Advance Reservations for Distributed Real-Time Workflows with Probabilistic Service Guarantees* SOCA 09 Submission - Additional material

Tommaso Cucinotta
Real-Time Systems Laboratory
Scuola Superiore Sant'Anna, Pisa, Italy
Email: cucinotta@sssup.it

Kleopatra Konstanteli and Theodora Varvarigou
Advanced Distributed Computing Laboratory
National Technical University of Athens, Athens, Greece
Email: {kkonst,dora}@telecom.ntua.gr

V. Formalization of the problem

B. Probabilistic Formalisation

1) Probabilistic response-time guarantees: The response-time constraints may be relaxed in a probabilistic sense, if, instead of relying on worst-case estimates for the computation requirements $\left\{c_{i,j}^{(a)}\right\}$, as well as the message sizes $\left\{m_i^{(a)}\right\}$, they are (more effectively, for multimedia) considered as non-completely known values, and modeled as stochastic variables. For the sake of simplicity, it is assumed that they are independent and identically distributed (i.i.d.), and that the provider has an estimate of a certain quantile of their distributions: $\Pr\left[c_{i,j}^{(a)} \leq C_{i,j}^{(a)}\right] \geq \alpha_i^{(a)}$, and $\Pr\left[m_i^{(a)} \leq M_i^{(a)}\right] \geq \beta_i^{(a)}$, with $\prod_{i \in \mathcal{A}^{(a)}} \alpha_i^{(a)} \beta_i^{(a)} \geq \phi^{(a)}$. Then, introducing the notation $c_i^{(a)} \triangleq \sum_{j \in \mathcal{H}} x_{i,j}^{(a)} c_{i,j}^{(a)}$, $C_i^{(a)} \triangleq \sum_{j \in \mathcal{H}} x_{i,j}^{(a)} C_{i,j}^{(a)}$ and $L_i^{(a)} \triangleq \sum_{s \in \mathcal{S}} y_{i,s}^{(a)} L_s$, then the probability that the response-time is $\leq R^{(a)}$ may be written as:

(*) The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7 under grant agreement n.214777 IRMOS – Interactive Realtime Multimedia Applications on Service Oriented Infrastructures.

$$
\begin{aligned}
&\Pr\left[\rho^{(a)} \le R^{(a)}\right] = \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(\rho_i^{(a)} + \frac{m_i^{(a)}}{b_i^{(a)}} + \sum_{s\in\mathcal{S}} y_{i,s}^{(a)} L_s\right) \le R^{(a)}\right] = \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(\rho_i^{(a)} + \frac{m_i^{(a)}}{b_i^{(a)}} + L_i^{(a)}\right) \le R^{(a)} \mid \bigwedge_{i\in\mathcal{A}^{(a)}} c_i^{(a)} \le C_i^{(a)} \wedge \bigwedge_{i\in\mathcal{A}^{(a)}} m_i^{(a)} \le M_i^{(a)}\right] \cdot \\
&\cdot \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[c_i^{(a)} \le C_i^{(a)}\right] \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[m_i^{(a)} \le M_i^{(a)}\right] + \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(\rho_i^{(a)} + \frac{m_i^{(a)}}{b_i^{(a)}} + L_i^{(a)}\right) \le R^{(a)} \mid \neg\left(\bigwedge_{i\in\mathcal{A}^{(a)}} c_i^{(a)} \le C_i^{(a)} \wedge \bigwedge_{i\in\mathcal{A}^{(a)}} m_i^{(a)} \le M_i^{(a)}\right)\right] \cdot \\
&\cdot \left(1 - \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[c_i^{(a)} \le C_i^{(a)}\right] \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[m_i^{(a)} \le M_i^{(a)}\right]\right) \ge \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(\rho_i^{(a)} + \frac{m_i^{(a)}}{b_i^{(a)}} + L_i^{(a)}\right) \le R^{(a)} \mid \bigwedge_{i\in\mathcal{A}^{(a)}} c_i^{(a)} \le C_i^{(a)} \wedge \bigwedge_{i\in\mathcal{A}^{(a)}} m_i^{(a)} \le M_i^{(a)}\right] \cdot \\
&\cdot \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[c_i^{(a)} \le C_i^{(a)}\right] \prod_{i\in\mathcal{A}^{(a)}} \Pr\left[m_i^{(a)} \le M_i^{(a)}\right] \ge \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(d_i^{(a)} + \frac{M_i^{(a)}}{b_i^{(a)}} + L_i^{(a)}\right) \le R^{(a)}\right] \prod_{i\in\mathcal{A}^{(a)}} \alpha_i^{(a)} \beta_i^{(a)} \ge \\
&\Pr\left[\sum_{i\in\mathcal{A}^{(a)}}\left(d_i^{(a)} + \frac{M_i^{(a)}}{b_i^{(a)}} + L_i^{(a)}\right) \le R^{(a)}\right] \phi^{(a)}
\end{aligned}
$$

And the expression inside the paper is obtained.

