Composing Thermostatically Controlled Loads to Determine the Reliability against Blackouts

Nils Müllner, Oliver Theel and Martin Fränzle
Carl von Ossietzky University of Oldenburg, Germany, and
OFFIS Institute for Computer Science
Email: nils.muellner|theel|fraenzle@informatik.uni-oldenburg.de

***Abstract*—Power grids are parallel systems in which consumers demand a shared resource independent of each other. A blackout occurs when the total demand increases or decreases too rapidly. This paper combines methods and concepts from three domains. The first one stems from estimating the power consumption based on thermostatically controlled loads via Markov chains. The second domain provides the composition of parallel systems enriched by intermediate lumping to construct a minimal aggregate transition model, in this case of a community of housings. The third domain provides reasoning about fault tolerance properties by introducing limiting window reliability as measure, suitable to account for the continuous risk of blackouts. Combined, the three methods and concepts allow to determine the risk of blackout of a community over time.**

I. Introduction

This paper demonstrates the practical application of combining parallel composition and intermediate lumping on thermostatically controlled loads (TCL). The first contribution is a method to determine the risk of blackout. The second contribution is to point out the potential of *independent* processes — in the context that processes do not influence each other — in contrast to mutually depending processes.

A. Related work

The TCL scenario is based on a temperature model introduced by Malhamè and Chong in 1985 [1]. Callaway applied their model in 2009 [2] to derive a discrete-time Markov chain (DTMC) to determine the temperature progress, extended by Koch et al. in 2011 [3]. Kamgarpour et al. contribute three suitable models to determine the risk of blackout in this context [4]. This paper contributes a compositional reasoning to this domain.

Markov chains, as introduced by Kemeny and Snell in 1969 [5],[1] can be lumped to reduce the state space. Lumping is a well-known technique of coalescing of states under an equivalence relation, based on the definition of probabilistic bisimulation by Larsen and Skou from 1989 [6], presented by Buchholz in 1994 [7]. Lumping has the potential to speed up model checking as discussed by Katoen et al. in 2007 [8]. Lumping transition models of *independent processes* was proposed by Hermanns and Katoen in 1999 [9] exemplarily for *a plain old telephone system*. Lumping is further implemented in popular tools like the Caesar/Aldebaraan Development Package (CADP) [10] to carry out formal verification and performance analysis with the non-stochastic process algebra LOTOS [11]. Lumping of independent processes is further addressed by Boudali et al. [12]. This paper contributes a novel application to this domain.

Determining fault tolerance measures of distributed systems has been introduced by Arora, Kulkarni et al. [13] in a deterministic context. Publications preceeding this paper focused on system decomposition [14] to allow for local lumping on the likely considerably smaller transition models of subsystems (i.e. *sub-Markov chains*). This paper comparatively points out the difference between systems comprising *dependent* and systems comprising *independent* processes.

B. Structure

The paper is organized as follows. Section II presents the system model and its conversion to a transition model. Section III explains how the aggregate transition model is constructed and applied to determine the risk of blackouts. Section IV provides a brief example to demonstrate the practical value of the approach. Section V concludes our work.

II. Model

The goal of this case study is to determine the risk of voltage peaks in a power grid that cause shutdowns. Such peaks occur when the accumulated load demanded by consumers changes too fast, i.e., when too many consumers simultaneously either increase or decrease their energy demand. This example considers the load to be caused by cooling systems controlled via thermostats.

A. The TCL model

Consider a set of homogeneous houses in a warm region. While the ambient temperature θ_a outside is constantly $32\,^\circ\mathrm{C}$, the desired set indoor temperature θ_s is $20\,^\circ\mathrm{C}$. A thermostat controls the cooling system in the house. It turns on when the temperature reaches the upper bound of the hysteresis $\delta = 0.5\,^\circ\mathrm{C}$, which is $20.5\,^\circ\mathrm{C}$, and it turns off when reaching the lower boundary which is $19.5\,^\circ\mathrm{C}$.

B. The deadband

Similar to defining safety to demarcate legal from illegal states when determining fault tolerance properties [14], temperature bands can be used to specify *comfort zones* in which the thermostat should operate. Consider that the thermostat has a latency of one time step and measures the temperature at discrete evenly distributed time points utilizing a *bang-bang*

[1] We refer to revised version of 1976.

control [15]. A *bang-bang* control is a simple on/off switch turning the cooling system on when it is too hot or turning it off when it is too cold. It is reasonable to specify the comfort zones according to how far the actual temperature deviates from θ_s. The system is in a *legal state* within interval $[19.5, 20.5]$, a switching should occur within a *reasonable interval*, that is, not too long after the deadband is left, and in an undesired state beyond that interval, when switching did not occur *timely* as depicted in Figure 1.

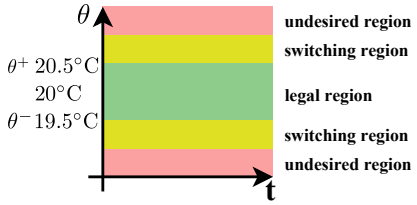

Figure 1. Specifying legal and undesired states and a witching region

This shows how the classification into fault, error and failure by Avižienis et al. [16] from the context of fault tolerance can be mapped onto temperature intervals.

C. Temperature dynamics

The following equation describes the temperature progress:

$$\theta(t+1) = \underbrace{a\theta(t)}_{i)} + \underbrace{(1-a)(\theta_a - m(t)R \cdot P)}_{ii)} + \underbrace{g(t)}_{iii)} \; [2, p.8] \quad (1)$$

The equation[2] reads as follows: the temperature in the next time step is $i)$ the temperature of the current time step plus $ii)$ the temperature progress depending on whether the thermostat is turned on or off plus $iii)$ some noise. Parameter a "governs the thermal characteristics of the thermal mass and is defined as $a = \exp(-h/CR)$" [2] with h being the duration of a time step measured in seconds, C being the thermal capacitance measured in $kWh/°\mathrm{C}$ and R being the thermal resistance measured in $°\mathrm{C}/kW$. The switch m is defined in [2, p.9] as follows:

$$m_i(t_{n+1}) = \begin{cases} 0, & \theta(t) < \theta_s - \delta = \theta_- \\ 1, & \theta(t) > \theta_s + \delta = \theta_+ \\ m(t) & \text{otherwise} \end{cases} \quad (2)$$

