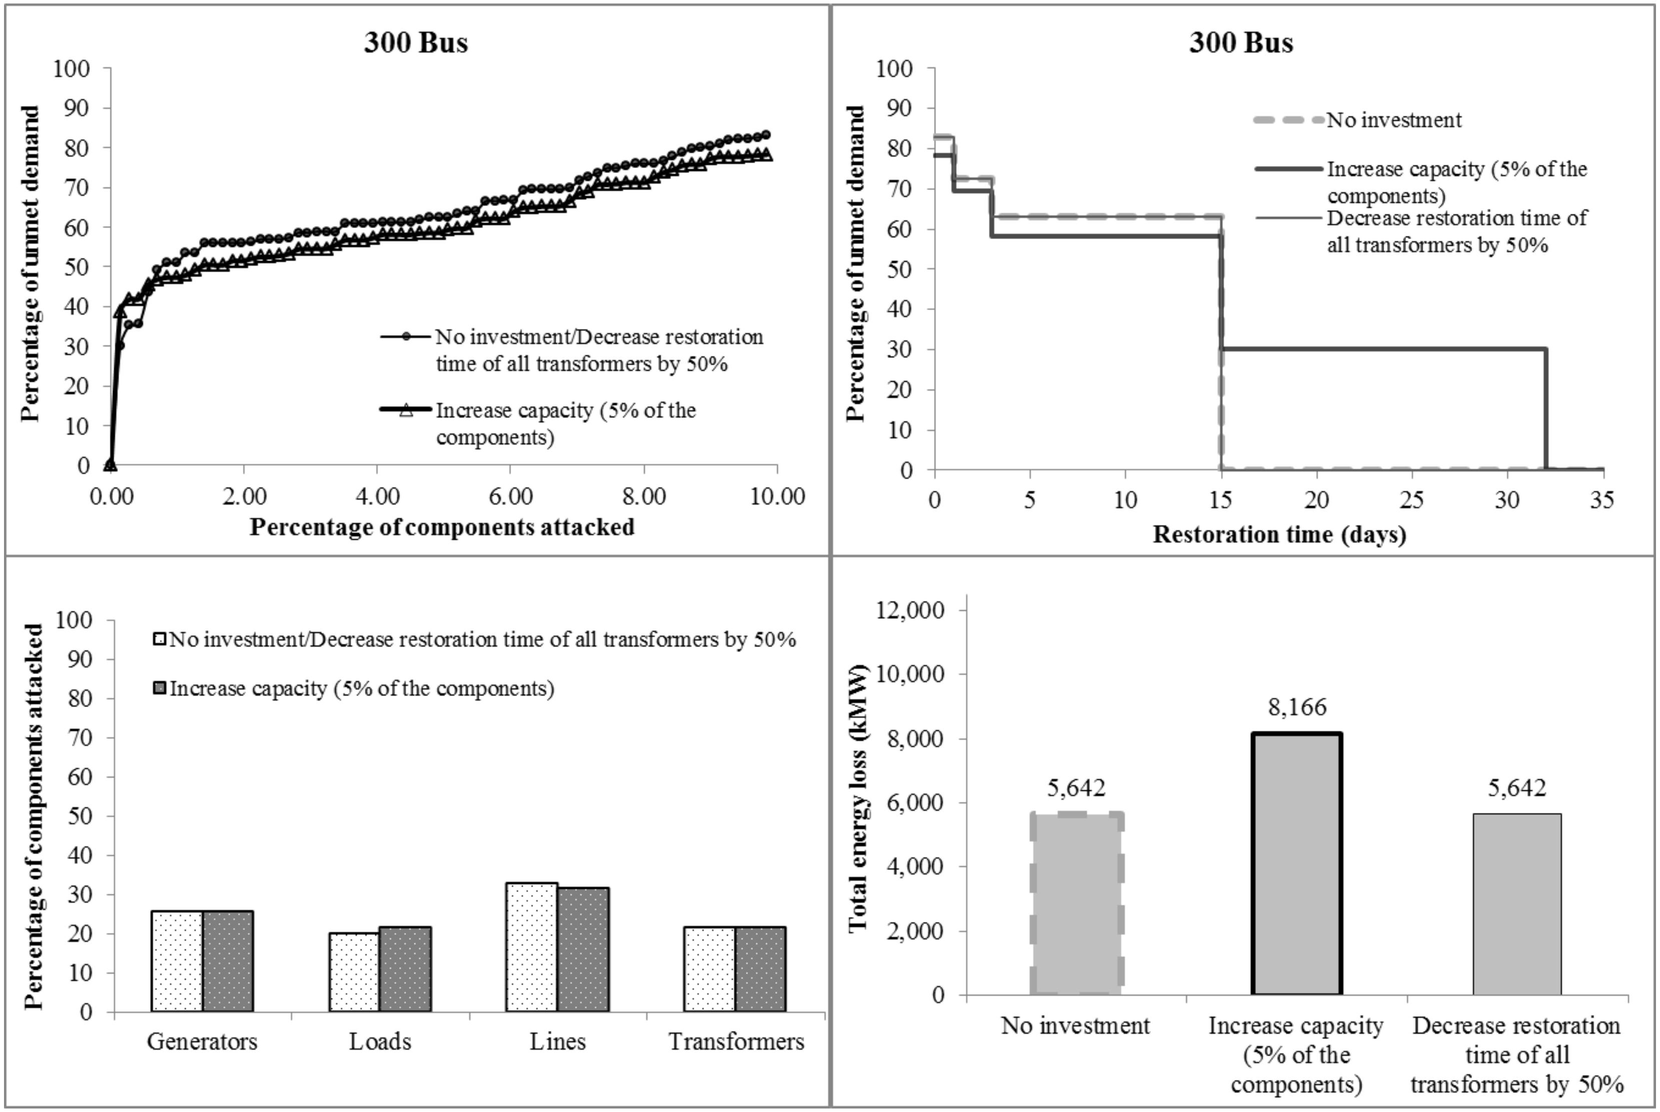

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 24-bus system)

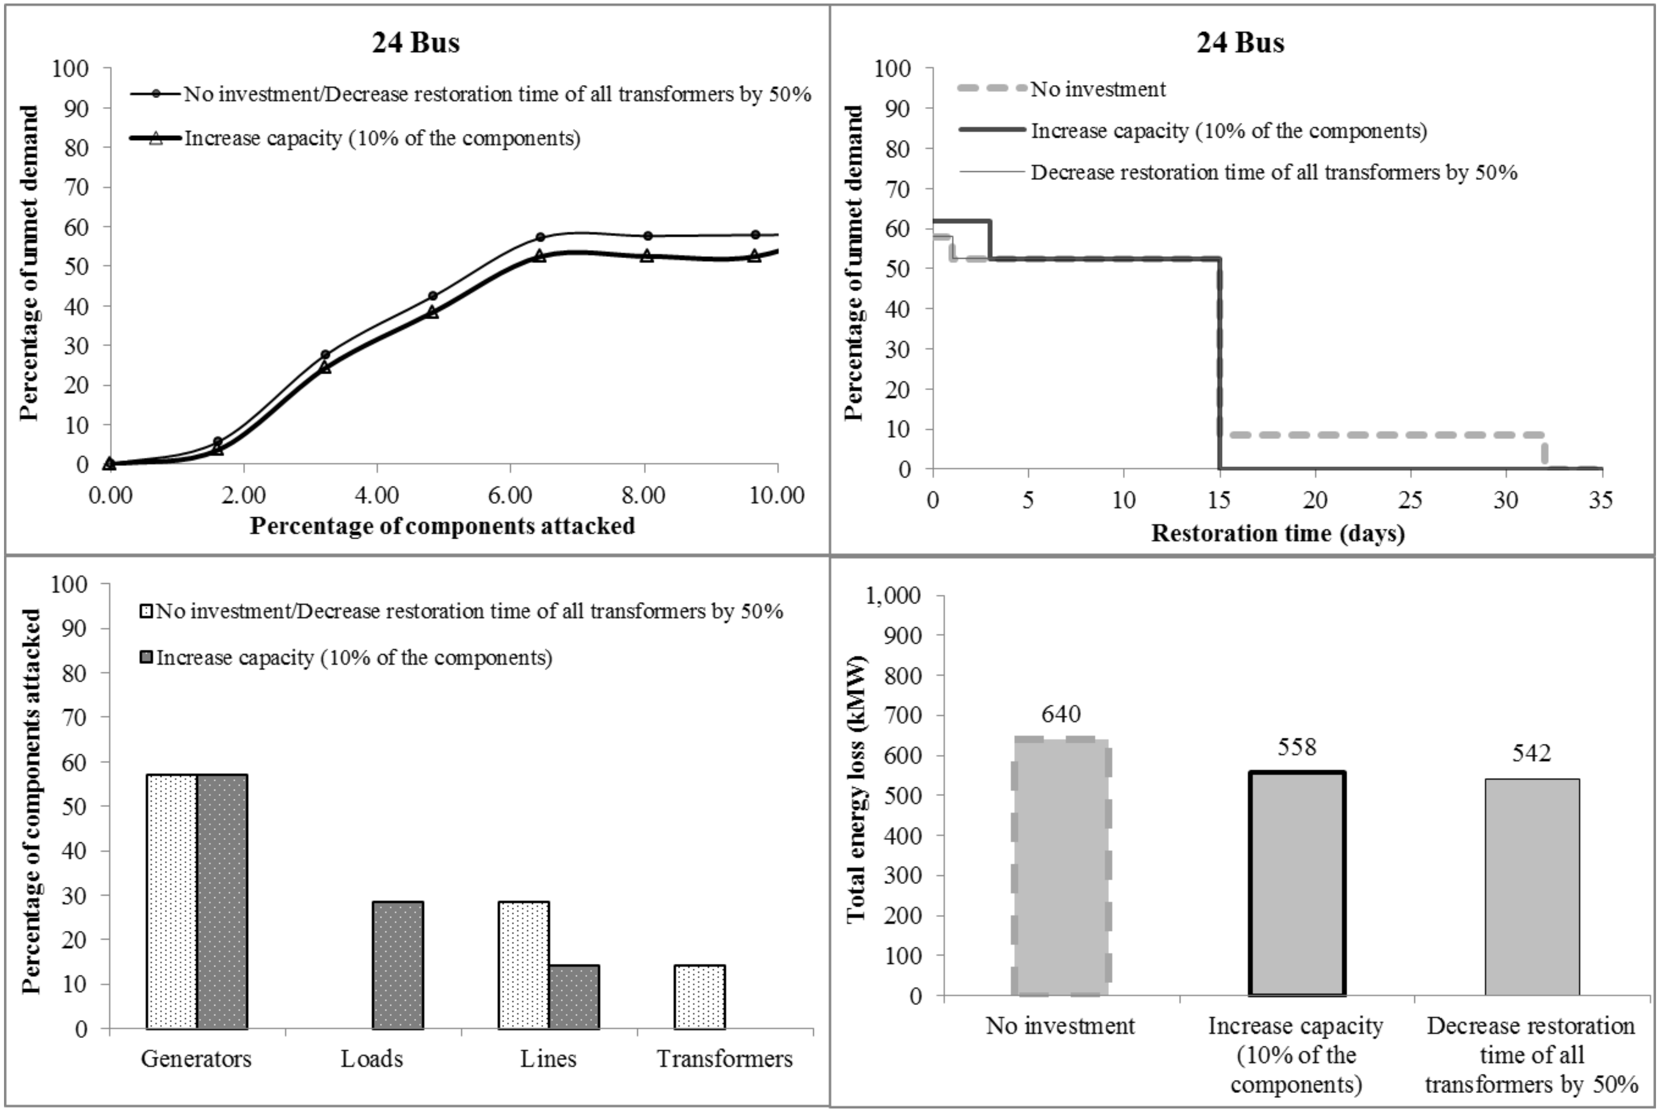

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 48-bus system)

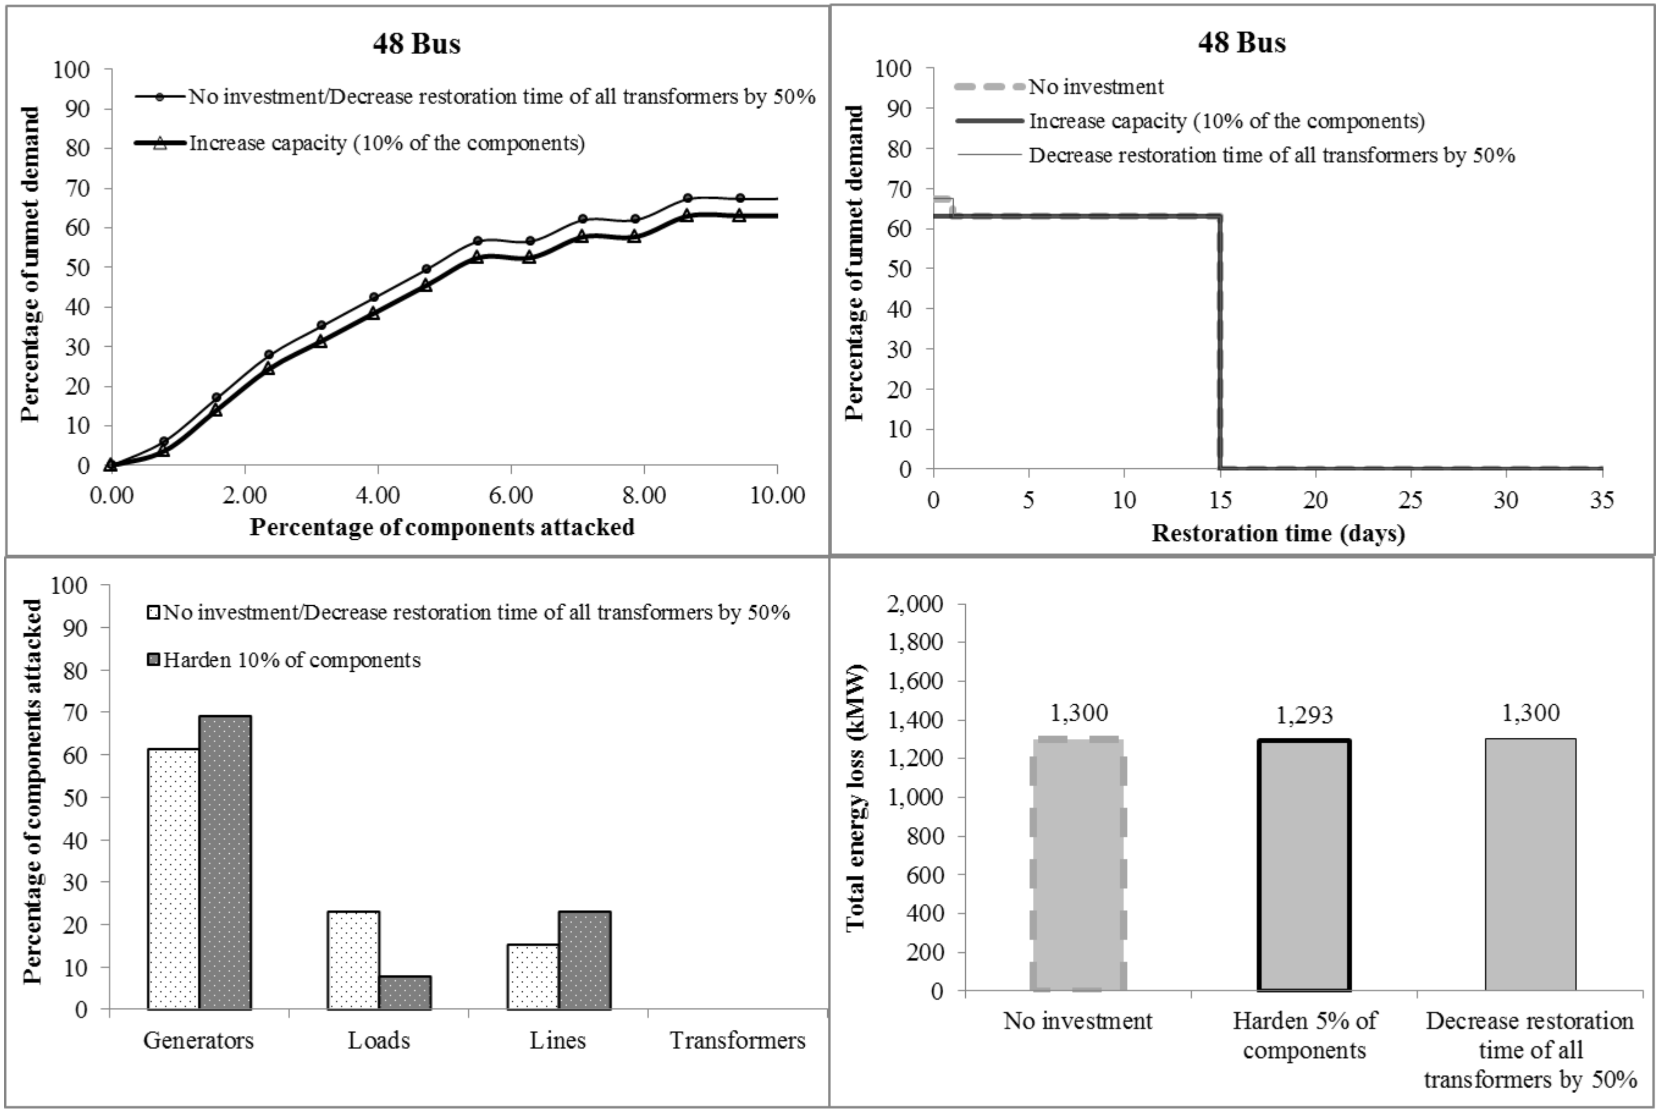

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 118-bus system)

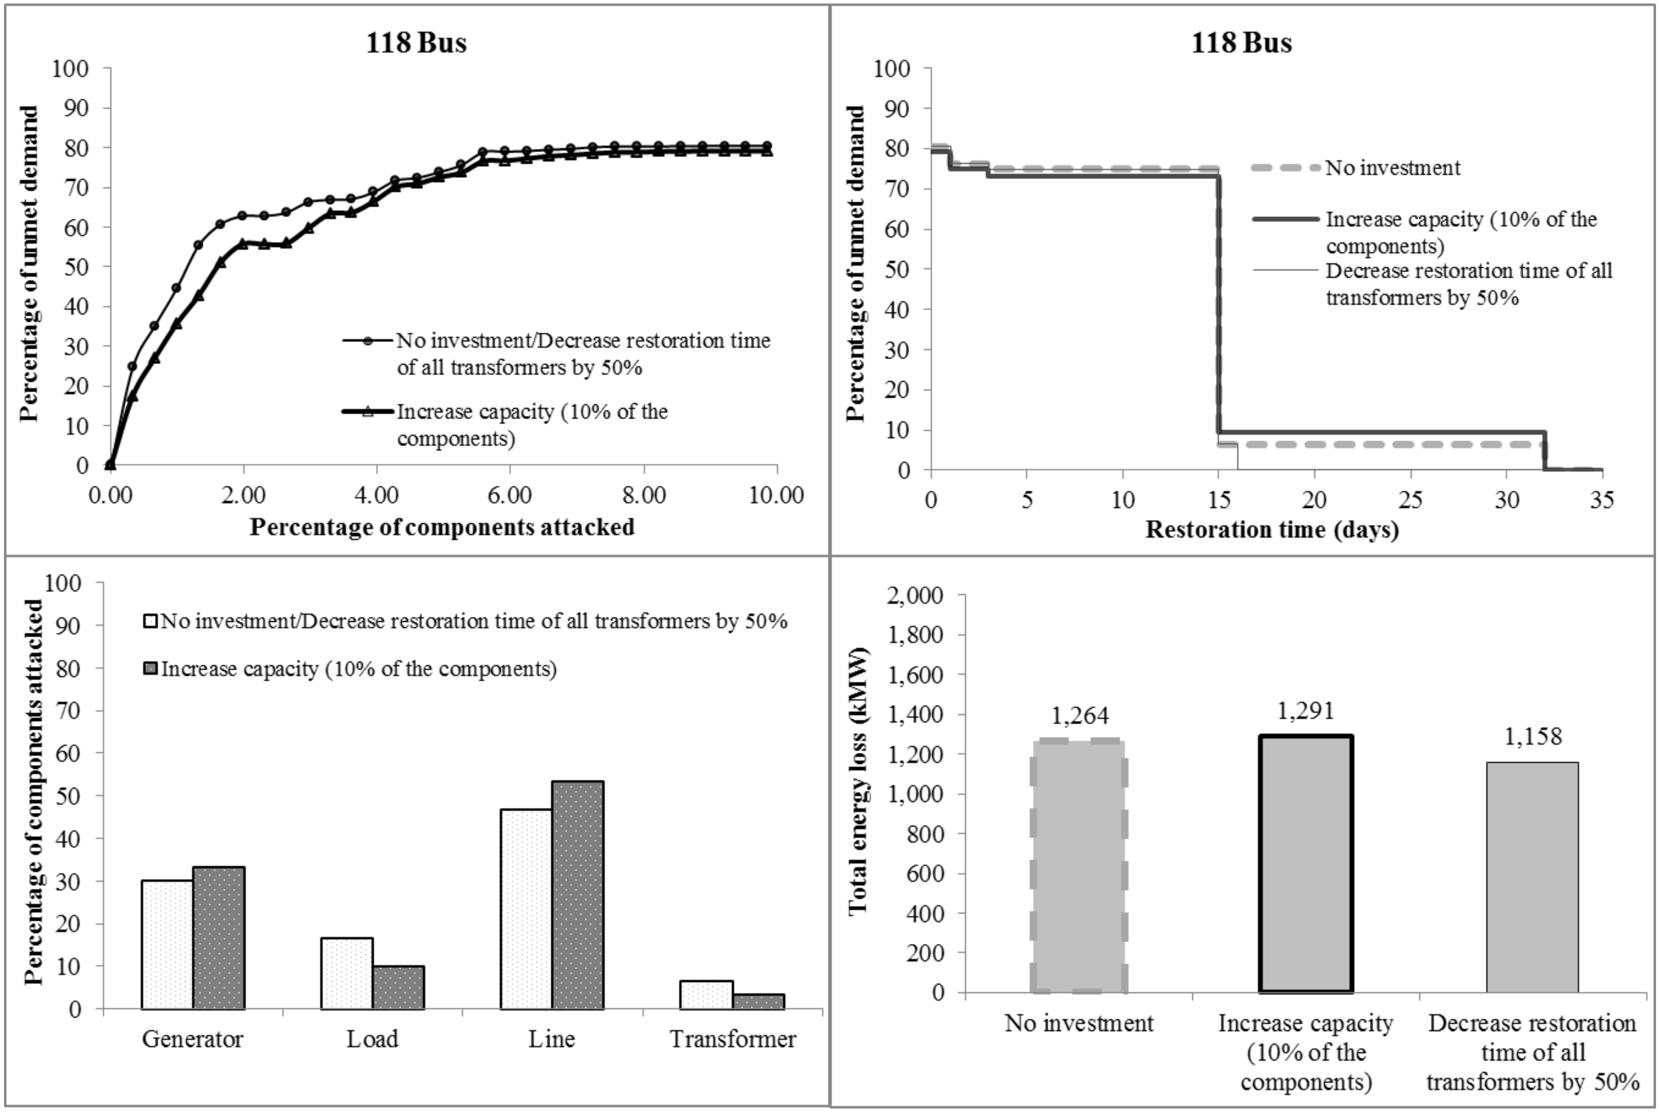

Appendix N. Increasing capacity versus decreasing restoration times of transformers (increasing capacity of the 10% of the components against a static attacker with no cascading knowledge, 300-bus system)

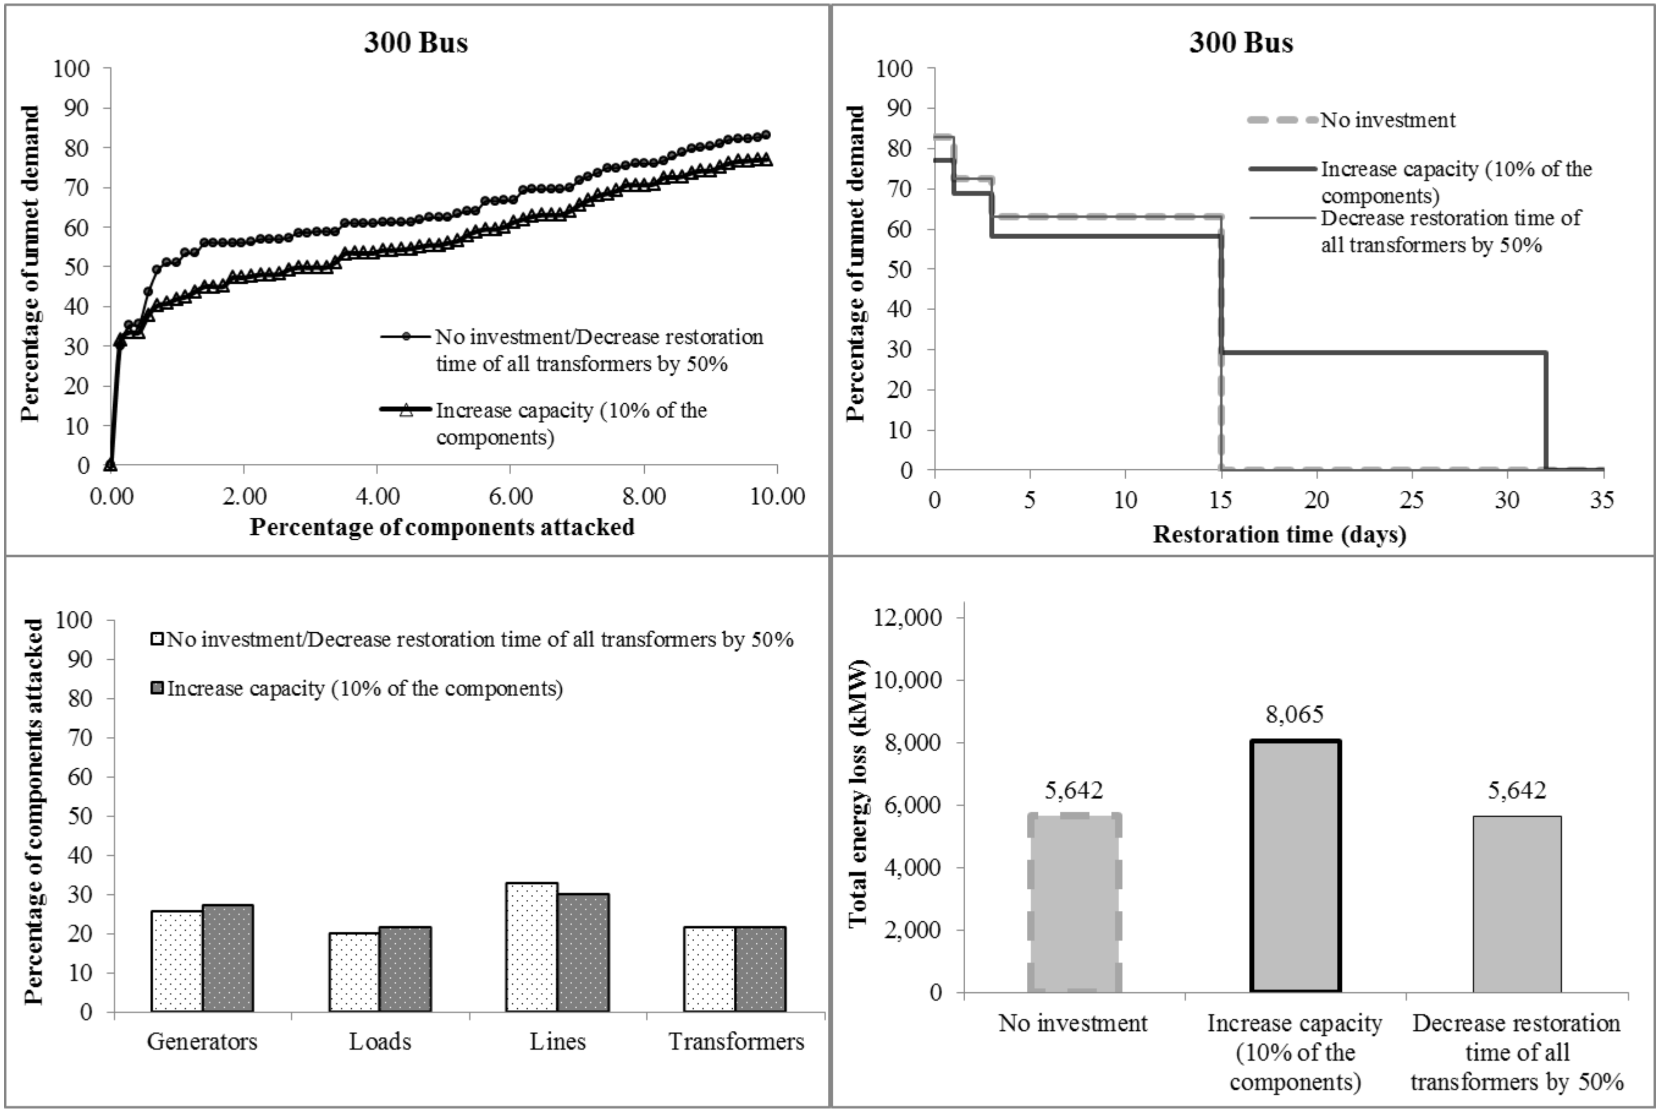

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a dynamic attacker, 24-bus system)

