

# Task Execution Time Modeling for Heterogeneous Computing Systems

Shoukat Ali<sup>†</sup>, Howard Jay Siegel<sup>†</sup>, Muthucumaru Maheswaran<sup>‡</sup>,  
Debra Hensgen<sup>◇</sup>, and Sahra Ali<sup>†</sup>

<sup>†</sup>Purdue University  
School of Electrical and Computer Engineering  
West Lafayette, IN 47907-1285 USA  
{alis@ecn., hj@}purdue.edu

<sup>‡</sup>Department of Computer Science  
University of Manitoba  
Winnipeg, MB R3T 2N2 Canada  
maheswar@cs.umanitoba.ca

<sup>◇</sup>OS Research and Evaluation  
OpenTV  
Mountain View, CA 94043 USA  
dhensgen@opentv.com

## Abstract

*A distributed heterogeneous computing (HC) system consists of diversely capable machines harnessed together to execute a set of tasks that vary in their computational requirements. Heuristics are needed to map (match and schedule) tasks onto machines in an HC system so as to optimize some figure of merit. This paper characterizes a simulated HC environment by using the expected execution times of the tasks that arrive in the system onto the different machines present in the system. This information is arranged in an “expected time to compute” (ETC) matrix as a model of the given HC system, where the entry( $i, j$ ) is the expected execution time of task  $i$  on machine  $j$ . This model is needed to simulate different HC environments to allow testing of relative performance of different mapping heuristics under different circumstances. In particular, the ETC model is used to express the heterogeneity among the runtimes of the tasks to be executed, and among the machines in the HC system. An existing range-based technique to generate ETC matrices is described. A coefficient-of-variation based technique to generate ETC matrices is proposed, and compared with the range-based technique. The coefficient-of-variation-based ETC generation method provides a greater control over the spread of values (i.e., heterogeneity) in any given row or column of the ETC matrix than the range-based method.*

## 1. Introduction

A distributed heterogeneous computing (HC) system consists of diversely capable machines harnessed together to execute a set of tasks that vary in their computational requirements. Heuristics are needed to map (match and

schedule) tasks onto machines in an HC system so as to optimize some figure of merit. The heuristics that match a task to a machine can vary in the information they use. For example, the current candidate task can be assigned to the machine that becomes available soonest (even if the task may take a much longer time to execute on that machine than elsewhere). In another approach, the task may be assigned to the machine where it executes fastest (but ignores when that machine becomes available). Or the current candidate task may be assigned to the machine that completes the task soonest, i.e., the machine which minimizes the sum of task execution time and the machine ready time, where machine ready time for a particular machine is the time when that machine becomes available after having executed the tasks previously assigned to it (e.g., [13]).

The discussion above should reveal that more sophisticated (and possibly wiser) approaches to the mapping problem require estimates of the execution times of all tasks (that can be expected to arrive for service) on all the machines present in the HC suite to make better mapping decisions. One aspect of the research on HC mapping heuristics explores the behavior of the heuristics in different HC environments. The ability to test the relative performance of different mapping heuristics under different circumstances necessitates that there be a framework for generating simulated execution times of all the tasks in the HC system on all the machines in the HC system. Such a framework would, in turn, require a quantification of heterogeneity to express the variability among the runtimes of the tasks to be executed, and among the capabilities of the machines in the HC system. The goal of this paper is to present a methodology for synthesizing simulated HC environments with quantifiable levels of task and machine heterogeneity. This paper characterizes the HC environments so that it will be easier for the researchers to describe the workload and the machines used in their simulations using a common scale.

Given a set of heuristics and a characterization of HC

---

This research was supported by the DARPA/ITO Quorum Program under the NPS subcontract numbers N62271-98-M-0217 and N62271-98-M-0448, and under the GSA subcontract number GS09K99BH0250. Some of the equipment used was donated by Intel.

environments, one can determine the best heuristic to use in a given environment for optimizing a given objective function. In addition to increasing one's understanding of the operation of different heuristics, this knowledge can help a working resource management system select which mapper to use for a given real HC environment.