Parameter P describes the energy transfer rate to or from the thermal mass measured in kW. The term $g(t)$ is a noise term. Table I shows the standard parameters used by Callaway [2]: Parameter η is required to describe the total power demand

Parameter	Meaning	Standard value	Unit
R	average thermal resistance	2	$°\mathrm{C}/kW$
C	average thermal capacitance	10	$kWh/°\mathrm{C}$
P	average energy transfer rate	14	kw
η	load efficiency	2.5	
θ_s	temperature set point	20	$°\mathrm{C}$
δ	thermostat hysteresis	0.5	$°\mathrm{C}$
θ_a	ambient temperature	32	$°\mathrm{C}$

Table I. MODEL PARAMETERS

[2]The equation has been adapted from [2, p.8] in the context of this paper. For instance, the original version uses w instead of g as noise term. To avoid ambiguity with the window size, Equation 1 shows $g_i(t_n)$ as noise term.

y in the nect time step: $y(t+1) = \sum_{i=1}^{N} \frac{1}{\eta} P \cdot m(t+1)$. The parameter η describes the efficiency "and can be interpreted as the coefficient of performance" [2].

D. Deterministic execution

To determine the influence of each single parameter, the system execution is evaluated at first without noise, that is, without part iii) of Equation 1. Then, the system is deterministic without probabilistic influence. The corresponding implementation in iSat [17] is provided online.[3] We first

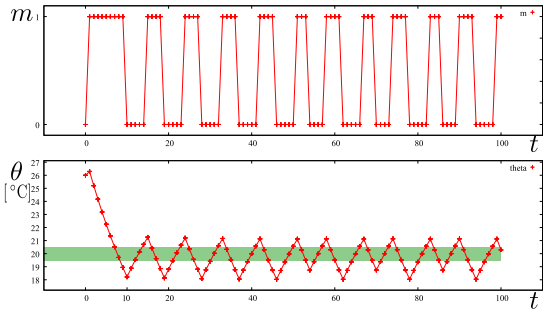

Figure 2. The TCL model executing with standard parameters

describe the initial behavior and then the behavior in the limit. The upper graph in Figure 2 shows the status of the switch and the lower graph shows the temperature evolving over time. Initially, the system detects that the temperature is too high and initiates cooling one time step later. At time step 8, it enters the deadband for the first time — at time step 7 it is just above the hysteresis — and continues cooling until time step 9. The system requires ten time steps to reach the lower boundary of the hysteresis for the first time, the switch is turned off for (alternatingly) three or four steps, the switch is turned on for (alternatingly) two or three steps, and the repetitive switching cycle shown in Figure 3 occurs the first time at time instant 44 and persists at least until time step 100. It even holds until time step 1000, not depicted in the graph, so that the assumption that the cyclic behavior is stable is justified. Each vertex is labeled

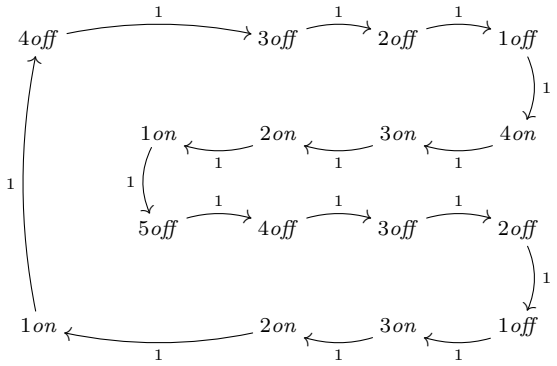

Figure 3. Repetitive cycle of a swicth remaining in a certain state

number status, referring to the number of steps the system will remain in the state. The deterministic setting without noise allows to understand how the single parameters influence the

[3]www.informatik.uni-oldenburg.de/~phoenix/isatcallaway.zip

equation. Therefore, we repeat the same setting but change each parameter, one at a time, amplifying it by a factor of ten compared to the standard parameters from Table I, except for the parameters altered in Figures 8 and 9, which are amplified by adding $10\,^{\circ}\mathrm{C}$ in Figure 8 and subtracting $10\,^{\circ}\mathrm{C}$ in Figure 9. When the isolation of the house via parameter

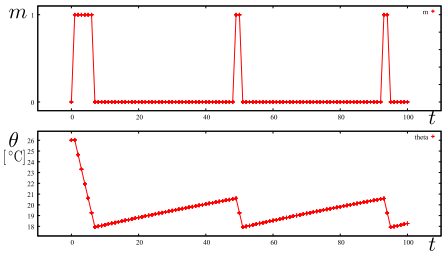

Figure 4. Temperature dynamics with standard parameters and altered average thermal resistance $R = 20\,^{\circ}\mathrm{C}/kW$

R is increased, it heats up at a far slower pace as shown in Figure 4. Furthermore, the cooling process is more efficient. The switching delay forces the system to even cool below the *safety* threshold as the temperature reaches below $18\,^{\circ}\mathrm{C}$. Figure 5 shows that amplifying the cooling power via P cools

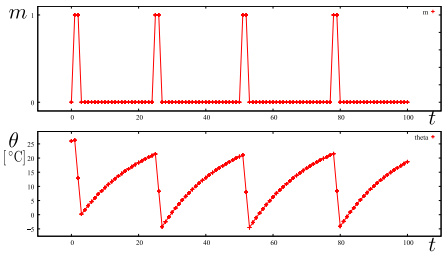

Figure 5. Temperature dynamics with standard parameters and altered average energy transfer rate $P = 140kw$

the house down rapidly. With the delay of one time step given, the cooling device freezes the house even below $0\,^{\circ}\mathrm{C}$. In case

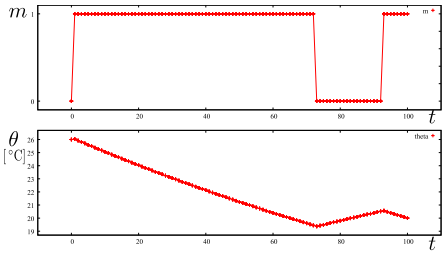

Figure 6. Temperature dynamics with standard parameters and altered average thermal capacitance $C = 100kWh/\,^{\circ}\mathrm{C}$

the thermal capacitance is increased — imagine for instance the house filled with a liquid instead of air — via parameter C, both cooling and heating phases are slowed down as shown in Figure 6. If the deadband is relaxed via parameter δ as shown in Figure 7, the cooling and heating phases take longer as well. Since it would be unreasonable to amplify the ambient temperature via θ_a beyond a certain point, $10\,^{\circ}\mathrm{C}$ are added instead of multiplying it by a factor of 10. As shown in Figure 8, the heating phases are shortened and the cooling phases are extended. The setting in depicted in Figure 9 lowers