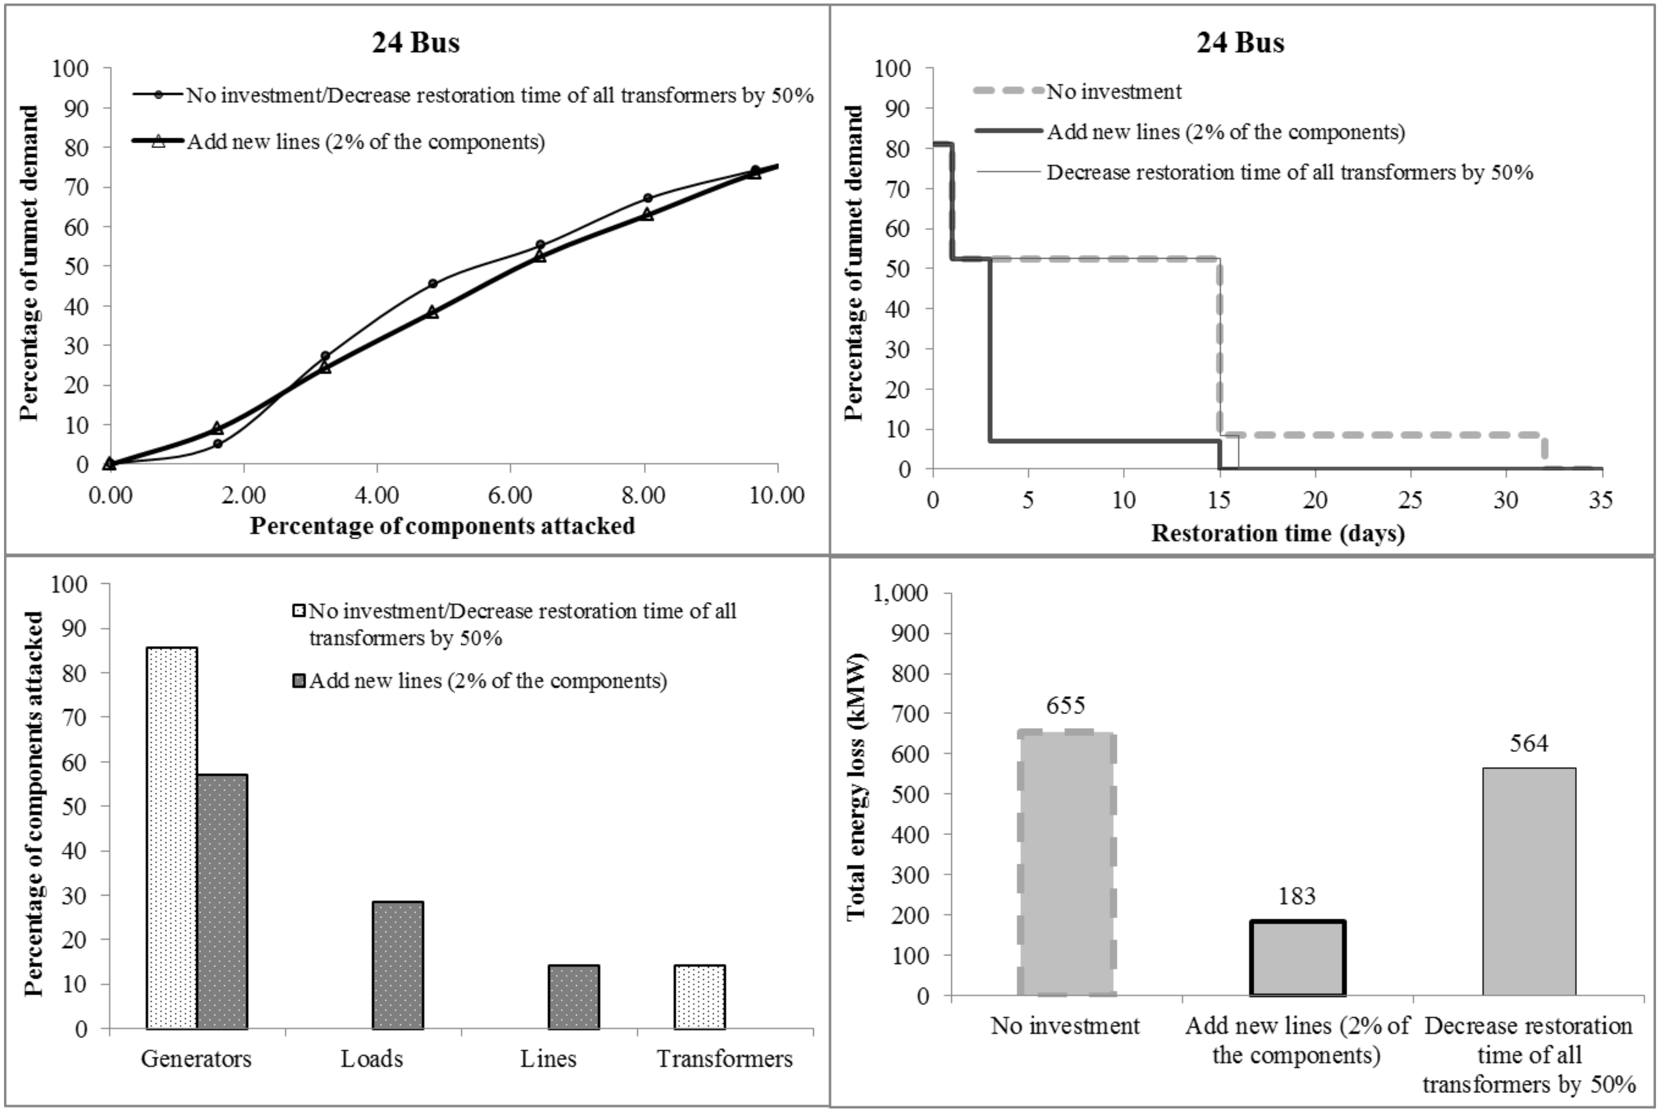

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a dynamic attacker, 48-bus system)

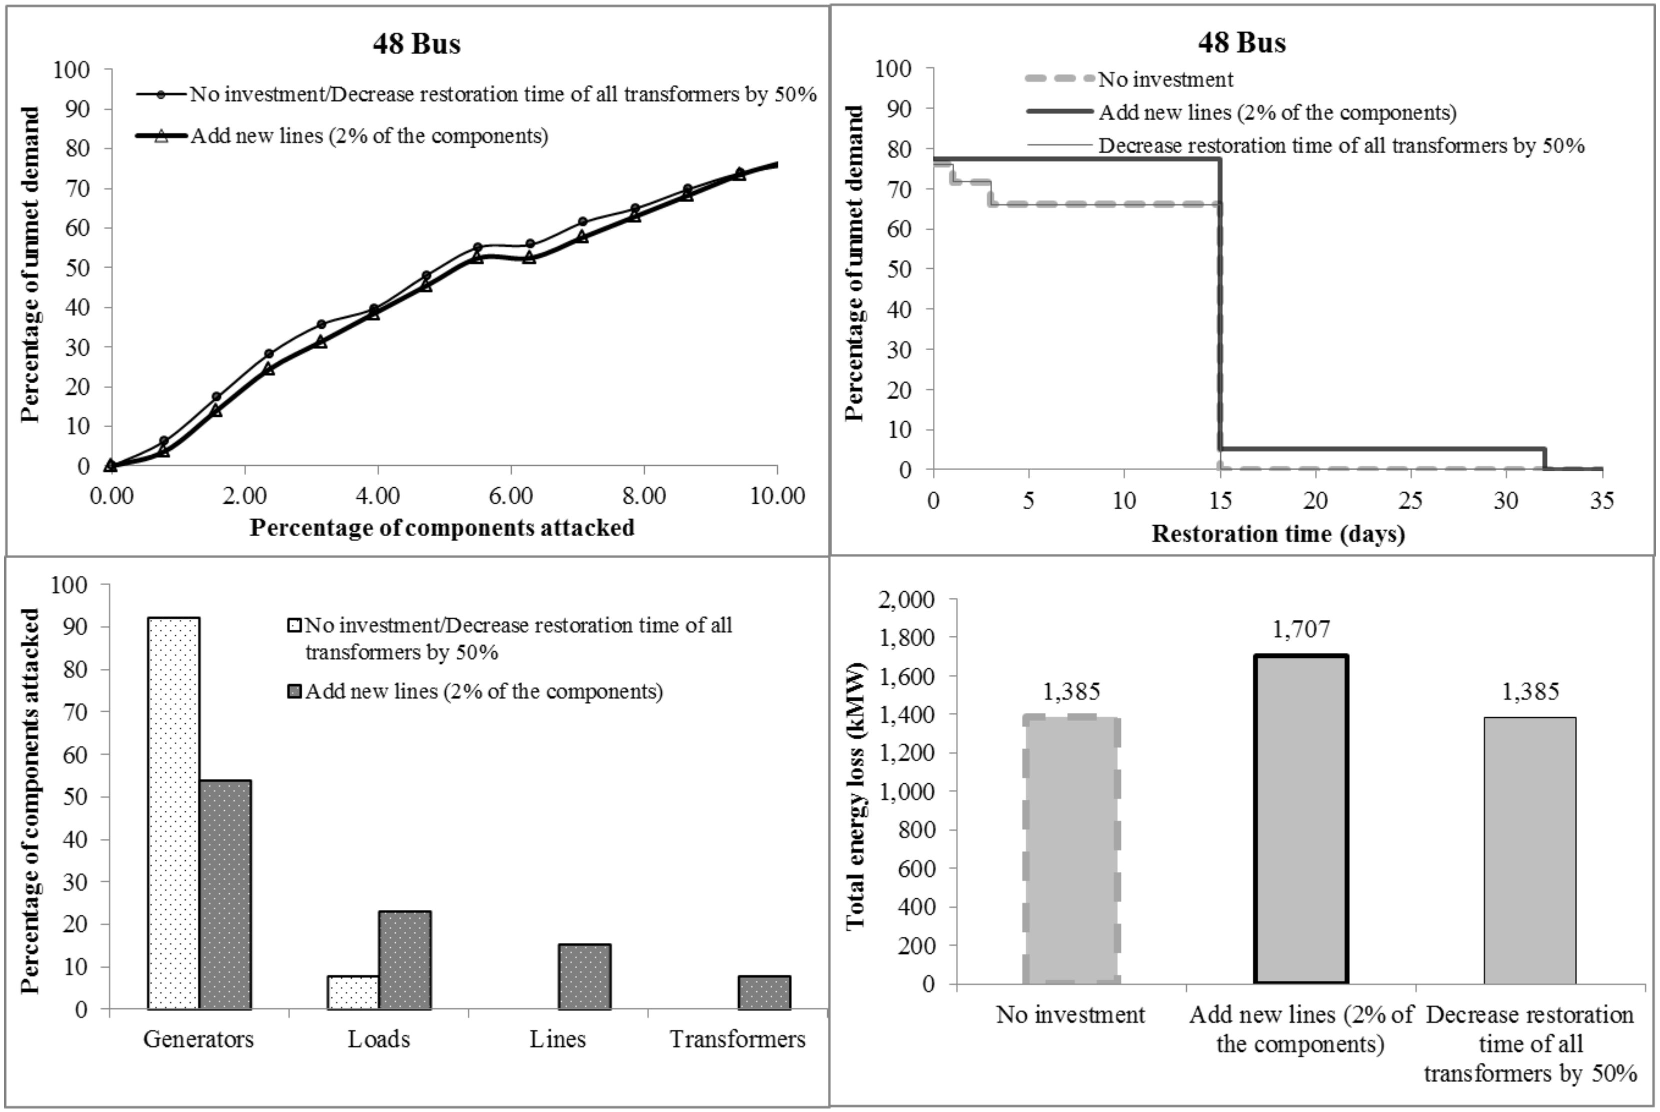

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a dynamic attacker, 118-bus system)

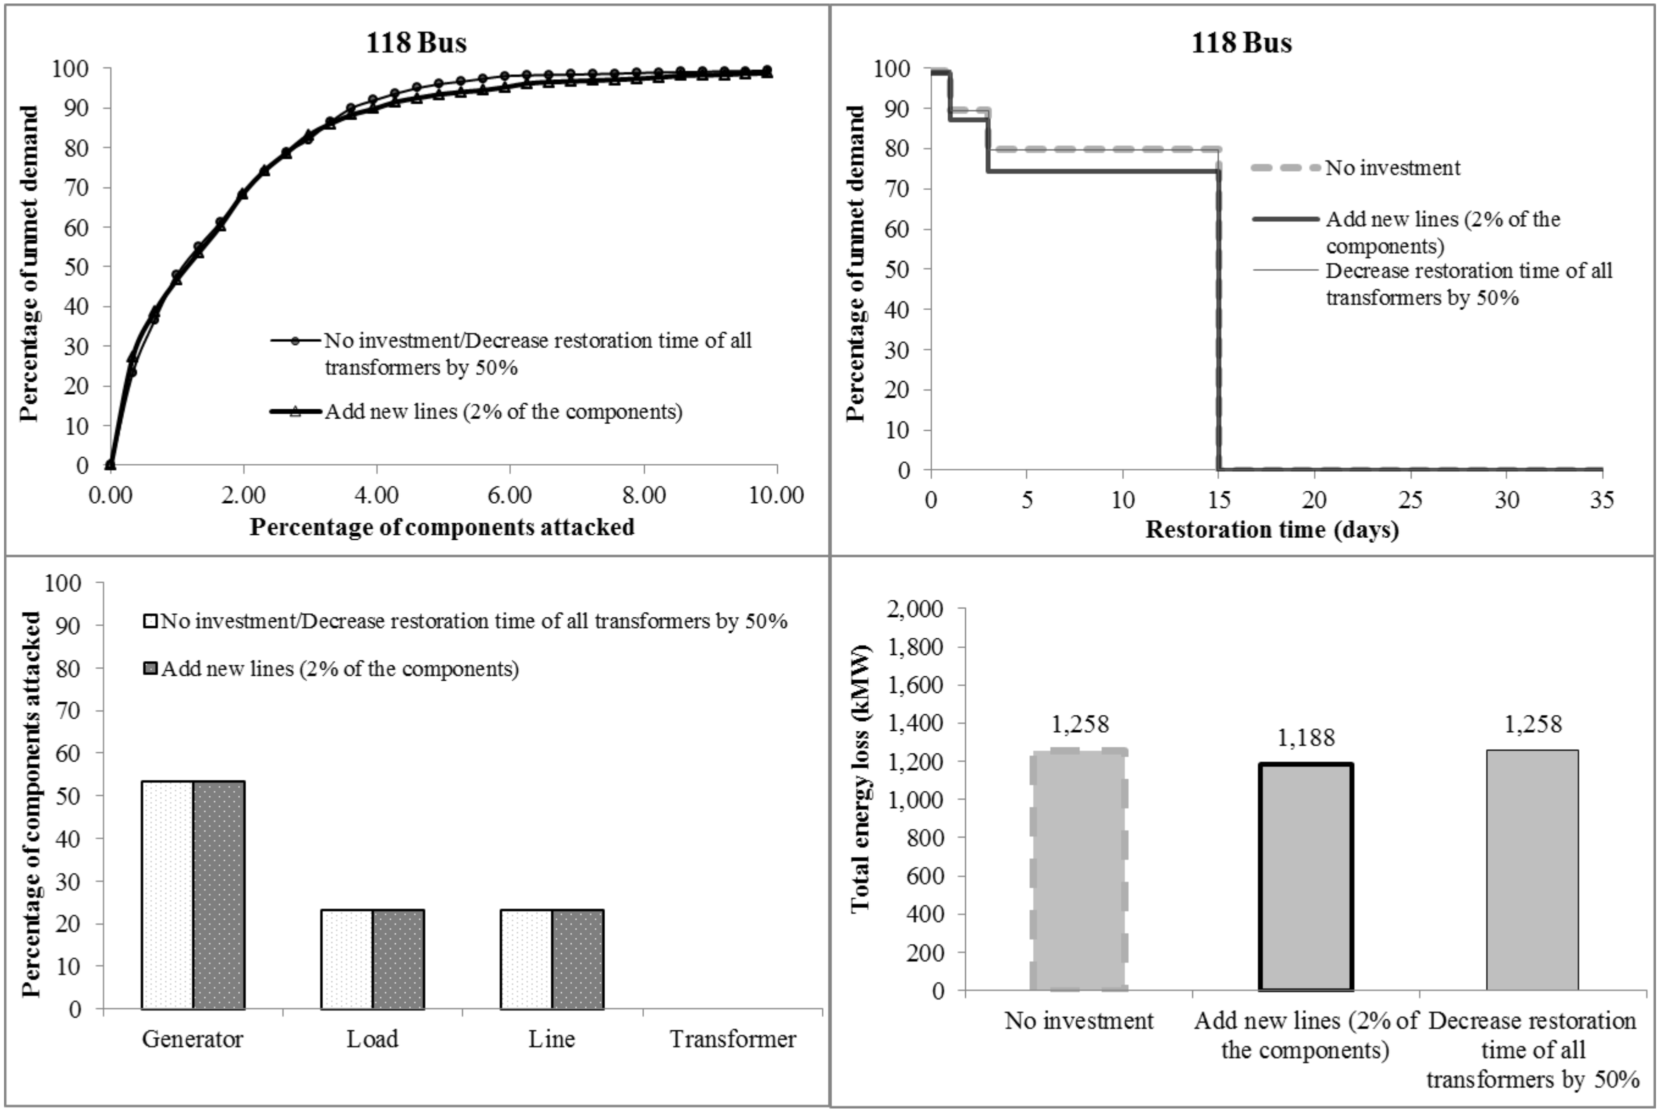

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a dynamic attacker, 300-bus system)

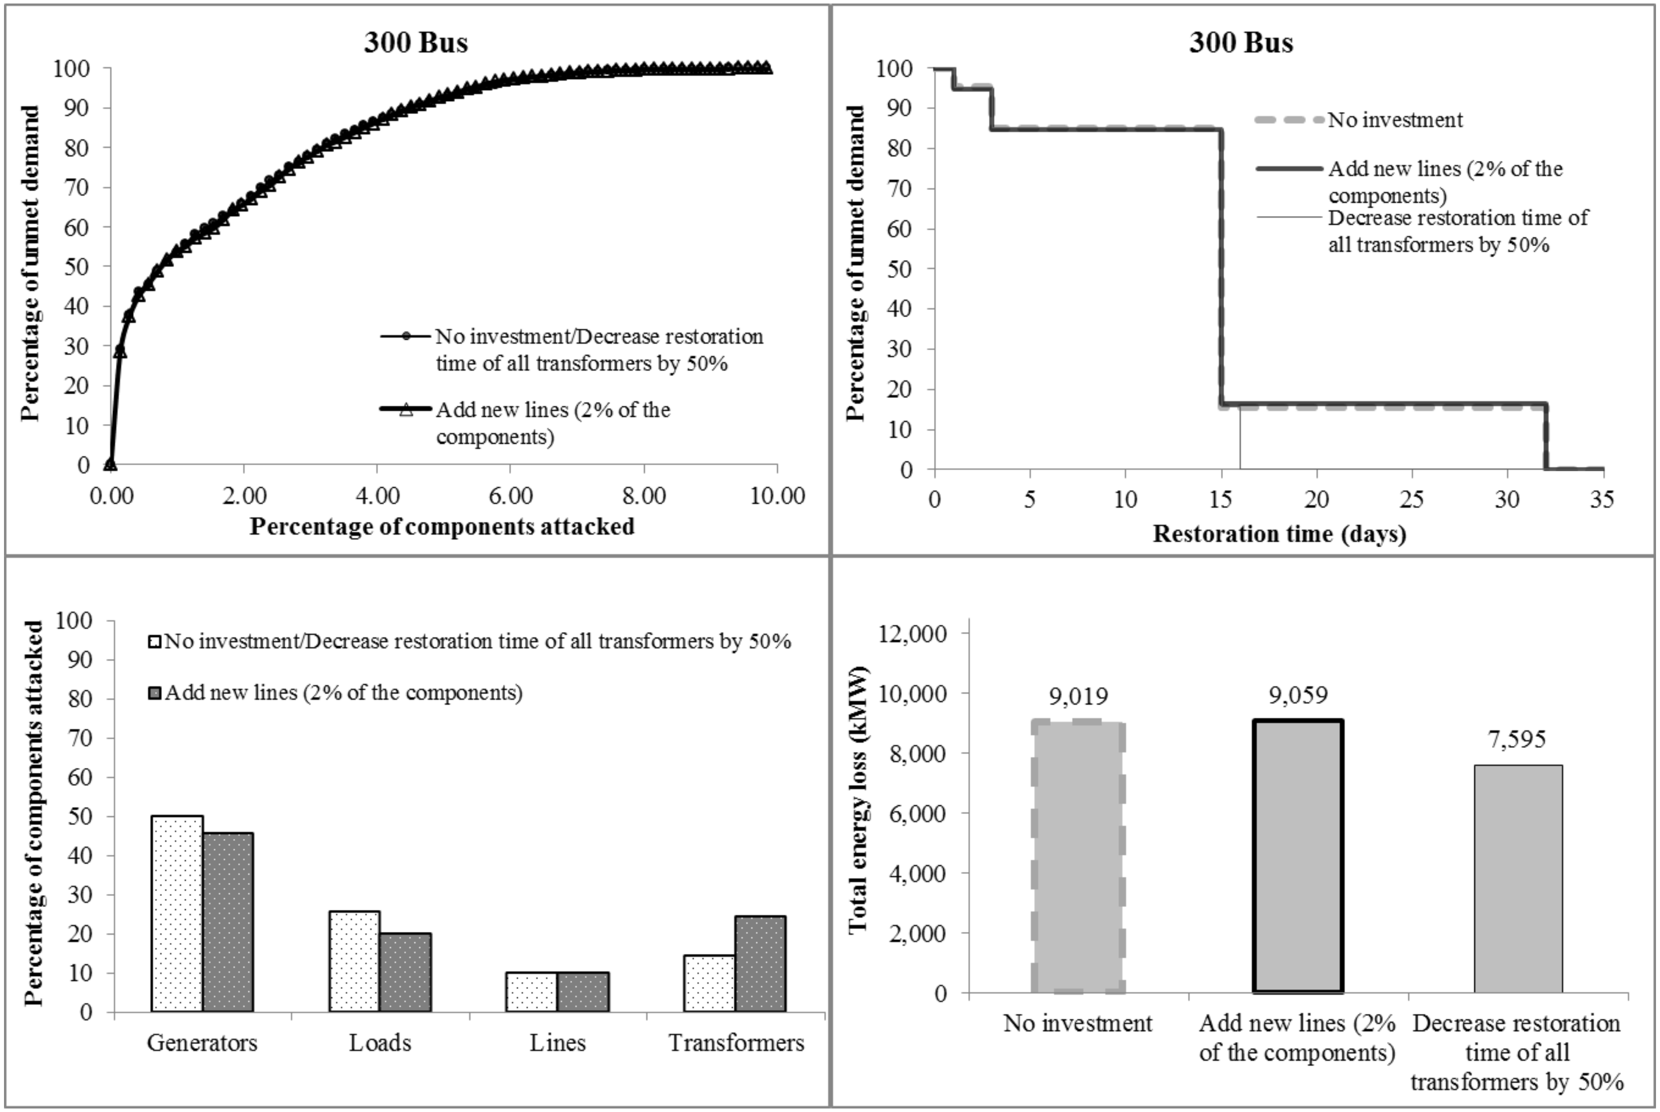

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a dynamic attacker, 24-bus system)

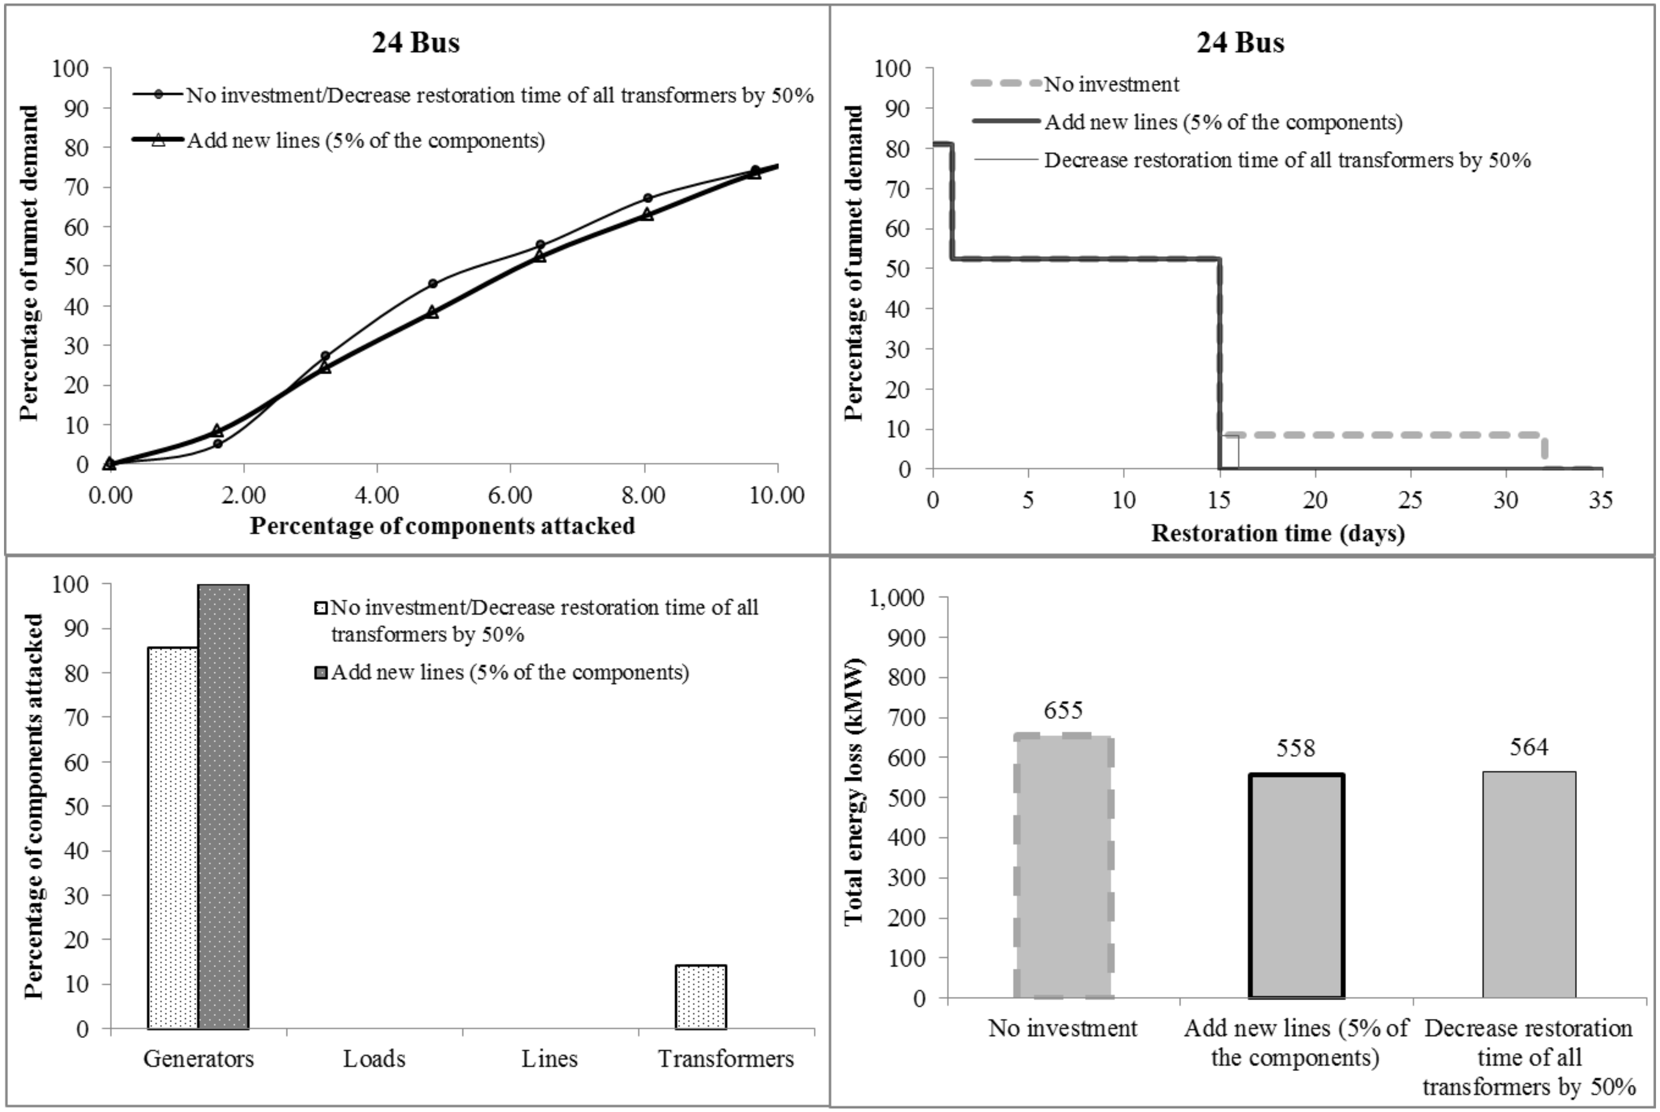

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a dynamic attacker, 48-bus system)