This research is part of a DARPA/ITO Quorum Program project called MSHN (pronounced "mission") (Management System for Heterogeneous Networks) [7]. MSHN is a collaborative research effort that includes the Naval Postgraduate School, NOEMIX, Purdue, and University of Southern California. It builds on SmartNet, an implemented scheduling framework and system for managing resources in an HC environment developed at NRaD [5]. The technical objective of the MSHN project is to design, prototype, and refine a distributed resource management system that leverages the heterogeneity of resources and tasks to deliver the requested qualities of service. The methodology developed here for generating simulated HC environments may be used to design, analyze and evaluate heuristics for the Scheduling Advisor component of the MSHN prototype.

The rest of this paper is organized as follows. A model for describing an HC system is presented in Section 2. Based on that model, two techniques for simulating an HC environment are described in Section 3. Section 4 briefly discusses analyzing the task execution time information from real life HC scenarios. Some related work is outlined in the Section 5.

## 2. Modeling Heterogeneity

To better evaluate the behavior of mapping heuristics, a model of the execution times of the tasks on the machines is needed so that the parameters of this model can be changed to investigate the performance of the heuristics under different HC systems and under different types of tasks to be mapped. One such model consists of an expected time to compute (ETC) matrix, where the entry  $(i, j)$  is the expected execution time of task  $i$  on machine  $j$ . The ETC matrix can be stored on the same machine where the mapper is stored, and contains the estimates for the expected execution times of a task on all machines, for all the tasks that are expected to arrive for service over a given interval of time. (Although stored with the mapper, the ETC information may be derived from other components of a resource management system (e.g., [7])). In an ETC matrix, the elements along a row indicate the estimates of the expected execution times of a given task on different machines, and those along a column give the estimates of the expected execution times of different tasks on a given machine.

The exact actual task execution times on all machines may not be known for all tasks because, for example, they might be a function of input data. What is typically assumed in the HC literature is that estimates of the expected

execution times of tasks on all machines are known (e.g., [6, 10, 12, 16]). These estimates could be built from task profiling and machine benchmarking, could be derived from the previous executions of a task on a machine, or could be provided by the user (e.g., [3, 6, 8, 14, 18]).

The ETC model presented here can be characterized by three parameters: machine heterogeneity, task heterogeneity, and consistency. The variation along a row is referred to as the machine heterogeneity; this is the degree to which the machine execution times vary for a given task [1]. A system's machine heterogeneity is based on a combination of the machine heterogeneities for all tasks (rows). A system comprised mainly of workstations of similar capabilities can be said to have "low" machine heterogeneity. A system consisting of diversely capable machines, e.g., a collection of SMP's, workstations, and supercomputers, may be said to have "high" machine heterogeneity.

Similarly, the variation along a column of an ETC matrix is referred to as the task heterogeneity; this is the degree to which the task execution times vary for a given machine [1]. A system's task heterogeneity is based on a combination of the task heterogeneities for all machines (columns). "High" task heterogeneity may occur when the computational needs of the tasks vary greatly, e.g., when both time-consuming simulations and fast compilations of small programs are performed. "Low" task heterogeneity may typically be seen in the jobs submitted by users solving problems of similar complexity (and hence have similar execution times on a given machine).

Based on the above idea, four categories were proposed for the ETC matrix in [1]: (a) high task heterogeneity and high machine heterogeneity, (b) high task heterogeneity and low machine heterogeneity, (c) low task heterogeneity and high machine heterogeneity, and (d) low task heterogeneity and low machine heterogeneity.

The ETC matrix can be further classified into two categories, consistent and inconsistent [1], which are orthogonal to the previous classifications. For a consistent ETC matrix, if a machine  $m_x$  has a lower execution time than a machine  $m_y$  for a task  $t_k$ , then the same is true for any task  $t_i$ . A consistent ETC matrix can be considered to represent an extreme case of low task heterogeneity and high machine heterogeneity. If machine heterogeneity is high enough, then the machines may be so much different from each other in their compute power that the differences in the computational requirements of the tasks (if low enough) will not matter in determining the relative order of execution times for a given task on the different machines (i.e., along a row). As a trivially extreme example, consider a system consisting of Intel Pentium III and Intel 286. The Pentium III will almost always run any given task from a certain set of tasks faster than the 286 provided the computational requirements of all tasks in the set are similar (i.e.,

low task heterogeneity), thereby giving rise to a consistent ETC matrix.