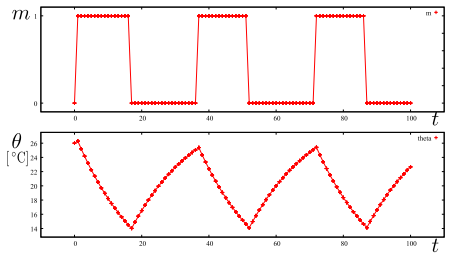

Figure 7. Temperature dynamics with standard parameters and altered thermal hysteresis $\delta = 5\,^{\circ}\mathrm{C}$

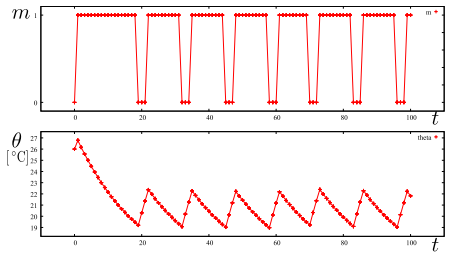

Figure 8. Temperature dynamics with standard parameters and altered temperature set point $\theta_a = 42\,^{\circ}\mathrm{C}$

the set point θ_s to $10\,^{\circ}\mathrm{C}$ which also shortens the heating phase and flattens the graph. Amplifying the load efficiency via η has

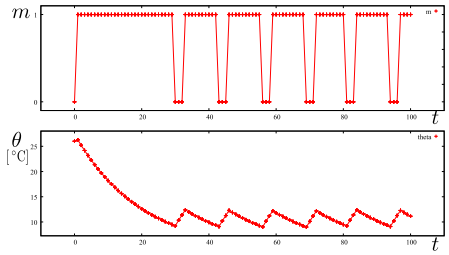

Figure 9. Temperature dynamics with standard parameters and altered ambient temperature $\theta_s = 10\,^{\circ}\mathrm{C}$

almost no effect and is therefore not depicted.

E. Adding noise

When the TCL model without noise is explored, then the system executes along one deterministic execution trace. By adding a general noise term, like the part $iii)$ in Equation 1, the transition model becomes Markovian. The execution traces then *spread over time* as exemplarily shown in Figure 10 by Koch et al. [3].

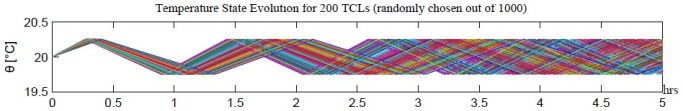

Figure 10. Temperature state evolution via simulation by Koch et al. [3, p.3]

In this setup, 200 households execute in parallel. Noise causes them to reach the deadband boundaries at different times. While all households are initially in the same state, their progress differs. After about three hours, all synchronicity is lost. One time step equals ten seconds in this example.

F. Binning

Limiting window availability (LWA) is the probability that a system provides a desired service (or satisfies a safety specification) within an alotted time window, considering that the stationary distribution initially holds [18]. The method to compute a window property like LWA is based on system and probabilistic influence to be translated into a DTMC. An intermediate step to finally acquire a DTMC from the TCL scenario is the discretization of the continuous temperature domain. A discretization in this context is commonly known as *binning* [2, 3]. The temperature domain is partitioned into — in this case equally sized — *bins*.

The probabilistic execution traces reach a bin with a certain probability in the next time step. The progress of each household along the temperature domains — one domain for m being off and one for m being on — can be formally be described with a DTMC as pictured in Figure 11.

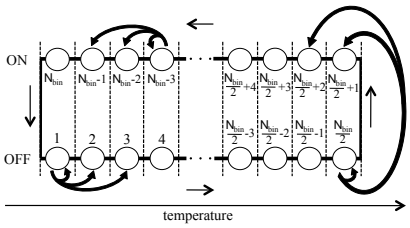

Figure 11. The state bin transition model by Koch et al. [3, p.2]

The Figure shows how the temperature domain is binned for both on and off states of m, and that the transition probabilities can be computed for each state tuple. The TCL example points out the limitations of deriving precise transition probabilities analytically. Transition probabilities are often derived via approximate methods like simulation or sampling. In this case, binning is an abstraction introducing an error. The coarser the bins are, the larger becomes the abstraction error. Soudjani and Abate [19, 20, 21] currently work on methods to compute the error that is introduced by the abstraction. Notably, they propose a method to *directly* compute the transition probabilities in a product chain of multiple housings, contrary to the sequential construction of the lumped product Markov chain that is discussed in this section.

The analytic methods proposed in the previous chapters rely on the quality of the provided probabilities. The discussion of power grids addressed that determining this quality is important. Furthermore, it showed, that safety can be formulated and a DTMC can be constructed to evaluate the safety over time.

III. Method

A. Population lumping

The example contains homogeneous housings with uniform parameters. This is not unrealistic, given the uniformity of communities in suburban areas. Each housing is modeled as a process. The goal is to construct one DTMC as surrogate transition model for one housing in the community. By multiplying it with the Kronecker product, the probability of too many houses within the population switching simultaneously can be computed. Lumping can be applied between each two Kronecker multiplications to minimize the product chain to a counting abstraction.

The complexity of the aggregate DTMC of all households depends on the number of households and the granularity of the applied binning. This is similar to the size of the state space being the product over all register domains in the previous examples. The size of the full product chain here is n^k with n bins per household and k households. Notably, each bin has to be accounted for twice: once for *on* and once for *off* mode as shown in Figure 11. In order to arrive at a tractable Markov chain, it is reasonable to select a binning according to the number of households such that the full product chain remains tractable.

Assume the following symbolic DTMC $\mathcal{D}_1$ for one housing as given. The states are labeled *number status*. Probability p_i is the probability that the temperature in a house remains in its current bin i for one time step and $1 - p_i$ is the probability that it progresses to the next temperature bin. The matrix is intentionally designed simple with only two bins and a sparse matrix to demonstrate lumpability.

↓ from/to →	1 on	2 on	1 off	2 off
1 on	p_1	$1-p_1$		
2 on		p_2	$1-p_2$	
1 off			p_3	$1-p_3$
2 off	$1-p_4$			p_4

Table II. Example symbolic DTMC for a surrogate housing $\mathcal{D}_1$

Consider the Markov chain to be irreducible and labeled $\mathcal{D}_1$. With the houses being mutually independent and executing Equation 1 in parallel, maximal parallel execution semantics apply and the Kronecker product $\otimes$ can be used as discussed in [22].