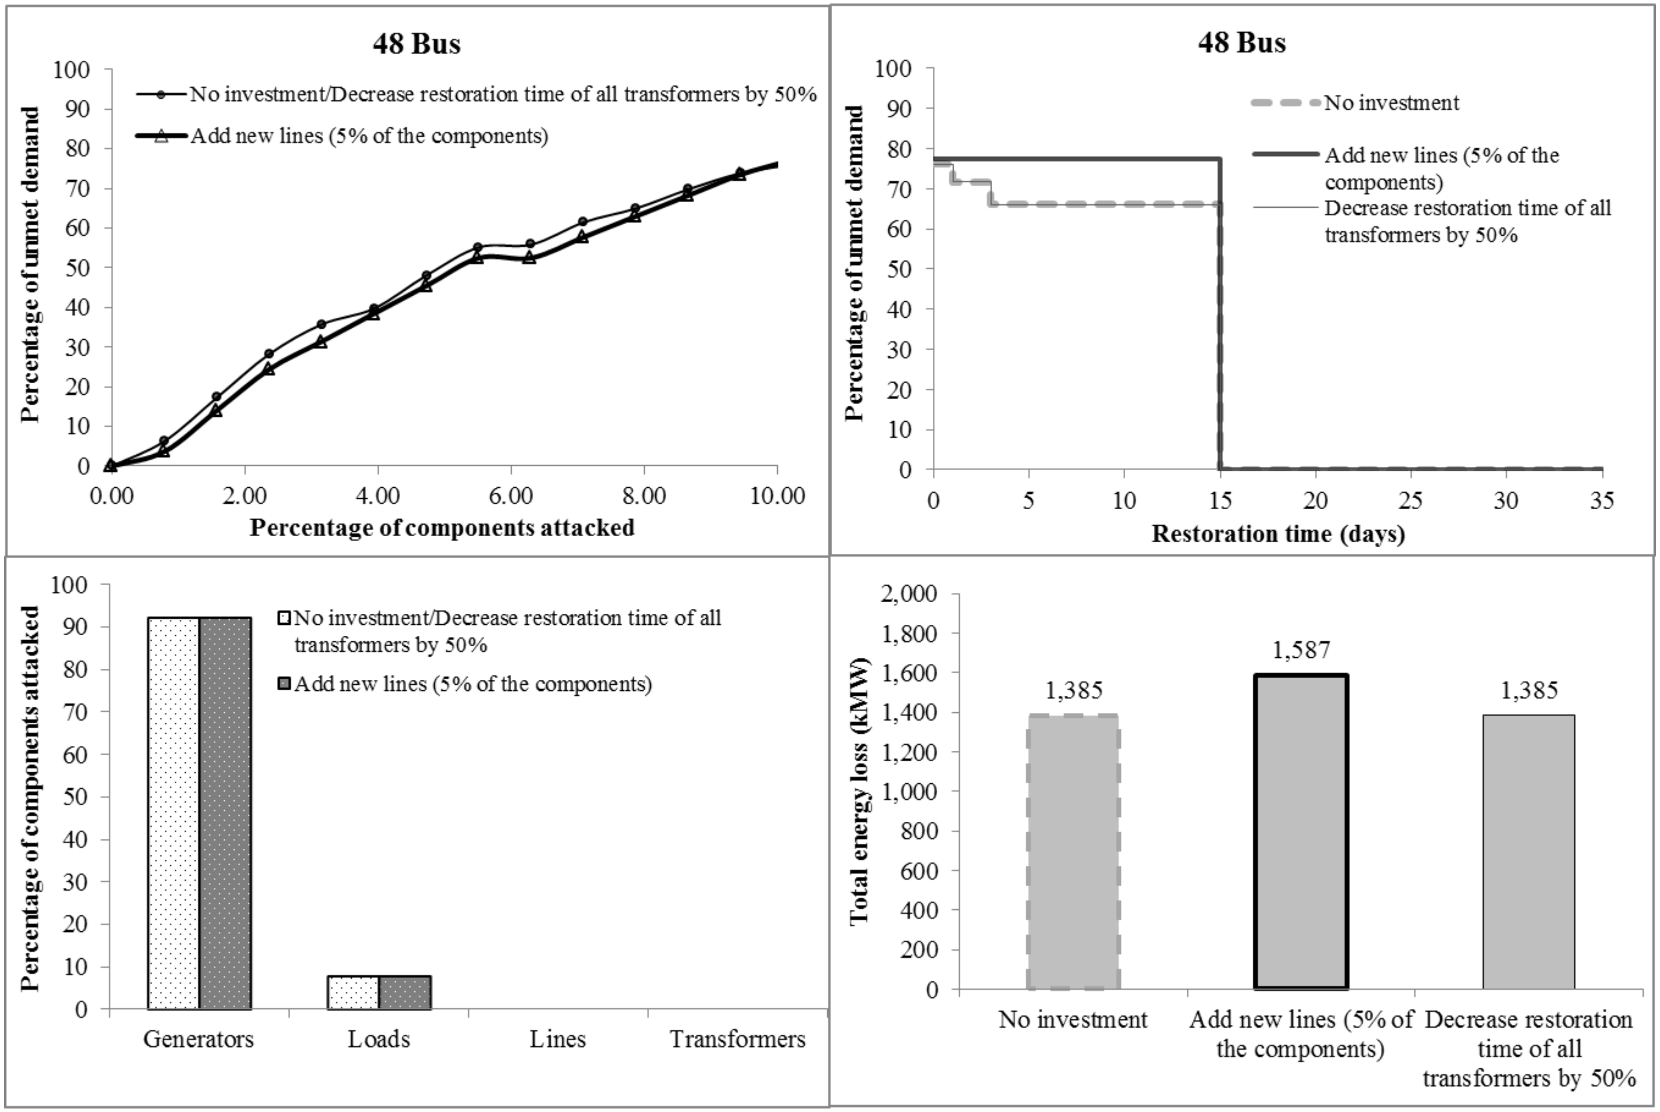

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a dynamic attacker, 118-bus system)

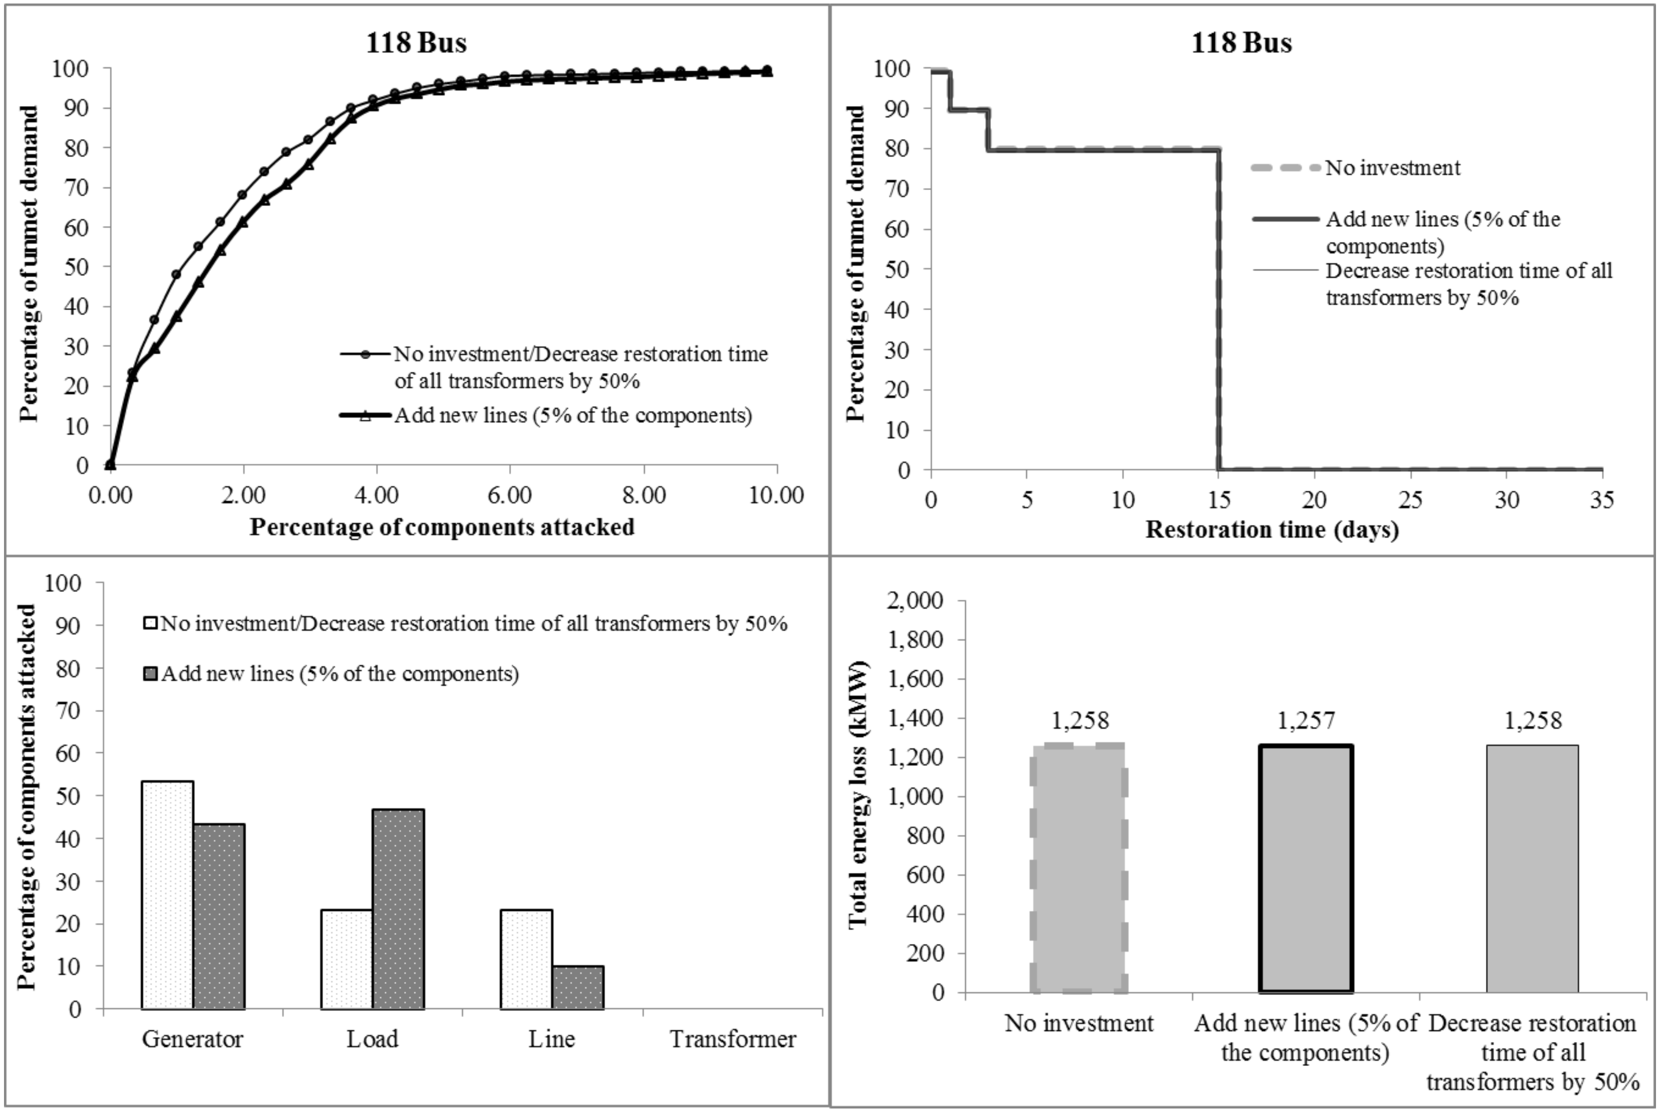

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a dynamic attacker, 300-bus system)

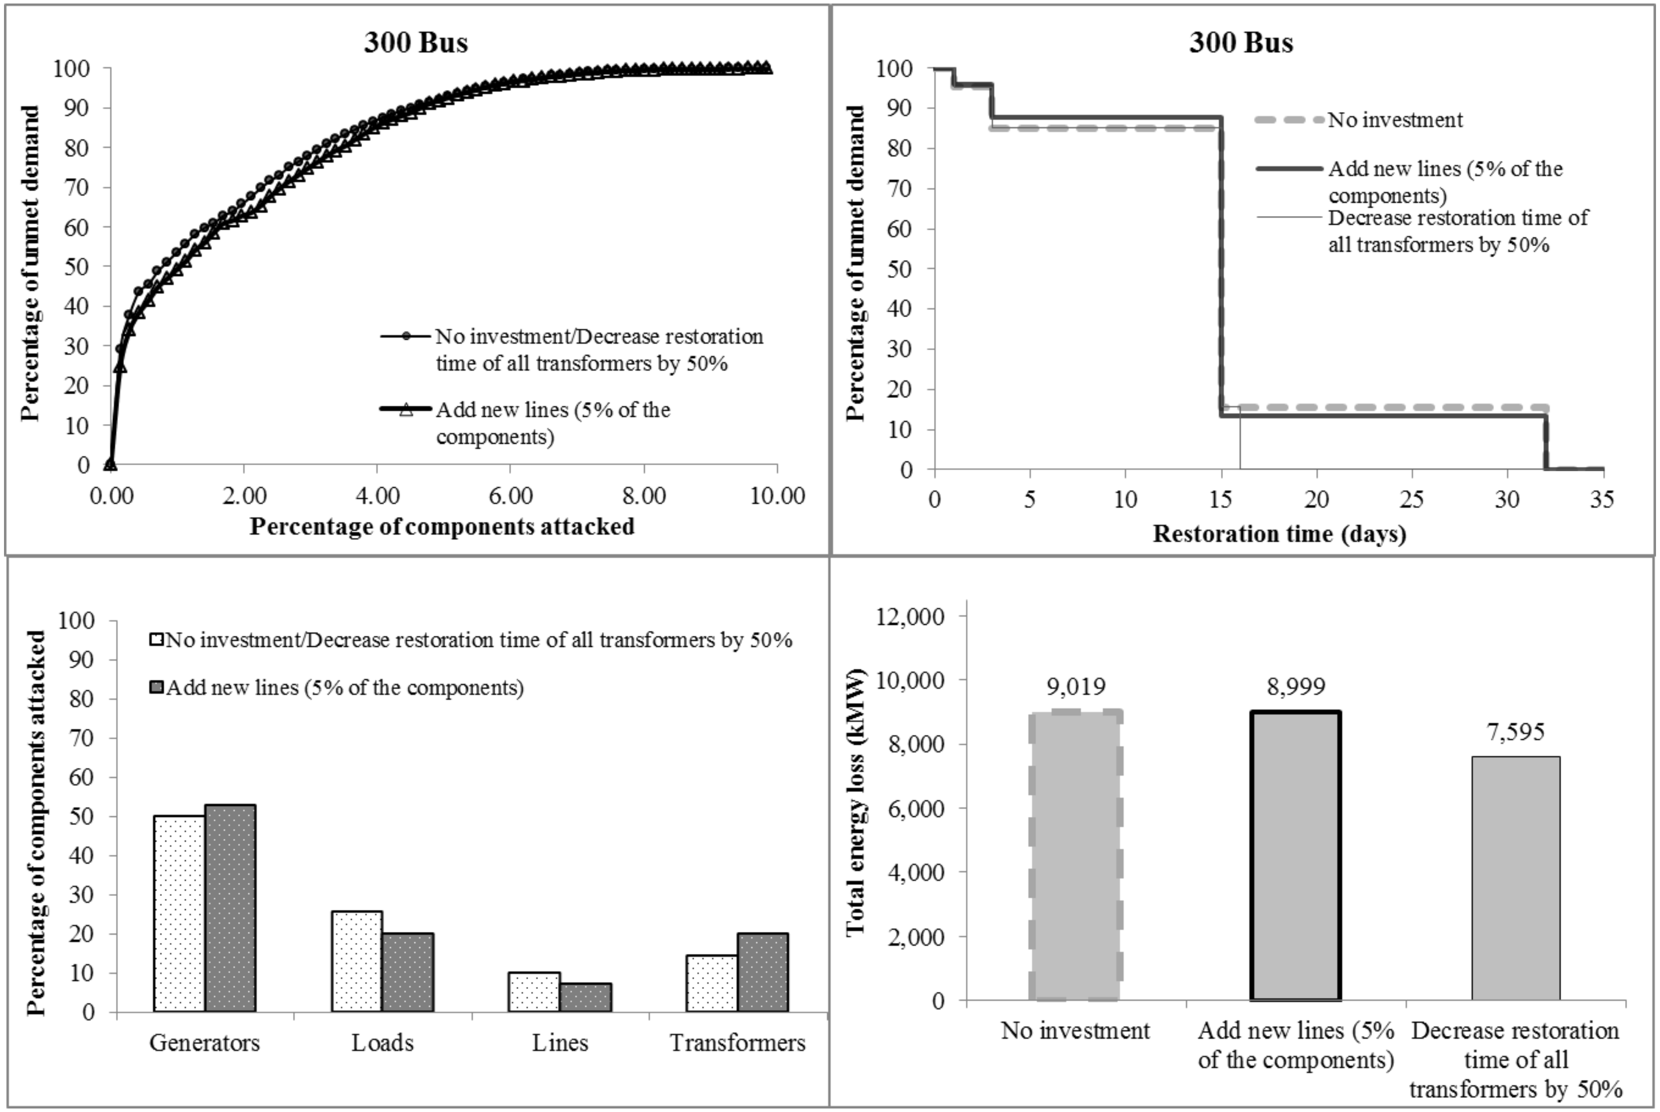

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a dynamic attacker, 24-bus system)

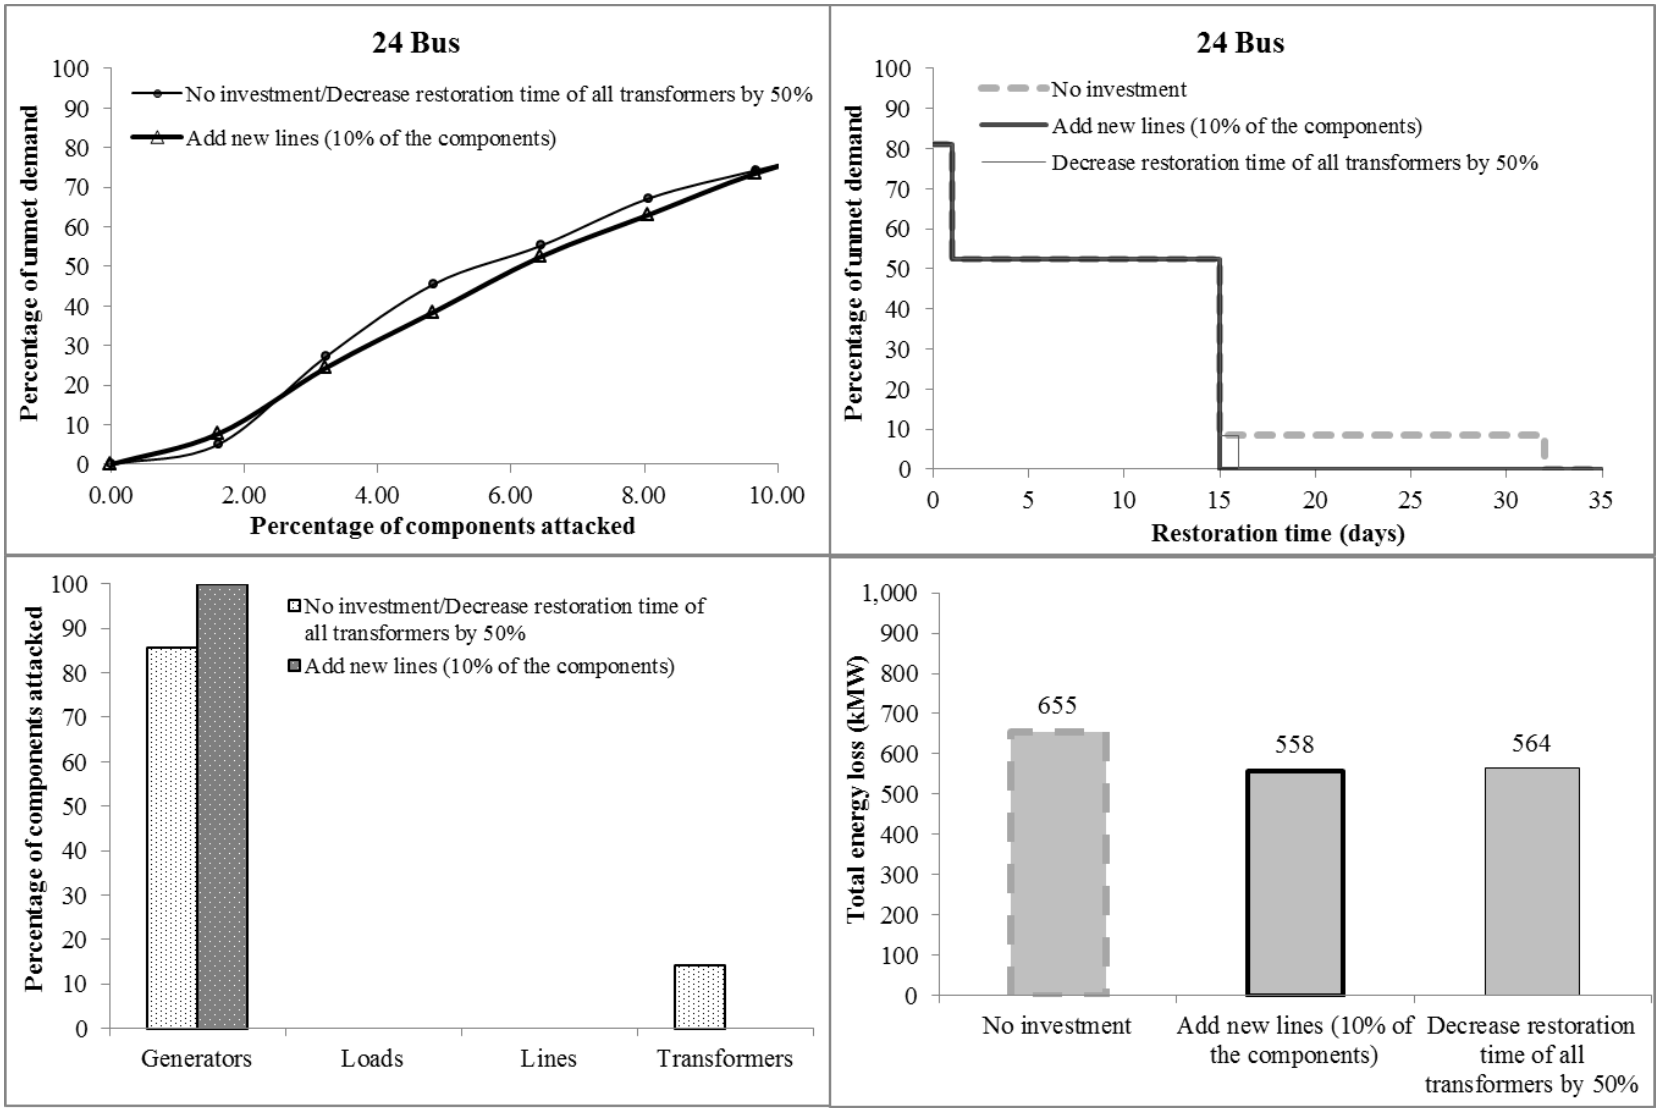

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a dynamic attacker, 48-bus system)

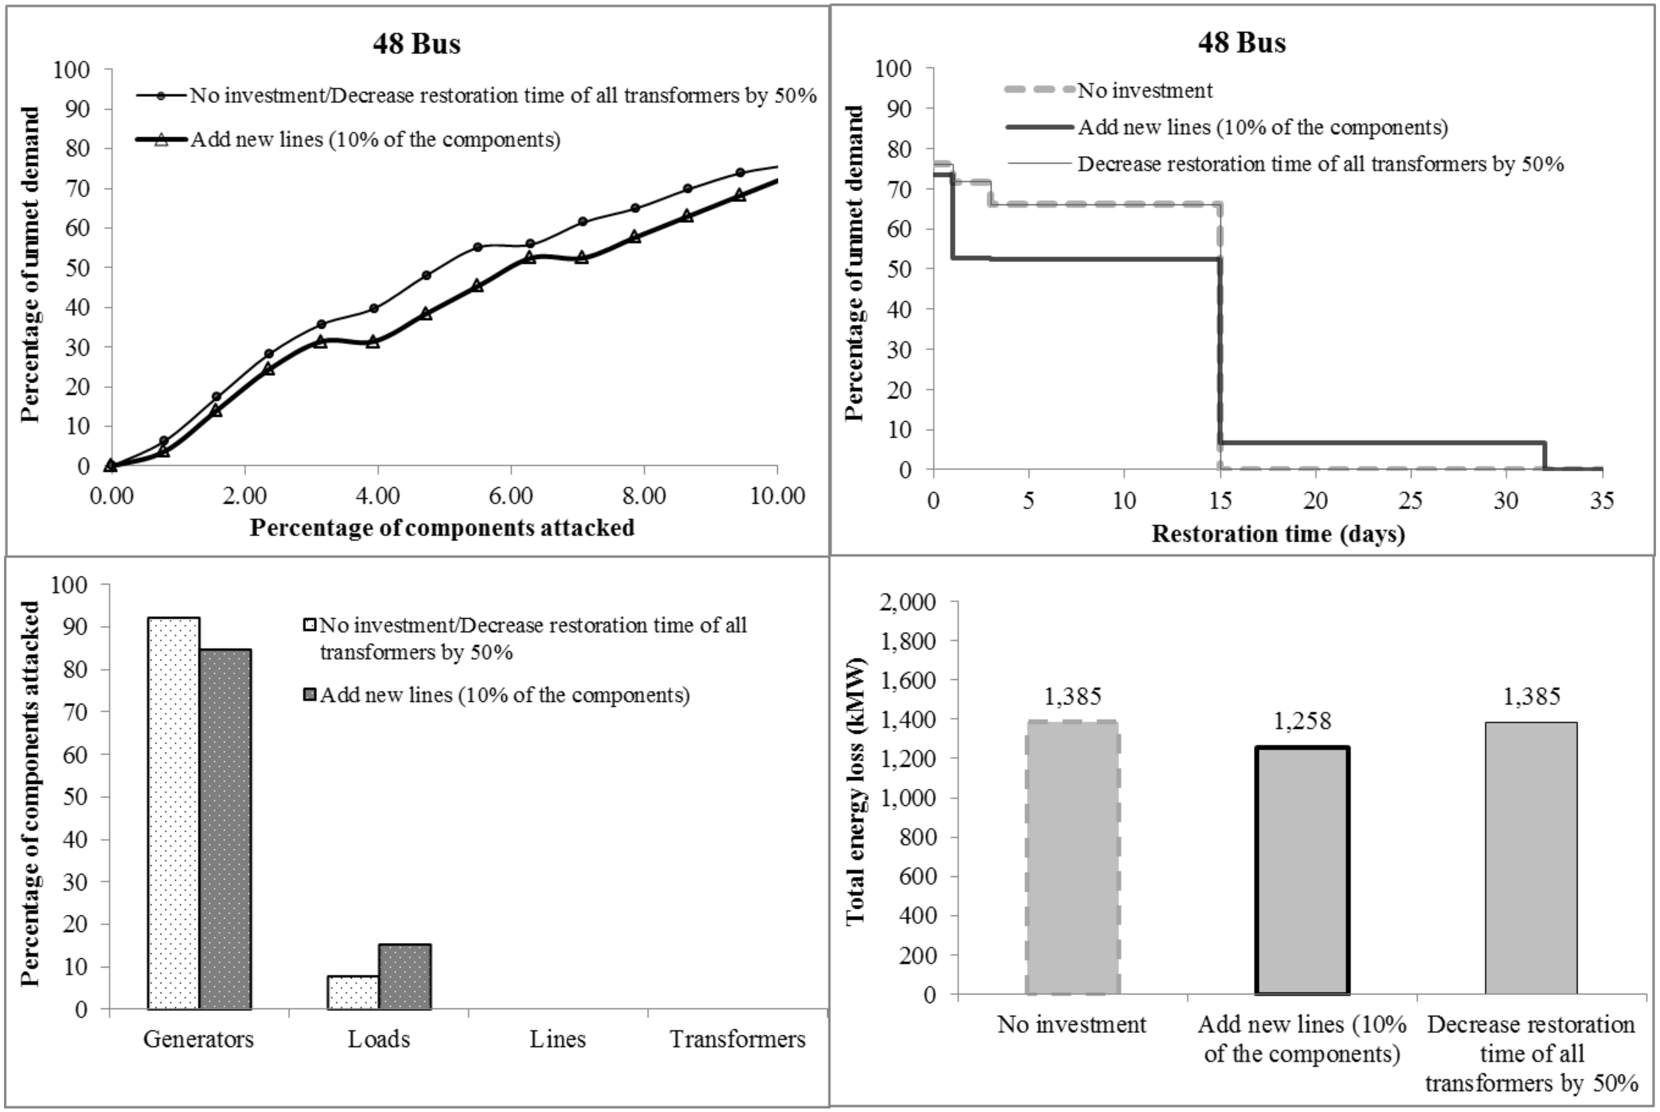

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a dynamic attacker, 118-bus system)

