Contact during exam [Faglig kontakt under eksamen]:
Bjarne E. Helvik (92667)

EXAM IN COURSE [EKSAMEN I EMNE]
TTM4120 Dependable Systems [Pålitelige systemer]

Monday [Mandag] 2011-05-30
09:00 – 13:00

The English version starts on page 2.

Den norske bokmålsutgaven starter på side 9.

Hjelpemidler:
D - No printed or handwritten material is allowed. Predefined simple calculator [Ingen trykte eller håndskrevne hjelpemidler tillatt. Forhåndsbestemt enkel kalkulator]

Sensur 2011-06-20

English version[1]

A company delivering services over the Internet has a fault tolerant ICT system at its operational site. A sketch of the system is shown in Figure 1. It consists of two routers that are homed to different ASes in the Internet, and both of them are connected to two computing clusters. The asymptotic availability of routers, links and computing clusters are A_r, A_l and A_c. From a performance point of view, it is sufficient that one router and one cluster is working for the site to deliver adequate service.

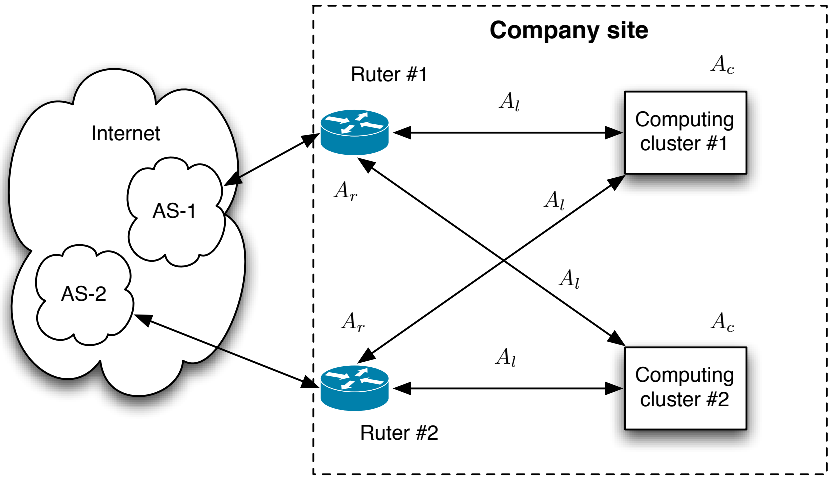

Figure 1: Operational site with fault tolerant ICT system.[Operasjonssenter med feiltolerant IKT system.]

The system has been operational for some time, and a dependability analysis is launched to improve its reliability and availability.

The routers have failed several times. Router #1 has failed eight times. The time to first failure and between the subsequent failures are 0.046, 0.027, 8.051, 0.265, 0.075, 0.024, 12.865 and 1.670 days. Down times after failures are negligible and the router is as new after it is restored/repaired after a failure, i.e the time to first failure and the times between failures have the same statistical properties.

a) Make a plot of the empirical reliability function based on these observations. What is the observed MTFF=MTBF in this case? What is the probability that the router will have a time between failures that is longer than the MTBF? Make the estimate directly from the observations.

[1] In case of divergence between the English and the Norwegian version, the English version prevails.

A similar problem is dealt with in exercise 1.

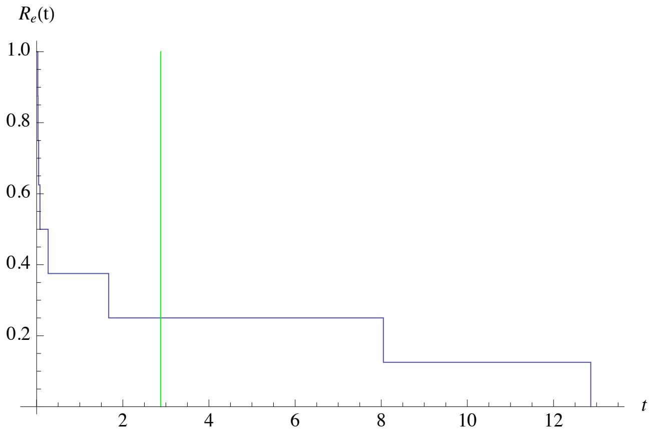

The observed MTFF=MTBF is 2.878 and R(MTFF) = 1/4

b) If a failure process Poissonian (e.g that of Router #1), what is the probability of observing a time between failures that is longer than the MTBF? Motivate the answer. If you should fit the failure data by a common distribution, which one would you choose? Motivate the answer, give the reliability function for the distribution and indicate ranges of parameter(s) if relevant. Let say you select a distribution with pdf f and parameters a and b, i.e., $f(t; a, b)$, *outline* the method you would use to obtain an estimator of the parameters. (It is required that you identify the method, present its principle and give a formal description of how the principle is applied. It is not expected that you derive a specific estimator or estimate parameters from the data.)

