Quantitative evaluation of availability measures of gas distribution networks

Laura Carnevali, Marco Paolieri, Fabio Tarani, Enrico Vicario
Dipartimento di Ingegneria dell'Informazione - Università di Firenze
{laura.carnevali, marco.paolieri, fabio.tarani, enrico.vicario}@unifi.it

ABSTRACT

Rising competition among gas distribution companies, growing availability of smart metering devices, and increasingly strict requirements on agreed service levels stimulate research on advanced modeling and solution techniques for quantitative evaluation of gas distribution networks. We propose a novel methodology for modeling and evaluation of the transient network behavior after a component failure.

The approach relies on a topological model of the fluid dynamics and a stochastic timed model of the actions started after a component failure. Fluid dynamic analysis evaluates the service level of end-users in each possible operating condition of the network, also supporting the derivation of stochastic parameters for the failure management model. In turn, such model is analyzed to evaluate the probability over time of the network operating conditions. Transient probabilities are then aggregated on the basis of the results of fluid dynamic analysis to derive availability measures. Special attention is paid to make the structure of the stochastic model independent of the network topology. To provide a proof of concept, the approach is exemplified on a small-sized network equipped with a backup pipe, evaluating for each end-user the transient probability of not being served after a component failure as well as the mean outage time. These measures comprise a valid ground for the evaluation of different failure management processes and the definition of demand-response strategies.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Modeling Techniques

General Terms

Performance, Algorithms

Keywords

Gas distribution networks, transient availability measures, Markov regenerative processes, stochastic state classes.

ValueTools'13, December 10 – 12 2013, Turin, Italy

1. INTRODUCTION

Quantitative evaluation of expected availability is taking growing relevance for the efficient operation of gas distribution networks (i.e., middle/low pressure networks), due to several drivers including competitive challenges raised by novel industrial organization of utilities, emphasis on homeland security and serviceability, automation capabilities offered by smart monitoring and actuation devices, increasing interest in demand-response control applications [23]. This motivates investigation in modeling and solution methods, both in the tactical perspective that supports decision during run-time operation, and in the strategic perspective that supports planning of topology, localization of sensing and actuation devices, preventive maintenance, and evaluation of sustainable service levels.

The subject of network reliability and availability has been widely investigated in electric and telecommunication systems, seeking for efficient infrastructures that facilitate interoperability among innovative technological solutions [27]. As notable examples, a real-time optimization model of demand response is solved in [11] to maximize the consumer utility while minimizing energy costs [11]; a Markov model of the actions taken in reaction to a failure in a telecommunication network is leveraged in [2] to assess survivability metrics with respect to distributed generation, demand response, and smart grids; hierarchical composition is exploited in [18] to derive availability and performability measures for telecommunication systems. Yet, various major differences characterize gas networks with respect to telecommunication and electric systems, with notable issues such as: the localization of the point of failure and network reconfiguration, which may involve much less automation and may result in a large variability of timings; the regulation of controllable inputs, which involves processes running on a much slower and more variable time scale; a lower level of network redundancy and a different perspective on the criticality of interruptions; the possibility to operate the network at different levels of pressure, trading the efficiency of operation for the resilience to transient faults.

Modeling targeted to natural gas networks is less explored. In particular, most of the literature on the analysis and simulation of gas networks focuses on the fluid-dynamics perspective, mainly oriented to assess flow rates and pressures across network elements [15, 12, 22]. Optimization of operations has been addressed in various aspects, notably to favor efficient integration within multi-carrier systems combining provisioning of electric and gas power [21, 19, 17, 20]. Stochastic modeling is applied in [5] to consider different rates of leakage that may occur in a pipe fault and thus predict the impact that this may have on pressures and flow rates across the network, supporting the planning of appropriate actions to mitigate risks. In [24] fluid-dynamic analysis of a section of a real gas network is repeated for different configurations of demand reflecting the statistics of usage in different day hours and seasons. The effects of sequential restoration and constrained network capacity are considered in [14] to support reliability assessment by deriving average measures of interruption rate and outage time experienced by end-users, exemplifying the approach on a small-sized gas network.

In this paper, we propose a method for modeling and evaluation of availability of middle/low pressure gas networks. To this end, we assume that changes of the operating conditions of the network due to daily/seasonal demand variations or demand-response mechanisms are independent of the actions taken in reaction to a failure of a network element.

In the modeling stage, this permits to separate the representation of the fluid dynamics from the stochastic characterization of the failure management steps. On the one hand, the fluid dynamics model follows a relatively conventional graph-theoretical representation of the network topology, supporting well known techniques for the evaluation of pressures and flow rates under a given configuration of components and parameters. On the other hand, the failure management model provides a representation of the different functional behaviors that may occur when a network component fails, with stochastic distributions depending on the failure localization and the consequent pressure regulation within the network. As a salient trait, stochastic modeling affords the use of non-Markovian temporal parameters that overcome the limits of validity induced by unbounded support and exponential distribution. Moreover, the model structure is independent of the network topology, which makes the model general and almost guarantees a constant level of complexity of the stochastic analysis.

In the evaluation stage, the evolution over time of the failure management actions is analyzed through transient analysis based on stochastic state classes and generalized Markov renewal equations, as proposed in [16]. This provides the probability over time of any feasible operating condition of the network after a failure. Such probabilities are then aggregated on the basis of the results of fluid dynamic analysis (identifying the service status of each end-user in each network condition), providing transient and average availability measures for end-users. The approach is experimented on a small example of the literature, providing a proof of concept of the overall methodology and paving the way for future developments concerned with the evaluation of different failure management processes and the definition of demand-response strategies.

The rest of the paper is organized in four sections. In Section 2, we present the fluid-dynamics model (Section 2.1), discuss a characterization of stochastic temporal parameters of the failure management actions (Section 2.2), and present the failure-management model (Section 2.3). In Section 3, we recall the salient aspects of the solution technique of [16] (Section 3.1) and discuss how transient and average availability measures are derived (Section 3.2). In Section 4, we exemplify the approach on a small-sized network of the literature (Section 4.1) and we present the obtained results (Section 4.2). Finally, conclusions are drawn in Section 5.

2. MODELING

A gas distribution network is a hybrid system combining the fluid dynamic behavior of gas pressure and flow rate with the temporal behavior of the actions that are taken to recover from the failure of a network element. To decouple the topological representation of the fluid dynamics from the stochastic temporal characterization of the failure management process, we assume that the operating conditions of the network are completely determined by given pressures and mass flow rates when a failure occurs, thus neglecting any change of the network operation during the steps of failure management. This results in a *fluid-dynamics model* and a *failure management model*.

