

A Note on Stable Flow-Equivalent Aggregation in Closed Networks

Giuliano Casale
College of William and Mary
Computer Science Department
23185-8795 Williamsburg, VA, USA
casale@cs.wm.edu
Office: +1-757-221-2029
Fax: +1-757-221-1717

Abstract

We introduce the Conditional Mean Value Analysis (CMVA) algorithm, an exact solution method for product-form load-dependent closed queueing networks that provides a numerically stable solution of models where the load-dependent Mean Value Analysis (MVA) is numerically unstable. Similarly to the MVA algorithm for constant-rate queues, CMVA performs operations in terms of mean quantities only, i.e., queue-lengths, throughput, response times. Numerical stability derives from a new version of the MVA arrival theorem for load-dependent models which is expressed in terms of mean queue-lengths instead of marginal probabilities. The formula is obtained by the analysis of the conditional state spaces which describe network equilibrium as seen by jobs during their residence times at queues. We also provide a generalization of CMVA to multiclass models that preserves the numerical stability property.

Keywords: Queueing Networks, Flow-Equivalent Aggregation, Product-Form, Numerical Stability

1 Introduction

The performance evaluation of distributed applications and computer systems is often performed by simple product-form queueing networks [1] which can be solved by efficient algorithms and are simple to parameterize [6]. Real systems with features such as finite capacity constraints, memory swapping, or high-variability in service times violate the underlying assumptions of product-form networks, but still can be approximated accurately by these models using the flow-equivalent aggregation technique [3,6] (also known as the Norton's theorem for queueing networks or the Chandy-Herzog-Woo method). In the flow-equivalent aggregation approach, the real system is modeled by a load-dependent queue, called flow-equivalent server, whose service rate equals in any feasible state of the network the observed throughput of the system; after this step, the

resulting load-dependent model is solved by the load-dependent Mean Value Analysis (MVA) algorithm [7,9]. Unfortunately, the load-dependent MVA often exhibits numerical difficulties under heavy-load conditions which eventually result in infeasible results such as negative throughputs [7]. The convolution method [2] can then be adopted, but it still does not grant numerical stability and requires the additional implementation of numerical scaling techniques to prevent floating-point range exceptions [2,5].

To simplify the analysis of flow aggregation models, we develop a new mean value analysis algorithm, called Conditional Mean Value Analysis (CMVA), which is numerically stable, simple to implement and generalizes effectively to multiclass models. These goals are achieved by studying the conditional marginal queue-length probabilities observed during job residence times at the flow equivalent server. The obtained characterization provides a new version of the load-dependent MVA arrival theorem [8] which involves only mean queue-lengths. As a result, the CMVA algorithm has a simple recursive structure that is very similar to the MVA for constant-rate models and enjoys similar numerical stability properties.

The organization of this work is as follows. After introducing required background in Section 2, we develop a conditional analysis of queue-lengths during job residence times in Section 3. The CMVA algorithm is introduced in Section 4 and generalized to the multiclass case in Section 5. In both sections we report numerical experiments on a large number of random models to validate the numerical stability of CMVA. Section 6 summarizes the contributions of this work.

2 Background

We consider a closed product-form queueing network composed by N single class jobs and M queues. The first $M - 1$ queues have constant-rate and mean service demand D_i , $1 \leq i \leq M - 1$; queue M is a flow-equivalent server with demand $D_M = 1$ and service rate function μ . The demands D_i , $1 \leq i \leq M - 1$, are the product between the mean service time per job visit S_i and the job visit ratio V_i to queue i for each visit to queue M assumed as reference. If present, the delay server has mean demand D_0 and is indexed by $i = 0$.