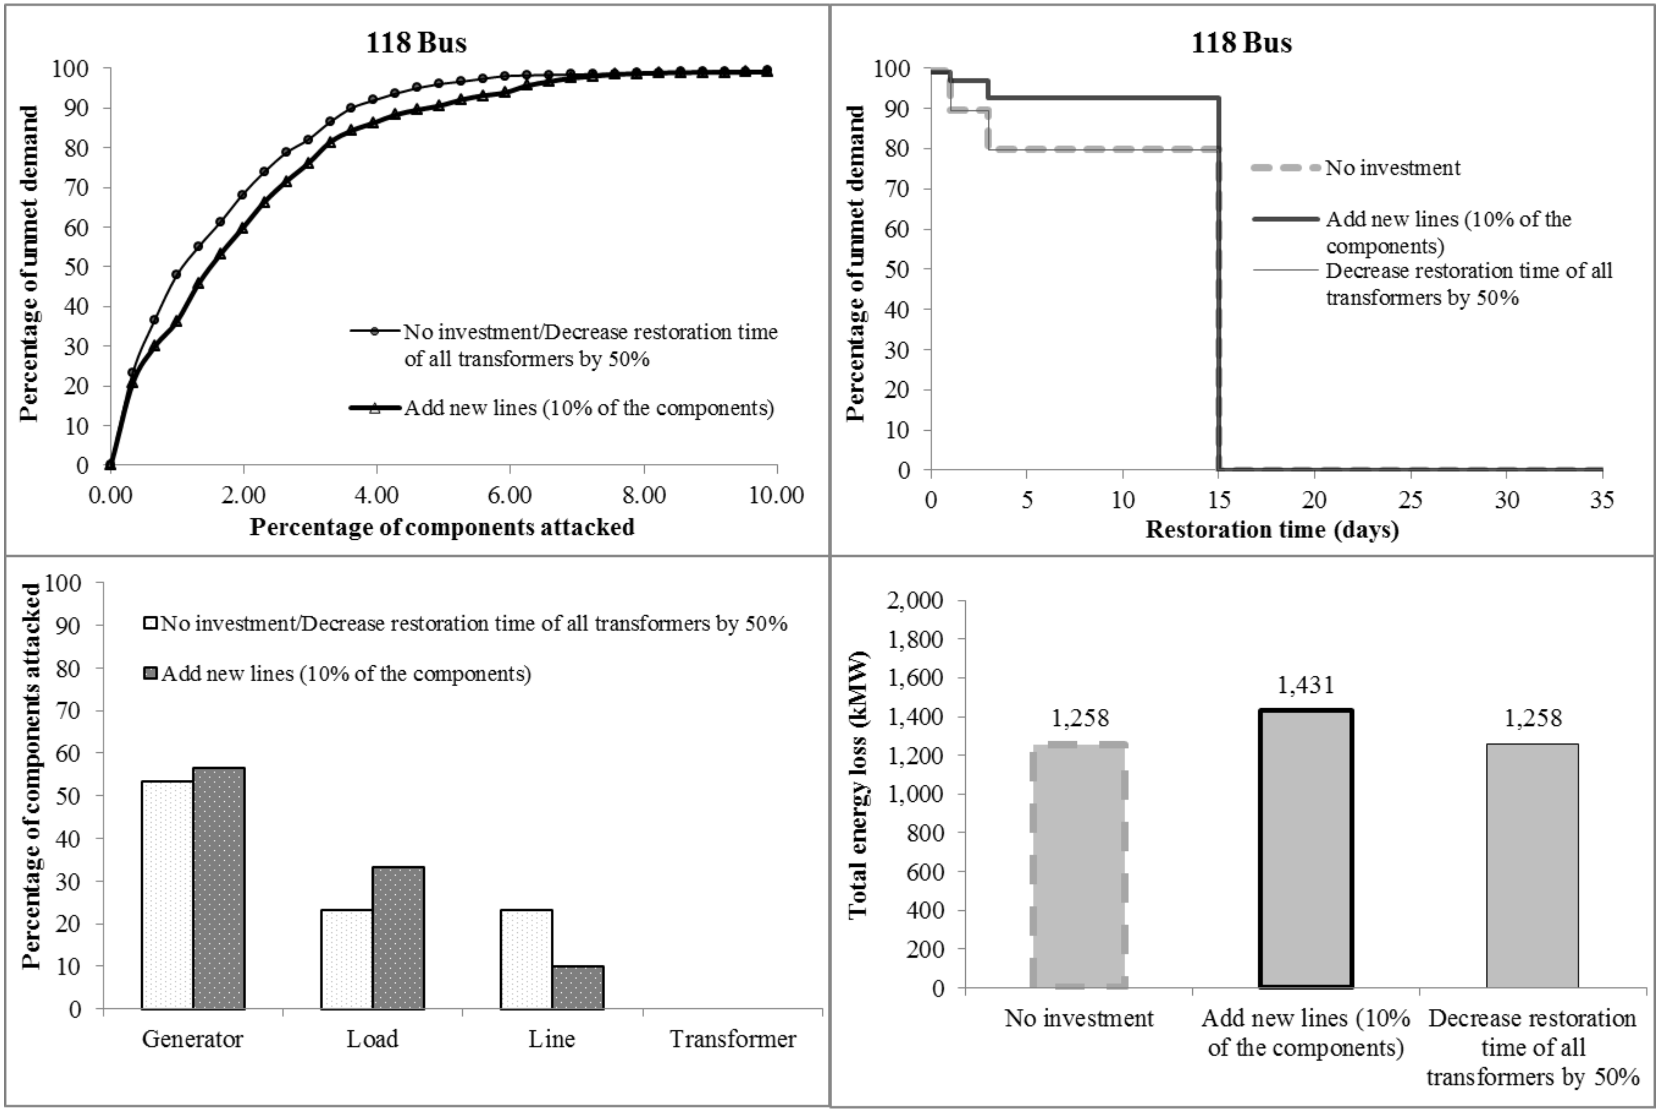

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a dynamic attacker, 300-bus system)

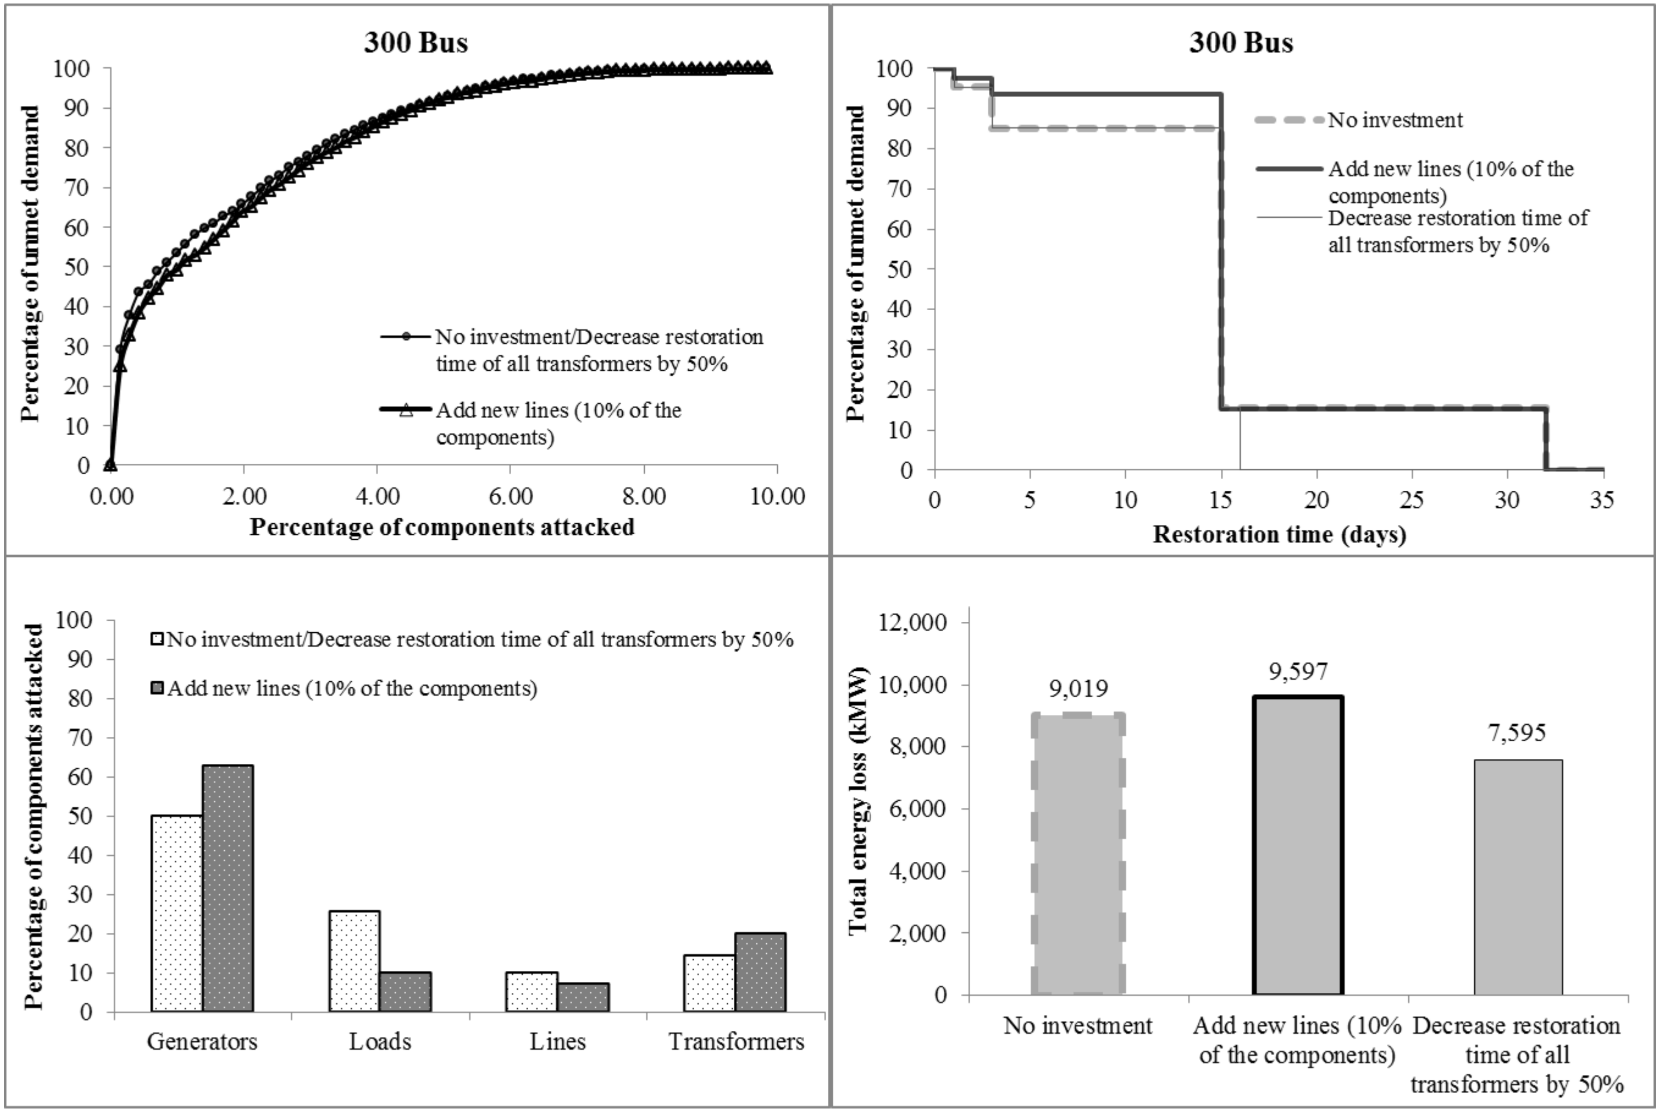

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 24-bus system)

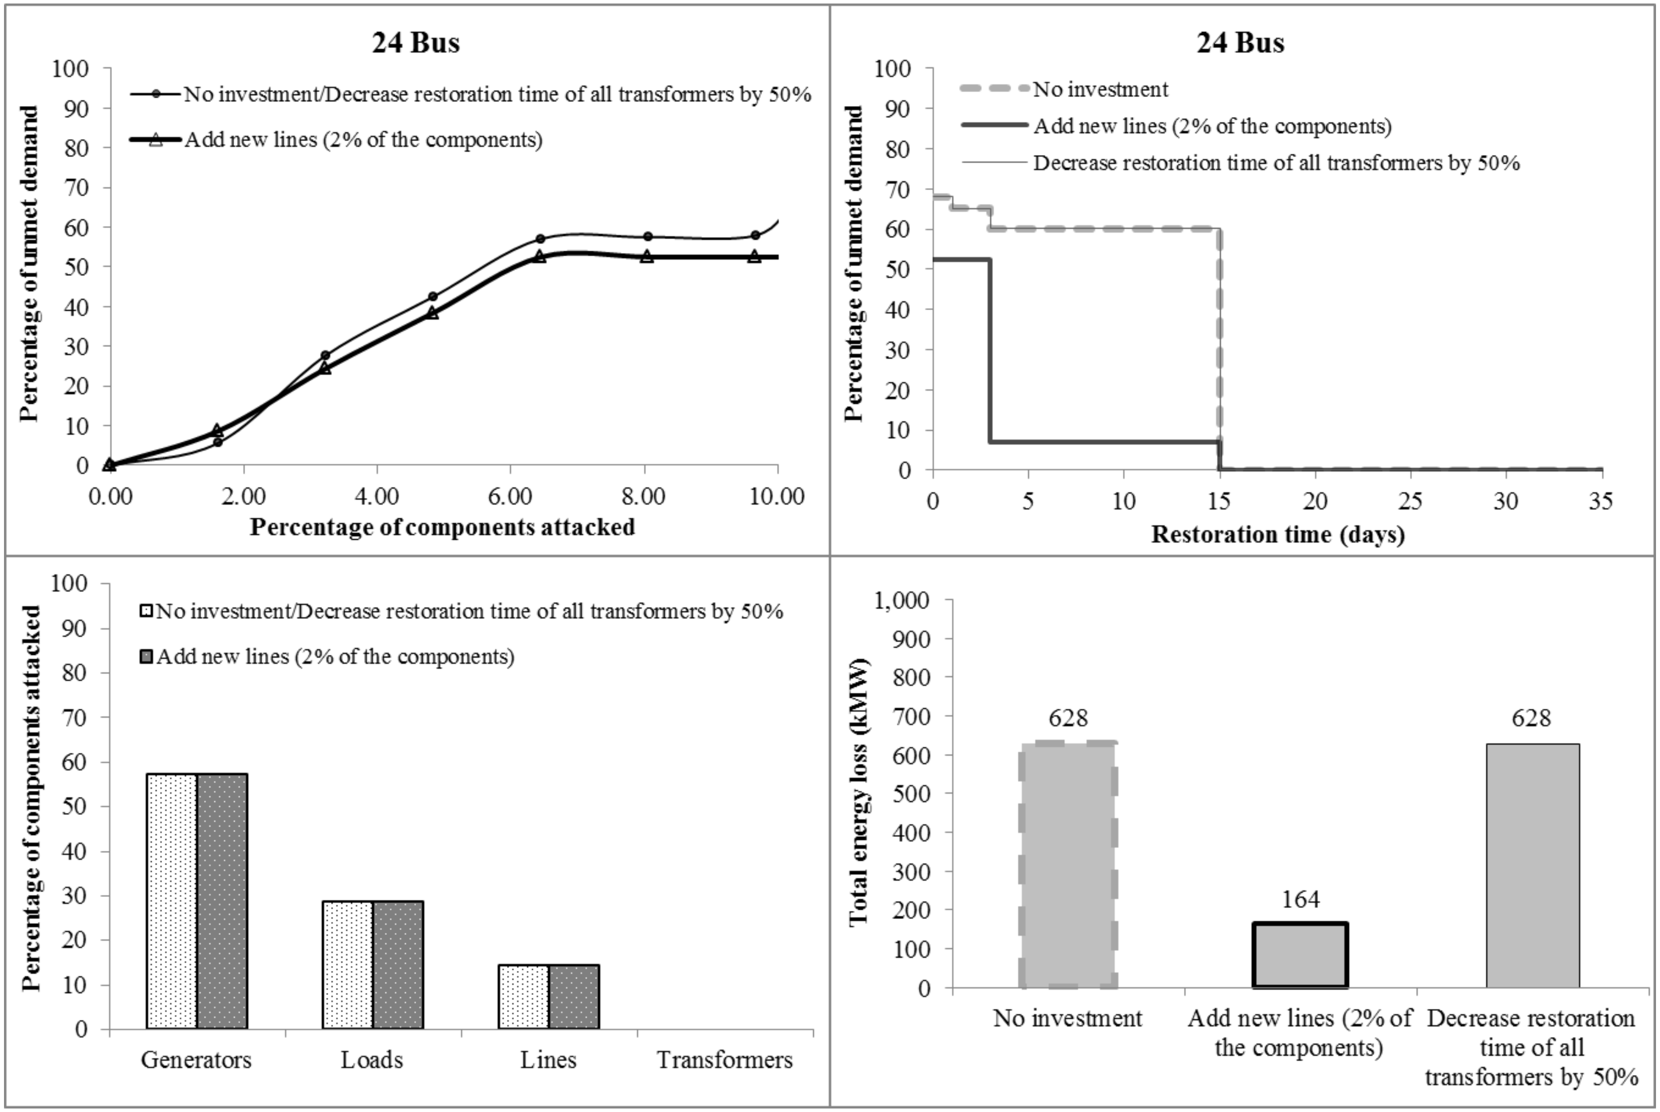

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 48-bus system)

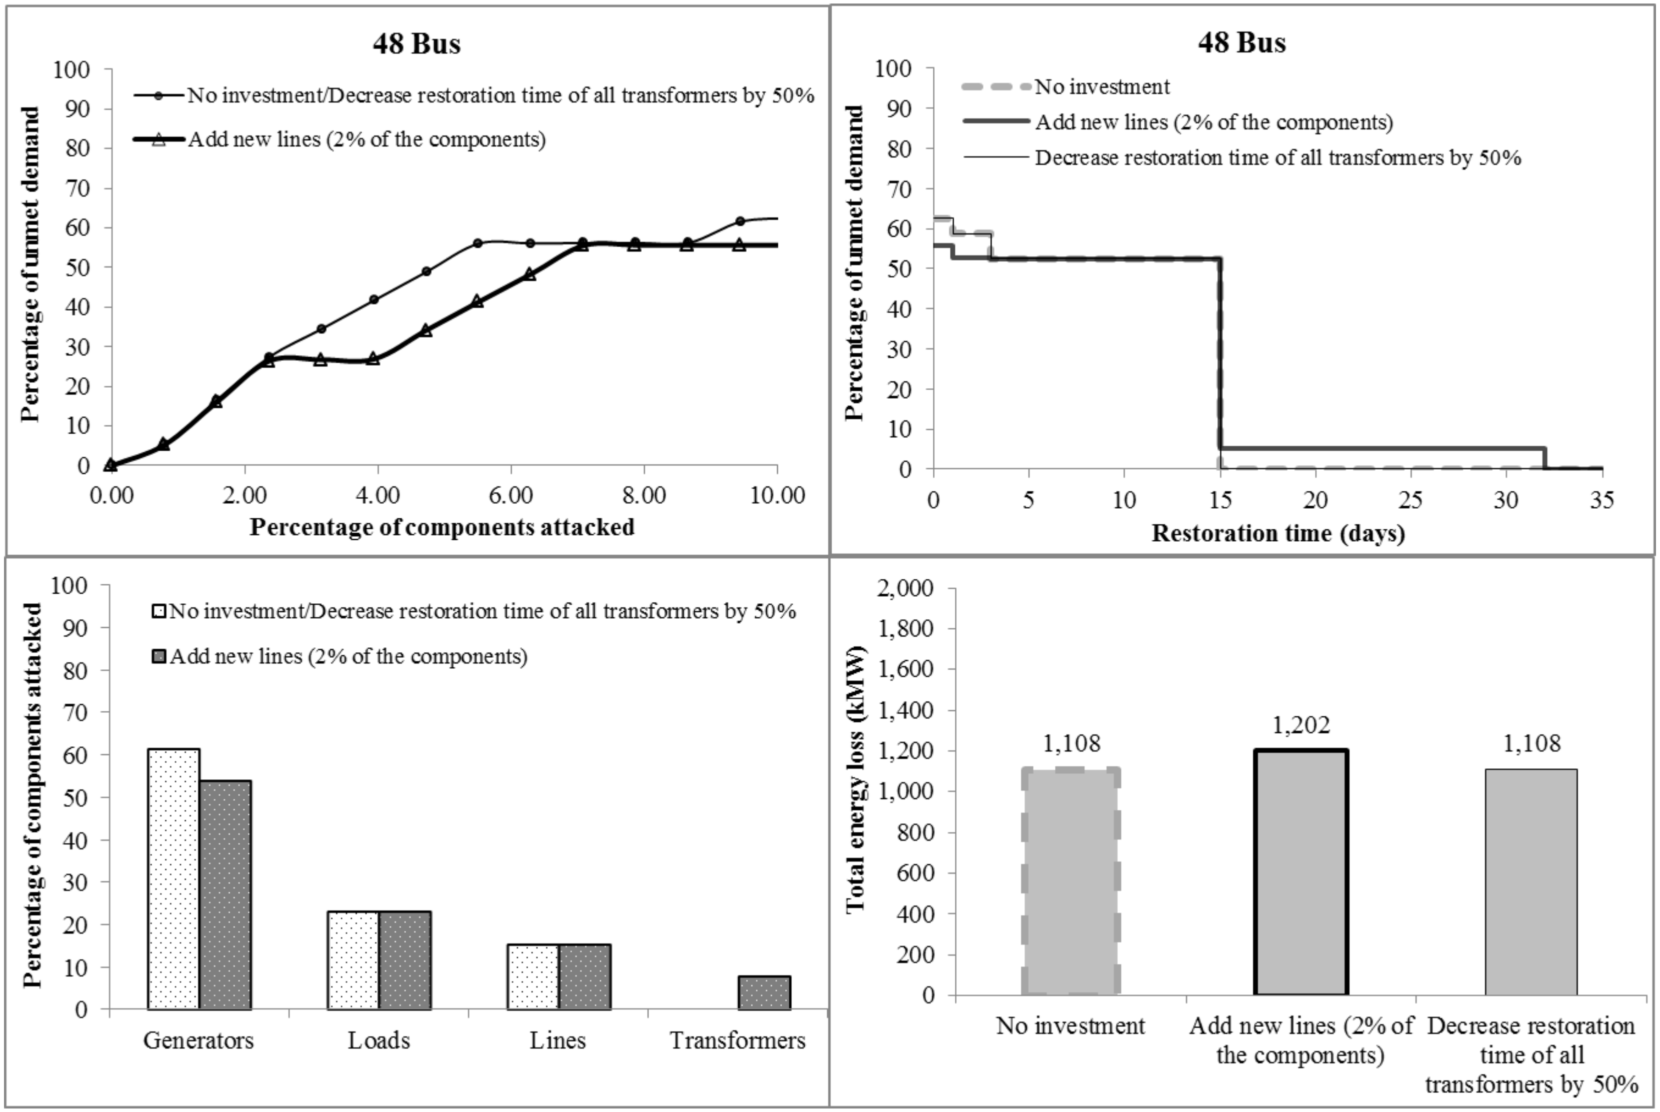

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 118-bus system)

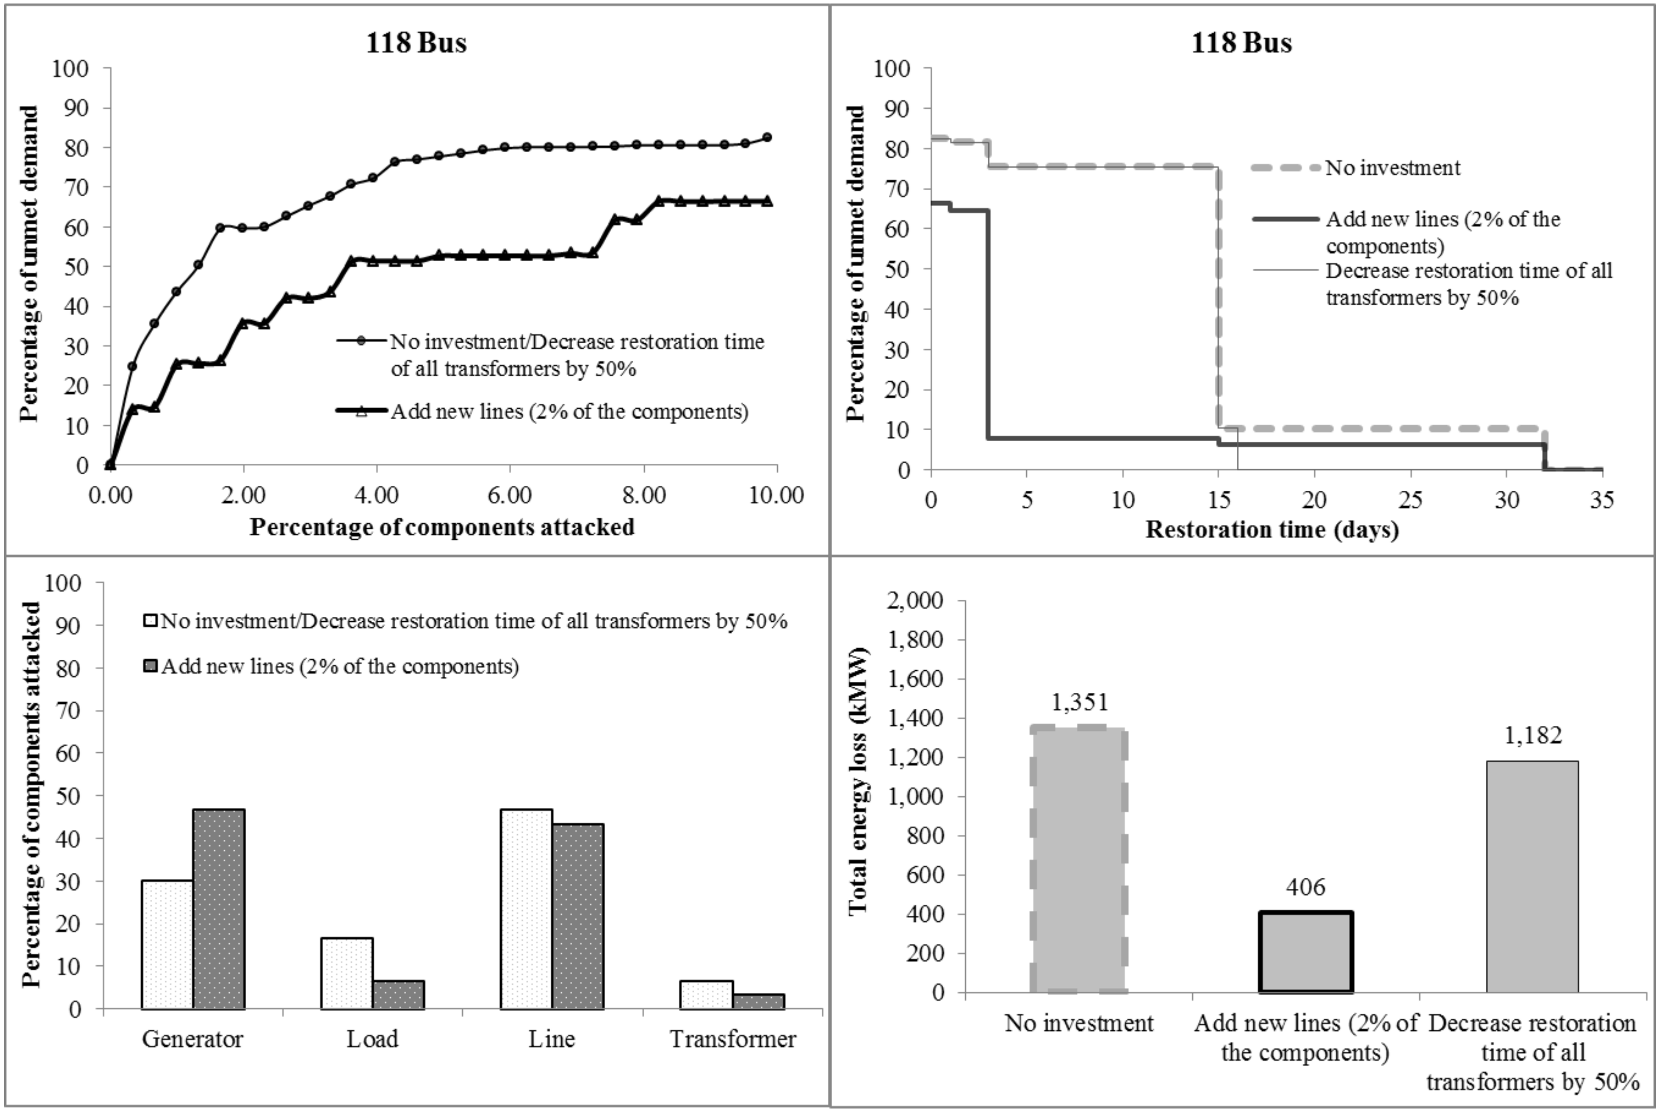

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 300-bus system)

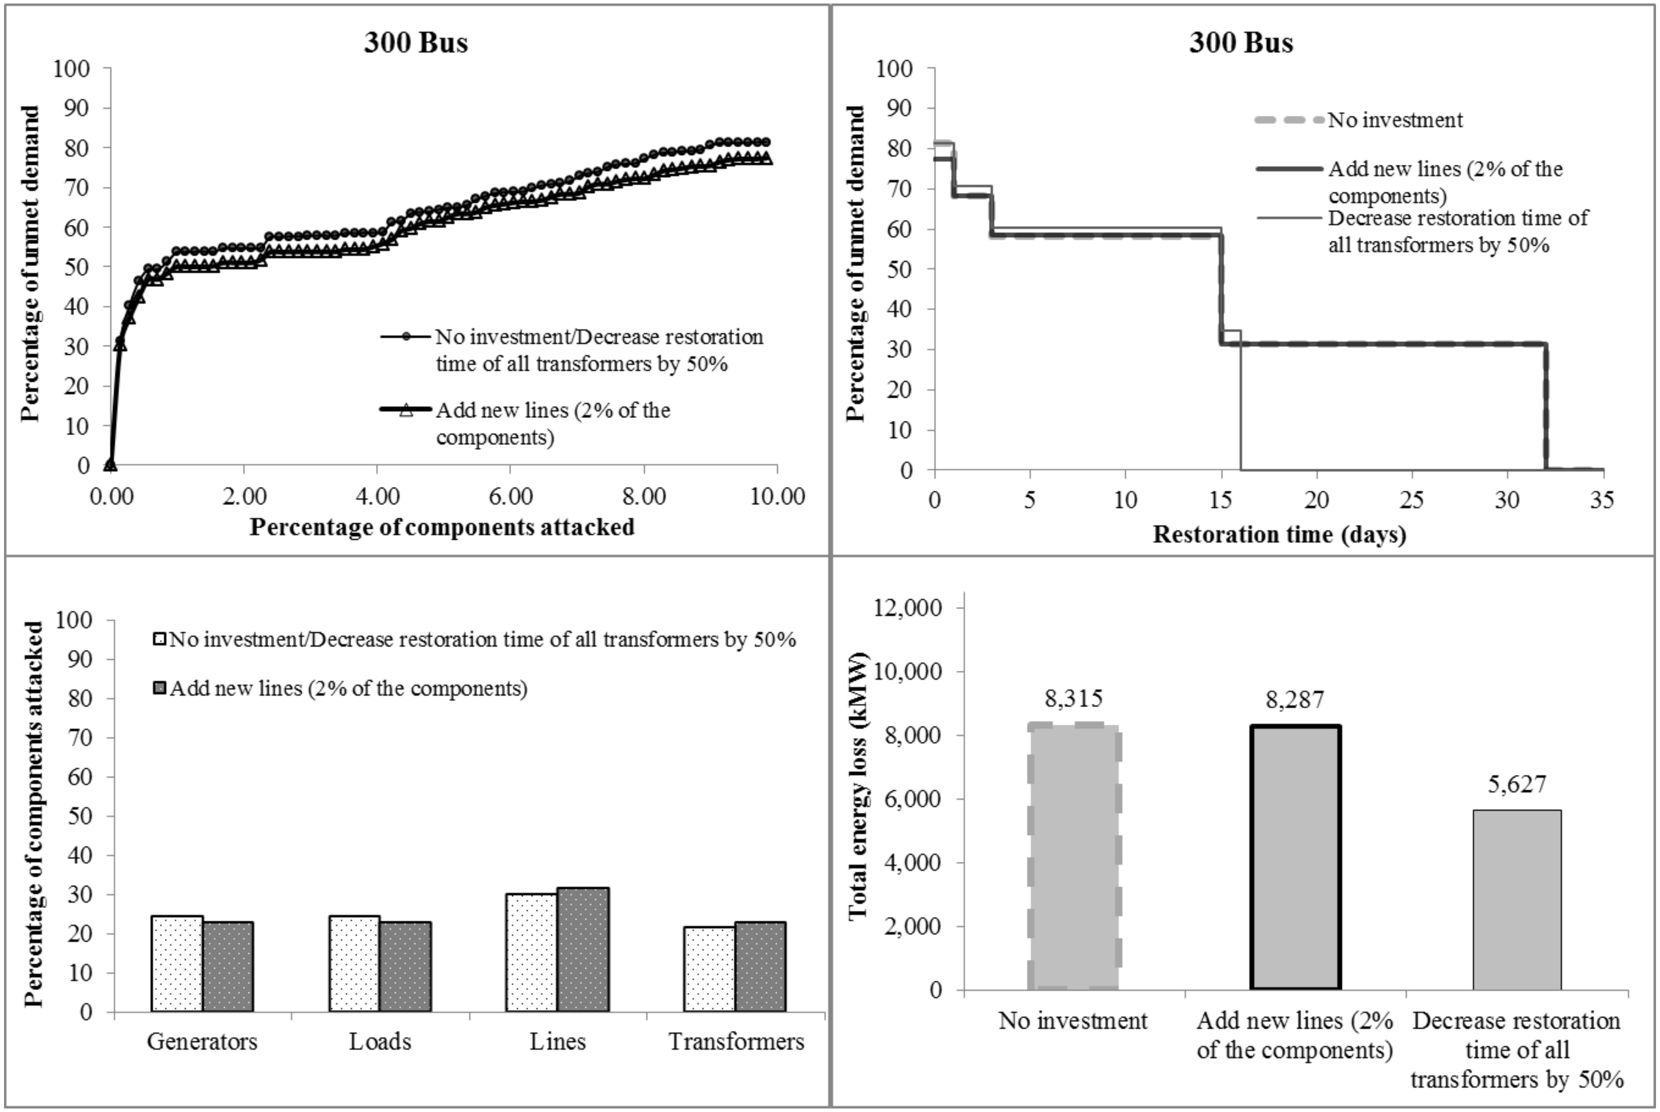

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 24-bus system)