A Poisson process has n.e.d. times between failures. Hence, $R(\text{MTBF}) = \exp(-\text{MTBF}/\text{MTBF}) = e^{-1} = 0.37$.

The observed distribution has a heavier tail, are more skewed and have a larger variance than we would anticipate from a Poisson distribution. Due to the few observations, it is hard to suggest which distribution that fits best, but a hyper-exponential $R(\tau) = \int_\tau^\infty (a\text{Exp}[-\lambda 1t] + (1-a)\text{Exp}[-\lambda 2t])dt = \frac{e^{-(\lambda 1+\lambda 2)\tau}\left(-(-1+a)e^{\lambda 1\tau}\lambda 1+ae^{\lambda 2\tau}\lambda 2\right)}{\lambda 1\lambda 2}$, a gamma distribution $R(\tau) = \int_\tau^\infty \lambda/\Gamma(\alpha)$ (λt)^(α-1) Exp$[-\lambda t]dt$ with shape parameter less than 1 or a Weibull distribution (which it is) $R(\tau) = Exp[-(\lambda t)^\alpha]$with shape parameter less than 1 may be used.

The method (suggested used in the syllabus and in gen. stat. theory) is maximum likelihood estimation (see exercise 1 and textbook). It find the values of the parameters that maximises the probability of the observed outcome. Formally: let $\{T_{FF_1}, T_{FF_2}, \ldots, T_{FF_n}\}$ be a random sample. The maximum likelihood estimator (MLE) $\hat{a}, \hat{b}$ of a, b is the values that gives the maximum of the likelihood function

$$L(T_{FF_1}, T_{FF_2}, \ldots, T_{FF_n}; a, b) = f(T_{FF_1}; a, b) \cdot f(T_{FF_2}; a, b) \cdot \ldots \cdot f(T_{FF_n}; a, b)$$

since the observations are i.i.d.

Let now $\{t_1, t_2, \ldots, t_n\}$ denote the observed values in the sample. A maximum likelihood estimate $\hat{a}, \hat{b}$ of a, b is the maximum of the likelihood of the sample. $\hat{a}, \hat{b} = \max_{a,b}\left(L(t_1, t_2, \ldots, t_n; a, b)\right) = \max_{a,b}\left(\prod_{i=1}^{n} f(t_i; a, b)\right)$

Doing this is usually easier when we take $L^*(\ldots; a, b) = \ln(L(\ldots; a, b))$. The two functions will have the same extrema since $\ln$ is i strictly increasing function. The extrema of $L^*(\ldots; a, b)$ are the values of a, b for which the first derivatives of $L^*(\ldots; a, b)$ equals 0, i.e.,

$$\frac{\mathrm{d}L^*(\ldots; a, b)}{\mathrm{d}a} = 0, \frac{\mathrm{d}L^*(\ldots; a, b)}{\mathrm{d}b} = 0$$

In addition we have to ensure that the values found yields the global maximum.

c) What is the criteria for a system have a series parallel structure? Does the site have a series parallel structure? Assume that routers, links and computer clusters fail and are repaired independently of each other. What technique can be used to enable a reliability block diagram analysis of the site? Use this technique and find the availability of the service from the site.

See lecture notes at page 67. A system is series parallel if it may be divided in a set of subsystems $\{S_1, \ldots, Sn\}$ which ha a series, parallel or k-out-of-n structure between the subsystems. The procedure must be (recursively) repeatable for each subsystem S_i until the subsystems constitutes the primitive (un dividable) elements of the system.

The structure of the site is not series parallel.

We may analyse the system by using the pivoting technique to transform the site into four structures heaving the series-parallel property that enable a simple calculation. In this case, it is for symmetry and an obviously always failed structure, necessary only to investigate two.