In standard operating conditions, gas enters the network through supply nodes and is withdrawn at load nodes (i.e., end-users). While the pressure at the supply nodes is kept at a constant value controlled by pressure-regulating devices, the pressure at the load nodes depends on flow patterns in the network. To guarantee the correct operation of the overall network as well as to meet commercial standards, the pressure at each load node is required to be greater than a certain threshold. According to this, whenever a network component fails, a set of actions are taken to restore the correct network behavior. It is worth noting that, while we consider failures of network pipes, the approach can be easily extended to encompass failures of other components (e.g., regulating devices, valves, etc).

A pipe failure is detected by people (because gas is smelled) or by some automated system. While in the former case a failure detection entails its localization, in the latter case the network is scanned to locate the leak (as a safety measure, a sub-network is shut down while it is scanned). Then, the network is reconfigured to guarantee safety at load nodes and allow the faulty pipe to be repaired (e.g., the nearest valves positioned upstream and downstream of the leak are closed to isolate the faulty section, while other sectioning valves are opened if needed to reduce the number of end-users disconnected from the network). After reconfiguration, the load nodes are distinguished into: *i) offline* nodes, which are disconnected due to the network topology (i.e., no path exists from any of the supply nodes to such load nodes); *ii) online served* nodes, which are connected with sufficient pressure to meet their respective requirements; and, *iii) online not served* nodes, which are connected but with pressure lower than the required threshold. During the pipe repair, the supplied pressure is raised by a certain amount ΔP so as to make all online nodes served. We arbitrarily divide the regulation operation into four steps, considering the online nodes that become served after a pressure increase of $0.25\ \Delta P$, $0.5\ \Delta P$, $0.75\ \Delta P$, and ΔP, respectively (the number of steps could be increased to obtain a finer grained evaluation without impairing the subsequent analyses). When the faulty pipe has been repaired, the network is brought back to the original configuration and the pressure in supply nodes is lowered, undoing the changes made during the reconfiguration and regulation phases, respectively.

2.1 Fluid-dynamics model

Given a set of boundary conditions, represented by pressure at supply nodes and mass flow rates withdrawn at load nodes, fluid-dynamic analysis is performed to assess the network state in terms of pressures at nodes and mass flow rates in pipes [10]. Specifically, two sets of equations are considered. Equation 1 states that the signed sum of mass flow rates that enter or exit from each node n (i.e., Q_{in} and Q_{nj}, respectively), must be equal to zero or to the withdrawn mass flow rate Q_n^w depending on whether n is a passive node (i.e., neither a supply nor a load node) or a load node, respectively, i.e.,:

$$\sum_{i \in I_n^{\text{ent}}} Q_{in} - \sum_{j \in I_n^{\text{ex}}} Q_{nj} = \begin{cases} 0 & \forall \text{ passive node } n, \\ Q_n^w & \forall \text{ load node } n, \end{cases} \quad (1)$$

where I_n^{ent} and I_n^{ex} are the sets of indexes of pipes that enter and exit from node n, respectively. This yields a set of N_q equations in M unknowns, where N_q is the number of passive/load nodes, M is the number of pipes in the network, the unknowns are the mass flow rates in the pipes, and the N_l mass flow rates withdrawn at load nodes are given as an input. Equation 2 states that the pressure loss per unit length in each pipe m is calculated through the Darcy-Weisbach formula:

$$\frac{\partial P_m}{\partial L_m} = \frac{f}{D_m} \cdot \frac{\rho V^2}{2} \qquad \forall \text{ network pipe } m, \quad (2)$$

where ρ is the gas density, V is the average gas velocity, D_m is the pipe diameter, and f is the Darcy friction factor calculated by means of the Colebrook equation for turbulent flows and the Poiseuille formula for laminar flows. This leads to a set of M equations in N_q unknowns, where the unknowns are the pressure values at load/passive nodes and the N_s pressure values at supply nodes are given as an input. Coupling Equations 1 and 2 leads to a non-linear system of $N_q + M$ equations in $N_q + M$ unknowns, which is solved through an iterative procedure based on the Newton-Raphson method.

For each possible pipe failure, fluid-dynamic analysis is repeated under different boundary conditions which refer to a specific step in the failure management process. More specifically, fluid-dynamic analysis is performed:

- after the automated detection of a failure, to assess whether each load node that does not belong to the faulty subnetwork is either offline or online and, in the latter case, to determine whether the pressure level meets the contractual requirements of the node (note that the nodes of the faulty subnetwork are definitely offline, as such subnetwork is shut down and isolated while the fault is being localized);
- after the network reconfiguration, to determine whether each load node is offline or online with sufficient or insufficient pressure level;
- after the network reconfiguration in the presence of online load nodes that are not able to withdraw the gas with sufficient pressure, to find out whether a pressure increase in some of the supply nodes can re-establish the proper service level for such nodes (i.e., an inverse fluid-dynamic analysis is carried out to evaluate the minimum pressure increase ΔP to be achieved at some supply nodes so that all online load nodes are served);
- after each step of pressure regulation, to discern the load nodes that are still not served, and the total mass flow rate required by the load nodes but not delivered at a sufficient pressure level.

2.2 Stochastic timed characterization of the process of failure management

Each step of the failure management process has an execution time characterized by a probability distribution, which can be derived from the results of fluid-dynamic analysis (i.e., pressure regulation and undo regulation), or the network topology and the failed pipe (i.e., failure localization), or practice-based considerations (i.e., automated/manual failure detection, network reconfiguration, pipe repair, and undo network reconfiguration). This requires a deep knowledge of the network operation and the failure management procedure (e.g., the working principles of some smart devices may notably reduce the time needed to disconnect a faulty pipe, the maximum time needed to repair a component may descend from contractual requirements established with third-party companies, etc). In the practice, it may be quite hard to get a good knowledge of such processes, as companies usually avoid disclosure of classified information. Therefore, we discuss stochastic characterization of a possible failure management procedure, defined on the basis of reasonable assumptions and the lessons learned so far in the research collaboration with Terranova [1]. While the issue is addressed for the completeness of discussion to provide a proof of concept of the approach, it is worth remarking that the focus of this research is not measuring/fitting stochastic temporal parameters of the process of failure management, but rather the methodology of modeling as well as the quantitative measures that can be evaluated through the subsequent stochastic analysis once such parameters are known.

Manual failure detection and localization. We assume that this step takes a time supported over $[0, \infty)$, as, in principle, a gas leak may be smelled immediately or not smelled at all. From a practical point of view, it does not appear likely to measure a statistics of the time elapsed between the occurrence of a gas leak and a notice from people, mainly because the time at which a pipe failed is actually unknown. Moreover, it may be hard to define a valid model as well, due to the many variables involved (e.g., population density, leak extent, wind speed and direction, etc) and thus an accurate estimation of variance and higher order moments is unlikely. Therefore, we assume that only the mean value μ_m is known and adopt an EXP distribution with rate $1/\mu_m$ (in the experiments, we arbitrarily assume $\mu_m = 8$ h).