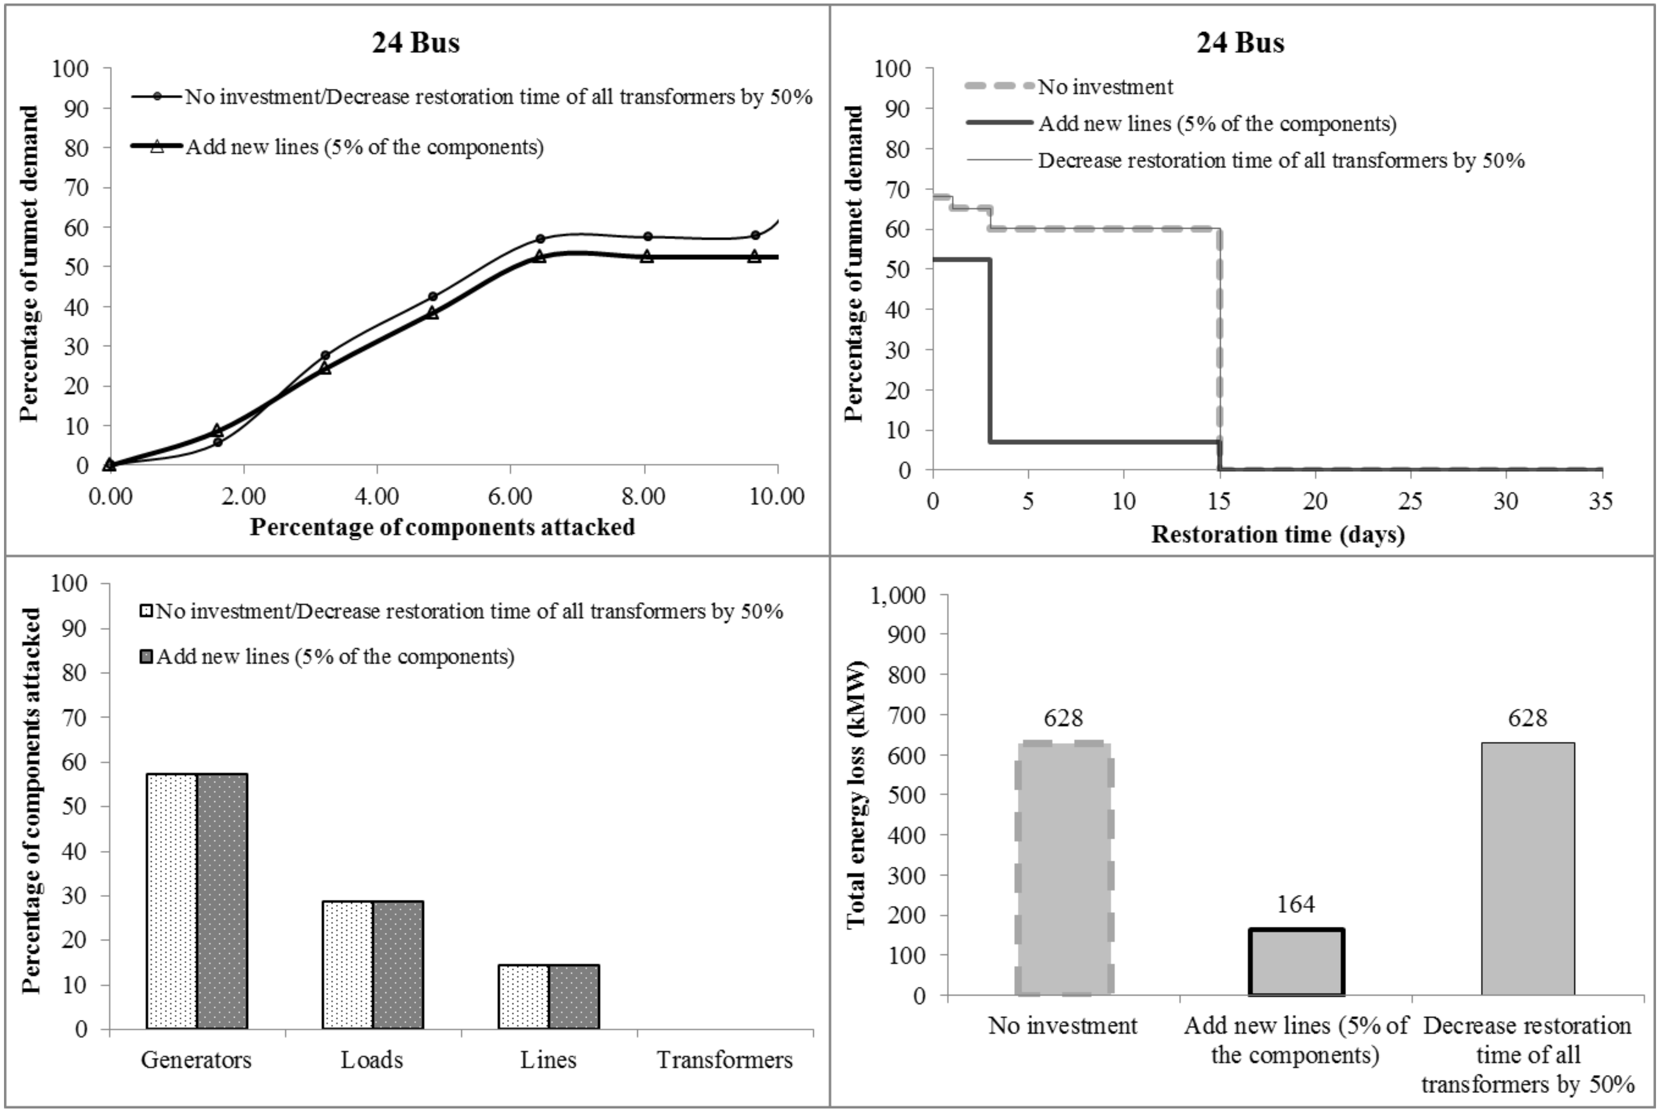

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 48-bus system)

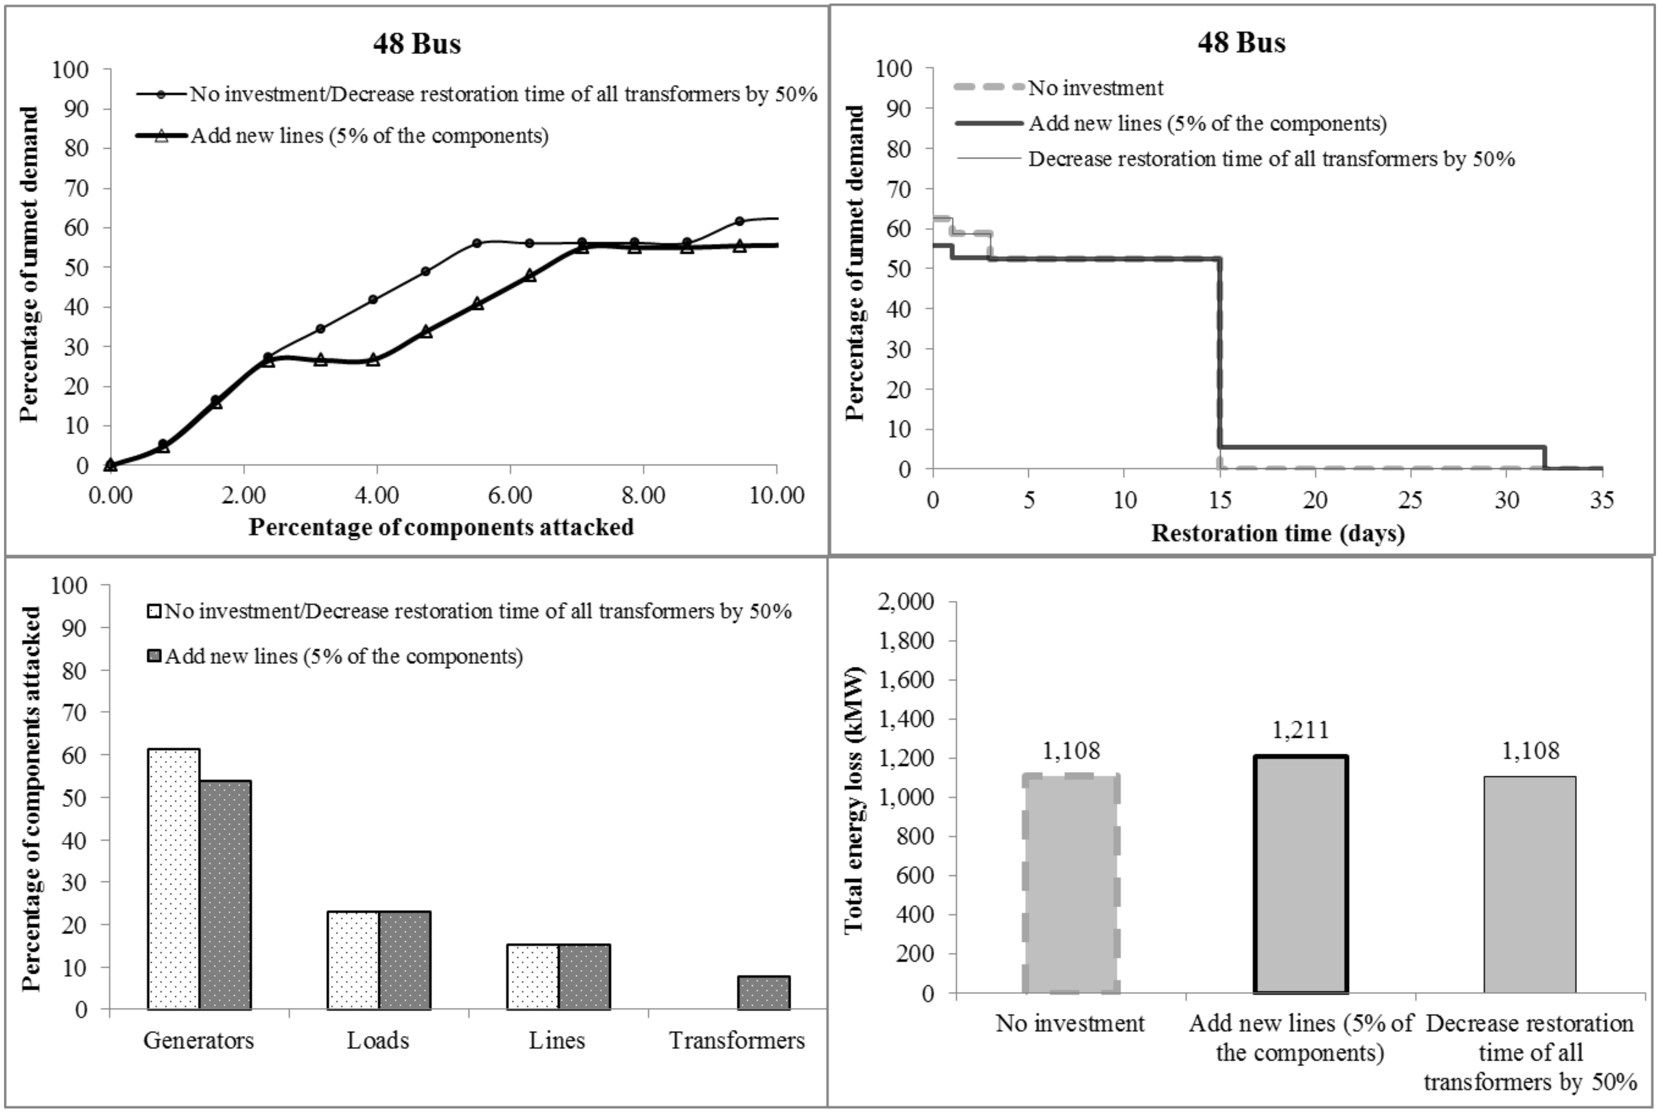

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 118-bus system)

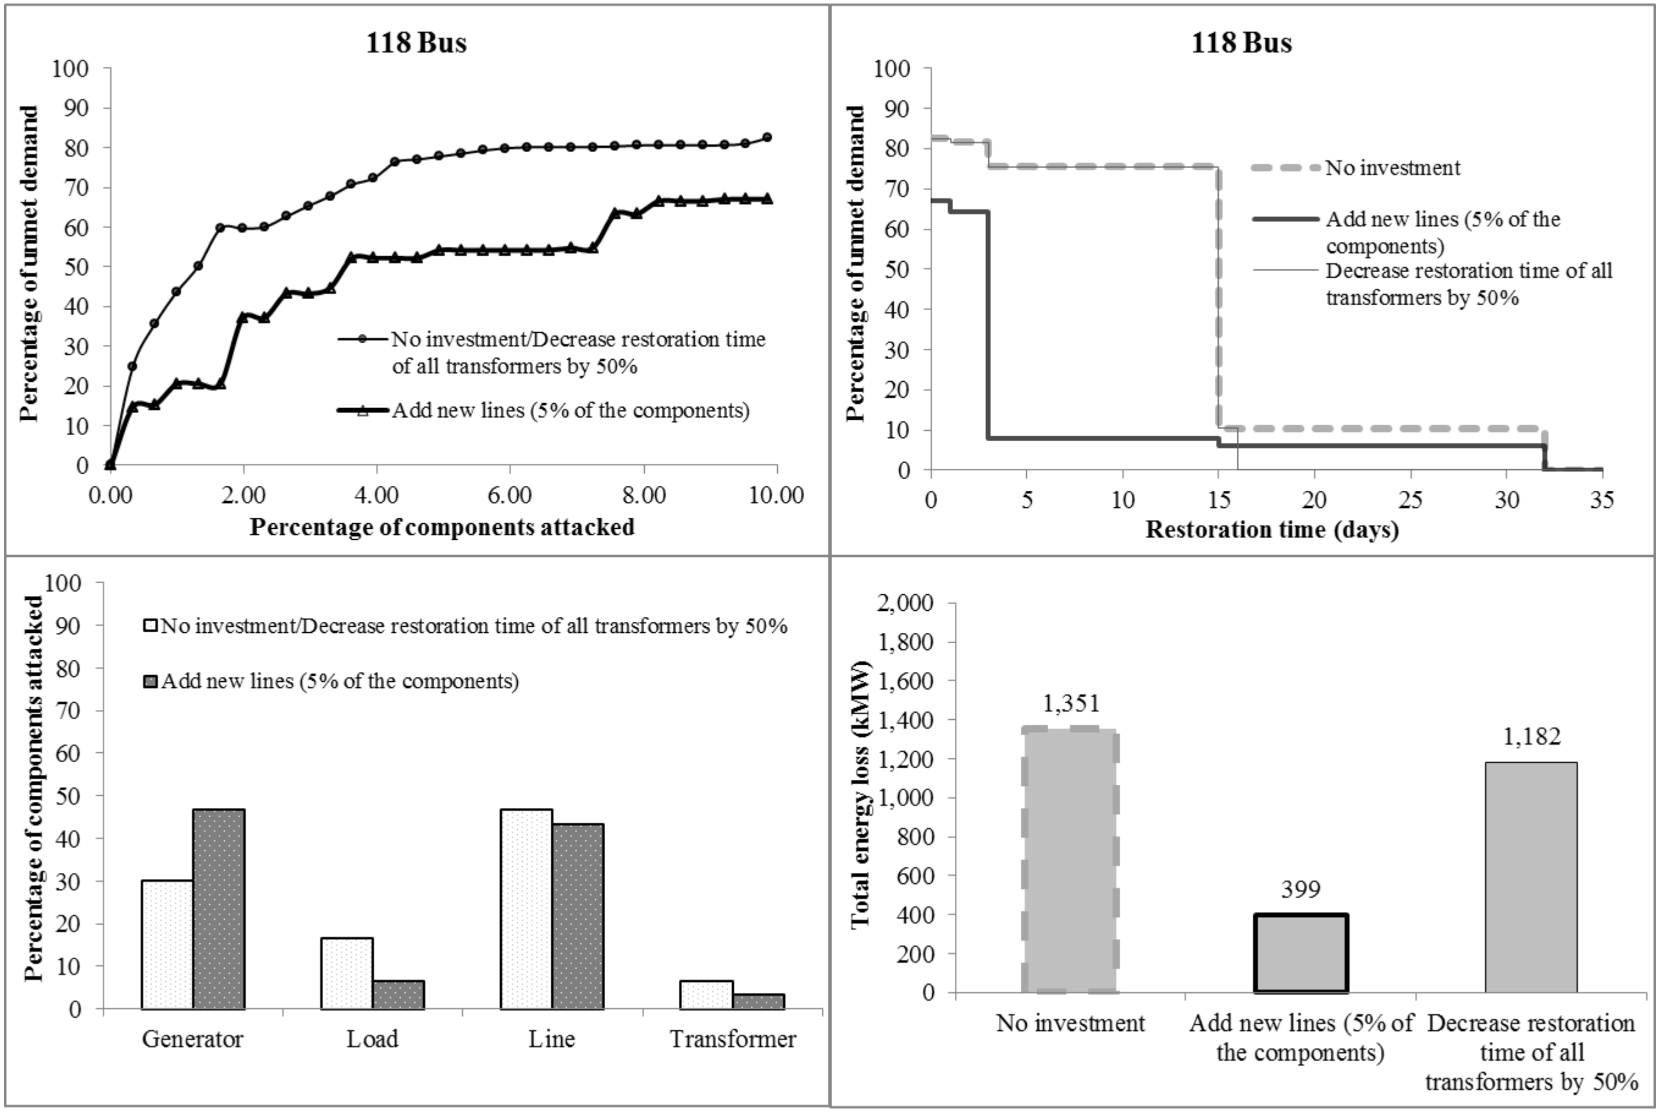

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 300-bus system)

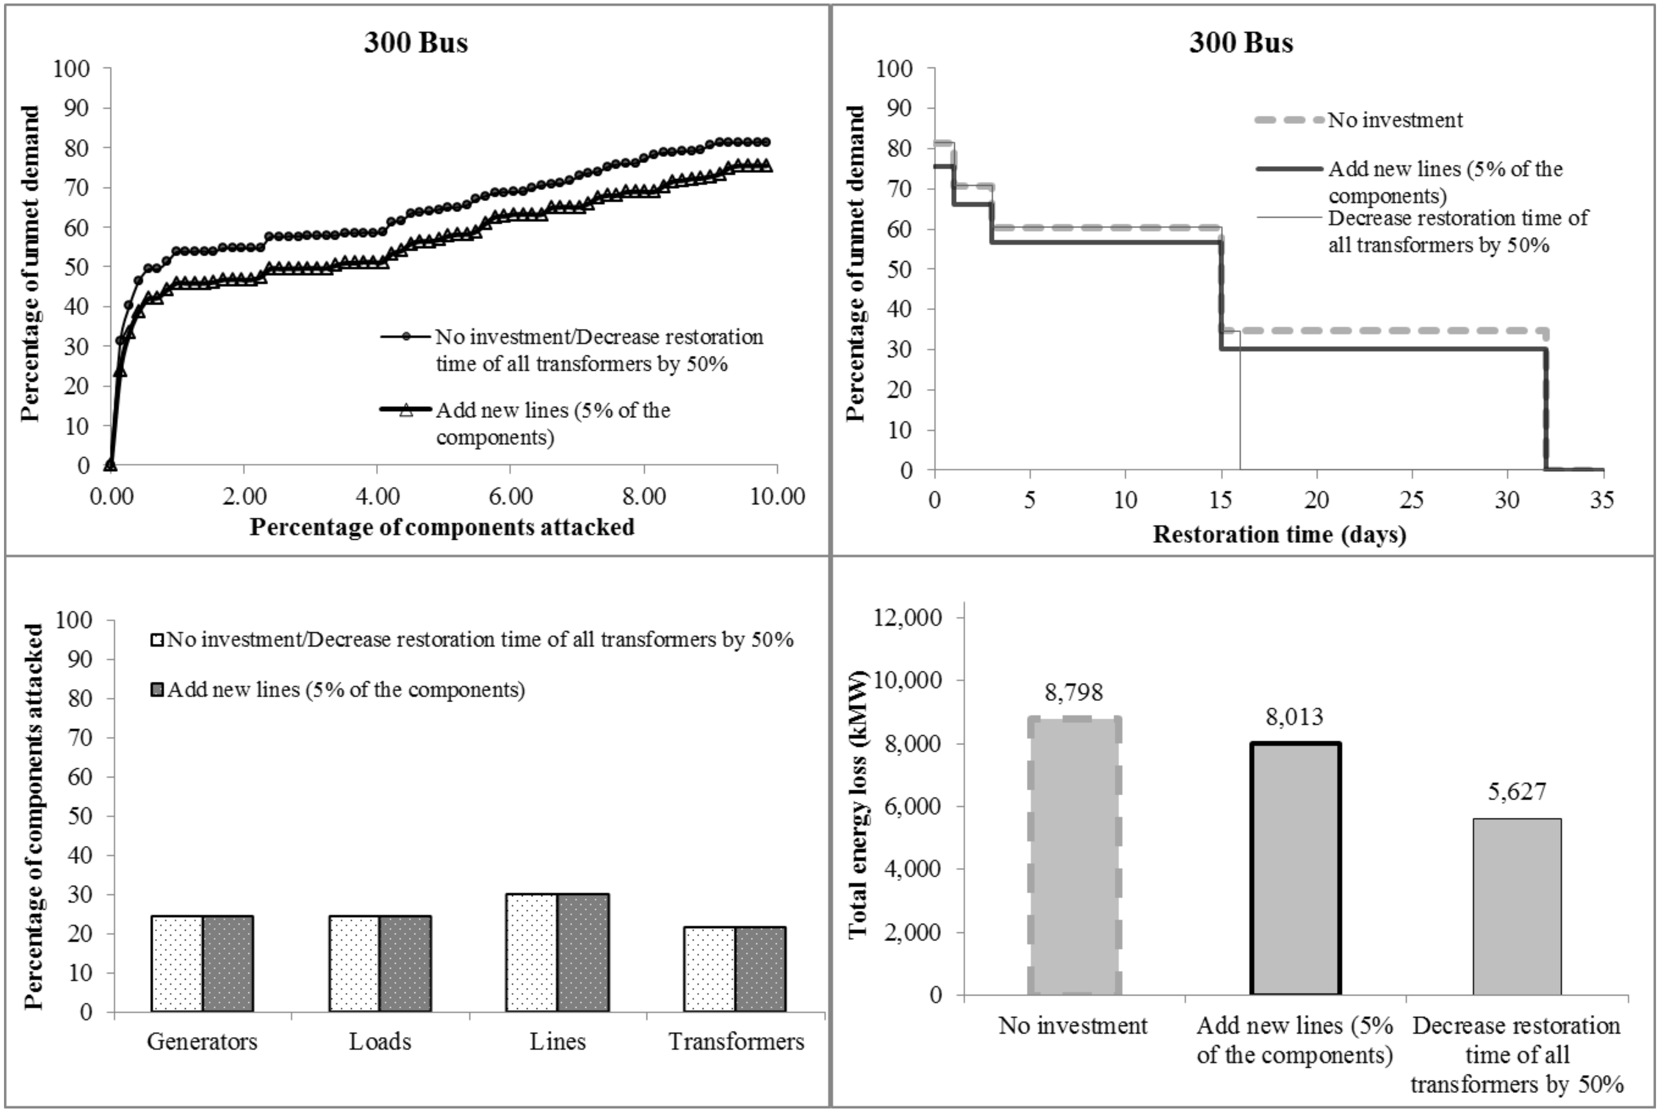

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 24-bus system)

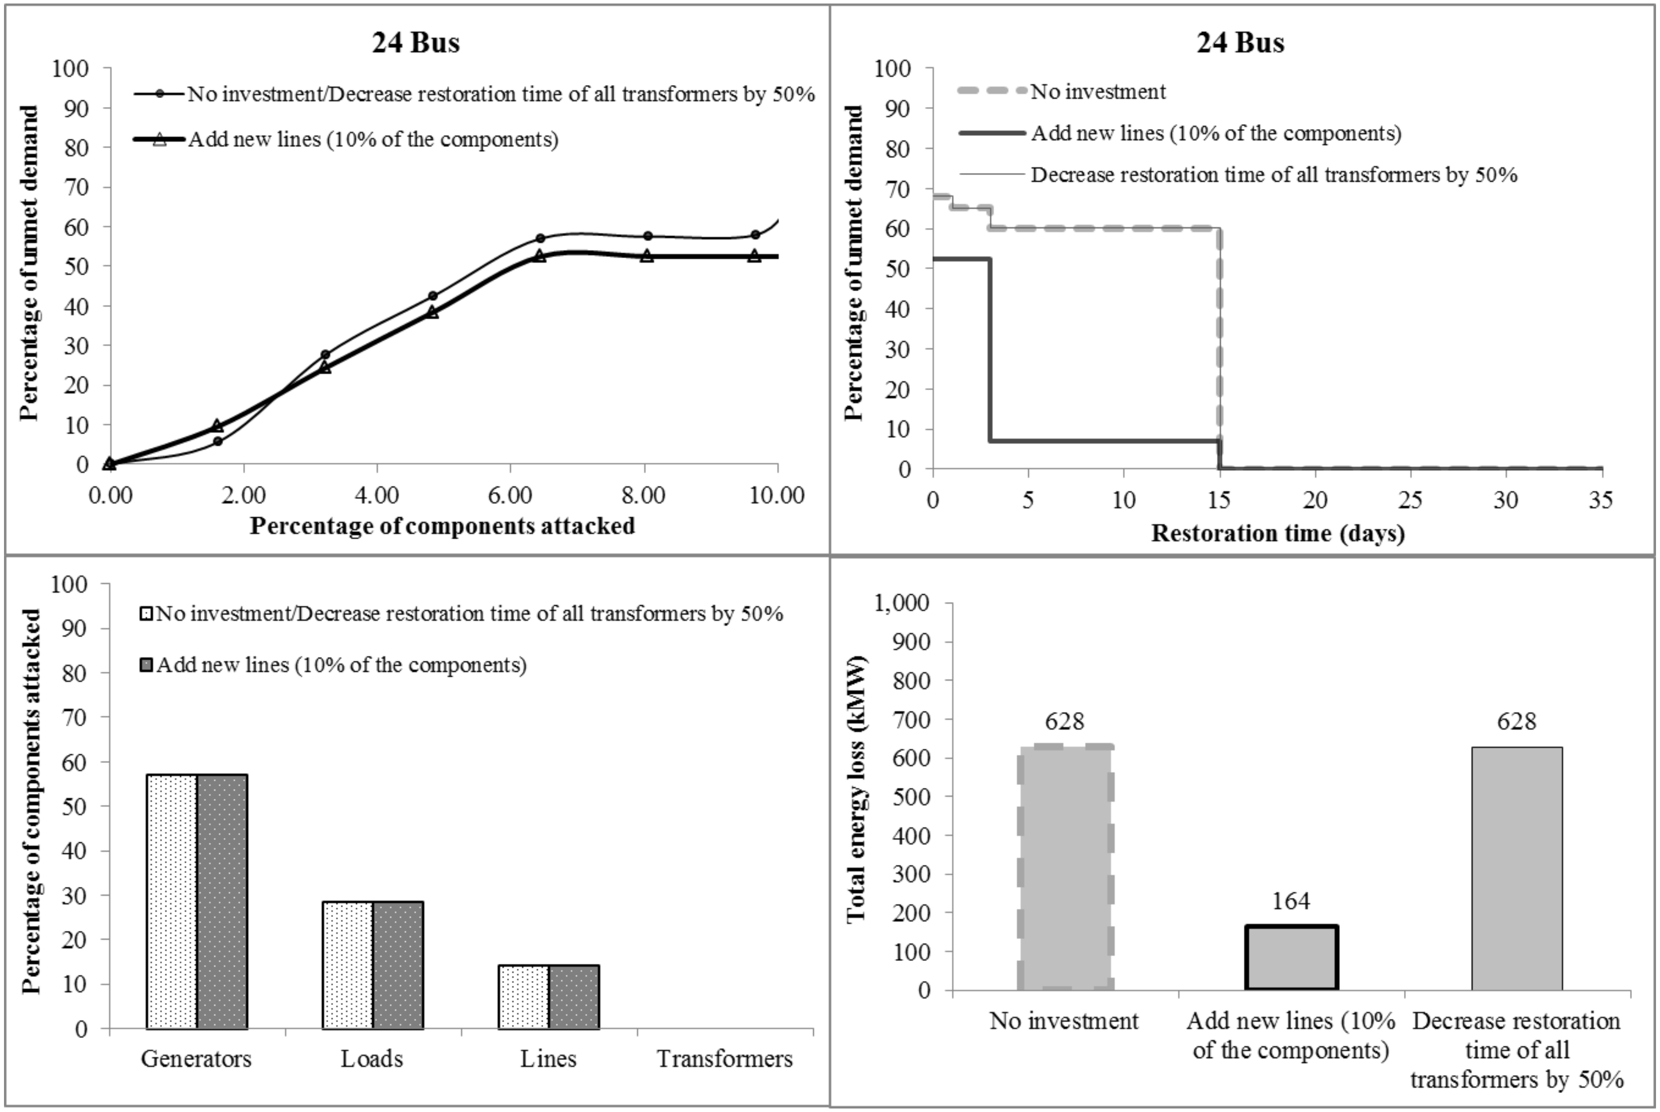

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 48-bus system)