Pivoting on the routers, we get the following diagrams and calculations (using notation from the mma package in the exercises) :

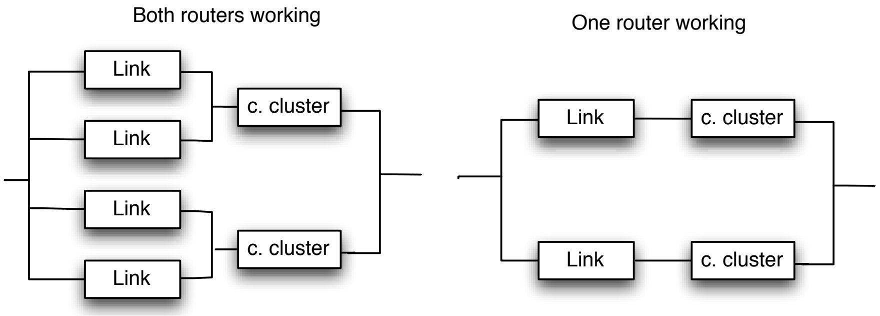

i	0	1	2	3	4	5	6	7	8	9
Φ_i	{ }	{r1}	{r2}	{l11}	{l12}	{l21}	{l22}	{c1}	{c2}	{r1, r2}
i	10	11	12	13	14	15	16	17	18	19
Φ_i	{r1, l11}	{r1, l12}	{r1, l21}	{r1, l22}	{r1, c1}	{r1, c2}	{r2, l11}	{r2, l12}	{r2, l21}	{r2, l22}
i	20	21	22	23	24	25	26	27	28	29
Φ_i	{r2, c1}	{r2, c2}	{l11, l12}	{l11, l21}	{l11, l22}	{l11, c1}	{l11, c2}	{l12, l21}	{l12, l22}	{l12, c1}
i	30	31	32	33	34	35	36			
Φ_i	{l12, c2}	{l21, l22}	{l21, c1}	{l21, c2}	{l22, c1}	{l22, c2}	{c1, c2}			

Table 1: Indexing of all operational modes with two or fewer failed network elements. ri indicates that router i has failed, ci indicates that cluster i has failed and lij indicates that the link between router i and cluster j has failed. (The third row is intentionally left blank for your use.) [Indeksering av alle operative modi med to eller færre feilte nettelementer. ri indikerer at router i har feilet, ci indikerer at klynge i har feilet og lij indikerer at lenken mellom ruter i og klynge j har feilet. (Den tredje raden er med hensikt latt være tom og er for ev. bruk under eksamen.)]

$R = \{\mathrm{Ar}, \mathrm{MUTr}\}; L = \{\mathrm{Al}, \mathrm{MUTl}\}; \mathrm{Cc} = \{\mathrm{Ac}, \mathrm{MUTc}\}\{\mathrm{Ac}, \mathrm{MUTc}\}$
Both routers ok: $\mathrm{A1} = ((L \sqcup L) \sqcap \mathrm{Cc}) \sqcup ((L \sqcup L) \sqcap \mathrm{Cc})$//First
$1 - (1 - \mathrm{Ac}\,(1 - (1 - \mathrm{Al})^2))^2$

one router ok: $\mathrm{A2} = (L \sqcap \mathrm{Cc}) \sqcup (L \sqcap \mathrm{Cc})$//First
$1 - (1 - \mathrm{AcAl})^2$

System
$A = \mathrm{A1Ar}\hat{}2 + 2 * \mathrm{A2Ar}(1 - \mathrm{Ar}) + 0(1 - \mathrm{Ar})\hat{}2$
$= 2\,(1 - (1 - \mathrm{AcAl})^2)\,(1 - \mathrm{Ar})\mathrm{Ar} + \left(1 - (1 - \mathrm{Ac}\,(1 - (1 - \mathrm{Al})^2))^2\right)\mathrm{Ar}^2$
$= -\mathrm{AcAlAr}\left(-4 + 2\mathrm{AlAr} + \mathrm{AcAl}\left(2 + \left(2 - 4\mathrm{Al} + \mathrm{Al}^2\right)\mathrm{Ar}\right)\right)$

The ICT system will have a number of operational modes depending on which network elements are working or not. A subset of these is listed in Table 1.

d) Show how an upper and lower bound of the availability of the services from the site may be obtained by considering only operational modes with two or fewer failed network elements. Obtain the expression for the upper bound with the same assumptions as in c) above.