Automated failure detection and localization. We consider a system that periodically compares cumulative mass balances in closed sub-networks (i.e., the integral over time of involved mass flow rates) so that the probability that a leak is detected depends on how long it has been active when a check is performed, yielding a quasi-triangular distribution. We assume that the mean value μ_a and variance σ_a^2 of such distribution are known, and we approximate it with an Erlang distribution having the same mean value and variance. In the experiments, we assume $\mu_a = 24$ h, $\sigma_a^2 = 96$ h^2, and an Erlang distribution with $k = 6$ and $\lambda = 0.25$.

We assume that a gas leak is located by means of a sensors-equipped vehicle, which scans the leak-affected sub-network by circulating on the streets above the pipes at a constant speed v_{loc}. We consider a time t_{locStart} to start the localization procedure and a time t_{displ} to displace the vehicle to another location to start scanning a new pipe (as a continuous path covering all the pipes only once may not exist). Once the vehicle path is established, the number of needed displacements N_d is derived and the min/max localization time is evaluated as the sum of the setup time (i.e., $t_{\text{locStart}} + N\ t_{\text{displ}}$) and the min/max scan time (easily calculated given the length of scanned pipes). As each point of each pipe has the same probability to suffer a leak, a uniform distribution over the identified min-max interval is assumed.

Network reconfiguration and undo reconfiguration. We arbitrarily assume that network reconfiguration and its undoing take a uniformly distributed time over $[1, 2]$ h. In a more accurate model, the network topology and the specific reconfiguration actions should be taken into account, e.g., the presence of manually/remotely operated valves, the length of the pipes and the vehicle used to drive across the network in case of manually operated valves.

Pressure regulation and undo regulation. The duration t_{reg} of a regulation step is easily calculated given the pressure increase rate r_{press} (allowed by the adopted regulating devices) and the pressure increase of each step (determined by the number of steps K and the total pressure increase $\Delta\ P$ derived by fluid-dynamic analysis). To allow a margin of laxity, we assume a uniform distribution over a 0.25 h-width domain centered around t_{reg}. The whole time needed for undoing the pressure regulation is easily obtained as $K\ t_{reg}$ and, in a similar manner to the previous case, we assume a uniform distribution supported over a 1 h-width domain centered around $K\ t_{reg}$.

Pipe repair. On the basis of conventional contractual requirements with end-users, we assume that the time to repair a pipe has a uniform distribution over $[24, 72]$ h.

2.3 Failure management model

The model is defined as a *stochastic Time Petri Net* (sTPN) [26, 6] extended with enabling and flush functions [7], which change the enabling condition of transitions and the token moves after each firing, augmenting the modeling convenience without restricting the expressivity of the model or disrupting stochastic analysis. We briefly recall sTPN syntax and semantics (Sections 2.3.1 and 2.3.2, respectively) and discuss the failure management model (Section 2.3.3).

2.3.1 STPN syntax

An sTPN is a tuple $\langle P; T; A^-; A^+; A^\bullet; m_0; EFT^s; LFT^s; \mathcal{F}; \mathcal{C}; E; L\rangle$. The elements $\langle P;\ T;\ A^-;\ A^+;\ A^\bullet;\ m_0;\ EFT^s; LFT^s; \mathcal{F}; \mathcal{C}\rangle$ comprise the model of *stochastic Time Petri Nets* (sTPN) [26, 6]. Specifically, P is a set of places; T is a set of transitions disjoint from P; $A^- \subseteq P \times T$, $A^+ \subseteq T \times P$, and $A^\bullet \subseteq P \times T$ are the sets of precondition, postcondition, and inhibitor arcs; $m_0 : P \to \mathbb{N}$ is the initial marking associating each place with an initial non-negative number of tokens; $EFT^s : T \to \mathbb{Q}_0^+$ and $LFT^s : T \to \mathbb{Q}_0^+ \cup \{\infty\}$ associate each transition $t \in T$ with a static *Earliest Firing Time* and a (possibly infinite) static *Latest Firing Time*, respectively, with $EFT^s(t) \leq LFT^s(t)\ \forall\ t \in T$; $\mathcal{C} : T \to \mathbb{R}^+$ associates each transition $t \in T$ with a weight; $\mathcal{F} : T \to F_t^s$ associates each transition $t \in T$ with a static Cumulative Distribution Function (CDF) supported over $[EFT^s(t), LFT^s(t)]$.

As typical in Petri Nets (PNs), a place p is called an *input*, an *output*, or an *inhibitor* place for a transition t if $\langle p, t\rangle \in A^-$, $\langle t, p\rangle \in A^+$, or $\langle p, t\rangle \in A^\bullet$, respectively. As usual in Stochastic Petri Nets (SPNs), a transition t is called *immediate* (IMM) if $[EFT^s(t), LFT^s(t)] = [0, 0]$ and *timed* otherwise; a timed transition t is said to be *exponential* (EXP) if $F_t^s(x) = 1 - e^{\lambda x}$ over $[0, \infty]$ for some rate $\lambda \in \mathbb{R}_0^+$ and *general* (GEN) otherwise; a GEN transition t is called *deterministic* (DET) if $EFT^s(t) = LFT^s(t) > 0$ and *distributed* otherwise, i.e., if $EFT^s(t) \neq LFT^s(t)$. For each distributed transition $t \in T$, we assume that its CDF F_t^s is absolutely continuous over its support $[EFT^s(t), LFT^s(t)]$ and, thus, that there exists a Probability Density Function (PDF) f_t^s such that $F_t^s(x) = \int_0^x f_t^s(y)dy$.

$E : T \to \{true, false\}^{\mathbb{N}^P}$ relates each transition with an *enabling function* $E(t) : \mathbb{N}^P \to \{true, false\}$ associating each marking $m : P \to \mathbb{N}$ with a boolean value. $L : T \to \mathcal{P}(P)^{\mathbb{N}^P}$ relates each transition to a *flush function* $L(t) : \mathbb{N}^P \to \mathcal{P}(P)$ associating each marking $m : P \to \mathbb{N}$ with a subset of P.

2.3.2 STPN semantics

The *state* of an sTPN is a pair $\langle m, \tau\rangle$ where marking $m : P \to \mathbb{N}$ associates each place with a non-negative number of tokens and $\tau : T \to \mathbb{R}_0^+$ associates each transition with a dynamic time-to-fire. A transition t is *enabled* by m if: *i)* each of its input places contains at least a token, i.e., $m(p) \geq 1\ \forall\ \langle p, t\rangle \in A^-$, *ii)* none of its inhibitor places contains any token, i.e., $m(p) = 0\ \forall\ \langle p, t\rangle \in A^\bullet$, and *iii)* its enabling function evaluates to true in m, i.e., $E(t)(m) = true$. An enabled transition t is *firable* in $s = \langle m, \tau\rangle$ if its time-to-fire is not higher than that of any other enabled transition, i.e., $\tau(t) \leq \tau(t')\ \forall\ t' \in T^e(m)$ where $T^e(m)$ is the set of transitions enabled by m. If multiple transitions are firable in s, one of them is selected according to the random switch determined by $\mathcal{C}$, i.e., $Prob\{t \text{ selected}\} = \mathcal{C}(t)/\sum_{t_i \in T_f} \mathcal{C}(t_i)$, where $T^f(s)$ is the set of transitions that are firable in s.