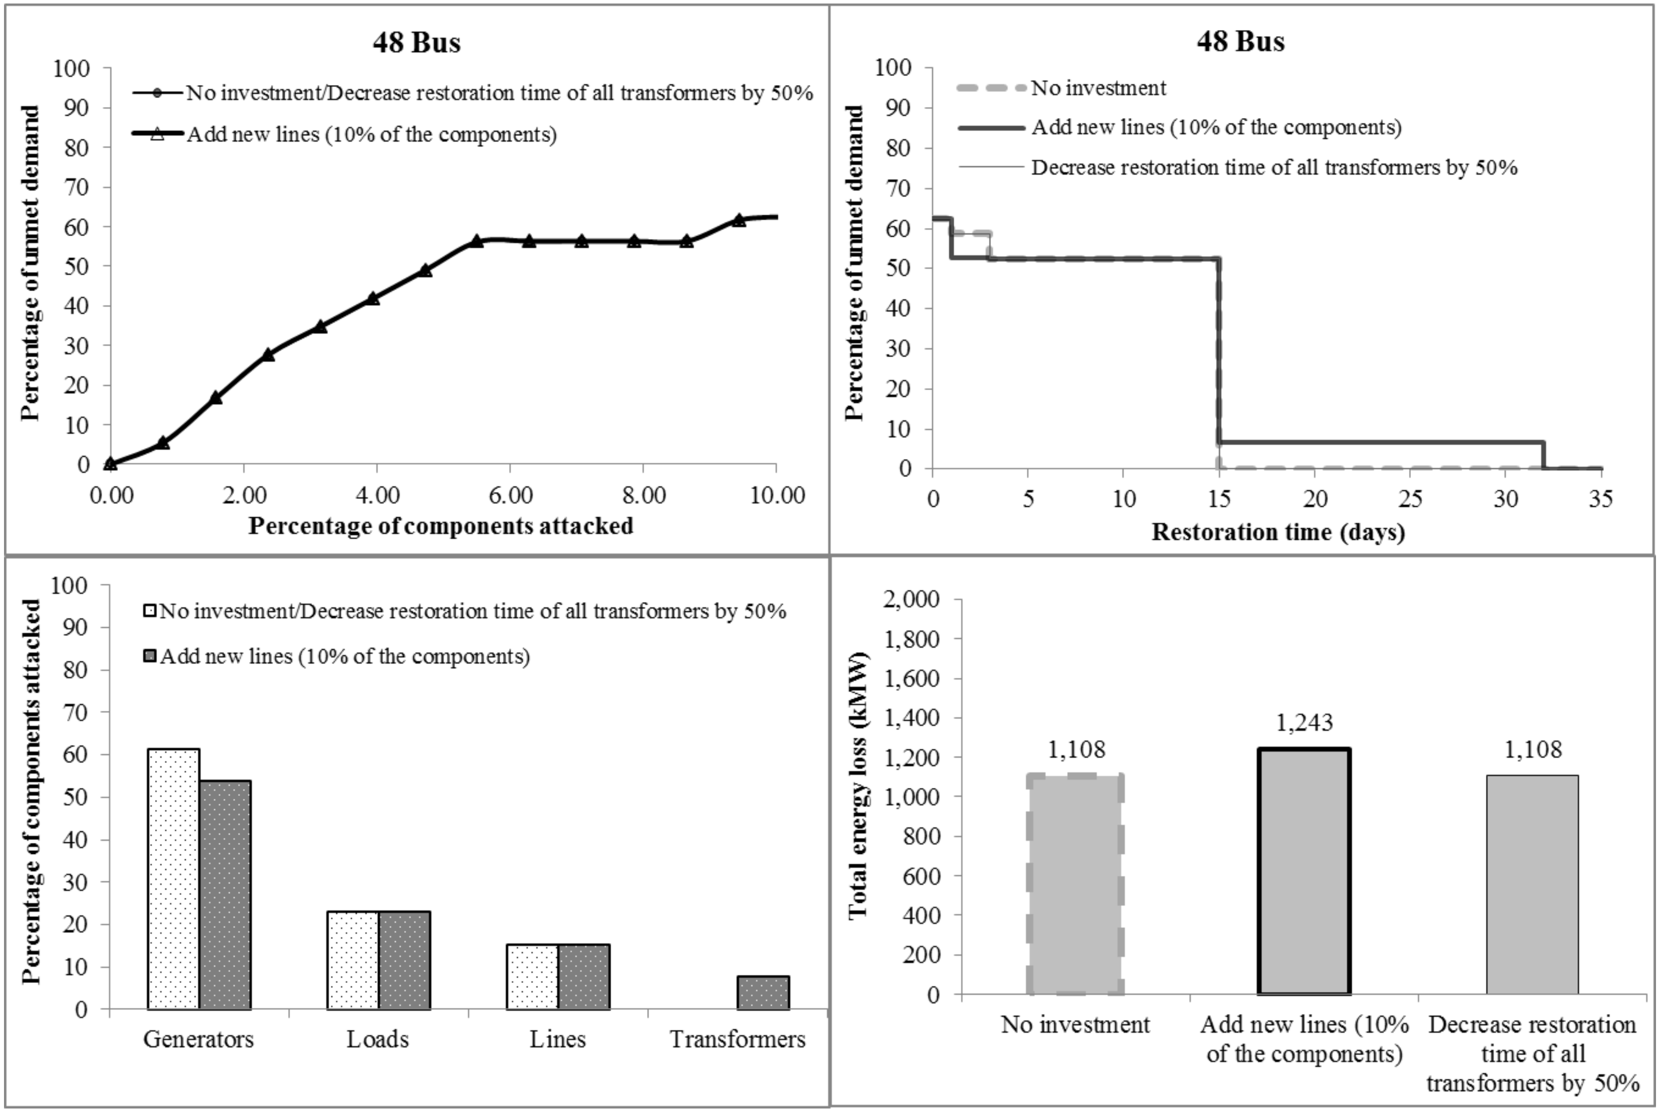

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 118-bus system)

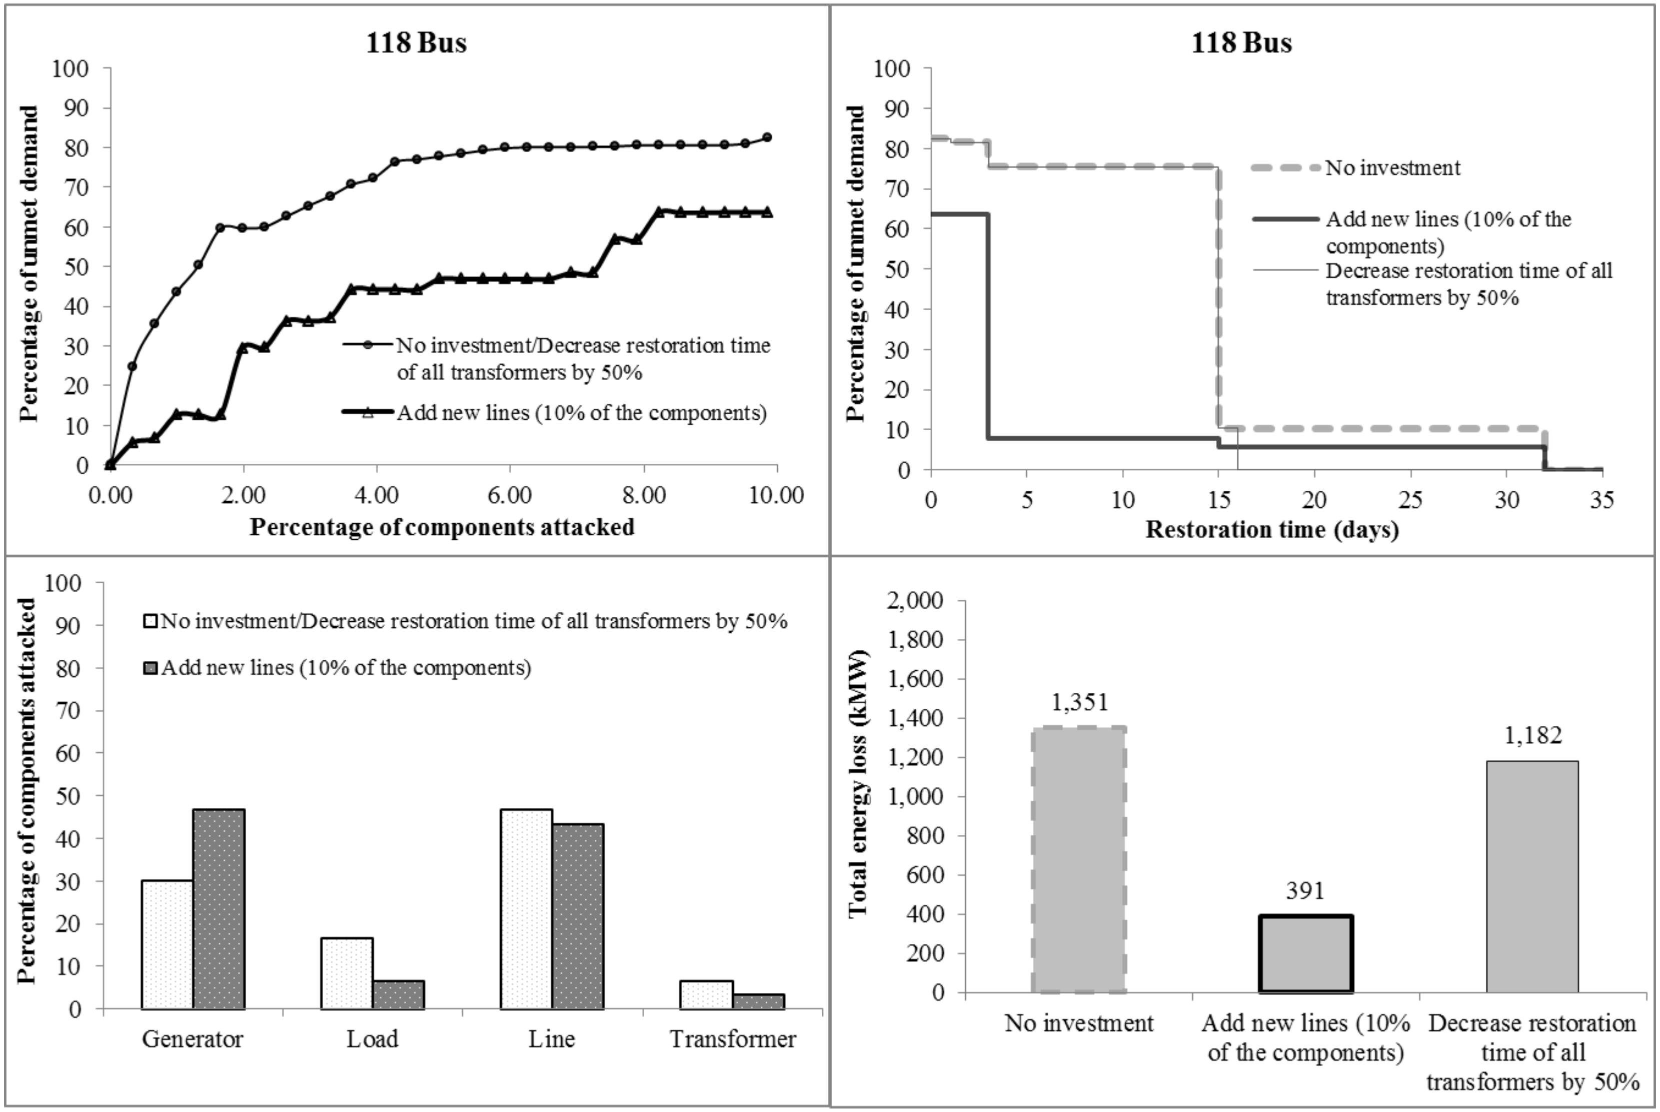

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 300-bus system)

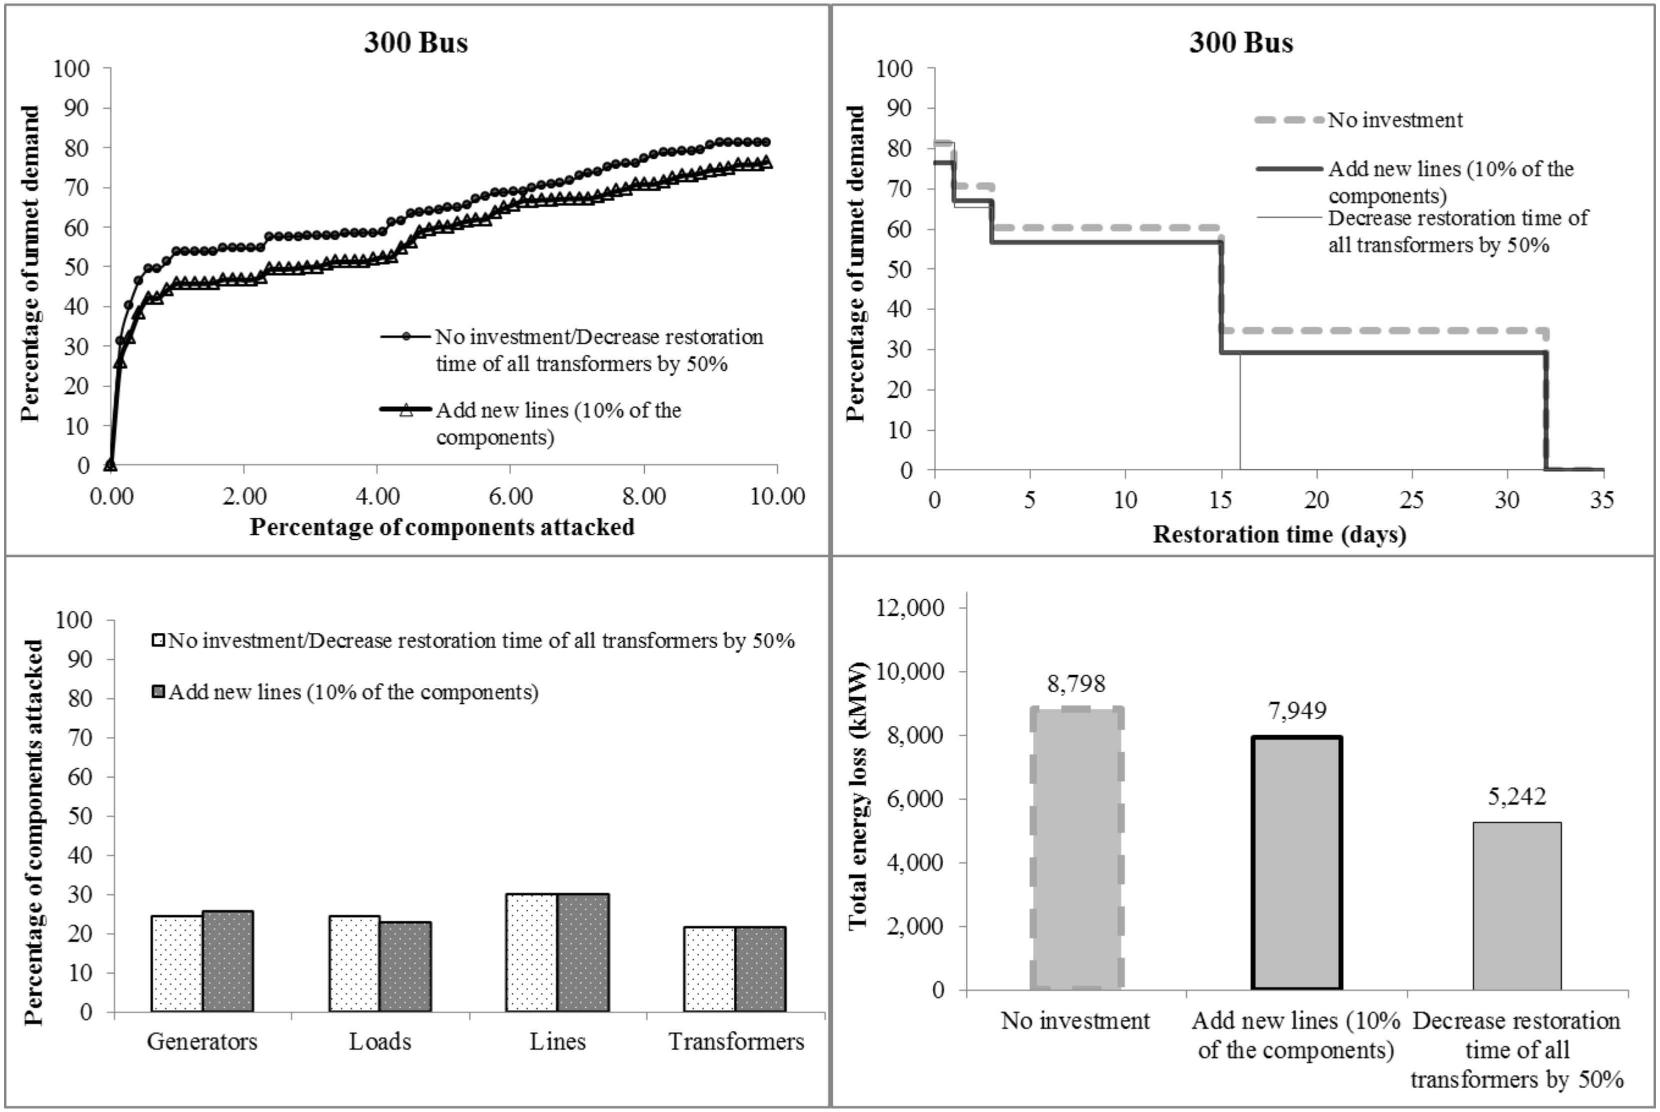

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers
(adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 24-bus system)

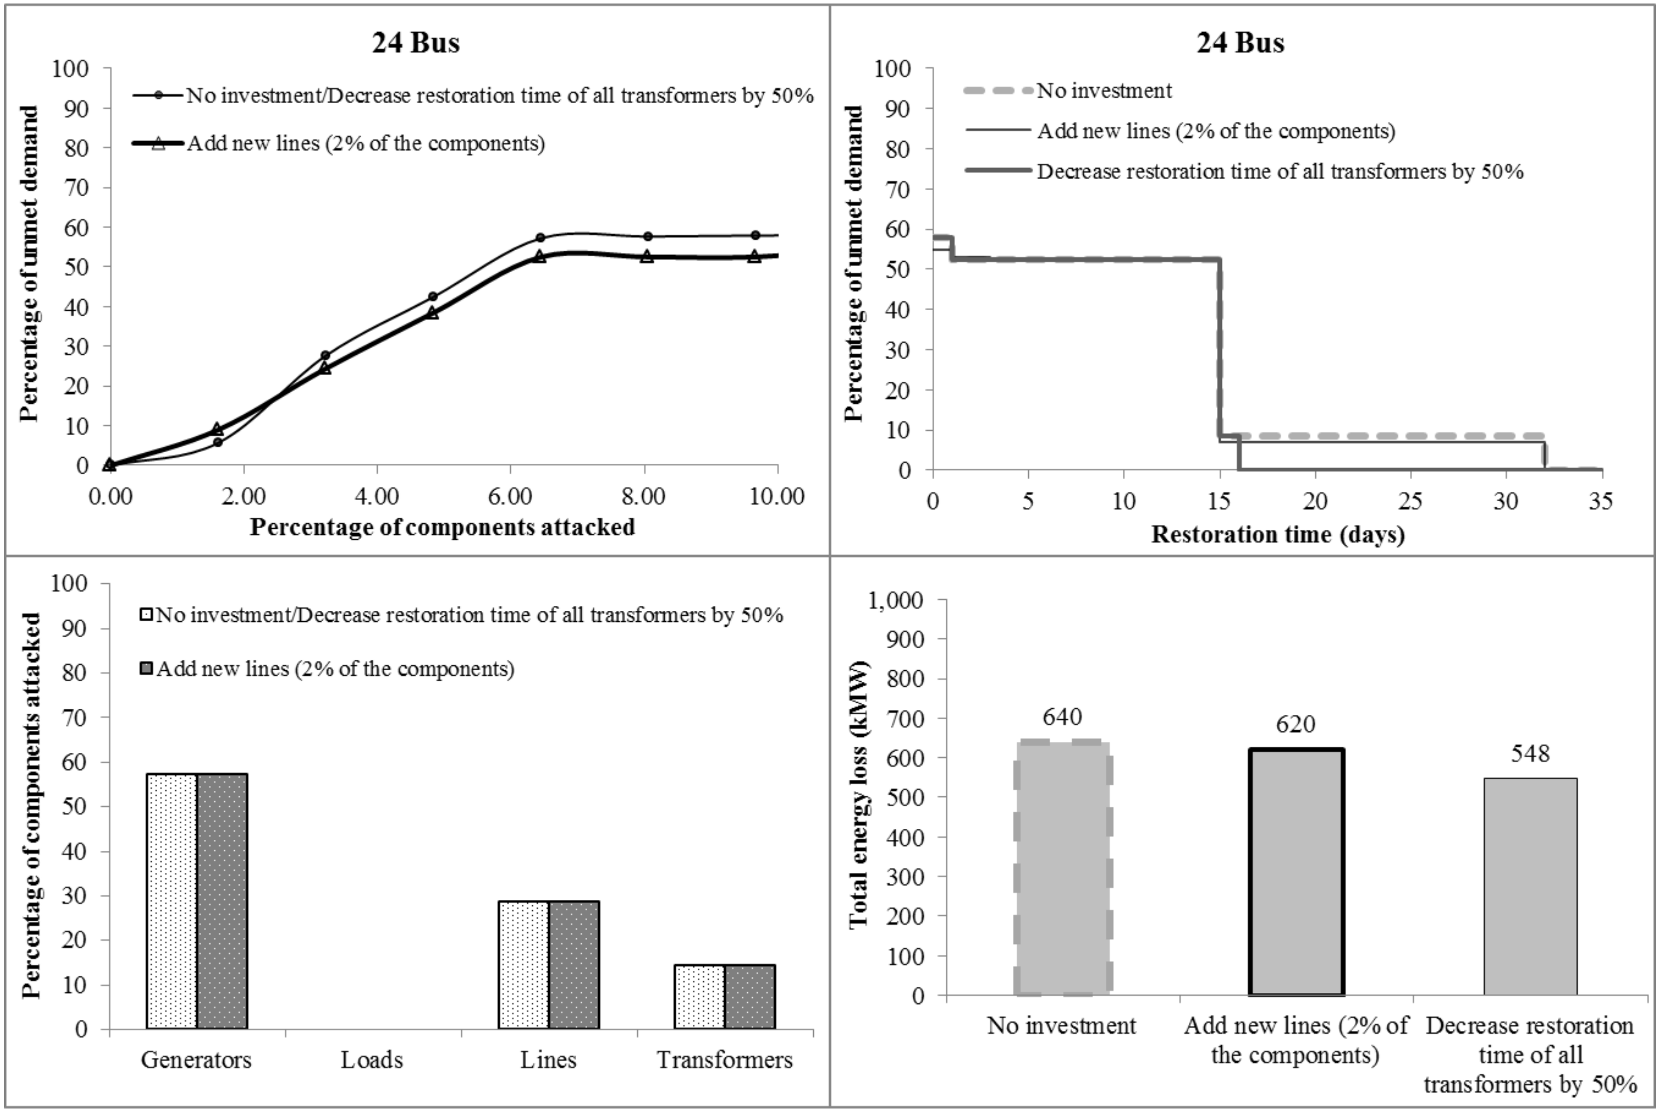

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 48-bus system)

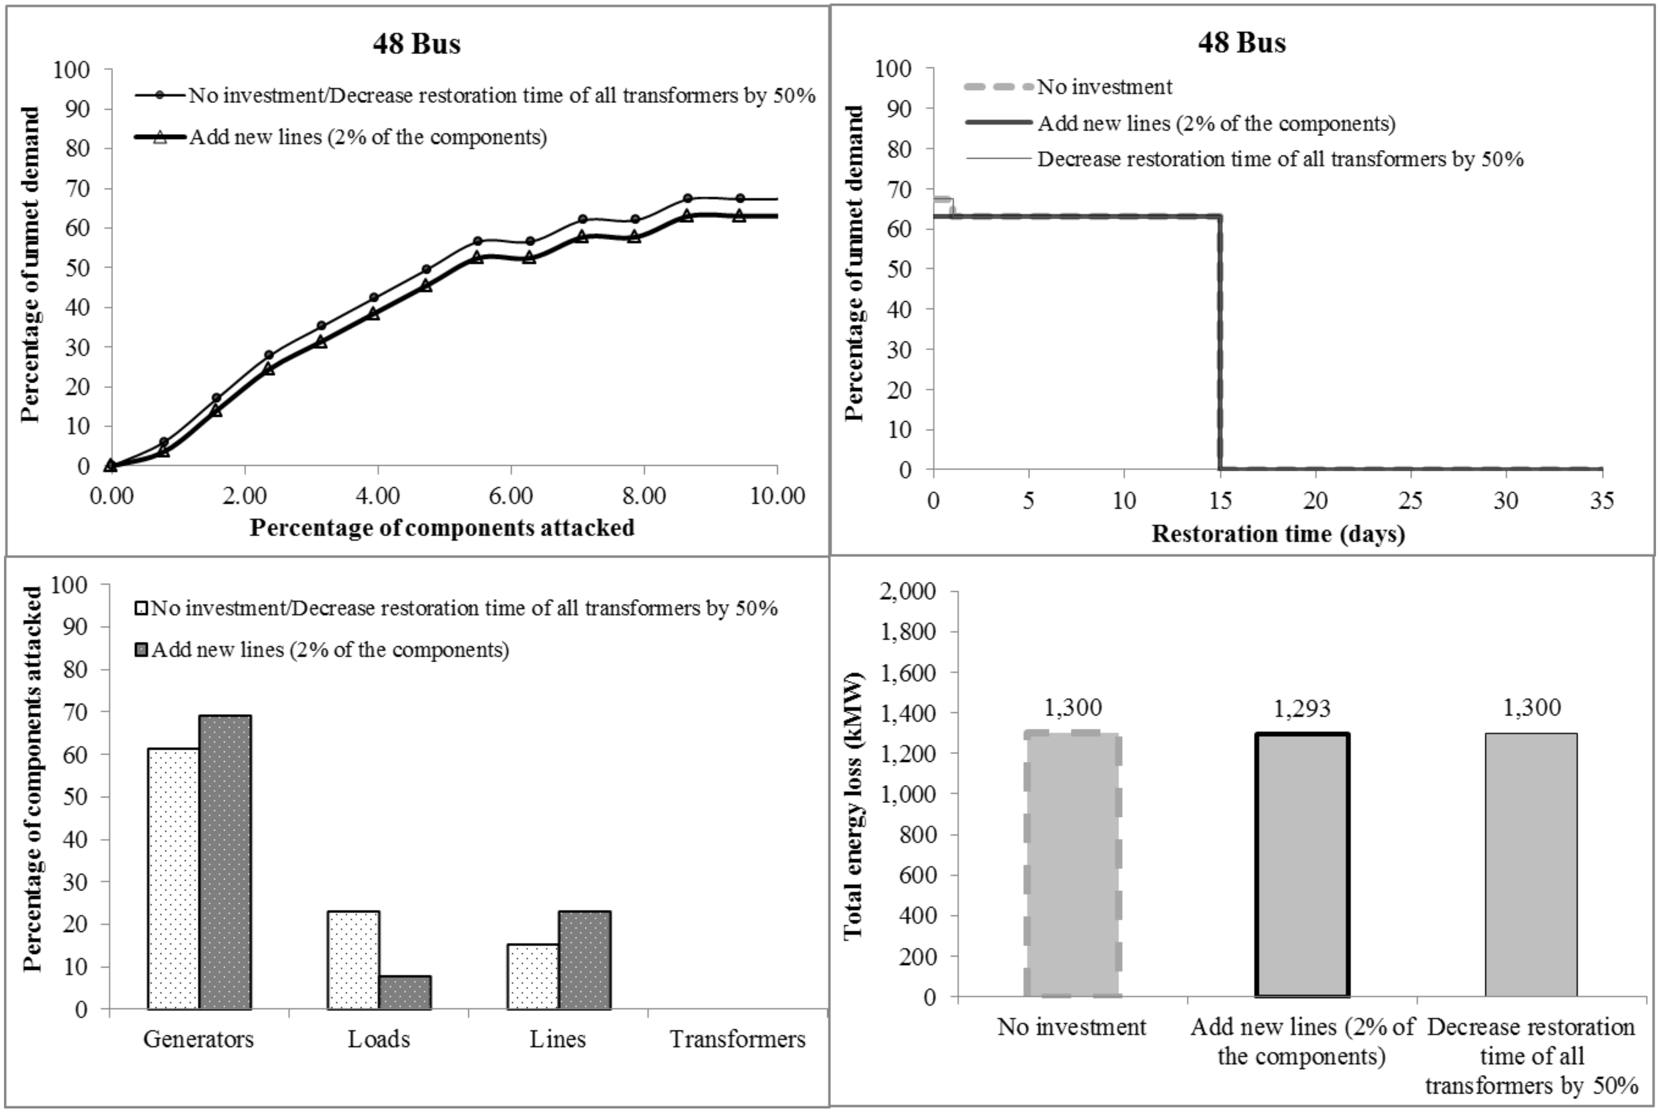

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 118-bus system)