2) Probabilistic availability guarantee: The probability for an application workflow to find enough available resources when actually activated, to be constrained to be higher than $\xi^{(a)}$, may be formalized as follows.

Let $E_j^{(a)}$ denote the event that services of application a deployed on host j be active. Let $U_{R,j}^{(a)}$ denote the overall CPU share requirements on host j due to $\mathcal{A}^{(a)}$, and $U_{A,j}^{(a)}$ denote the stochastic variable representing the available computation power on j when $\mathcal{A}^{(a)}$ is activated. Similarly, let $B_{R,s}^{(a)}$ denote the overall network bandwidth requirements on subnet s due to $\mathcal{A}^{(a)}$, and $B_{A,s}^{(a)}$ denote the stochastic variable representing the available network bandwidth when $\mathcal{A}^{(a)}$ is activated. Also, let $\pi_i^{(a)} \triangleq r^{(a)} \frac{\sum_{j\in\mathcal{H}} x_{i,j}^{(a)} C_{i,j}^{(a)}}{u_i^{(a)}}$ denote the probability of activation of service i in application a, $\pi_{i,j}^{(a)} \triangleq 1 - (1 - \pi_i^{(a)}) x_{i,j}^{(a)}$ denote a variable with a value of $\pi_i^{(a)}$ if i is allocated on j and 1 otherwise, and $\pi_{i,\{s\}}^{(a)} \triangleq 1 - (1 - \pi_i^{(a)}) y_{i,s}^{(a)}$ denote a variable with a value of $\pi_i^{(a)}$ if i is allocated within subnet s and 1 otherwise. Finally, let $u_i^{(a)} \triangleq \frac{\sum_{j\in\mathcal{H}} x_{i,j} c_{i,j}^{(a)}}{d_i^{(a)}}$ be denote a variable with the computing requirements of $\tau_i^{(a)}$.

Then, assuming that all the $E_j^{(a)}$ events are independent among each other, the probability for an application pipeline $\mathcal{A}^{(a)}$ to find enough available resources, if activated at any time $t_k \in I^{(a)}$, may be formalized as:

$$
\begin{aligned}
& \prod_{j\in\mathcal{H}} \Pr\left[U_{R,\,j}^{(a)} \leq U_{A,\,j}^{(a)} \mid E_j^{(a)}\right] \prod_{s\in\mathcal{S}} \Pr\left[B_{R,\,s}^{(a)} \leq B_{A,\,s}^{(a)} \mid E_j^{(a)}\right] \\
\geq\; & \prod_{j\in\mathcal{H}} \sum_{\mathcal{B}\subset\mathcal{A}(t_k)\setminus\{a\}} \Pr\left[U_{R,\,j}^{(a)} \leq U_{A,\,j}^{(a)} \mid E_j^{(a)} \bigwedge_{b\in\mathcal{B}} E_j^{(b)} \wedge \bigwedge_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \overline{E_j^{(b)}}\right] \cdot \Pr\left[\bigwedge_{b\in\mathcal{B}} E_j^{(b)} \wedge \bigwedge_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \overline{E_j^{(b)}}\right] \cdot \\
\cdot\; & \prod_{s\in\mathcal{S}} \sum_{\mathcal{B}\subset\mathcal{A}(t_k)\setminus\{a\}} \Pr\left[B_{R,\,s}^{(a)} \leq B_{A,\,s}^{(a)} \mid E_j^{(a)} \bigwedge_{b\in\mathcal{B}} E_j^{(b)} \wedge \bigwedge_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \overline{E_j^{(b)}}\right] \cdot \Pr\left[\bigwedge_{b\in\mathcal{B}} E_j^{(b)} \wedge \bigwedge_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \overline{E_j^{(b)}}\right] = \\
=\; & \prod_{j\in\mathcal{H}} \sum_{\mathcal{B}\subset\mathcal{A}(t_k)\setminus\{a\}} \Pr\left[\sum_{i\in\mathcal{A}^{(a)}} x_{i,\,j}^{(a)} u_i^{(a)} \leq U_j - \sum_{b\in\mathcal{B}} \sum_{i\in\mathcal{A}^{(b)}} x_{i,\,j}^{(b)} u_i^{(b)}\right] \cdot \\
\cdot\; & \prod_{b\in\mathcal{B}} \prod_{i\in\mathcal{A}^{(b)}} \pi_{i,\,j}^{(b)} \prod_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \prod_{i\in\mathcal{A}^{(b)}} \overline{\pi_{i,\,j}^{(b)}} \cdot \\
\cdot\; & \prod_{s\in\mathcal{S}} \sum_{\mathcal{B}\subset\mathcal{A}\setminus\{a\}} \Pr\left[\sum_{i\in\mathcal{A}^{(a)}\setminus\left\{n^{(a)}\right\}} y_{i,\,s}^{(a)} b_i^{(a)} \leq B_s - \sum_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B},\, i\in\mathcal{A}^{(b)}\setminus\left\{n^{(b)}\right\}} y_{i,\,s}^{(b)} b_i^{(b)}\right] \cdot \\
\cdot\; & \prod_{b\in\mathcal{B}} \prod_{i\in\mathcal{A}^{(b)}} \pi_{i,\,\{s\}}^{(b)} \prod_{b\in\mathcal{A}(t_k)\setminus\{a\}\setminus\mathcal{B}} \prod_{i\in\mathcal{A}^{(b)}} \overline{\pi_{i,\,\{s\}}^{(b)}}
\end{aligned}
$$