The state of an sTPN evolves depending on the times-to-fire sampled by transitions and the resolution of random switches according to the weights of transitions. Specifically, when a transition t fires, the state $s = \langle m, \tau\rangle$ is replaced by a new state $s' = \langle m', \tau'\rangle$. Marking m' is derived from m by: *i)* removing a token from each input place of t and assigning zero tokens to the places belonging to the set $L(t)(m)$ identified by the value of the flush function of t in m, which yields an intermediate marking m_{tmp}, and *ii)* adding a token to each output place of t, which finally yields m', i.e.,

$$
\begin{aligned}
m_{tmp}(p) &= \begin{cases} m(p) - 1 & \forall\ p\ .\ \langle p, t\rangle \in A^- \\ & \quad \wedge\ p \notin L(t)(m), \\ 0 & \forall\ p\ .\ p \in L(t)(m), \\ m(p) & \text{otherwise} \end{cases} \\
m'(p) &= \begin{cases} m_{tmp}(p) + 1 & \forall\ p\ .\ \langle t, p\rangle \in A^+, \\ m_{tmp} & \text{otherwise} \end{cases}
\end{aligned}
\tag{3}
$$

Transitions that are enabled both by m_{tmp} and by m' are said *persistent*, while those that are enabled by m' but not by m_{tmp} or m are said *newly-enabled*. If the fired transition t is still enabled after its own firing, it is always regarded as newly enabled [3, 25]. For any transition t_p that is persistent after the firing of t, the time-to-fire is reduced by the time elapsed in the previous state s (which is equal to the time-to-fire of t measured at the entrance in s):

$$\tau'(t_p) = \tau(t_p) - \tau(t). \tag{4}$$

For any transition t_n that is newly-enabled after the firing of t, the time-to-fire takes a random value sampled in the static firing interval according to the static CDF $F^s_{t_n}$:

$$EFT^s(t_n) \leq \tau'(t_n) \leq LFT^s(t_n),\quad Prob\{\tau'(t_n) \leq x\} = F^s_{t_n}(x). \tag{5}$$

2.3.3 The model structure

The failure management model accounts for sequencing and timing constraints of the actions taken in reaction to the failure of a network pipe. As a characterizing trait, the structure of the model as well as the distribution of various of its temporal parameters are independent of both the network topology and the failed pipe. This almost guarantees a constant level of complexity in the subsequent stochastic analysis based on the theory of stochastic state classes, which largely depends on the degree of concurrency among the active timers of the model and on the length of behaviors in which timers overlap their activities [6].

As shown in Figure 1, the IMM transition *fail* models the occurrence of a pipe failure. Its output places $p0$ and $p2$ are chained with *i)* the sequence of GEN transitions *autoDet* and *loc*, which model the automated failure detection and the subsequent failure localization, respectively, and *ii)* the EXP transition *manualDetLoc*, which accounts for the manual failure detection and localization.

Flush functions are associated with *manualDetLoc* and *loc* to account for the fact that the automated detection or the subsequent localization of a failure is stopped as soon as the failure is detected (and located) by report, and viceversa. According to this, $L(manualDetLoc)(m) = \{p0, p1\}\ \forall\ m : P \to \mathbb{N}$ s.t. $m(p2) = 1$, which flushes places $p0$ and $p1$ whenever *manualDetLoc* fires. Note that, in a similar manner, $L(loc)(m) = \{p2\}\ \forall\ m : P \to \mathbb{N}$ s.t. $m(p1) = 1$, which empties place $p2$ whenever *loc* fires.

When either *loc* or *manualDetLoc* fires, a token arrives in $p3$, enabling the GEN transition *reconfig* which models the network reconfiguration. When *reconfig* fires, a token arrives in both $p4$ and $p6$, which enables *i)* the four chained GEN transitions *reg1* through *reg4* representing the steps of pressure regulation, and *ii)* the GEN transition *rep* modeling the repair operation of the failed pipe. When both *rep* and *reg4* have fired and, thus, both places $p5$ and $p10$ contain a token, the IMM transition *undo* becomes enabled and fires. This deposits a token in both places $p11$ and $p13$, enabling the concurrent GEN transitions *undoReconfig* and *undoReg*, which model the restoration of the network configuration and pressure level, respectively. When both these transitions have fired, both places $p12$ and $p14$ contain a token. This enables transition *end*, whose firing models the completion of the procedure of failure management.

3. EVALUATION

As a relevant trait, the failure management model includes multiple concurrent GEN transitions, which goes beyond the limits of the so called enabling restriction and motivates the use of the solution technique proposed in [16] to perform transient stochastic analysis. Such theory of analysis supports the evaluation of the outage time of each load node, both in the transient regime and on the average. We recall here the salient traits of the analysis technique (Section 3.1), referring the reader to [16] for more details. We also discuss in detail the evaluated availability measures (Section 3.2).

3.1 Quantitative transient analysis

The solution technique of [16] supports the transient analysis of models with multiple concurrent GEN transitions, which has an underlying Generalized Semi-Markov Process (GSMP) with equal-speed timers [13, 9]. The state of the underlying GSMP is sampled after each transition firing and an additional timer called τ_{age} is maintained in order to account for the absolute elapsed time. This identifies an embedded Discrete Time Markov Chain (DTMC) called *transient stochastic graph* whose states are named *transient stochastic state classes* (transient classes for short).

Each transient class is made of a marking plus the joint support and PDF of τ_{age} and the times-to-fire of the enabled transitions. The marginal PDF of τ_{age} permits to derive the PDF of the absolute time at which a transient class can be entered, enabling the evaluation of continuous-time transient probabilities of reachable markings within a given time horizon, provided that the number of transient classes that can be reached within that time interval is either bounded or can be truncated under the assumption of some approximation threshold on the total unallocated probability.

The complexity of the solution technique can be reduced by applying the regenerative analysis of [16] in the case that the model underlies a Markov Regenerative Process (MRP) that always reaches a regeneration point, i.e., a state where the future behavior is independent from the past behavior through which it has been reached. The regenerative approach limits transient analysis to the first regeneration epoch and repeats it from every regenerative point, supporting the derivation of the local and global kernels that characterize the behavior of the MRP [8, 9, 4] and enabling the evaluation of the transient probabilities of reachable markings at any time through the numerical integration of generalized Markov renewal equations.