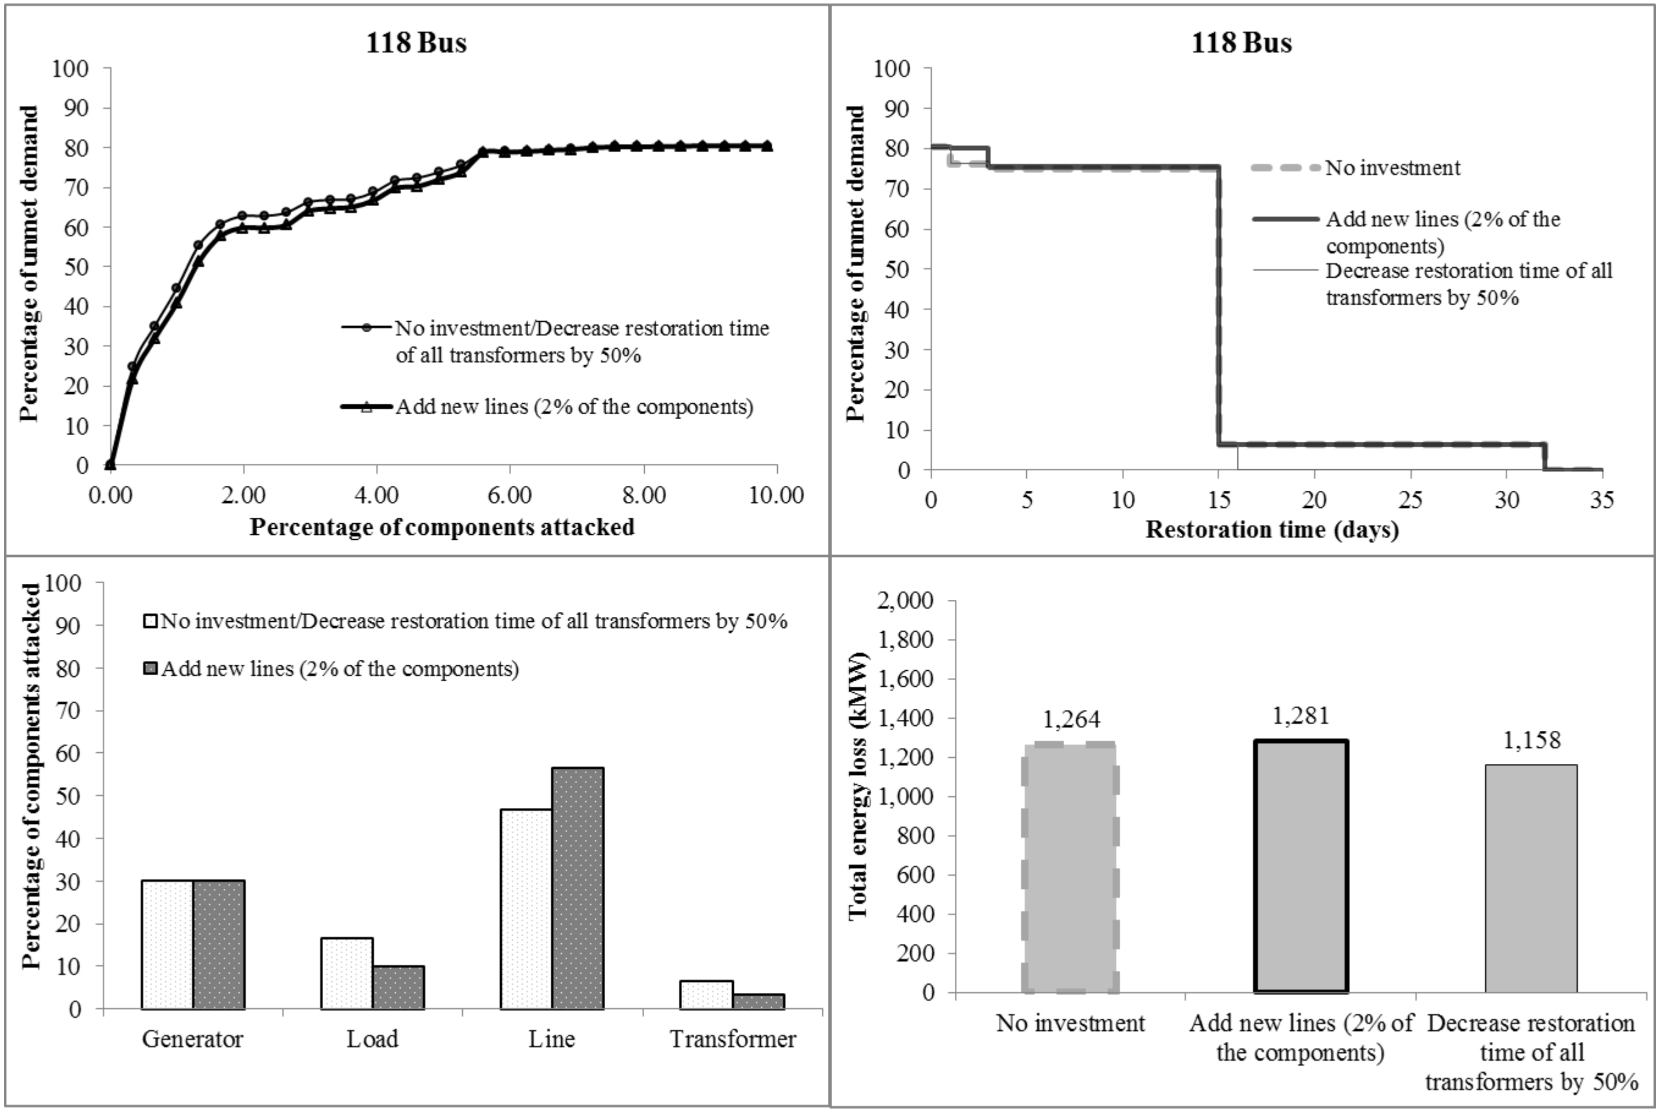

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 2% of the components against a static attacker with no cascading knowledge, 300-bus system)

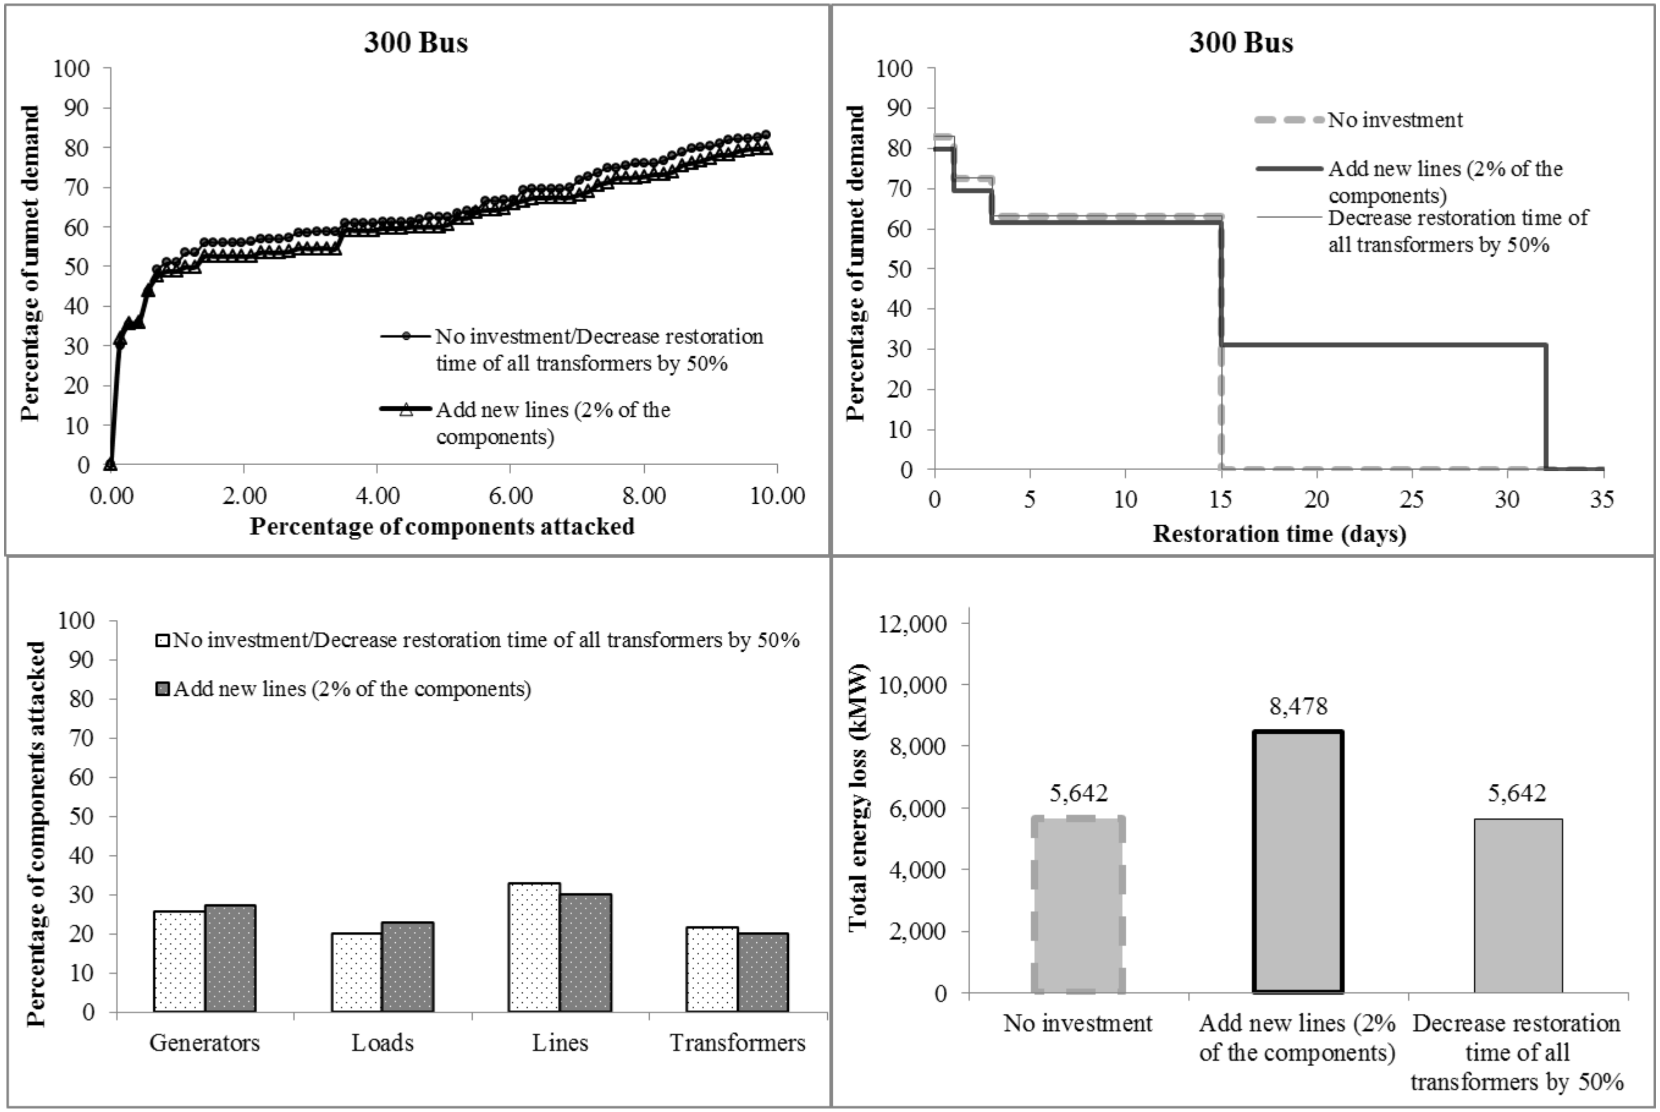

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 24-bus system)

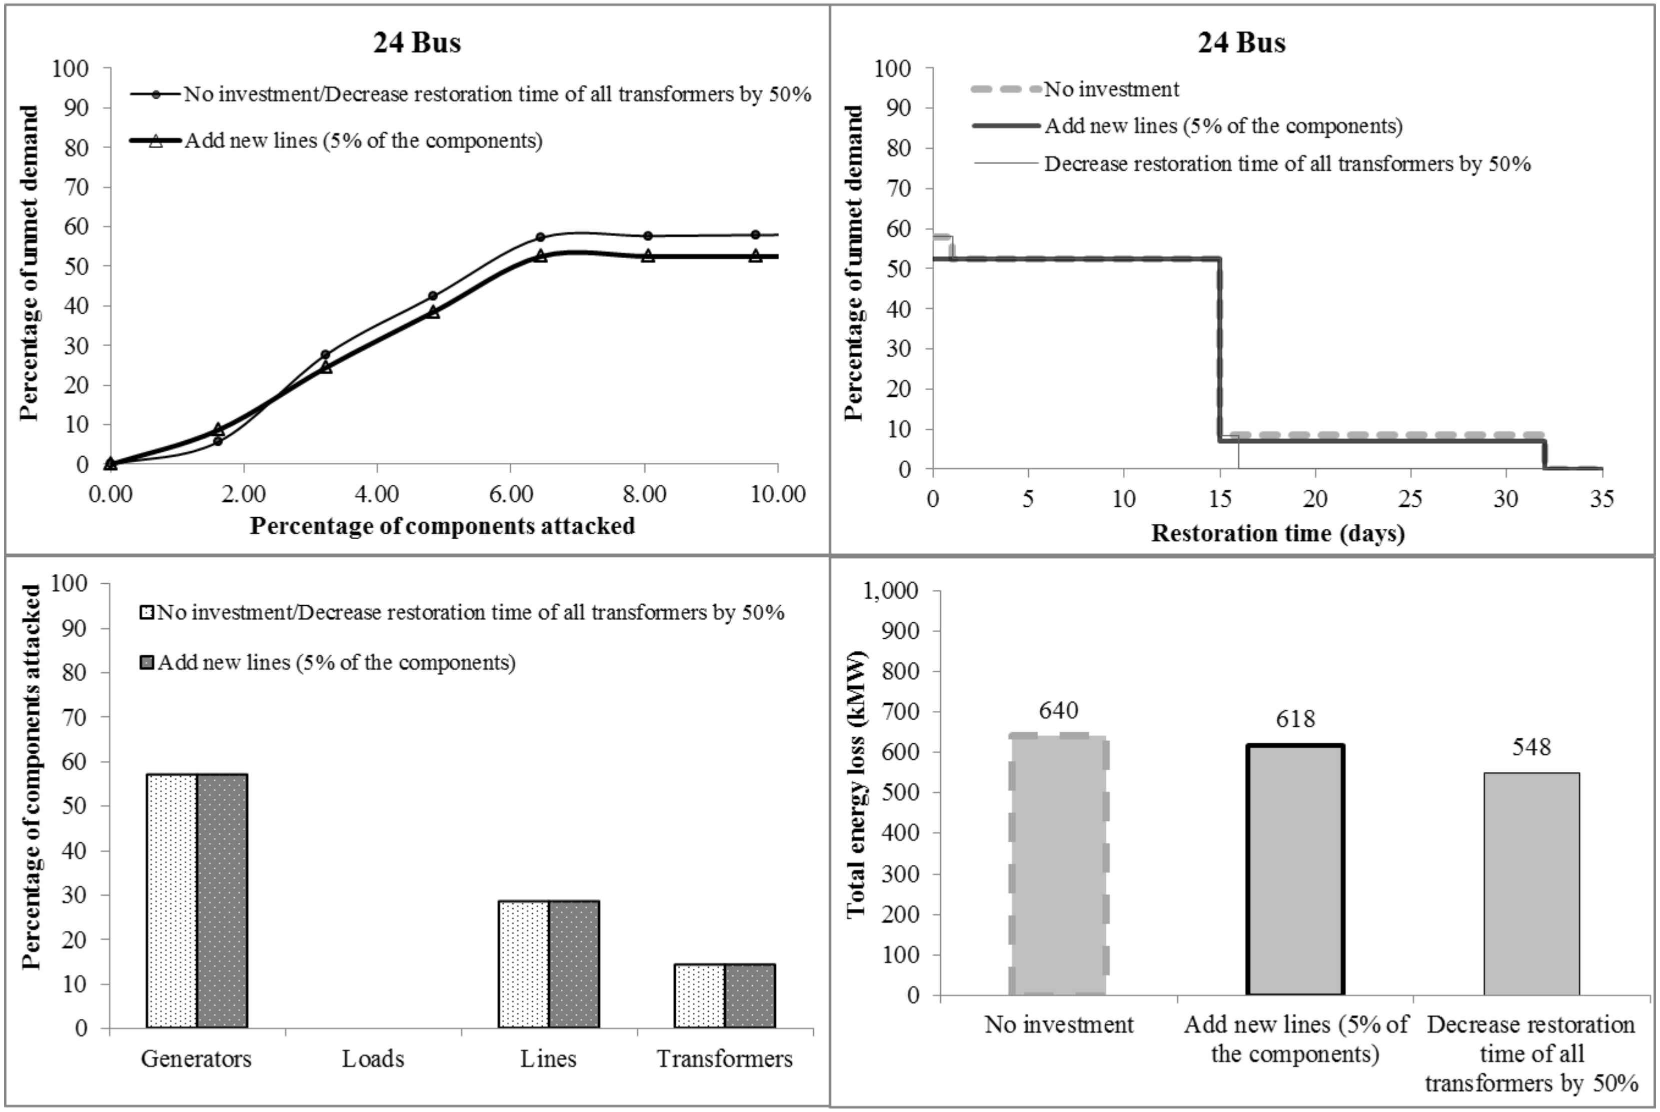

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 48-bus system)

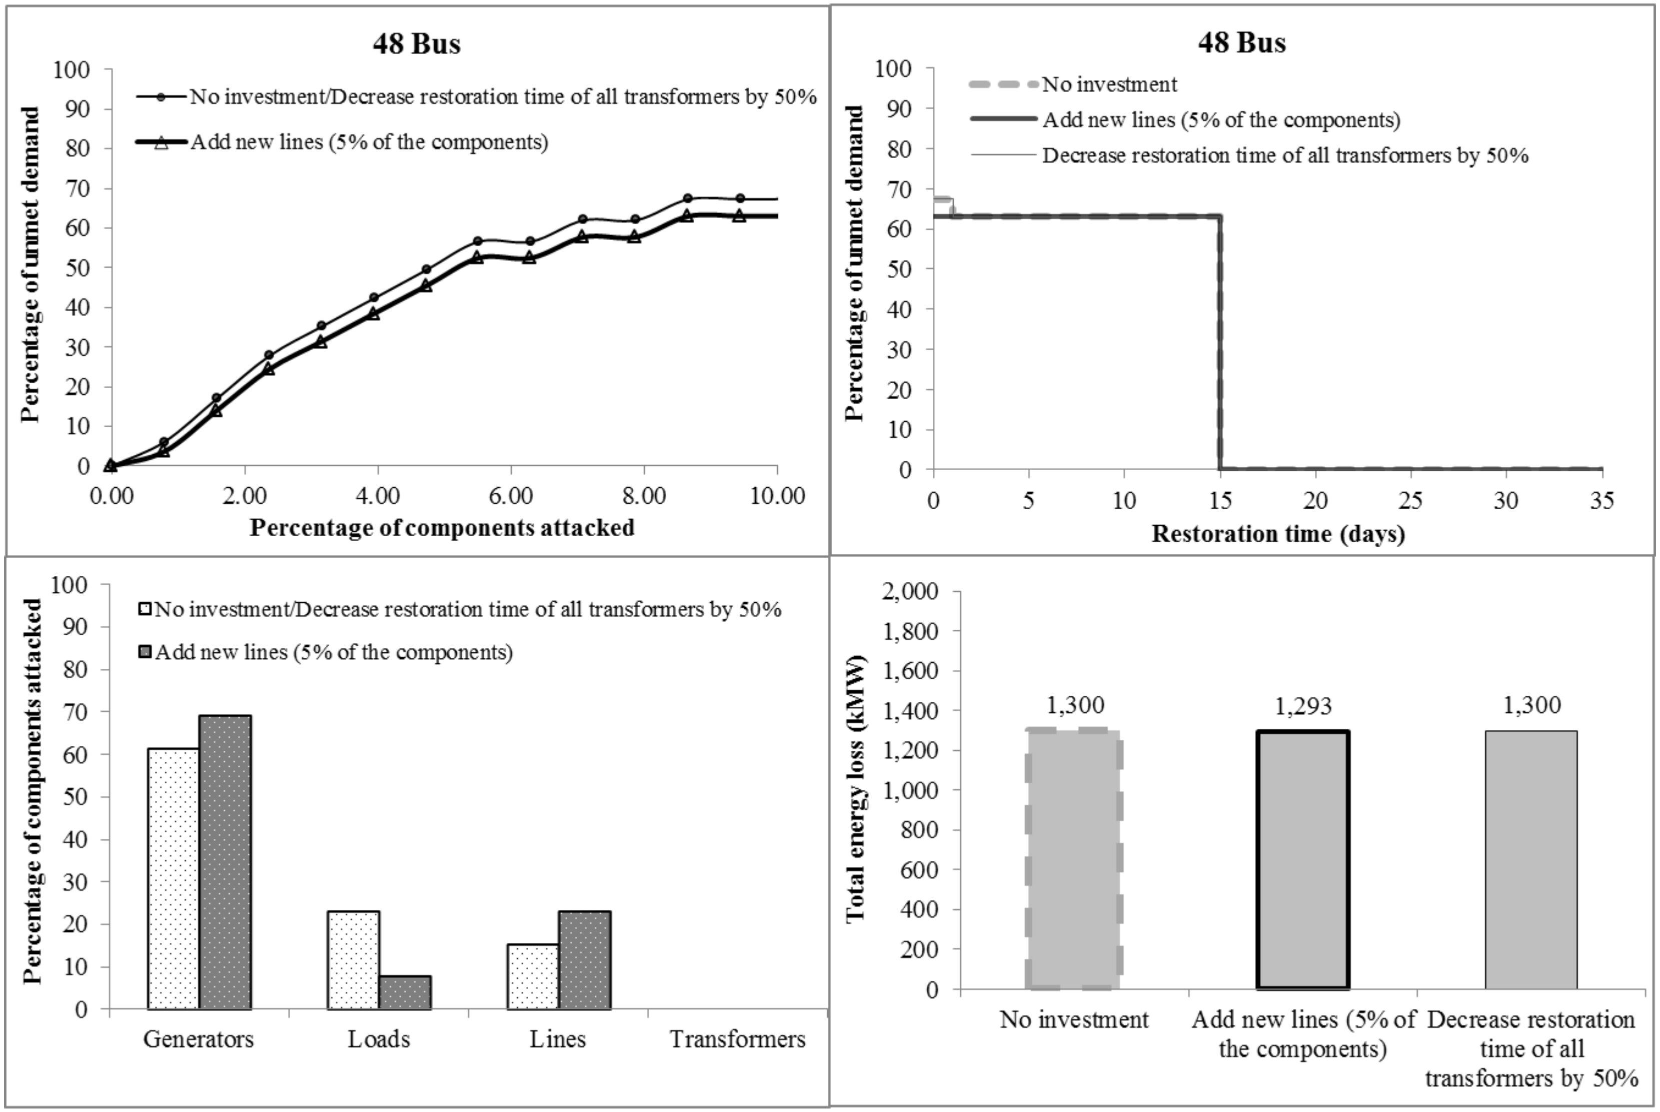

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 118-bus system)

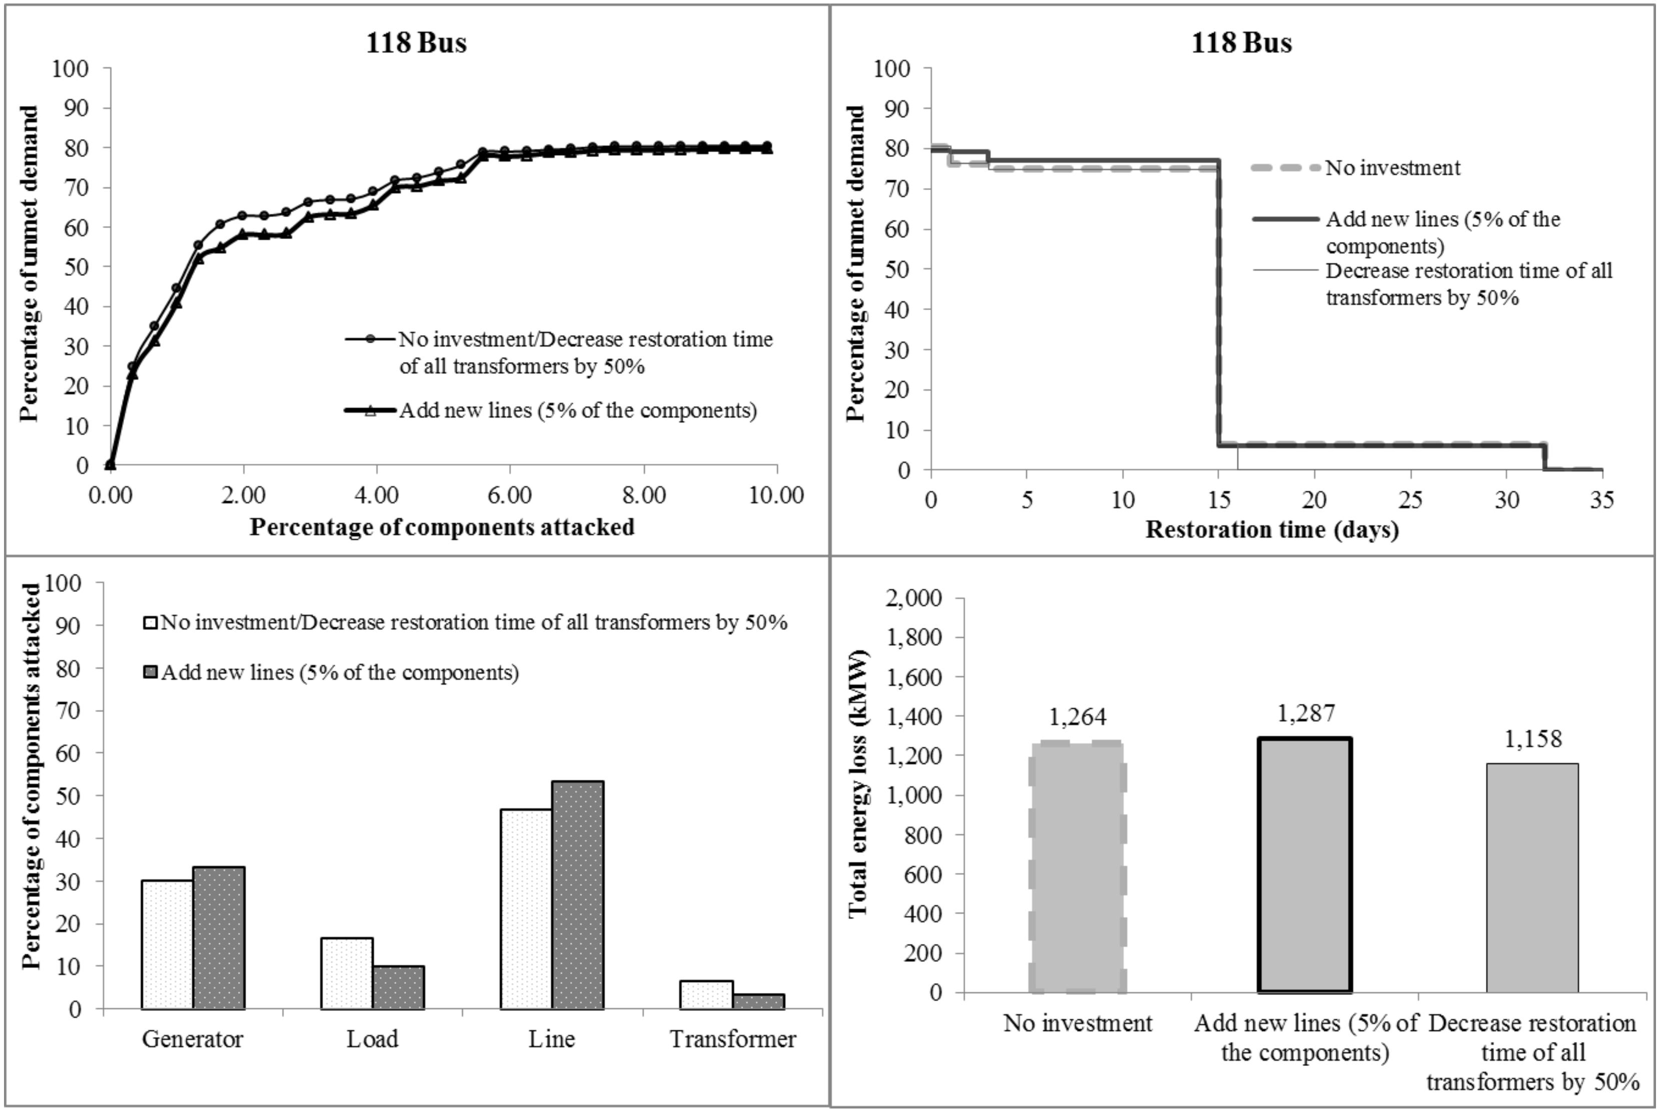

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 5% of the components against a static attacker with no cascading knowledge, 300-bus system)