Let $\vec{n} = (n_0, n_1, \dots, n_M)$, $\sum_{i=0}^M n_i = N$, $n_i \geq 0$, be a feasible allocation of jobs across the network. The equilibrium probability of state \vec{n} is given by [1]

$$P[\vec{n} | N, t] = \frac{\prod_{i=0}^M D_i^{n_i}}{G(N, t) \prod_{n=1}^{n_M} \mu(t + n - 1)}, \quad (1)$$

where the normalizing constant $G(N, t)$ assures that probabilities sum to one and the *rate shift* t allows to select the feasible service rates of the flow-equivalent server as a range of values of the function $\mu(\cdot)$. That is, if $t = c$ the flow-equivalent server uses the rate $\mu(c + n_M - 1)$ when its current queue-length is n_M and for $t = 1$ it therefore operates with the standard rates $\mu(1), \mu(2), \dots, \mu(N)$. The following quantities derived from (1) are also considered throughout the derivations, see [2,7,9] for proofs. The marginal probability that

$j \geq 1$ jobs are enqueued at station i is

$$P[n_i = j | N, t] = \begin{cases} D_i X(N, t) P[n_i = j - 1 | N - 1, t], & i \leq i \leq M - 1, \\ \frac{D_M}{\mu(j + t - 1)} X(N, t) P[n_i = j - 1 | N - 1, t], & i = M, \end{cases} \quad (2)$$

in which $X(N, t) = G(N - 1, t)/G(N, t)$ denotes network throughput as measured at queue M assumed as reference for job completions. The probability that queue i is idle is instead

$$P[n_i = 0 | N, t] = G^{-i}(N, t)/G(N, t) = 1 - \sum_{j=1}^N P[n_i = j | N, t] = 1 - U_i(N, t), \quad (3)$$

where $U_i(N, t) = \sum_{j=1}^N P[n_i = j | N, t]$, $0 \leq i \leq M$, is queue i utilization and $G^{-i}(N, t)$ is the normalizing constant of a model with queue i less.

2.1 Numerical Instability in the Load-Dependent MVA

For a model with population N , the MVA algorithm recursively computes (2) from the solution of a model with $N - 1$ jobs by first determining the throughput $X(N, t)$ as [7]

$$X(N, t) = \frac{N}{D_0 + \sum_{i=1}^{M-1} D_i \sum_{j=1}^N j P[n_i = j - 1 | N - 1, t] + D_M \sum_{j=1}^N j P[n_M = j - 1 | N - 1, t] / \mu(t + j - 1)},$$

and then computing the marginal queue-length probabilities for the model with N jobs by (2)-(3). At the end of the computation, mean response times are immediately computed by Little's Law [6] as $R_i(N, t) = Q_i(N, t)/X(N, t)$, $0 \leq i \leq M$, where the mean queue-lengths $Q_i(N, t)$ are

$$Q_i(N, t) = \begin{cases} D_0 X(N, t), & i = 0, \\ \sum_{j=1}^N j P[n_i = j | N, t], & 1 \leq i \leq M. \end{cases} \quad (4)$$

The source of numerical instability of the load-dependent MVA is equation (3). If computed for utilization values $U_i(N, t)$ close to one, (3) can return negative values due to round-off errors. If the sign of $P[n_i = 0 | N, t]$ is negative, also the computation of the other probabilities by (2) can produce macroscopic errors which typically result in negative throughputs, see [7] for an illustrative example.

Numerical instability can be solved for constant-rate queues by using mean queue-lengths instead of marginal probabilities. However, no recurrence relation for mean queue-lengths of load-dependent queues is available in the literature and therefore the load-dependent MVA remains numerically unstable. In the next section, we overcome this limitation by a queue-length recurrence for flow-equivalent servers.

3 Conditional Queue-Lengths During Residence Times

Throughout this section, we study quantities defined on the probabilities $P[\bar{n} | n_i \geq 1, N, t]$ describing the network equilibrium behavior conditioned on a job residing at queue i . Zahorjan shows early in [10] the

analytical relation between the state probability $P[\vec{n} | N, t]$ and the conditional probability $P[\vec{n} | n_i \geq 1, N, t]$. The result is used to provide insights on the BCMP theorem [1], but has never found application to the computational solution of queueing network models. We fill this gap by deriving from this theory a recurrence relation for load-dependent queue-lengths that is similar to the simple queue-length recursion in the constant-rate MVA [8]. We begin by proving the following result.