It is worth noting that, in the approach of [16], regeneration points can be easily identified as the transient classes

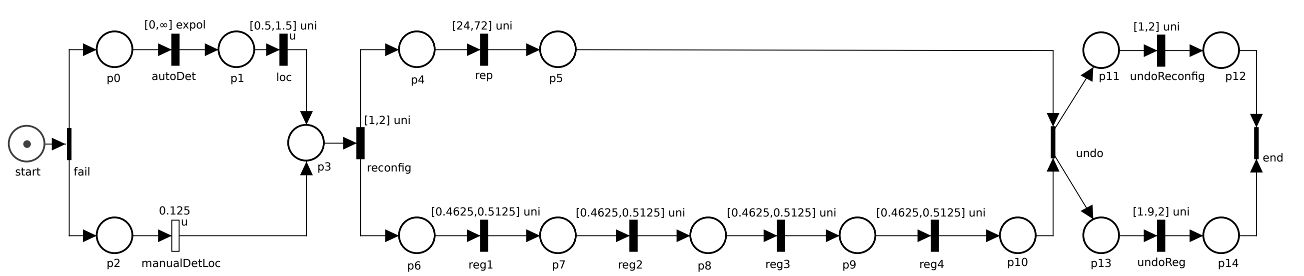

Figure 1: The sTPN specification of the failure management model. IMM, EXP, and GEN transitions are represented by thin bars, thick empty bars, and thick black bars, respectively. The distributions associated with timed transition refer to the example analyzed in Section 4.

where each times-to-fire is either *i)* newly-enabled, or *ii)* exponentially distributed, or *iii)* deterministic, or *iv)* bounded to take a deterministic delay with respect to a time-to-fire satisfying any of the previous conditions. According to this, the failure management model defined in Section 2.3 is guaranteed to reach a regeneration point at the completion of the following operations: automated/manual detection & localization of a failure, network reconfiguration, and repair & regulation. Therefore, transient evaluation through the regenerative analysis technique of [16] is allowed.

3.2 Evaluated measures

Availability measures are derived by leveraging and combining the results of fluid dynamic and stochastic analysis. On the one hand, for each pipe that can suffer a gas leak, regenerative analysis of the corresponding failure management model provides the probability over time of any marking that can be reached within a given time bound. A reachable marking comprises a logical state of the failure management model and corresponds to a specific operating condition of the network, e.g., in the failure management model of Figure 1, marking $p1$ $p2$ corresponds to the condition in which a fault has been automatically detected and it is being localized, marking $p4$ $p6$ corresponds to the condition in which the network has been reconfigured and the first step of pressure regulation is ongoing, etc. According to this, stochastic analysis permits to derive the transient probability of any operating condition of the network after the failure of one of its components.

On the other hand, fluid dynamic analysis is performed for each operating condition of the network (under the corresponding boundary conditions) in order to derive the service level of each load node (i.e., offline, online with insufficient pressure level, online with sufficient pressure level). According to this, the sum of the transient probability of network conditions in which a node experiences the same service level yields the transient probability that such node receives that service level, e.g., if a load node is online but with an insufficient pressure level starting from the network reconfiguration until the end of the third step of pressure regulation, then the probability over time that it is not served can be evaluated as the sum of the transient probability of any marking such that $p6$, or $p7$, or $p8$ contains a token (see the model of Figure 1). In doing so, the results of fluid dynamic analysis comprise a criterion to aggregate transient probabilities evaluated by stochastic analysis with the aim to derive availability measures for each load node.

In addition, the failure management model can be augmented with an absorbing place for each transition that represents the completion of a step of the failure management process, as shown in Figure 2. In doing so, for each added absorbing place p_a, the sum of the transient probability of any marking such that p_a contains a token comprises the CDF of the completion time of the input transition of p_a (measured since the entrance in the initial state). This enables the evaluation of the CDF of the completion time of any step of the failure management process, also permitting to derive average availability measures such as the mean value. For instance, the CDF of the completion time of the first step of pressure regulation can be evaluated as the sum of the transient probability of any marking such that place $reg1\_done$ contains a token.

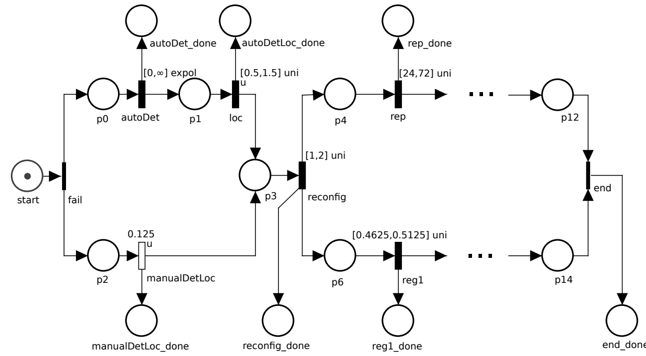

Figure 2: A fragment of a variant of the sTPN model shown in Figure 1, where an absorbing output place is added to each transition that represents the completion of a step of the failure management process.

The distribution of the time elapsed along a sequence of state transitions in the failure management model can be computed from the transient stochastic graph. In particular, a DET transition can be enabled in the initial state of the sequence with time-to-fire higher than the maximum sequence duration. Such transition, always enabled but never firable, keeps track of the time elapsed along the sequence of states without affecting the model behavior; its marginal PDF in the last state of the sequence provides the PDF of the sequence duration by a linear transformation, and it allows to compute moments such as the mean time.

To compute the mean time between a given pair of states m_1 and m_2, all the sequences from m_1 to m_2 need to be considered in the transient stochastic graph: the average time of each sequence is weighted by the probability of starting its execution (provided by the transient stochastic graph). Moreover, if regenerations can occur before the start of the sequence, the number of sequences to be analyzed can be reduced by leveraging the results of Markov renewal theory.

4. AN EXAMPLE

We illustrate here the gas distribution network considered in the experimental validation (Section 4.1) and we discuss the obtained results (Section 4.2).

4.1 Experimental setting

Figure 3 shows a topological representation of the gas distribution network analyzed in the experiments, which is identical to that presented in [14]. The network is made of a supply node, four load nodes marked as A through D, and eight pipes numbered from 4 to 12 (note that pipes numbered from 1 to 3 in [14] are not considered here as they are part of a high-pressure transmission network). The gas is provided by the supply node, while the sectioning valve belonging to pipe 9 is kept closed in ordinary operating conditions. According to this, the gas is supplied radially, so that load nodes A and B are served by pipes 4 and 7 (which comprise the *upper branch*), while load nodes C and D are served by pipes 5 and 11 (which comprise the *lower branch*).