An indicator function is introduced for whether the service i delivered or not. In this case it is most convenient to let the function be one when the service has failed, i.e.,

$$I(\Phi_i) = \begin{cases} 1 & \text{Service failed} \\ 0 & \text{Service working} \end{cases}$$

We see that $I(\Phi_9) = I(\Phi_{36}) = 1$, otherwise it is 0 in the modes listed in Table 1. We denote the set of modes listed in Table 1 the dominant modes D. A lower bound on the availability is obtained when we assume that the service has failed in all modes outside the dominant. Similarly an upper bound is obtained be assuming the the service is working for all modes in the outside dominant modes. Hence,

$$\begin{aligned} 1 - \sum_{i \in D} P(\Phi_i) I(\Phi_i) - \sum_{i \notin D} P(\Phi_i) &= 1 - \sum_{i \in D} P(\Phi_i) I(\Phi_i) - (1 - \sum_{i \in D} P(\Phi_i)) \\ &= \sum_{i \in D} P(\Phi_i)(1 - I(\Phi_i)) \\ &\leq A \leq 1 - \sum_{i \in D} P(\Phi_i) I(\Phi_i) \end{aligned}$$

It is seen that $P(\Phi_9) = (1 - A_r)^2 A_c^2 A_l^4$ and $P(\Phi_{36}) = (1 - A_c)^2 A_r^2 A_l^4$. Hence, $A \leq 1 - [A_l^4((1 - A_r)^2 A_c^2 + (1 - A_c)^2 A_r^2)]$.

In the remaining questions we concentrate on a computing cluster. It consists of n processors operating in load sharing. $k \leq n$ out of these are needed to provide the service with sufficient quality. In the configuration in Figure 1, one cluster is active taking the entire load, the other one is a stand-by. The intensity of permanent, i.e. hardware, failures are negligible and we may assume that *all* failures are transient or due to logical faults. When a processor fails, it is restarted in a negative exponential time with expectation μ^{-1}. Processors restart independently of each other and all restarts are i.i.d. The failure intensity per processor is λ. When a processor fails, with probability $(1 - c)$, *all* processors in the same cluster stop and have to be restarted.

- **e)** Establish a Markov model (state diagram) which may be used to determine the asymptotic availability, A_c, of the computing cluster. Let p_i be the probability that the cluster has i failed processors. Establish a set of equations that may be used to find $p_i, i = 0, \ldots, n$ and A_c for general n. (The equations shall not be solved.)

The state diagram

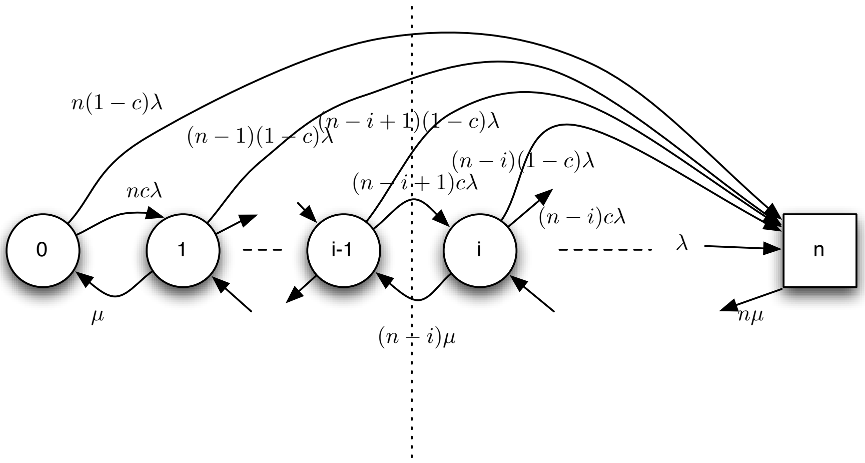

$A_c = \sum_{i=0}^{n-k} p_i$.

$\sum_{i=0}^{n} p_i = 1$.

Establishes a balance equation across all sections between states as illustrated in the diagram

$p_i i\mu = p_{i-1}(n-i+1)\lambda + \sum_{j=0}^{i-2} p_j (n-j)(1-c)\lambda, \quad i = 2, ..., n$ and $p_1\mu = p_0 n\lambda$.

In Figure 2 the unavailabilities of the computer cluster, i.e., $U_{6,n} = 1 - A_c$, is obtained from the equations found in e) for $n = 7, \ldots, 14$, in the case where $k = 6$ when $\mu = 1$ and $\lambda = 0.001$. The two following models of the coverage factor c are considered:

1. $c = 0.992$, i.e., a constant probability that a processor failure does not cause the entire cluster to fail.