Theorem 1 (Mean Queue-Length During Job Residence Times). *Let $Q_i^{cond}(n_i \geq 1, N, t)$ be the mean queue-length at queue i as seen by a resident job at the same queue. Then,*

$$Q_i(N, t) = U_i(N, t)[1 + Q_i^{cond}(n_i \geq 1, N, t)]. \quad (5)$$

Proof. The marginal probability $P[n_i = j | N, t]$ and its conditional counterpart $P[n_i = j | n_i \geq 1, N, t]$ are related by definition to each other by

$$P[n_i = j | n_i \geq 1, N, t] = \frac{P[n_i = j | N, t]}{P[n_i \geq 1 | N, t]}, \quad j \geq 1, \quad (6)$$

which is a re-normalization of $P[n_i = j | N, t]$, $j \geq 1$, over the conditional state space $\{\vec{n} | N, n_i \geq 1\}$. Noting that $P[n_i \geq 1 | N, t] = U_i(N, t)$ we have immediately

$$P[n_i = j | N, t] = U_i(N, t)P[n_i = j | n_i \geq 1, N, t], \quad (7)$$

and the computation of the mean queue-length at i becomes

$$Q_i(N, t) = \sum_{j=1}^N jP[n_i = j | N, t] = U_i(N, t) \sum_{j=1}^N jP[n_i = j | n_i \geq 1, N, t], \quad 1 \leq i \leq M \quad (8)$$

Because the event $n_i = 0$ is not allowed in the conditional state space $\{\vec{n} | N, n_i \geq 1\}$, we can insert the normalization condition $\sum_{j=1}^N P[n_i = j | n_i \geq 1, N, t] = 1$ in (8) and perform the additional step

$$Q_i(N, t) = U_i(N, t)[1 + \sum_{j=1}^N (j-1)P[n_i = j | n_i \geq 1, N, t]]. \quad (9)$$

The result follows noting that the right-hand side summation is the queue-length $Q_i^{cond}(n_i \geq 1, N, t)$. \square

From the MVA arrival theorem for constant-rate queues [8], the last result immediately implies $Q_i^{cond}(n_i \geq 1, N, t) = Q_i(N-1, t)$, for the constant-rate queues $1 \leq i \leq M-1$. We now investigate the conditional queue-length of the flow-equivalent server.

Theorem 2 (Mean Queue-Length at the Flow-Equivalent Server). *The conditional mean queue-length of the flow-equivalent server M is given by*

$$Q_M^{cond}(n_M \geq 1, N, t) = Q_M(N-1, t+1). \quad (10)$$

Proof. Summing (2) over all feasible states \vec{n} of the underlying Markov process it is

$$G(N, t) = \sum_{\vec{n}} \frac{\prod_{i=0}^M D_i^{n_i}}{\prod_{n=1}^{n_M} \mu(t+n-1)} = \sum_{\vec{n}: n_i=0} \frac{\prod_{i=0}^M D_i^{n_i}}{\prod_{n=1}^{n_M} \mu(t+n-1)} + \sum_{\vec{n}: n_i \geq 1} \frac{\prod_{i=0}^M D_i^{n_i}}{\prod_{n=1}^{n_M} \mu(t+n-1)}.$$

Note that the summation for $n_i = 0$ is the normalizing constant $G^{-i}(N, t)$. Moreover, isolating in the second summation the common term $D_i/\mu(t)$, it can be verified, according to the given definitions, that the remaining terms sum to the normalizing constant $G(N - 1, t + 1)$. Thus, we get

$$G(N, t) = G^{-i}(N, t) + \left(\frac{D_i}{\mu(t)} \right) G(N - 1, t + 1). \quad (11)$$

Considering the normalizing constant sensitivity formula [4]

$$D_i \frac{\partial G(N, t)}{\partial D_i} = Q_i(N, t) G(N, t).$$

and imposing the equality of derivatives of the two sides of (11), we find

$$Q_i(N, t) = \left(\frac{D_i}{\mu(t)} \right) \frac{G(N - 1, t + 1)}{G(N, t)} [1 + Q_i(N - 1, t + 1)]. \quad (12)$$

where we exploited that $G^{-i}(N, t)$ is not function of D_i and therefore its derivative is equal to zero. The result now follows easily from (11), (3) and (5) since

$$Q_i(N, t) = \left(1 - \frac{G^{-i}(N - 1, t + 1)}{G(N, t)} \right) [1 + Q_i(N - 1, t + 1)] = U_i(N, t) [1 + Q_i(N - 1, t + 1)].$$

□

The last proof does not depend on the number of load-dependent queues in the model and holds also with multiple load-dependent queues provided that a different rate shift index t is used for each queue.