With respect to [14], the operating parameters of the network components have been chosen so as to experience different degrees of network unavailability following different

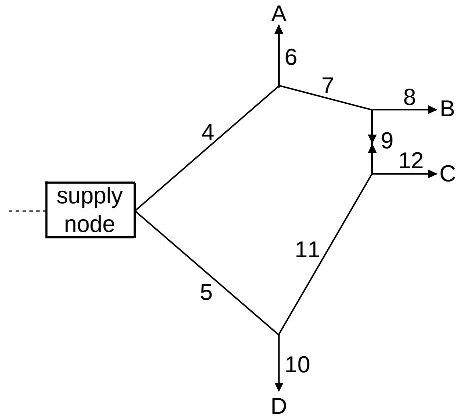

Figure 3: A gas distribution network.

pipe failures. Specifically, Table 1 reports the mass flaw rates withdrawn at load nodes and the pressure at the supply node (which are considered as input values for fluid dynamic analysis) as well as the minimum pressure required by each load node. During the regular operation of the network, the pressure in each load node is greater than the corresponding pressure threshold, so that all nodes are properly served.

Node	**Flow rate** (m^3/h)	**Setup pressure** (bar)	**Pressure threshold** (bar)
supply	-	3.5	-
A	200	-	3.0
B	200	-	3.0
C	150	-	3.0
D	200	-	3.0

Table 1: Operating parameters of the nodes of the gas distribution network shown in Figure 3 in ordinary operating conditions.

When a pipe failure is automatically detected on either the upper or the lower branch (the valve belonging to pipe 9 is closed), the gas supply to the entire branch is interrupted to safely search for the leak. Hence, all the load nodes belonging to the shut off branch become offline. Conversely, when a pipe failure is detected (and located) by human report, all load nodes remain online as the branch shut off is not performed. Once the failure is detected and located (either manually or automatically), two scenarios are distinguished depending on whether the failed pipe belongs to the ring or not, respectively.

- If the failed pipe belongs to the ring, then it is excluded from the network through a reconfiguration operation and the valve belonging to pipe 9 is opened. This permits to identify a path from the supply node to each load node so that all of them are online again. If some online load node experiences an insufficient pressure level, fluid dynamics calculations can be performed to determine the minimum pressure increase ΔP to be achieved at the supply node so that all online load nodes return to be properly served.
- If the failed pipe does not belong to the ring, then only the load node directly connected with that pipe remains offline and becomes served again at the end of the failure management process. In this special case, the activities of network reconfiguration, pressure regulation, undo reconfiguration, and undo regulation are skipped. Therefore, in the corresponding failure management model, transitions *reconfig*, *reg1* through *reg4*, *undoReconfig* and *undoReg* are IMM.

We consider failures of pipes 4 through 12, except for failures of pipe 9 which are not taken into account. In fact, as the sectioning valve belonging to pipe 9 is closed in ordinary operating conditions, its failures as well as failures of pipe 9 cannot be detected through the approach described in this paper, and would appear only during the step of network reconfiguration after the occurrence of a gas leak in another pipe. This envisages a different problem concerned with failures of spare components, which is not addressed in the present experiments.

For failures of pipes 4 through 8 and 10 through 12, the temporal parameters that characterize the failure management process can be derived as discussed in Sections 2.1 and 2.2. Specifically, given the length of each pipe reported in [14] and here shown in Table 2, the localization time is characterized by assuming a start-up time $t_{locStart} = 0.5$ h, a vehicle speed $v_{loc} = 5$ km/h, and a time $t_{displ} = 0.5$ h to move the vehicle from a pipe to another not contiguous pipe. For failures of pipes belonging to the ring, the pressure regulation time is characterized by assuming a pressure increase rate $r_{press} = 0.2$ bar/h and by using the results of the fluid dynamic analysis which provides the minimum pressure increase ΔP to be actuated at the supply node so that all online nodes are served (also shown in Table 2). Tables 3 and 4 finally show the obtained stochastic parameters.

pipe	**length** (km)	ΔP (bar)
4	5	0.63
5	5	0.39
6	3	-
7	4	0.2
8	5	-
10	3	-
11	5	0.12
12	3	-

Table 2: For the network shown in Figure 3, the length of pipes and the minimum pressure increase ΔP relative to failures of pipes on the ring.

pipe	**localiz.**	**regulation step**	**undo regulation**
4	$[0.5, 1.5]$	$[0.7625, 0.8125]$	$[3.05, 3.25]$
5	$[0.5, 1.5]$	$[0.4625, 0.5125]$	$[1.85, 2.05]$
6	$[3.6, 4.2]$	-	-
7	$[1.5, 2.1]$	$[0.225, 0.275]$	$[0.9, 1.1]$
8	$[2.1, 3.1]$	-	-
10	$[3.6, 4.2]$	-	-
11	$[1.5, 2.5]$	$[0.125, 0.175]$	$[0.5, 0.7]$
12	$[2.5, 3.1]$	-	-

Table 3: For the network of Figure 3, supports of the uniform distributions of the failure management model that depend on the network topology and the faulty pipe (times are expressed in hours).

automated detection	manual detection & localization	network reconfig.	pipe repair
Erlang $(6, 0.25)$	EXP (0.125)	uniform $[1, 2]$	uniform $[24, 72]$

Table 4: For the network shown in Figure 3, distributions of the failure management model that are independent of both the network topology and the faulty pipe (times are expressed in hours).

4.2 Experimental results

Without loss of generality, from now on we consider the case of a failure of pipe 5 as an example for discussion. Note that we deliberately focus on a failure of a pipe that belongs to the network ring, as such failures will leave more load nodes not served than failures of radial pipes.

As illustrated in Section 2.1, fluid dynamic analysis is repeated under different boundary conditions to assess the service level experienced by each load node after each step of the failure management process that changes either the network topology or the pressure at the supply node, i.e., the steps of automated failure detection, network reconfiguration, and pressure regulation (note that it is not necessary to evaluate service levels after the concurrent steps of undoing network reconfiguration and undoing pressure regulation, as they restore the ordinary operation conditions). Table 5 reports such results for failures of pipe 5.

failure management step	online served nodes	online not served nodes	offline nodes
automated detection	A,B	-	C,D
network reconfiguration	-	A,B,C,D	-
pressure regulation step 1	A	B,C,D	-
pressure regulation step 2	A	B,C,D	-
pressure regulation step 3	A,B	C,D	-
pressure regulation step 4	A,B,C,D	-	-

Table 5: Service level of each load node after each step of failure management that changes either the network topology or the pressure at the supply node, referred to a failure of pipe 5.

Stochastic analysis yields the transient probability of each logical state of the model, which corresponds to a specific operating condition of the network. Such probabilities are then aggregated on the basis of the results of fluid dynamic analysis reported in Table 5, which assess the service level received by each load node in each operating condition and thus permit to derive availability measures for end-users. Figure 4 shows the transient probability that nodes A, B, C, and D are not served after a failure of pipe 5.