In inconsistent ETC matrices, the relationships among the task computational requirements and machine capabilities are such that no structure as that in the consistent case is enforced. Inconsistent ETC matrices occur in practice when: (1) there is a variety of different machine architectures in the HC suite (e.g., parallel machines, superscalars, workstations), and (2) there is a variety of different computational needs among the tasks (e.g., readily parallelizable tasks, difficult to parallelize tasks, tasks that are floating point intensive, simple text formatting tasks). Thus, the way in which a task's needs correspond to a machine's capabilities may differ for each possible pairing of tasks to machines.

A combination of these two cases, which may be more realistic in many environments, is the partially-consistent ETC matrix, which is an inconsistent matrix with a consistent sub-matrix [2, 13]. This sub-matrix can be composed of any subset of rows and any subset of columns. As an example, in a given partially-consistent ETC matrix, 50% of the tasks and 25% of the machines may define a consistent sub-matrix.

Even though no structure is enforced on an inconsistent ETC matrix, a given ETC matrix generated to be inconsistent may have the structure of a partially consistent ETC matrix. In this sense, partially-consistent ETC matrices are a special case of inconsistent ETC matrices. Similarly, consistent ETC matrices are special cases of inconsistent and partially-consistent ETC matrices.

It should be noted that this classification scheme is used for generating ETC matrices. Later in this paper, it will be shown how these three cases differ in generation process. If one is given an ETC matrix, and is asked to classify it among these three classes, it will be called a consistent ETC matrix only if it is fully consistent. It will be called inconsistent if it is not consistent.

Often an inconsistent ETC matrix will have some partial consistency in it. For example, a trivial case of partial-consistency always exists; for any two machines in the HC suite, *at least* 50% of the tasks will show consistent execution times.

### 3. Generating the ETC Matrices

#### 3.1. Range Based ETC Matrix Generation

Any method for generating the ETC matrices will require that heterogeneity be defined mathematically. In the range-based ETC generation technique, the heterogeneity of a set of execution time values is quantified by the range of the execution times [2, 13]. The procedures given in this section for generating the ETC matrices produce inconsistent ETC matrices. It is shown later in this section how consistent and

- (1) for  $i$  from 0 to  $(t - 1)$
- (2)      $\tau[i] = U(1, R_{task})$
- (3)     for  $j$  from 0 to  $(m - 1)$
- (4)          $e[i, j] = \tau[i] \times U(1, R_{mach})$
- (5)     endfor
- (6) endfor

**Figure 1. The range-based method for generating ETC matrices.**

partially-consistent ETC matrices could be obtained from the inconsistent ETC matrices.

Assume  $\underline{m}$  is the total number of machines in the HC suite, and  $\underline{t}$  is the total number of tasks expected to be serviced by the HC system over a given interval of time. Let  $U(a, b)$  be a number sampled from a uniform distribution with a range from  $a$  to  $b$ . (Each invocation of  $U(a, b)$  returns a new sample.) Let  $R_{task}$  and  $R_{mach}$  be numbers representing task heterogeneity and machine heterogeneity, respectively, such that higher values for  $R_{task}$  and  $R_{mach}$  represent higher heterogeneities. Then an ETC matrix  $e[0..(t - 1), 0..(m - 1)]$ , for a given task heterogeneity and a given machine heterogeneity, can be generated by the range-based method given in Figure 1, where  $e[i, j]$  is the estimated expected execution time for the task  $i$  on the machine  $j$ .

As shown in Figure 1, each iteration of the outer **for** loop samples a uniform distribution with a range from 1 to  $R_{task}$  to generate one value for a vector  $\underline{\tau}$ . For each element of  $\underline{\tau}$  thus generated, the  $m$  iterations of the inner **for** loop (Line 3) generate one row of the ETC matrix. For the  $i$ -th iteration of the outer **for** loop, each iteration of the inner **for** loop produces one element of the ETC matrix by multiplying  $\tau[i]$  with a random number sampled from a uniform distribution ranging from 1 to  $R_{mach}$ .