The product Markov chain of two houses with uniform parameters is the Kronecker product of two Markov chains $\mathcal{D}_1$. It calculates to $\mathcal{D}_1 \otimes \mathcal{D}_1$ and is labeled $\mathcal{D}_2$ — the index in $\mathcal{D}_i$ refers to the number of households — shown in Table IV on the next page. Empty quadrants are omitted according to the scheme shown in Table III, in which the black cells represent the omitted zero values.

↓ from/to →	first quarter	second quarter	third quarter	fourth quarter
first quarter				
second quarter				
third quarter				
fourth quarter				

Table III. Omission scheme for lumping the DTMC in Table IV

Lumping is conducted as described in [18] to reduce the DTMC shown in Table IV to the DTMC shown in Table V. The state lumping follows the schematics shown in Figure 12 which describes the equivalence classes. It also shows the symmetry of the equivalence classes in the state space: The states mirrored at the diagonal are pairwise bisimilar. States $\langle 1\text{on}, 2\text{on}\rangle$ and $\langle 2\text{on}, 1\text{on}\rangle$ become state $\langle \mathbb{1}\text{on}, 2\text{on}\rangle$. The other equivalence classes are labeled analogously. This process can be repeated for k uniform households, that is, their respective transition models, until $\mathcal{D}'_k$ is composedly constructed.

B. Computing the complexity with enumerative combinatorics

Enumerative combinatorics provide the means to compute the number of states the lumped aggregate DTMC com-

↓first quarter row	first quarter column			
↓ from/to →	⟨1on, 1on⟩	⟨1on, 2on⟩	⟨1on, 1off⟩	⟨1on, 2off⟩
⟨1on, 1on⟩	p_1^2	$p_1 \cdot (1-p_1)$		
⟨1on, 2on⟩		$p_1 \cdot p_2$	$p_1 \cdot (1-p_2)$	
⟨1on, 1off⟩			$p_1 \cdot p_3$	$p_1 \cdot (1-p_3)$
⟨1on, 2off⟩	$p_1 \cdot (1-p_4)$			$p_1 \cdot p_4$

↓first quarter row	second quarter column			
↓ from/to →	⟨2on, 1on⟩	⟨2on, 2on⟩	⟨2on, 1off⟩	⟨2on, 2off⟩
⟨1on, 1on⟩	$p_1 \cdot (1-p_1)$	$(1-p_1)^2$		
⟨1on, 2on⟩		$(1-p_1) \cdot p_2$	$(1-p_1) \cdot (1-p_2)$	
⟨1on, 1off⟩			$(1-p_1) \cdot p_3$	$(1-p_1) \cdot (1-p_3)$
⟨1on, 2off⟩	$(1-p_1) \cdot (1-p_4)$			$(1-p_1) \cdot p_4$

↓second quarter row	second quarter column			
↓ from/to →	⟨2on, 1on⟩	⟨2on, 2on⟩	⟨2on, 1off⟩	⟨2on, 2off⟩
⟨2on, 1on⟩	$p_2 \cdot p_1$	$p_2 \cdot (1-p_1)$		
⟨2on, 2on⟩		p_2^2	$(1-p_2) \cdot p_2$	
⟨2on, 1off⟩			$p_2 \cdot p_3$	$p_2 \cdot (1-p_3)$
⟨2on, 2off⟩	$p_2 \cdot (1-p_4)$			$p_2 \cdot p_4$

↓second quarter row	third quarter column			
↓ from/to →	⟨1off, 1on⟩	⟨1off, 2on⟩	⟨1off, 1off⟩	⟨1off, 2off⟩
⟨2on, 1on⟩	$(1-p_2) \cdot 1$	$(1-p_2) \cdot (1-p_1)$		
⟨2on, 2on⟩		$p_2 \cdot (1-p_2)$	$(1-p_2)^2$	
⟨2on, 1off⟩			$(1-p_2) \cdot p_3$	$(1-p_2) \cdot (1-p_3)$
⟨2on, 2off⟩	$(1-p_2) \cdot (1-p_4)$			$(1-p_2) \cdot p_4$

↓third quarter row	third quarter column			
↓ from/to →	⟨1off, 1on⟩	⟨1off, 2on⟩	⟨1off, 1off⟩	⟨1off, 2off⟩
⟨1off, 1on⟩	$p_3 \cdot p_1$	$p_3 \cdot (1-p_1)$		
⟨1off, 2on⟩		$p_3 \cdot p_2$	$p_3 \cdot (1-p_2)$	
⟨1off, 1off⟩			p_3^2	$p_3 \cdot (1-p_3)$
⟨1off, 2off⟩	$p_3 \cdot (1-p_4)$			$p_3 \cdot p_4$

↓third quarter row	fourth quarter column			
↓ from/to →	⟨2off, 1on⟩	⟨2off, 2on⟩	⟨2off, 1off⟩	⟨2off, 2off⟩
⟨1off, 1on⟩	$p_3 \cdot (1-p_1)$	$(1-p_3) \cdot (1-p_1)$		
⟨1off, 2on⟩		$(1-p_3) \cdot p_2$	$(1-p_3) \cdot (1-p_2)$	
⟨1off, 1off⟩			$(1-p_3) \cdot p_3$	$(1-p_3)^2$
⟨1off, 2off⟩	$(1-p_3) \cdot (1-p_4)$			$(1-p_3) \cdot p_4$

↓fourth quarter row	fourth quarter column			
↓ from/to →	⟨1on, 1on⟩	⟨1on, 2on⟩	⟨1on, 1off⟩	⟨1on, 2off⟩
⟨2off, 1on⟩	$(1-p_4) \cdot p_1$	$(1-p_4) \cdot (1-p_1)$		
⟨2off, 2on⟩		$(1-p_4) \cdot p_2$	$(1-p_4) \cdot (1-p_2)$	
⟨2off, 1off⟩			$(1-p_4) \cdot p_3$	$(1-p_4) \cdot (1-p_3)$
⟨2off, 2off⟩	$(1-p_4)^2$			$(1-p_4) \cdot p_4$