- Node A experiences regular service during the phase of automated/manual failure detection and localization; after the network reconfiguration, node A is still online but with an insufficient pressure level, and it returns to be properly served at the completion of the first step of pressure regulation. The probability over time that node A is not served is initially zero, has a peak of 0.0558 in a neighborhood of 2.6 h, and then decreases up to being in the order of $5 \cdot 10^{-3}$ from time 35 h on.
- Node B is not properly served since the network reconfiguration until the completion of the third step of pressure regulation, respectively. Thus, the probability over time that node B is not served has a similar trend to the probability of node A, almost coincident in the interval $[0, 2]$ h, but it reaches a higher peak equal to 0.1585 at around 3.5 h.
- Conversely, nodes C and D are offline or not adequately served since the completion of automated or manual failure detection until the end of the fourth step of pressure regulation. Thus, their probability of being not served starts increasing from time 0 with a higher slope than the probability of nodes A and B; it has a peak equal to 0.3405 in a neighborhood of time 3.8 h; then, it progressively decreases up to being lower than $5 \cdot 10^{-3}$ from time 35 h on.

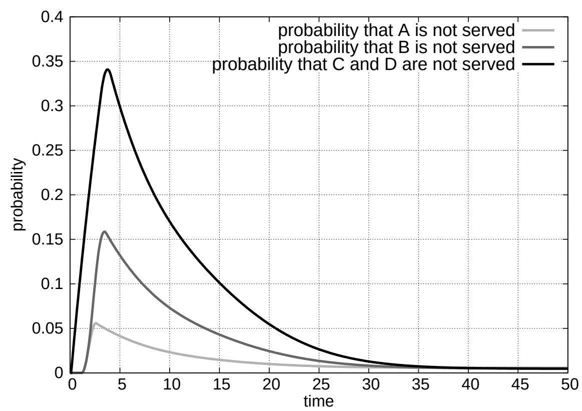

Figure 4: Probability that nodes A, B, C, and D are not served after a failure of pipe 5 (times are expressed in hours).

Table 6 shows the average availability measures obtained for load nodes A through D. Note that, if a statistics of failures is known, average availability measures for end-users could be derived over a long period of time.

- Given that node A is not served since the completion of the network reconfiguration until the end of the first regulation step, its average outage time is equal to the mean value of the stochastic distribution associated with the execution time of the first regulation step. As the latter is a uniform distribution supported over $[0.4625, 0.5125]$ h, the average outage time experienced by node A is equal to 0.4875 h.
- As node B is not served since the completion of the network reconfiguration until the end of the third regulation step, its average outage time is equal to the mean value of the stochastic distribution associated with the sum of the execution times of the first two regulation steps. Since all regulation steps have an execution time uniformly distributed between 0.4625 and 0.5125 h, and given that the mean value of the sum of two uniformly distributed random variables is equal to the sum of their mean values, the average outage time experienced by node B is equal to 0.9750 h.

- Given that nodes C and D are not served since the manual or automated failure detection until the completion of the last regulation step, their average outage time can be evaluated using an additional DET transition as illustrated in Section 3.2. This permits to derive the probability distribution of the outage time of C and D (shown in Figure 5), which has a mean value equal to 3.5318 h.

load node	**mean outage time experienced after a failure of pipe 5**
A	0.4875
B	0.9750
C	3.5318
D	3.5318

Table 6: Average availability measures a failure of pipe 5 (times are expressed in hours).

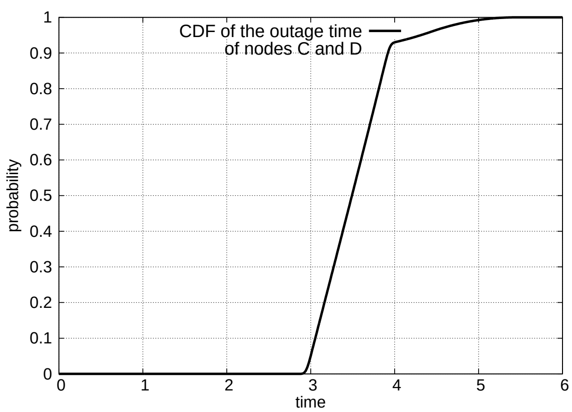

Figure 5: CDF of the time during which nodes C and D are not served.

As illustrated in Section 3.2, additional absorbing places can be used to derive the completion time CDF of any operation performed in reaction to a pipe failure. As a significant example, Figure 6 shows the completion time CDF of the overall failure management process, which has a mean value equal to 58.8866 h. Specifically, it was evaluated as the sum of the transient probability of any marking such that place *end_done* contains a token.

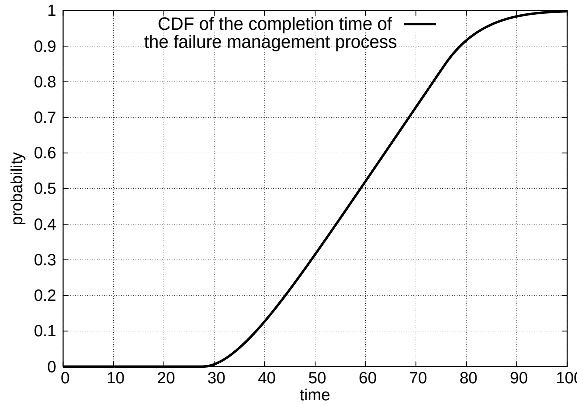

Figure 6: Completion time distribution of the failure management process.

5. CONCLUSIONS

We have presented the initial results of an on-going research project co-funded by Terranova [1], a company that provides innovative software solutions for support and automation of tele-metering and tele-management processes in gas, water, and electricity distribution networks. As a part of this collaboration, we have developed an approach for modeling and evaluation of the transient behavior of a gas distribution network after a pipe failure, with the intent to derive availability measures for each end-user. To this end, we have assumed that changes of the operating conditions of the network due to demand variations or demand-response applications occur in isolation, so as to be separated in time from the actions taken to recover from a pipe failure. This permits to decompose the system into a topological model of the fluid dynamics and a stochastic model of the failure management actions. As a characterizing trait, the outcomes of fluid dynamic analysis are used both to derive a stochastic characterization for parameters of the failure management model and to aggregate the results of stochastic analysis, providing transient and average availability measures for each end-user of the gas network. This comprises a valid basis for the evaluation of quality of service metrics such as the System Average Interruption Duration Index (SAIDI), and opens the way to the evaluation of different failure management processes and demand-response strategies.

It is worth remarking that the stochastic model includes concurrently enabled temporal parameters that have a non-Markovian distribution over a bounded support, which motivates the analysis through the method of transient stochastic state classes [16]. Moreover, while the approach has been experimented on a small-sized network of the literature to provide a proof of concept of the overall approach, the complexity of stochastic analysis is substantially independent of the network topology, thus appearing suitable to afford the analysis of cases of real complexity.