4 Conditional Mean Value Analysis (CMVA)

The practical consequence of Theorem 2 is the definition of a new mean value analysis algorithm that can be employed to solve flow-equivalent aggregation models. We first introduce a revised version of the load-dependent MVA arrival theorem that does not require marginal probabilities.

Theorem 3 (Revised Mean-Value Arrival Theorem.). *The mean response time of station i is*

$$R_i(N, t) = \begin{cases} D_0, & i = 0 \\ D_i[1 + Q_i(N - 1, t)], & 1 \leq i \leq M - 1, \\ D_M(N, t)[1 + Q_M(N - 1, t + 1)], & i = M, \end{cases} \quad (13)$$

where $D_M(N, t) = U_M(N, t)/X(N, t)$ is the mean service demand of the flow-equivalent server. This is equal to $D_M(1, t) = D_M/\mu(t)$ for $N = 1$ and for larger population is computed recursively as

$$D_M(N, t) = \frac{X(N - 1, t)}{X(N - 1, t + 1)} D_M(N - 1, t), \quad N \geq 2. \quad (14)$$

CONDITIONAL MEAN VALUE ANALYSIS ALGORITHM (CMVA)

INPUT: Service demands D_i , $0 \leq i \leq M$, population N , flow-equivalent server rates $\mu(k)$, $1 \leq k \leq N$.

Let $Q_i(0, t) = 0$ for $t = 1, \dots, N + 1$.

for $n = 1$ **to** N

for $t = 1$ **to** $N - n + 1$

 Compute $D_M(n, t)$ using (14) if $n \geq 2$, otherwise set $D_M(1, t) = D_M/\mu(t)$ if $n = 1$.

 Compute $R_i(n, t)$ using (13), for $i = 0, \dots, M$.

 Compute by Little's Law the throughput $X(n, t) = N / \sum_{i=0}^M R_i(n, t)$.

 Compute $Q_i(n, t) = \begin{cases} D_i X(n, t)[1 + Q_i(n - 1, t)], & 1 \leq i \leq M - 1, \\ D_M(n, t) X(n, t)[1 + Q_M(n - 1, t + 1)], & i = M, \end{cases}$

end for

end for

OUTPUT: Throughput $X(N) \equiv X(N, t = 1)$. Mean queue-lengths $Q_i(N) \equiv Q_i(N, t = 1)$, $i = 1, \dots, M$.

Figure 1: Pseudo-code of the CMVA algorithm for stable flow-equivalent aggregation.

Proof. The cases $i \neq M$ are well-known [6], we need only to prove the relation for the case $i = M$. From (11) and (3) we know that

$$U_M(N, t) = \left(\frac{D_M}{\mu(t)} \right) \frac{G(N - 1, t + 1)}{G(N, t)},$$

which implies by the Utilization Law [6] and the usual definition $X(N, t) = G(N - 1, t)/G(N, t)$ that

$$D_M(N, t) = \frac{U_M(N, t)}{X(N, t)} = \left(\frac{D_M}{\mu(t)} \right) \frac{G(N - 1, t + 1)}{G(N - 1, t)}. \quad (15)$$

We have then

$$D_M(N, t) = \left(\frac{D_M}{\mu(t)} \right) \frac{G(N - 2, t + 1)}{G(N - 2, t)} \frac{G(N - 1, t + 1)}{G(N - 1, t)} \frac{G(N - 2, t)}{G(N - 2, t + 1)},$$

which we rewrite by (15) as

$$D_M(N, t) = \frac{G(N - 2, t)}{G(N - 1, t)} \frac{G(N - 1, t + 1)}{G(N - 2, t + 1)} D_M(N - 1, t) = \frac{X(N - 1, t)}{X(N - 1, t + 1)} D_M(N - 1, t)$$

that concludes the proof. \square

The algorithmic application of Theorem 3 is the Conditional Mean Value Analysis (CMVA) algorithm illustrated in Figure 1. In the case of models with constant-rate queues only, the algorithm reduces to the MVA algorithm. Note that the number of operations performed by CMVA is analogous to the load-dependent MVA, and because of the double recursion on the population and on the rate shifts it grows as $O(MN^2)$. Similarly, the space requirements grows as $O(MN)$ as in the load-dependent MVA.