↓fourth quarter row	first quarter column			
↓ from/to →	⟨2off, 1on⟩	⟨2off, 2on⟩	⟨2off, 1off⟩	⟨2off, 2off⟩
⟨2off, 1on⟩	$p_4 \cdot p_1$	$p_4 \cdot (1-p_1)$		
⟨2off, 2on⟩		$p_4 \cdot p_2$	$p_4 \cdot (1-p_2)$	
⟨2off, 1off⟩			$p_4 \cdot p_3$	$p_4 \cdot (1-p_3)$
⟨2off, 2off⟩	$p_4 \cdot (1-p_4)$			p_4^2

Table IV. EXAMPLE TCL DTMC COMPOSITION $\mathcal{D}_2$, 16 STATES, 64 TRANSITIONS

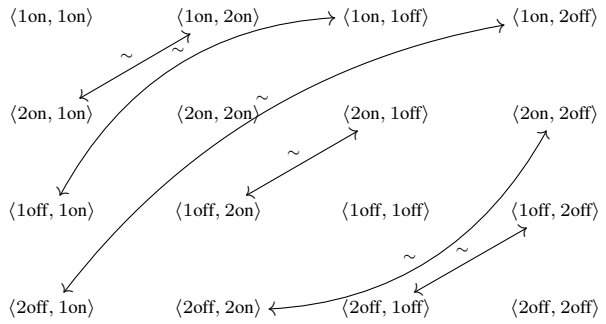

Figure 12. Lumping scheme showing which states are bisimilar

↓ from/to →	⟨1on, 1on⟩	⟨1on, 2on⟩	⟨1on, 𝟙off⟩	⟨1on, 2off⟩	⟨2on, 2on⟩	⟨2on, 𝟙off⟩	⟨2on, 2off⟩
⟨1on, 1on⟩	p_1^2	$2 \cdot p_1 \cdot (1-p_1)$			$(1-p_1)^2$		
⟨1on, 2on⟩		$p_1 \cdot p_2$	$p_1 \cdot (1-p_2)$		$(1-p_1) \cdot p_2$	$(1-p_1) \cdot (1-p_2)$	
⟨1on, 𝟙off⟩			$p_1 \cdot p_3$	$p_1 \cdot (1-p_3)$		$(1-p_1) \cdot p_3$	$(1-p_1) \cdot (1-p_3)$
⟨1on, 2off⟩	$p_1 \cdot (1-p_4)$	$(1-p_1) \cdot (1-p_4)$		$p_1 \cdot p_4$			$(1-p_1) \cdot p_4$

↓ from/to →	⟨1on, 2on⟩	⟨1on, 𝟙off⟩	⟨1on, 2off⟩	⟨2on, 2on⟩	⟨2on, 𝟙off⟩	⟨2on, 2off⟩	⟨𝟙off, 𝟙off⟩	⟨𝟙off, 2off⟩
⟨2on, 2on⟩				p_2^2	$2 \cdot (1-p_2) \cdot p_2$		$(1-p_2)^2$	
⟨2on, 𝟙off⟩		$p_2 \cdot p_3$	$p - 2 \cdot (1-p_3)$			$(1-p_2) \cdot (1-p_3)$	$(1-p_2) \cdot p_3$	
⟨2on, 2off⟩	$p_1 \cdot (1-p_4)$	$(1-p_1) \cdot (1-p_4)$				$p_2 \cdot p_4$		$(1-p_2) \cdot p_4$

↓ from/to →	⟨1on, 1on⟩	⟨1on, 𝟙off⟩	⟨1on, 2off⟩	⟨𝟙off, 𝟙off⟩	⟨𝟙off, 2off⟩	⟨2off, 2off⟩
⟨𝟙off, 𝟙off⟩				p_3^2	$2 \cdot (1-p_3) \cdot p_3$	$(1-p_3)^2$
⟨𝟙off, 2off⟩		$p_3 \cdot (1-p_4)$	$(1-p_3) \cdot (1-p_4)$		$p_3 \cdot p_4$	$(1-p_3) \cdot p_4$
⟨2off, 2off⟩	$(1-p_4)^2$		$2 \cdot (1-p_4) \cdot p_4$			p_4^2

Table V. LUMPED DTMC $\mathcal{D}_2'$, TEN STATES, 36 TRANSITIONS

prises. The state space explosion without lumping draws a state space according to *variation with repetition*. Therefore, there are $|\mathcal{S}| = n^k$ states when considering k houses and n bins. The successive lumping arrives at a state space of $|\mathcal{S}'| = \left(\binom{n}{k}\right) = \binom{n+k-1}{k} = \frac{(n+k-1)!}{(n-1)!\cdot k!}$ — the *multiset (rising) binomial coefficient* [23] — by *combination with repetition*. Figures 13(a) and 13(b) compare both state space explosions when adding more uniform households on the x-axis. They show that lumping *dampens* the explosion tremendously. The largest DTMC before the final lumping step in compositional lumping in this context is $\left(\binom{n}{k-1}\right) \cdot n$. Figure 13(a) compares the *initial* explosions up to ten households, while Figure 13(b) computes the scalability for up to 100 households. The figures demonstrate that instead of the exponential state space explosion depicted in the red graphs, the size of the DTMC increases almost linearly with lumping, depicted in the blue graphs. Compared to unlumped multiplication, the graph under

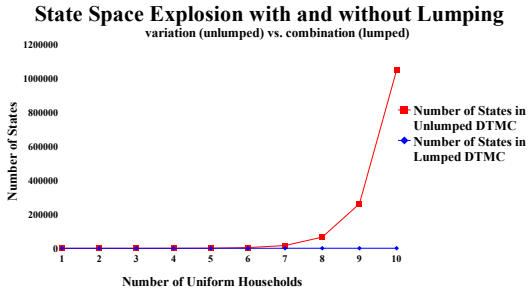

(a) Dampening the state space explosion in the first ten steps

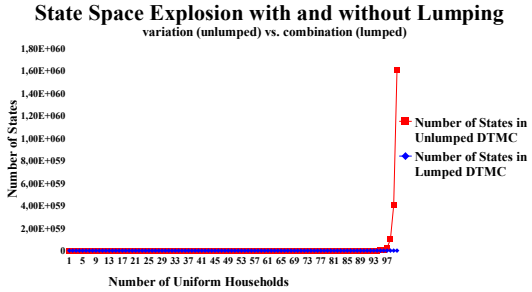

(b) Dampening the state space explosion in the first 100 steps

Figure 13. Dampening the state space explosion

application of lumping almost coincides with the x-axis. Both complexities are computed with *enumerative combinatorics*, that is, variation and combination with repetition. The tractability of the DTMC depends on the available computing power. Even with lumping and perfectly homogeneous households, $\mathcal{S}'$ contains $176,851$ states for 100 housings and the proposed binning. Yet, compared to approximately $1.61 \cdot 10^{60}$ states, sequential composition and lumping are obviously preferable.