The proposed approach is open to many extensions that we are presently developing in collaboration with Terranova. As a remarkable aspect, we are working to relax the assumption that the actions executed to recover from a failure do not overlap with variations of the operating conditions of the network. To this end, we are analyzing several factors that may result in changes of the operation mode, including: drift of loading conditions (e.g., in reaction to seasonal/daily weather variations or due to demand-response processes); drift of physical parameters of the network components; discontinuities in the feed from the higher pressure level; planned operations on network actuators; unplanned failure and recovery of the network components.

Such an extension requires that the model of the failure management actions is recurrently rejuvenated with up-to-date results of the fluid-dynamic analysis, paving the way for the realization of a predictive analysis engine able to support on-line monitoring of gas distribution networks. In the theoretical perspective, this will raise various issues of scientific relevance, mainly concerned with the derivation of a time horizon within which the operating conditions of the network must be sampled at least once to maintain a reliable estimate on the future behavior. In the applicative perspective, this will lay the foundations for the development of innovative applications able to support multiple activities including the scheduling of measurement actions of the network operation, planning of ordinary maintenance/substitution operations of components subject to wear and tear, arranging of vehicle routes for regular scan operations in search of gas leaks, and triggering of escalation policies whenever a potentially critical condition is detected.

Acknowledgment

We thank Terranova for advice in comprehending the issues of gas distribution networks, and Regione Toscana for support within the programme "POR CRO FSE 2007-2013" under the specific project ERNESTO. We also thank Massimo Nocentini for his contribution in the experimentation.

6. REFERENCES

[1] http://www.terranovasoftware.eu/eu.

[2] A. Avritzer, S. Suresh, D. S. Menasché, R. M. M. Leão, E. de Souza e Silva, M. C. Diniz, K. Trivedi, L. Happe, and A. Koziolek. Survivability models for the assessment of smart grid distribution automation network designs. In *Proc. of the 4th ACM/SPEC Int. Conf. on Performance Engineering*, ICPE '13, pages 241–252, New York, NY, USA, 2013. ACM.

[3] B. Berthomieu and M. Diaz. Modeling and Verification of Time Dependent Systems Using Time Petri Nets. *IEEE Trans. on SW Eng.*, 17(3):259–273, March 1991.

[4] A. Bobbio and M. Telek. Markov regenerative SPN with non-overlapping activity cycles. *Int. Computer Performance and Dependability Symp. - IPDS95*, pages 124–133, 1995.

[5] A. J. Brito, A. T. de Almeida, and C. M. Mota. A multicriteria model for risk sorting of natural gas pipelines based on ELECTRE TRI integrating Utility Theory . *European Journal of Operational Research*, 200(3):812 – 821, 2010.

[6] L. Carnevali, L. Grassi, and E. Vicario. State-Density Functions over DBM Domains in the Analysis of Non-Markovian Models. *IEEE Trans. on SW Eng.*, 35(2):178–194, 2009.

[7] L. Carnevali, M. Paolieri, K. Tadano, and E. Vicario. Towards the quantitative evaluation of phased maintenance procedures using non-markovian regenerative analysis. In *Proc.* 10^{th} *European Workshop on Perf. Eng. (EPEW13)*, Sept. 2013.

[8] H. Choi, V. G. Kulkarni, and K. S. Trivedi. Markov regenerative stochastic petri nets. *Perform. Eval.*, 20(1-3):337–357, 1994.

[9] G. Ciardo, R. German, and C. Lindemann. A characterization of the stochastic process underlying a stochastic Petri net. *IEEE Trans. on SW Eng.*, 20(7):506–515, 1994.

[10] C. Colebrook. Turbulent flow in pipes, with particular reference to the transition region between smooth and rough pipe laws. *Journal of the Institution of Civil Engineers (London).*

[11] A. Conejo, J. Morales, and L. Baringo. Real-time demand response model. *IEEE Trans. on Smart Grid*, 1(3):236–242, 2010.

[12] A. Costa, J. d. Medeiros, and F. Pessoa. Steady-state modeling and simulation of pipeline networks for compressible fluids. *Brazilian Journal of Chemical Engineering*, 15:344 – 357, 12 1998.

[13] P. W. Glynn. A gsmp formalism for discrete-event systems. *Proceedings of the IEEE*, 77:14–23, 1989.

[14] A. Helseth and A. T. Holen. Reliability modeling of gas and electric power distribution systems; similarities and differences. In *Proc. of* 9^{th} *Int. Conf. on Probabilistic Methods Applied to Power Systems*, June 2006.

[15] A. Herrán-González, J. D. L. Cruz, B. D. Andrés-Toro, and J. Risco-Martín. Modeling and simulation of a gas distribution pipeline network. *Applied Mathematical Modelling*, 33(3):1584 – 1600, 2009.

[16] A. Horváth, M. Paolieri, L. Ridi, and E. Vicario. Transient analysis of non-Markovian models using stochastic state classes. *Performance Evaluation*, 2012.

[17] G. Koeppel and G. Andersson. Reliability modeling of multi-carrier energy systems. *Energy*, 34(3):235 – 244, 2009.

[18] M. Lanus, L. Yin, and K. Trivedi. Hierarchical composition and aggregation of state-based availability and performability models. *Reliability, IEEE Transactions on*, 52(1):44–52, 2003.

[19] T. Li, M. Eremia, and M. Shahidehpour. Interdependency of natural gas network and power system security. *IEEE Trans. on Power Systems*, 23(4):1817–1824, 2008.

[20] A. Martinez-Mares and C. Fuerte-Esquivel. Integrated energy flow analysis in natural gas and electricity coupled systems. In *North American Power Symposium (NAPS), 2011*, pages 1–7, 2011.

[21] J. Munoz, N. Jimenez-Redondo, J. Perez-Ruiz, and J. Barquin. Natural gas network modeling for power systems reliability studies. In *Power Tech Conf. Proc., IEEE Bologna*, volume 4, pages 8 pp. Vol.4–, 2003.

[22] B. R. Munson, D. F. Young, T. H. Okiishi, and W. W. Huebsch. *Fundamentals of Fluid Mechanics.* Wiley, 2010.

[23] Smart Grids Task Force of the European Commission. Mission and work programme. Technical report, 2012.

[24] J. Szoplik. *The Gas Transportation in a Pipeline Network, Advances in Natural Gas Technology.* Dr. Hamid Al-Megren (Ed.), ISBN: 978-953-51-0507-7, InTech, 2012.

[25] E. Vicario. Static analysis and dynamic steering of time dependent systems using Time Petri Nets. *IEEE Trans. on SW Eng.*, 27(1):728–748, August 2001.

[26] E. Vicario, L. Sassoli, and L. Carnevali. Using stochastic state classes in quantitative evaluation of dense-time reactive systems. *IEEE Trans. SW Eng.*, 35(5):703–719, 2009.

[27] Y. Yan, Y. Qian, H. Sharif, and D. Tipper. A survey on smart grid communication infrastructures: Motivations, requirements and challenges. *IEEE Communications Surveys Tutorials*, 15(1):5–20, 2013.