The CMVA provides a numerically stable solution of load-dependent models since avoids the use of equation (3) that is numerically unstable. Furthermore, the algorithm does not involve any subtraction which prevents round-off instabilities similar to those arising in the evaluation of (3). In order to test the effectiveness of the stabilization in all cases, we performed the following two experiments:

Case Study 1. The network has $M \in [2, 10]$ queues. The demands and service rates are drawn with a uniform distribution in $[0, 1]$. The CMVA results are compared with the exact global balance solution of the underlying Markov process for all populations $1 \leq N \leq 100$. The total number of evaluated models in this experiment is 10000.

Case Study 2. This case is similar to Case Study 1, but the evaluated models are chosen by first generating a large number of random models and then selecting among them 1000 models for which the load-dependent MVA returns incorrect results. The CMVA is run on these model to verify that it can correctly solve models that the load-dependent MVA cannot solve in a numerically stable fashion.

In all experiments the CMVA results were always identical to the exact global balance solution on more than ten decimal figures. The evaluation has been performed both for mean throughput and mean queue-lengths and this is sufficient to grant the accurate computation by Little's Law and the Utilization Law of other performance mean indexes of interest (response times, utilizations), see [6].

5 Generalization of CMVA to Multiclass Models

For the sake of completeness, we discuss the multiclass extension of the result, despite multiclass flow aggregation is seldom used because of the high, often prohibitive, computational requirements [6].

We first consider the case where the service rate $\mu(n)$ is class-independent, a simplifying assumption that is usually taken to grant computational feasibility. In this case, n is intended as the total number of jobs at queue M , regardless of their service class. The derivation of the arrival theorem in this case has no significant differences from that of the single class case and for a model with C classes, population vector $\vec{N} = (N_1, N_2, \dots, N_C)$, and service demands $D_{i,c}$, $0 \leq i \leq M$, $1 \leq c \leq C$, leads to the response time formula

$$R_{i,c}(\vec{N}, t) = \begin{cases} D_{0,c}, & i = 0 \\ D_{i,c}[1 + Q_i(\vec{N} - 1_c, t)], & 1 \leq i \leq M - 1, \\ D_{M,c}(\vec{N}, t)[1 + \sum_{c'=1}^C Q_{M,c'}(\vec{N} - 1_c, t + 1)], & i = M, \end{cases} \quad (16)$$

where the additional c and c' indexes specify to which service class a measure refers to, and the load-dependent demand of the flow-equivalent server now satisfies the recursion

$$D_{M,c}(\vec{N}, t) = \frac{X_c(\vec{N} - 1_c, t)}{X_c(\vec{N} - 1_c, t + 1)} D_{M,c}(\vec{N} - 1_c, t),$$

where $D_{M,c}(\vec{N}, t) = D_{M,c}/\mu(t)$ if $N_c = 1$. The multiclass CMVA is then obtained by replacing in the algorithm pseudo-code the population n with a population vector \vec{n} and introducing a new inner loop to process the indexes of all C classes. The resulting computational complexities are thus $O(MR(\sum_{c=1}^C N_c)^{R+1})$ time and $O(MR(\sum_{c=1}^C N_c)^R)$ space.

A second important form of load-dependence involves the use of class-dependent service rates. When n_c , $1 \leq c \leq C$, jobs of class c are enqueued at the flow-equivalent server, its service rate for that class is $\mu_c(n_c)$. In order to account for the increased detail level on the service rates, it is sufficient to replace in the CMVA algorithm the rate shift index t by a *rate shift vector* $\vec{t} = (t_1, \dots, t_c, \dots, t_C)$ and compute response times as

$$R_{i,c}(\vec{N}, \vec{t}) = \begin{cases} D_{0,c}, & i = 0 \\ D_{i,c}[1 + Q_i(\vec{N} - 1_c, \vec{t})], & 1 \leq i \leq M - 1, \\ D_{M,c}(\vec{N}, \vec{t})[1 + \sum_{c'=1}^C Q_{M,c'}(\vec{N} - 1_c, \vec{t} + 1_c)], & i = M, \end{cases} \quad (17)$$