C. Control destroys bisimilarity

The sequential application of composition and lumping hinges on the mutual independence of the processes. *Control strategies* can prioritize housings to distribute limited resources, for instance when a limited amount of energy faces more demand than it can satisfy. In that case, processes lose

their independence. The demand by one prioritized process can delay the satisfaction of another process.

For instance, assume that in the above example of $\mathcal{D}_2$ in Table IV, one house constantly has a higher priority than another one. Further, assume that the power grid cannot tolerate both thermostats switching simultaneously from on to off or vice versa. In case both thermostats want to switch, the thermostat with the lower priority must wait exactly one time step. This adds two novel states to the system and replaces transitions accordingly as shown in Table VI.

↓ from/to →	⟨2on, 2on⟩	⟨1off, 2on⟩	⟨2on, 1off⟩	⟨1off, 3on⟩
⟨2on, 2on⟩	p_2^2	$(1-p_2)\cdot p_2$	$p_2\cdot(1-p_2)$	$(1-p_2)^2$

↓ from/to →	⟨2off, 2off⟩	⟨1on, 2off⟩	⟨2off, 1on⟩	⟨1on, 3off⟩
⟨2off, 2off⟩	p_4^2	$(1-p_4)\cdot p_4$	$p_4\cdot(1-p_4)$	$(1-p_4)^2$

↓ from/to →	⟨1off, 1off⟩	⟨2off, 1off⟩	⟨1on, 1on⟩	⟨2on, 1on⟩
⟨1off, 3on⟩	p_3	$1-p_3$		
⟨1on, 3off⟩			p_1	$1-p_1$

Table VI. PRIORITIZED TCL DTMC

In case the transition probabilities are not equal — $p_1 \neq p_2 \wedge p_1 \neq p_3 \wedge p_1 \neq p_4 \wedge p_2 \neq p_3 \wedge p_2 \neq p_4 \wedge p_3 \neq p_4$ — the DTMC becomes irreducible. For instance, the states $\langle 1\text{on}, 2\text{on}\rangle$ and $\langle 2\text{on}, 1\text{on}\rangle$ are then no longer probabilistic bisimilar as their outgoing transition probabilities would not coincide anymore. Although the processes do not propagate values to one another, thus excluding fault propagation: they depend on each other by sharing a mutual resource. When that resource is controlled, bisimilarity can be destroyed.

This paragraph demonstrated how sequential composition and lumping can be executed and pointed out, that the absence of fault propagation does not necessarily imply independence of the processes. The example introduced control to destroy bisimulation among non-communicating processes. Next, a small numerical example computes the probability for a small community to suffer from a black out.

IV. EXAMPLE

A. Interleaving application of the Kronecker product $\otimes$ and lumping

Consider a set of 1000 households. For the sake of argument, we assume the coarsest possible binning, yielding one bin for *on* and one for *off* mode. Acquiring transition probabilities is not the scope of this paper. We consider the following values as being provided: the probability to remain in the *on* bin is 0.9 and the probability to remain in the *off* bin is 0.8.

The full product chain without lumping contains $|\mathcal{S}| = 2^{1000} = 1,0715 \cdot 10^{301}$ states. When lumping is applied after each composition — which is a *counting abstraction* [24, p.195] —, the resulting DTMC contains only $|\mathcal{S}'| = 1001$ states: one state in which all thermostats are *off*, one in which only one thermostat is *on* and so on until one state in which all 1000 thermostats are *on*. Its computation took about 50 minutes on a Intel(R) Core(TM) i5-3317U CPU at 1.7 GHz equipped with 8GB DDR3 SODIMM with MatLab. The source code is provided provided online.[4] A graphical representation of the lumped product DTMC is shown in Figure 14.

[4] www.informatik.uni-oldenburg.de/~phoenix/matlabcallaway.zip

Figure 14. 1000 housings TCL power grid

Notably, there are no zero-probability transitions. The transitions in the *blue* areas are just very close to zero. The figure shows in the top row, in which all housings are *off*, a steep maximum at 100 housings simultaneously switching *on*. The bottom row, in which all housings are *on*, shows a shallower distribution with the maximum at 800 housings, indicating that about 200 housings simultaneously switch *off*.

With each housing being added, the matrix grows. Hence, each further addition takes longer than the previous one. The graph in Figure 15 indirectly shows how the computation time of adding further housings increases with each housing.

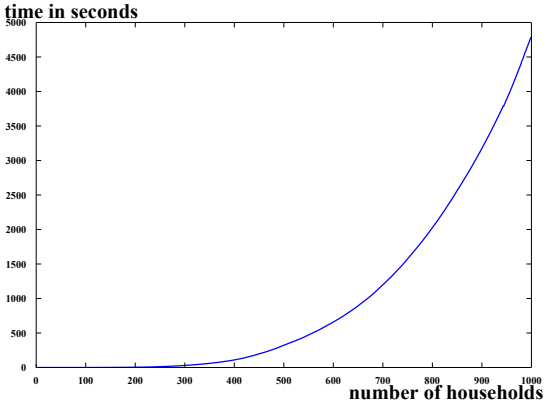

Figure 15. Time consumption to compute 1000 housings TCL power grid

Compared to decomposing hierarchical systems and otherwise mutually depending processes as discussed in [14], composing mutually independent processes is rather simple. With the households being homogeneous, one surrogate DTMC can be multiplied with the Kronecker product over and over again with the resulting matrix being lumped after each iteration. The example in this section used the coarsest possible matrix — in which the θ-domain of the temperature is not partitioned — to make a point: Writing a script to compose independent processes can be as trivial as in this case. Then, generating matrices containing thousands of states automatically is just a matter of time. The DTMC shown in Figure 14 was generated in less than an hour on a tablet PC with the specifications given

above. Contrary, constructing hierarchically structured systems is not as easy. The example in [14] contained only $|\mathcal{M}| = 648$ states and $|\mathcal{M}'| = 324$ states. Its computation by hand took about two weeks. The number of states is not a good indicator to reason about scalability. Instead, the effort that is required to construct a DTMC to compute the desired measure can should used.