In the range-based ETC generation, it is possible to obtain high task heterogeneity low machine heterogeneity ETC matrices with characteristics similar to that of low task heterogeneity high machine heterogeneity ETC matrices if  $R_{task} = R_{mach}$ . In realistic HC systems, the variation that tasks show in their computational needs is generally larger than the variation that machines show in their capabilities. Therefore it is assumed here that requirements of high heterogeneity tasks are likely to be more "heterogeneous" than the capabilities of high heterogeneity machines (i.e.,  $R_{task} \gg R_{mach}$ ). However, for the ETC matrices generated here, low heterogeneity in both machines and tasks is assumed to be same. Table 1 shows typical values for  $R_{task}$  and  $R_{mach}$  for low and high heterogeneities. Tables 2 through 5 show four ETC matrices generated by the range-based method. The execution time values in Table 2 are

**Table 1. Suggested values for  $R_{task}$  and  $R_{mach}$  for a realistic HC system for high heterogeneity and low heterogeneity.**

	high	low
task	$10^5$	$10^1$
machine	$10^2$	$10^1$

much higher than the execution time values in Table 5. The difference in the values between these two tables would be reduced if the range for the low task heterogeneity was changed to  $10^3$  to  $10^4$  instead of 1 to 10.

With the range-based method, low task heterogeneity high machine heterogeneity ETC matrices tend to have high heterogeneity for both tasks and machines, due to method used for generation. For example, in Table 5, original  $\tau$  vector values were selected from 1 to 10. When each entry is multiplied by a number from 1 to 100 for high machine heterogeneity this generates a task heterogeneity comparable to machine heterogeneity. It is shown later in Section 3.2 how to produce low task heterogeneity high machine heterogeneity ETC matrices which do show low task heterogeneity.

### 3.2. Coefficient-of-Variation Based ETC Matrix Generation

A modification of the procedure in Figure 1 defines the coefficient of variation,  $V$ , of execution time values as a measure of heterogeneity (instead of the range of execution time values). The coefficient of variation of a set of values is a better measure of the dispersion in the values than the standard deviation because it expresses the standard deviation as a percentage of the mean of the values [11]. Let  $\underline{\sigma}$  and  $\underline{\mu}$  be the standard deviation and mean, respectively, of a set of execution time values. Then  $V = \sigma/\mu$ . The coefficient-of-variation-based ETC generation method provides a greater control over spread of the execution time values (i.e., heterogeneity) in any given row or column of the ETC matrix than the range-based method.

The coefficient-of-variation-based (CVB) ETC generation method works as follows. A task vector,  $q$ , of expected execution times with the desired task heterogeneity must be generated. Essentially,  $q[i]$  is the execution time of task  $i$  on an “average” machine in the HC suite. For example, if the HC suite consists of an IBM SP/2, an Alpha server, and a Sun SPARC 5 workstation, then  $q$  would represent estimated execution times of the tasks on the Alpha server.

To generate  $q$ , two input parameters are needed:  $\mu_{task}$

and  $V_{task}$ . The input parameter,  $\mu_{task}$  is used to set the average of the values in  $q$ . The input parameter  $V_{task}$  is the desired coefficient of variation of the values in  $q$ . The value of  $V_{task}$  quantifies task heterogeneity, and is larger for higher task heterogeneity. Each element of the task vector  $q$  is then used to produce one row of the ETC matrix such that the desired coefficient of variation of values in each row is  $V_{mach}$ , another input parameter. The value of  $V_{mach}$  quantifies machine heterogeneity, and is larger for higher machine heterogeneity. Thus  $\mu_{task}$ ,  $V_{task}$ , and  $V_{mach}$  are the three input parameters for the CVB ETC generation method.

A direct approach to simulating HC environments should use the probability distribution that is empirically found to represent closely the distribution of task execution times. However, no standard benchmarks for HC systems are currently available. Therefore, this research uses a distribution which, though not necessarily reflective of an actual HC scenario, is flexible enough to be adapted to one. Such a distribution should not produce negative values of task execution times (e.g., ruling out Gaussian distribution), and should have a variable coefficient of variation (e.g., ruling out exponential distribution).

The gamma distribution is a good choice for the CVB ETC generation method because, with proper constraints on its characteristic parameters, it can approximate two other probability distributions, namely the Erlang-k and Gaussian (without the negative values) [11, 15]. The fact that it can approximate these two other distributions is helpful because this increases the chances that the simulated ETC matrices could be synthesized closer to some real life HC environment.

The uniform distribution can also be used but is not as flexible as the gamma distribution for two reasons: (1) it does not approximate any other distribution, and (2) the characteristic parameters of a uniform distribution cannot take all real values (explained later in the Section 3.3).