The first inequality is due to the non-complete expansion of the conditioned probability rule ($\Pr[A] = \Pr[A \wedge (B \vee \neg B)] = \Pr[A \mid B]\Pr[B] + \Pr[A \mid \neg B]\Pr[\neg B] \geq \Pr[A \mid B]\Pr[B]$).

The quantities inside the $\Pr[\cdot]$ operators in the expression above are not stochastic anymore, but they are a function of the problem variables. Therefore, depending on the values of the problem variables, the corresponding probabilities in the above formulas have a value of 1 if the condition is satisfied and 0 otherwise. Therefore, it is possible to introduce additional boolean problem variables $v_{\mathcal{B}}^i$ and $w_{\mathcal{B}}^s$ for the purpose of encoding whether or not such conditions are met or not, on the available CPU and network bandwidths, respectively.

In order for a boolean variable v to encode whether or not an inequality $e \geq 0$ is satisfied by the other problem variables, we use the following template:

$$
\begin{aligned}
e &\geq K(v-1) \\
e &\leq Kv - \epsilon
\end{aligned}
$$

with a sufficiently large constant K and a sufficiently small constant ϵ. In fact, if $v = 1$, then the first constraint mandates $e \geq 0$, while if $v = 0$, then it says $e \geq -K$ which, for a sufficiently large K does not limit at all the possible values of e. For the same reason, if $v = 1$, then the second constraint is ineffective, while if $v = 0$, then the second constraint mandates that $e \leq -\epsilon$, which for a sufficiently small constant ϵ amounts to requiring $e < 0$. From a dual perspective, if $e \geq 0$, then the second constraint mandates $v = 1$ and the first one is ineffective, while if $e < 0$, then the first constraint mandates $v = 0$ and the second one is ineffective.

Therefore, in the long probability expression computed above, substituting the $v_{\mathcal{B}\cup\{a\}}^i$ variables in place of the first probability and the $w_{\mathcal{B}\cup\{a\}}^s$ in place of the second one, an expression $P(\mathcal{A}(t_k))$ is obtained which is function of the considered time-instant t_k, actually of the applications with advance reservations in place at that time $\mathcal{A}(t_k)$. Now, the unconditioned probability of finding all needed resources available (event $E^{(a)}$), considering all of the time-instants $t_k \in I^{(a)}$, reasoning with the grouped time-slices $I_h \in \mathcal{G}^{(a)}$, is easily formalized as follows:

$$\begin{aligned}
\Pr\left[E^{(a)}\right] &= \sum_{t_k \in \mathcal{I}^{(a)}} \Pr\left[E^{(a)} \mid t_k\right] \Pr\left[\mathcal{A}^{(a)}\, activated\, in\, t_k\right] = \\
&= \sum_{I_h \in \mathcal{G}^{(a)}} \Pr\left[E^{(a)} \mid I_h\right] \Pr\left[\mathcal{A}^{(a)}\, activated\, in\, I_h\right] = \\
&= \sum_{I_h \in \mathcal{G}^{(a)}} P(\mathcal{A}(\min I_h)) \frac{tn_h}{f^{(a)} - s^{(a)}},
\end{aligned}$$

which corresponds to the formalization explicited in the paper.