B. The risk of blackout – limiting window reliability

The probability that the system blacks out is the accumulated transition probability of too many houses switching on or off simultaneously. For instance, if the system blacks out with 1000 simultaneous houses switching, the probability for a blackout computes as $pr(\overrightarrow{0, 1000}) \cdot pr_\Omega(\langle 0 \rangle) + pr(\overrightarrow{1000, 0}) \cdot pr_\Omega(\langle 1000 \rangle)$. The index here refers to the number of simultaneous switches necessary to cause a blackout. When the system breaks down for even 999 simultaneous switches, the probability for blackout computes as $pr(\overrightarrow{0, 1000}) \cdot pr_\Omega(\langle 0 \rangle) + pr(\overrightarrow{1000, 0}) \cdot pr_\Omega(\langle 1000 \rangle) + pr(\overrightarrow{0, 999}) \cdot pr_\Omega(\langle 0 \rangle) + pr(\overrightarrow{999, 0}) \cdot pr_\Omega(\langle 999 \rangle) + pr(\overrightarrow{1, 1000}) \cdot pr_\Omega(\langle 1 \rangle) + pr(\overrightarrow{1000, 1}) \cdot pr_\Omega(\langle 1000 \rangle)$ and so forth. Figure 16(a) shows the stationary distribution that is required to compute the probability to crash. Figure 16(b) shows the probability to crash according to the required number of simultaneous switches that are required for the system to blackout.

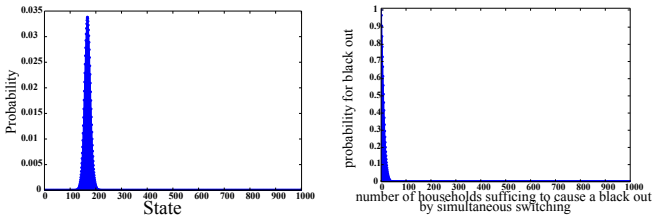

(a) Stationary distribution

(b) Risk per time step to black out depending on the number of simultaneously switching households

Figure 16. Determining the risk to crash

In this scenario, we are interested in the probability that closure is not violated within a given time window. The *limiting window reliability* is a probability distribution on first stopping times. In contrast to LWA [18], it is a probability on stopping times of taking *a wrong* transition, violating closure, while LWA measures the probability of taking *a right* transition to achieve convergence. The limiting window reliability, which is the chance to *survive* a given time window w, is simply computed by $(1 - pr_i)^w$ with respect to the critical number of simultaneous accumulated switches i.

With the limiting window reliability distribution, the ongoing risk of eventually suffering from a blackout can be computed. Figure 17 shows how the probability for each *safety predicate* — that is, either one house suffices to cause a blackout, or two, or three ... — converges to 1 over time. The x-axis showing the *number of households sufficing to cause a blackout* is cropped at 100 but extends to 1000. With fewer houses required to cause a blackout, the probability for a blackout increases at a faster pace. The figure shows that more than about 60 houses are required in the predicate to pose a thread for the community to survive the first 100 time steps from the limit onwards.

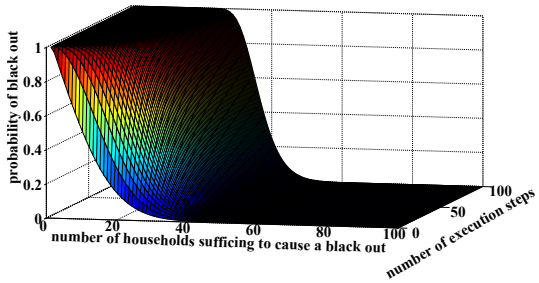

Figure 17. Limiting window reliability over 100 time steps

Figure 18 shows the same plot with the reliability being encoded in color for a larger time scale. This perspective shows that i) the demarcation between unreliable (dark red) and reliable (dark blue) is sharp (white) and ii) providing a safety threshold of even less than 100 houses suffices to provide for a *high* reliability for a time window of $10,000$ computation steps. Notably, the critical number that the graph converges to in the limit, is the total number or housings. What seems like counting to infinity twice — the limiting window reliability starts in the limit and then converges to the limit again — opens an important discussion. Similar to *limiting reliability* as discussed by Trivedi [25, p.321], the limiting window reliability for a limiting window is zero, too. In the limit, the red area extends to 1000 households.

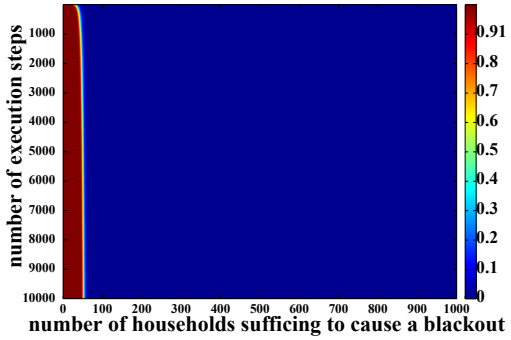

Figure 18. Limiting window reliability over 10,000 time steps

C. The result

The example demonstrated how a transition model and safety specifications can be derived from a real world system and how they can be utilized in the context of this paper. With an initial probability distribution, the risk of eventually taking a transition in which too many households simultaneously switch in the same direction can be computed. Contrary to LWA that computes the probability that the system reaches the legal states within a time window, the *window reliability* computes the probability that a system leaves the legal states in a time window. Contrary to the *limiting reliability* [25, p.321], it is reasonable in this context to compute the limiting *window* reliability. Consider the system to be initially supported by

a vent and an energy buffer to compensate for voltage peaks until it converges *sufficiently close* to its stationary distribution. From thereon, the limiting window reliability determines the probability with which the system *survives* a desired time window by summing up all relevant transition probabilities over that time.

V. Conclusion

This paper provided a practical case study to discuss important aspects. It highlighted

1) that determining probabilistic inputs is crucial to acquire realistic results with the presented methods,

2) that a DTMC and a safety predicate can be derived for some real world systems,

3) that absence of fault propagation does not automatically imply process independence and bisimilarity, and

4) that the notion of limiting/instantaneous window properties can be easily adapted to suit other contexts.

It furthermore demonstrated how synthesizing a transition model benefits from the processes being independent. The constructed DTMC accounts for thousand processes and was generated in less than an hour. On the contrary, the transition models of hierarchic systems cannot be constructed as easy.