where

$$D_{M,c}(\vec{N}, \vec{t}) = \frac{X_c(\vec{N} - 1_c, \vec{t})}{X_c(\vec{N} - 1_c, \vec{t} + 1_c)} D_{M,c}(\vec{N} - 1_c, \vec{t}),$$

and $D_{M,c}(\vec{N}, \vec{t}) = D_{M,c}/\mu_c(t_c)$ if $N_c = 1$. The pseudo-code in Figure 1 is immediately modified by replacing the loop on t with a loop on all feasible vectors \vec{t} . Because of this modification, the computation cost of the resulting algorithm is $O(MR(\sum_{c=1}^C N_c)^{2R+1})$ time and $O(MR(\sum_{c=1}^C N_c)^R)$ space.

A third form of load-dependence is also considered in the literature, i.e., service rates dependent on the current population $\vec{n}_M = (n_{M,1}, n_{M,2}, \dots, n_{M,C})$ at the flow-equivalent server M . This case is known to be prohibitively expensive because an additional nested recursion is required to evaluate each state of the underlying Markov process [9]. Because of the different structure imposed by the nested recursion, the mean value analysis does find application to this case [9] which is therefore outside the scope of this work.

We remark that the experimental results for the multiclass models yield results that are qualitatively similar to the single class case. The experimentation for the multiclass algorithms has been performed on 10000 models with class populations $N_c \leq 25$ and up to $C = 4$ service classes; global balance on larger models is computationally infeasible and therefore it would be impossible to check the accuracy of the multiclass CMVA results in these cases. All service rates are generated with a uniform distribution in $[0, 1]$. In all evaluated models, both queue-lengths and throughputs computed by CMVA were identical on more than ten decimal figures to the exact solutions computed by global balance. This confirms the numerical stability of the CMVA algorithm both on class-independent and class-dependent models.

6 Conclusion

Numerical instabilities arising in the solution of flow-equivalent aggregation problems can significantly limit the applicability of this methodology. We have presented the Conditional Mean Value Analysis (CMVA) algorithm to solve in a numerically stable fashion flow-equivalent aggregation models. The result relies on the queue-length formula (5) which relates the mean queue-length of a flow-equivalent server to the conditional queue-length as seen by a job during its residence time at that server. A revised MVA arrival theorem for load-dependent models deriving from (5) has been presented. We have extended CMVA to the multiclass case both for models with class-independent and class-dependent service rates. By extensive validation we have found that for all class of models the CMVA algorithm always provides accurate solutions.

Acknowledgement

The author thanks Gianfranco Balbo for his useful comments on an preliminary version of this work.

References

- [1] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios, "Open, closed, and mixed networks of queues with different classes of customers.," *JACM*, 22(2):248–260, 1975.
- [2] J. P. Buzen, "Computational algorithms for closed queueing networks with exponential servers," *Comm. ACM*, 16(9):527–531, 1973.
- [3] K.M. Chandy, U. Herzog, L. Woo, "Parametric analysis of queueing networks," *IBM J. Res. Dev.*, 19(1):36–42, 1975.
- [4] E. de Sousa e Silva and R. R. Muntz, "Simple relationships among moments of queue lengths in product form queueing networks," *IEEE Trans. Comp.*, 37(9):1125–1129, 1988.
- [5] S. Lam, "Dynamic scaling and growth behavior of queueing network normalization constants," *JACM*, 29(2):492–513, 1982.
- [6] E. D. Lazowska, J. Zahorjan, G. S. Graham, K. C. Sevcik, *Quantitative System Performance*. Prentice-Hall, 1984.
- [7] M. Reiser, "Mean-value analysis and convolution method for queue-dependent servers in closed queueing networks," *PEVA*, 1:7–18, 1981.
- [8] M. Reiser, S.S. Lavenberg, "Mean-value analysis of closed multichain queueing networks," *JACM*, 27(2):312–322, 1980.
- [9] C. H. Sauer, "Computational algorithms for state-dependent queueing networks," *ACM Trans. Comp. Sys.*, 1:67–92, Feb. 1983.
- [10] J. Zahorjan, "The distribution of network states during residence times in product form queueing networks," *PEVA*, 4:99–104, 1984.