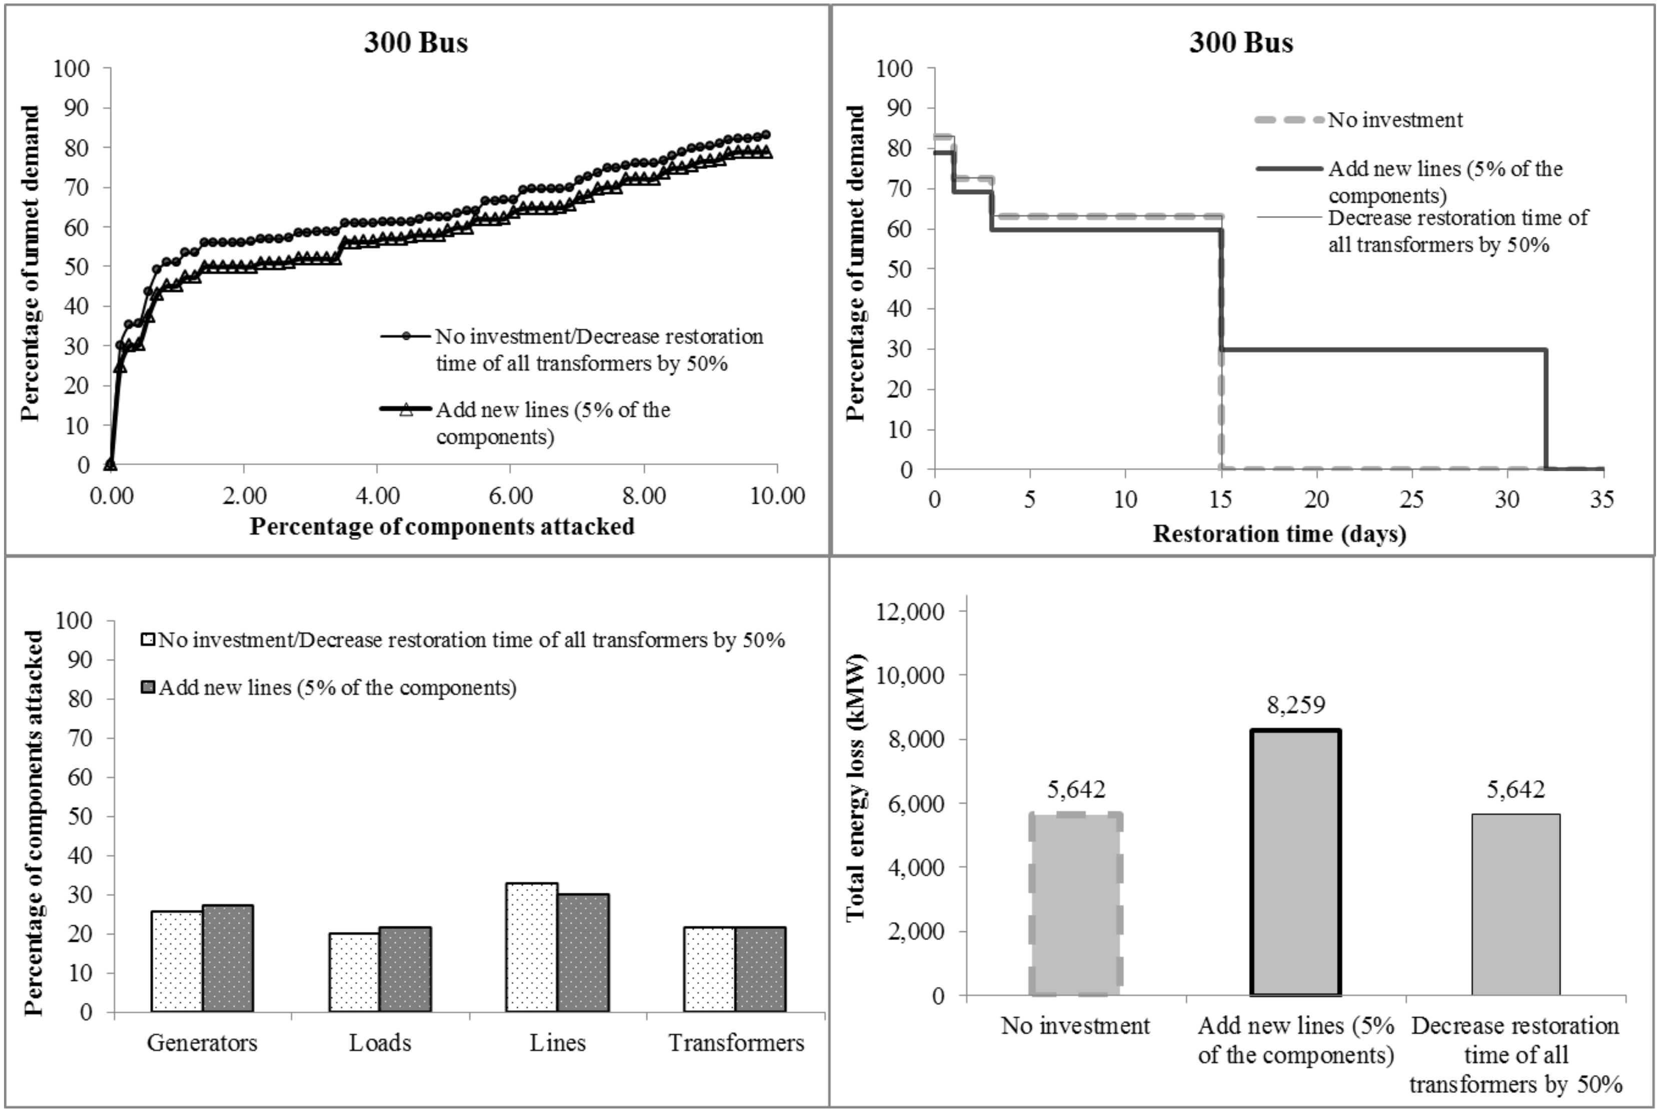

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 24-bus system)

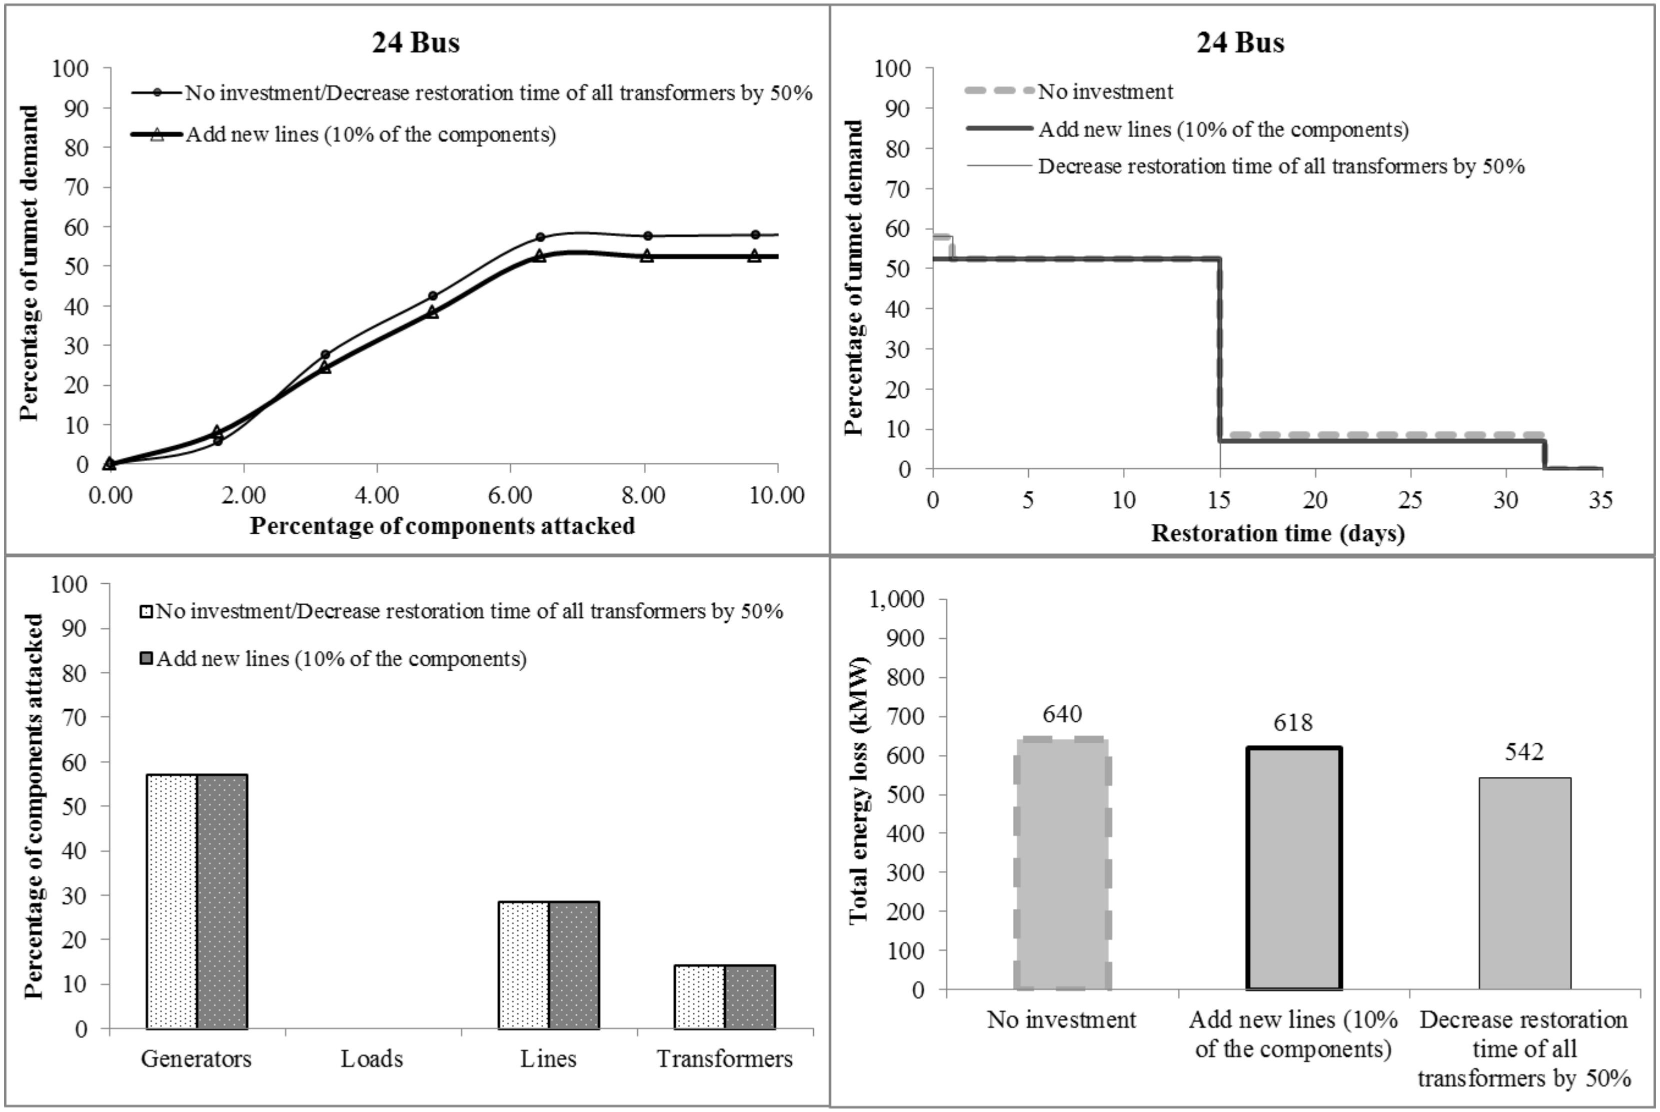

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 48-bus system)

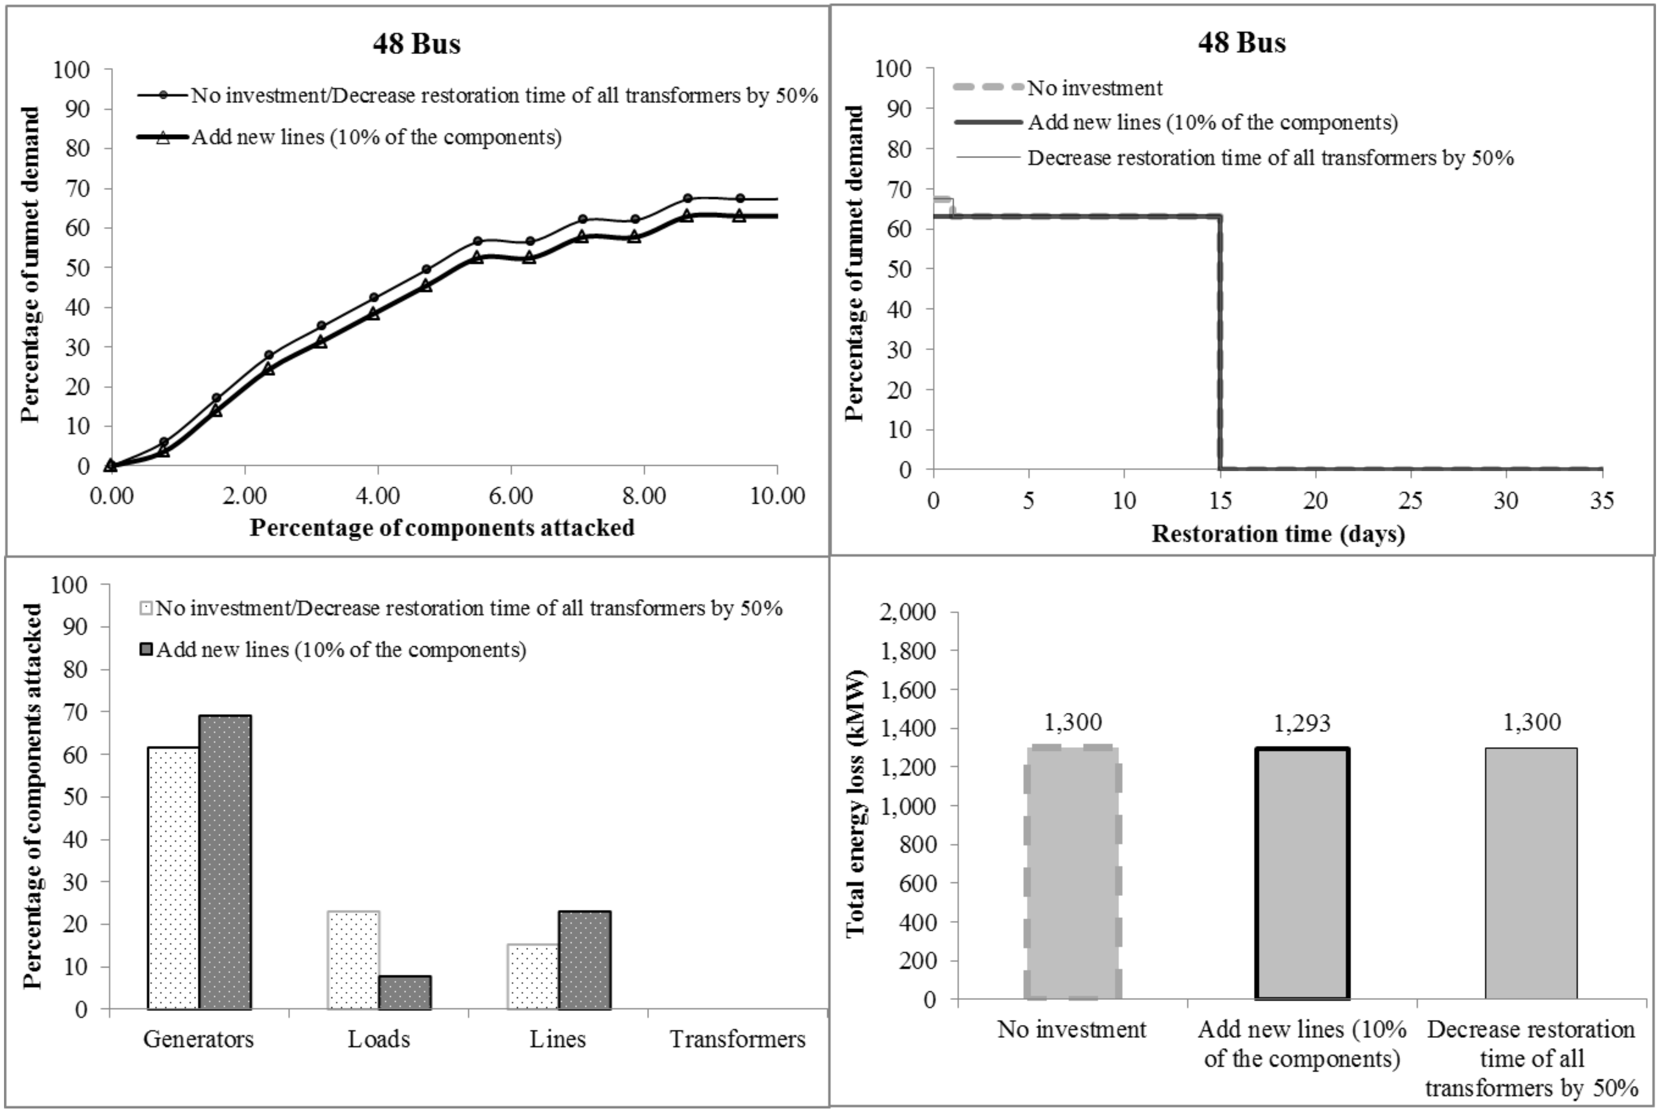

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 118-bus system)

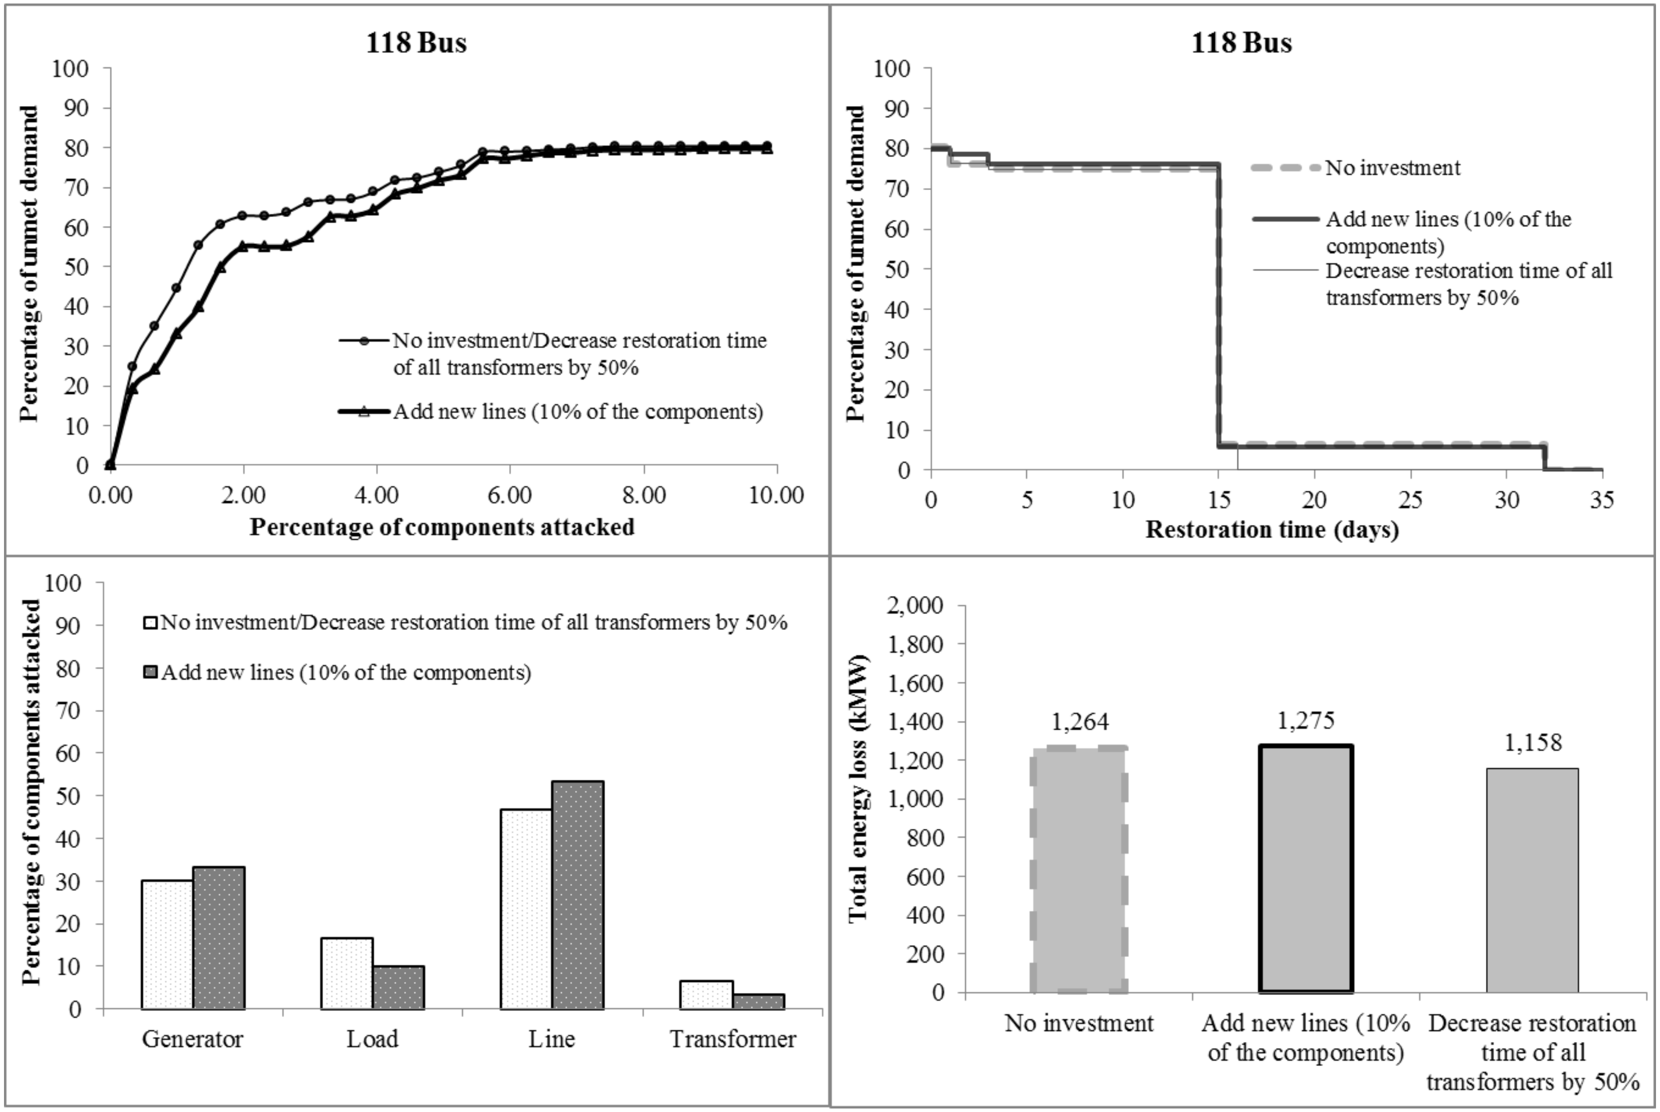

Appendix O. Adding new transmission lines versus decreasing restoration times of transformers (adding new lines of the 10% of the components against a static attacker with no cascading knowledge, 300-bus system)

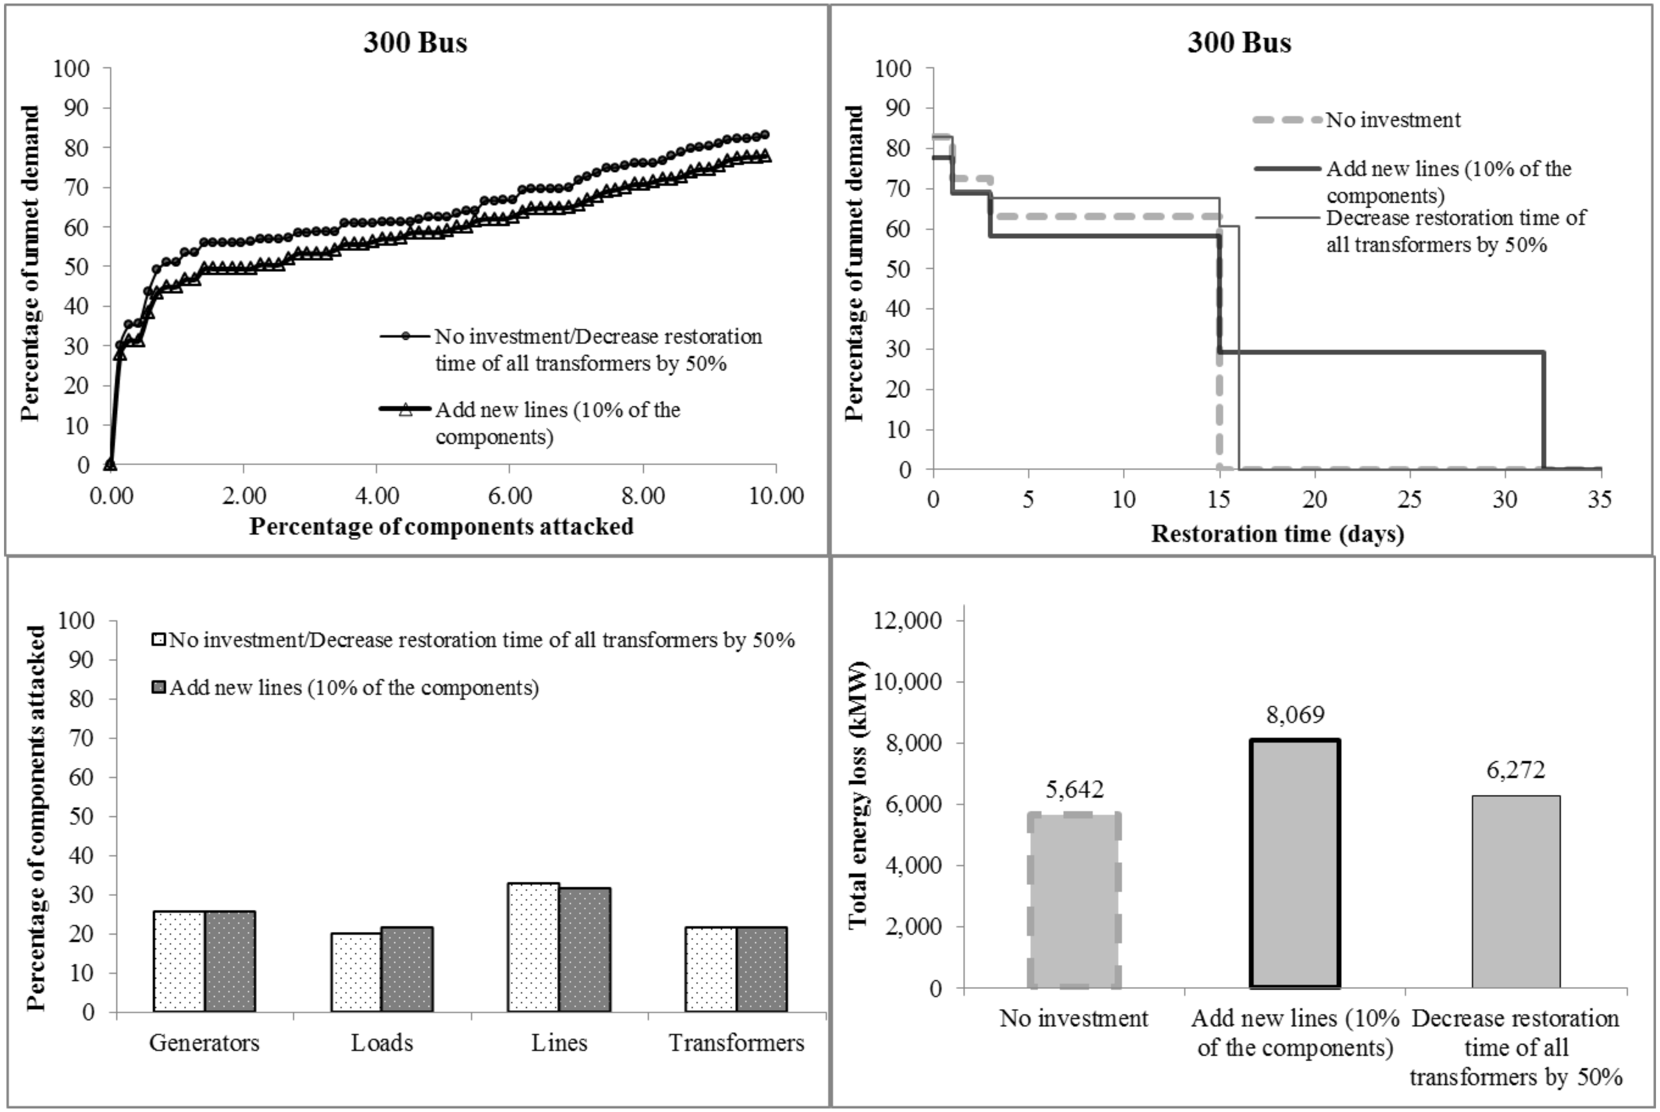
300 Bus
Percentage of unmet demand
Percentage of components attacked
No investment/Decrease restoration time of all transformers by 50%
Add new lines (10% of the components)
300 Bus
Percentage of unmet demand
No investment
Add new lines (10% of the components)
Decrease restoration time of all transformers by 50%
Restoration time (days)
Percentage of components attacked
No investment/Decrease restoration time of all transformers by 50%
Add new lines (10% of the components)
Generators
Loads
Lines
Transformers
Total energy loss (kMW)
No investment
Add new lines (10% of the components)
Decrease restoration time of all transformers by 50%