Acknowledgment

This work was partly supported by the German Research Council under grant SFB/TR 14 *AVACS*, the EU Commission under grant FP7-ICT-2009-257005 *MoVeS* and by the funding initiative *Niedersächsisches Vorab* of the Volkswagen Foundation and the Ministry of Science and Culture of Lower Saxony (as part of the *Interdisciplinary Research Center on Critical Systems Engineering for Socio-Technical Systems*).

References

[1] Roland Malhamè and Chee-Yee Chong. Electric-load Model Synthesis by Diffusion Approximation of a High-order Hybrid-state Stochastic-system. *IEEE Transactions on Automatic Control*, 30:854 – 860, 1985.

[2] Duncan S. Callaway. Tapping the Energy Storage Potential in Electric Loads to Deliver Load Following and Regulation, with Application to Wind Energy. *Energy Conversion and Management*, 50:1389 – 1400, May 2009.

[3] Stephan Koch, Johanna L. Mathieu, and Duncan S. Callaway. Modeling and Control of Aggregated Heterogeneous Thermostatically Controlled Loads for Ancillary Services. In *Proceedings of the 17th Power Systems Computation Conference*, Stockholm, Sweden, 2011.

[4] Maryam Kamgarpour, Christian Ellen, Sadegh Esmaeil Zadeh Soudjani, Sebastian Gerwinn, Johanna L. Mathieux, Nils Müllner, Alessandro Abate, Duncan S. Callaway, Martin Fränzle, and John Lygeros. Modeling Options for Demand Side Participation of Thermostatically Controlled Loads. In *Proceedings of the IREP Symposium-Bulk Power System Dynamics and Control -IX (IREP), August 25-30, 2013, Rethymnon, Greece*, 2013.

[5] John G. Kemeny and James L. Snell. *Finite Markov Chains*. University Series in Undergraduate Mathematics. New York, NY, USA, 2, 1976 edition, 1976.

[6] Kim G. Larsen and Arne Skou. Bisimulation Through Probabilistic Testing. In *Conference Record of the 16th ACM Symposium on Principles of Programming Languages (POPL1989*, pages 344 – 352, 1989.

[7] Peter Buchholz. Exact and Ordinary Lumpability in Finite Markov Chains. *Journal of Applied Probability*, 31(1):59–75, 1994.

[8] Joost-Pieter Katoen, Tim Kemna, Ivan S. Zapreev, and David N. Jansen. Bisimulation Minimisation Mostly Speeds up Probabilistic Model Checking. In *Proceedings of the 13th international conference on Tools and algorithms for the construction and analysis of systems*, TACAS'07, pages 87–101, Berlin, Heidelberg, 2007. Springer-Verlag.

[9] Holger Hermanns and Joost-Pieter Katoen. Automated Compositional Markov Chain Generation for a Plain-Old Telephone System. In *SCIENCE OF COMPUTER PROGRAMMING*, pages 97–127, 1999.

[10] Hubert Garavel, Frédéric Lang, and Radu Mateescu. An overview of CADP 2001. Research Report RT-0254, INRIA, 2001.

[11] LOTOS - A Formal Description Technique Based on the Temporal Ordering of Observational Behaviour. Standard, September 1989. Information Processing Systems, Open Systems Interconnection.

[12] Hichem Boudali, Pepijn Crouzen, and Mariëlle Stoelinga. A Rigorous, Compositional, and Extensible Framework for Dynamic Fault Tree Analysis. *IEEE Trans. Dependable Sec. Comput.*, 7(2):128–143, 2010.

[13] Sandeep S. Kulkarni. *Component Based Design of Fault-Tolerance*. PhD thesis, 1999. Advisors: Anish Arora, Mukesh Singhal, Ten-Hwang Lai.

[14] Nils Müllner, Oliver Theel, and Martin Fränzle. Combining Decomposition and Reduction for the State Space Analysis of Self-Stabilizing Systems. In *Journal of Computer and System Sciences (JCSS)*, volume 79, pages 1113 – 1125. Elsevier Science Publishers B. V., November 2013. The paper is an extended version of [22].

[15] L. M. Sonneborn and F. S. van Vleck. The Bang-bang Principle for Linear Control Systems. *Journal of the Society for Industrial and Applied Mathematics (SIAM)*, 2:151–159, 1964.

[16] Algirdas Avižienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic Concepts and Taxonomy of Dependable and Secure Computing. *IEEE Transactions on Dependable and Secure Computing*, 1:11 – 33, 2004.

[17] Martin Fränzle, Christian Herde, Tino Teige, Stefan Ratschan, and Tobias Schubert. Efficient Solving of Large Non-linear Arithmetic Constraint Systems with Complex Boolean Structure. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 1(3-4):209 – 236, 2007.

[18] Nils Müllner and Oliver Theel. The Degree of Masking Fault Tolerance vs. Temporal Redundancy. In *Proceedings of the 25th IEEE Workshops of the International Conference on Advanced Information Networking and Applications (WAINA2011)*, Track "The Seventh International Symposium on Frontiers of Information Systems and Network Applications (FINA2011)", pages 21 – 28, Biopolis, Singapore, 2011. IEEE Computer Society Press.

[19] Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate. Adaptive and Sequential Gridding Procedures for the Abstraction and the Verification of Stochastic Processes. 2013. Submitted for review to the Society for Industrial and Applied Mathematics (SIAM).

[20] Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate. Probabilistic Reachability Computation for Mixed Deterministic-Stochastic Processes. 2013. unpublished draft.

[21] Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate. Aggregation of Thermostatically Controlled Loads by Formal Abstractions. *Prodceedings of the Eurpoean Control Conference (ECC2013)*, 2013. submitted for review.

[22] Nils Müllner, Oliver Theel, and Martin Fränzle. Combining Decomposition and Reduction for State Space Analysis of a Self-Stabilizing System. In *Proceedings of the 26th IEEE International Conference on Advanced Information Networking and Applications (AINA2012)*, pages 936 – 943, Fukuoka-shi, Fukuoka, Japan, March 2012. IEEE Computer Society Press. Best Paper Award.

[23] William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, January 1968.

[24] Xiang Fu, Tevfik Bultan, and Jianwen Su. Formal Verification of E-Services and Workflows. In *Proc. ESSW*, pages 188–202. Springer-Verlag, 2002.

[25] Kishor S. Trivedi. *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*. John Wiley & Sons, second edition, 2002.