2. $c = 0.999 - 0.001n$, i.e., the probability that a processor failure does not cause the entire cluster to fail decreases linearly with the number of processors in the cluster.

Six processors are sufficient to provide the services from the site. In the configuration in Figure 1 each computing cluster has $n = 7$ processors, i.e. one processor provides the spare capacity in the cluster in case of failure. A new configuration is considered, where the two clusters are combined and all processors are operated as one load-shared cluster of $n = 14$ processors. The objective of this is to be able to tolerate more failures.

f) What is the number of failed processors that can be tolerated in the two configurations? Will the configuration of Figure 1 or the new single cluster configuration provide the highest availability? Answer the question for both the case with coverage model 1 and the case with coverage model 2 presented above.

The current configuration can tolerate 3 failed processors, two in one cluster and one in the operative. (Of course more failures may be “tolerated in the std-by cluster, in total up to eight, but all but one in the standby.) The new configuration may tolerate 8 failures.

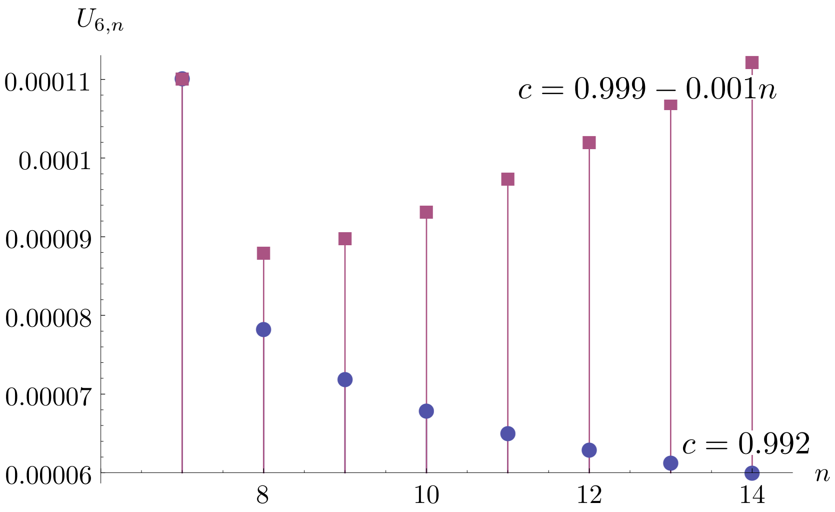

Figure 2: Unavailability of a computing cluster for a range of spare capacities and two coverage models. [Utilgjengelighet av en datamaskinklynge for ulik reservekapasitet og to dekningsmodeller.]

The two clusters operate independently. Hence, the availability will in both cases be $1 - U_{6,7}^2 = 1 - 0.00011^2 = 1 - 1.21 * 10^{-8}$, which is higher than for the new configuration in both cases.

g) Assume that we have ideal fault handling, i.e., $c = 1$. Are there any dependencies between the processors in this case? Under this assumption, will the configuration of Figure 1 or the new single cluster configuration provide the higher availability? Hint: it is not necessary to perform numerical calculations to answer the question.

They are independent, since the fail independently and are restored independently after failures.

Let i be the number of failed processors in cluster #1 and j be the number of failed processors in cluster #2. The probability of this is $p_i p_j$. This probability is the same whether they are merged into one cluster or not. Hence, we see that

$$A_{old} = \sum_{(i \leq 1 \wedge j \leq 7) \vee (i \leq 7 \wedge j \leq 1)} p_i p_j < A_{new} = \sum_{j+i \leq 8} p_i p_j$$

since $\{(i \leq 1 \wedge j \leq 7) \vee (i \leq 7 \wedge j \leq 1)\} \subset \{j + i \leq 8\}$. Under the the independence assumption the new single cluster configuration is best.

The brute force way to show this is to compute $A_c(n) = \sum_{i=6}^{n} \binom{n}{i} (\frac{\mu}{\mu+\lambda})^i (\frac{\lambda}{\mu+\lambda})^{n-i}$ and compare $A_c(14) \approx 1 - 10^{-16}$ with $1 - (1 - A_c(7))^2 \approx 1 - 4.4\ 10^{-10}$.