The gamma distribution [11, 15] is defined in terms of characteristic shape parameter,  $\underline{\alpha}$ , and scale parameter,  $\underline{\beta}$ . The characteristic parameters of the gamma distribution can be fixed to generate different distributions. For example, if  $\alpha$  is fixed to be an integer, then the gamma distribution becomes an Erlang-k distribution. If  $\alpha$  is large enough, then the gamma distribution approaches a Gaussian distribution (but still does not return negative values for task execution times).

Figures 2(a) and 2(b) show how a gamma density function changes with the shape parameter  $\alpha$ . When the shape parameter increases from two to eight, the shape of the distribution changes from a curve biased to the left to a more balanced bell-like curve. Figures 2(a), 2(c) and 2(d) show

**Table 2. A high task heterogeneity low machine heterogeneity matrix generated by the range-based method using  $R_{task}$  and  $R_{mach}$  values of Table 1.**

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$
$t_1$	333304	375636	198220	190694	395173	258818	376568
$t_2$	442658	400648	346423	181600	289558	323546	380792
$t_3$	75696	103564	438703	129944	67881	194194	425543
$t_4$	194421	392810	582168	248073	178060	267439	611144
$t_5$	466164	424736	503137	325183	193326	241520	506642
$t_6$	665071	687676	578668	919104	795367	390558	758117
$t_7$	177445	227254	72944	139111	236971	325137	347456
$t_8$	32584	55086	127709	51743	100393	196190	270979
$t_9$	311589	568804	148140	583456	209847	108797	270100
$t_{10}$	314271	113525	448233	201645	274328	248473	170176
$t_{11}$	272632	268320	264038	140247	110338	29620	69011
$t_{12}$	489327	393071	225777	71622	243056	445419	213477

**Table 3. A high task heterogeneity high machine heterogeneity matrix generated by the range-based method using  $R_{task}$  and  $R_{mach}$  values of Table 1.**


	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$
$t_1$	2425808	3478227	719442	2378978	408142	2966676	2890219
$t_2$	2322703	2175934	228056	3456054	6717002	5122744	3660354
$t_3$	1254234	3182830	4408801	5347545	4582239	6124228	5343661
$t_4$	227811	419597	13972	297165	438317	23374	135871
$t_5$	6477669	5619369	707470	8380933	4693277	8496507	7279100
$t_6$	1113545	1642662	303302	244439	1280736	541067	792149
$t_7$	2860617	161413	2814518	2102684	8218122	7493882	2945193
$t_8$	1744479	623574	1516988	5518507	2023691	3527522	1181276
$t_9$	6274527	1022174	3303746	7318486	7274181	6957782	2145689
$t_{10}$	1025604	694016	169297	193669	1009294	1117123	690846
$t_{11}$	2390362	1552226	2955480	4198336	1641012	3072991	3262071
$t_{12}$	96699	882914	63054	199175	894968	248324	297691

**Table 4. A low task heterogeneity low machine heterogeneity matrix generated by the range-based method using  $R_{task}$  and  $R_{mach}$  values of Table 1.**


	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$
$t_1$	22	21	6	16	15	24	13
$t_2$	7	46	5	28	45	43	31
$t_3$	64	83	45	23	58	50	38
$t_4$	53	56	26	42	53	9	58
$t_5$	11	12	14	7	8	3	14
$t_6$	33	31	46	25	23	39	10
$t_7$	24	11	17	14	25	35	4
$t_8$	20	17	23	4	3	18	20
$t_9$	13	28	14	7	34	6	29
$t_{10}$	2	5	7	7	6	3	7
$t_{11}$	16	37	23	22	23	12	44
$t_{12}$	8	66	47	11	47	55	56

**Table 5. A low task heterogeneity high machine heterogeneity matrix generated by the range-based method using  $R_{task}$  and  $R_{mach}$  values of Table 1.**


	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$
$t_1$	440	762	319	532	151	652	308
$t_2$	459	205	457	92	92	379	60
$t_3$	499	263	92	152	75	18	128
$t_4$	421	362	347	194	241	481	391
$t_5$	276	636	136	355	338	324	255
$t_6$	89	139	37	67	9	53	139
$t_7$	404	521	54	295	257	208	539
$t_8$	49	114	279	22	93	39	36
$t_9$	59	35	184	262	145	287	277
$t_{10}$	7	235	44	81	330	56	78
$t_{11}$	716	601	75	689	299	144	457
$t_{12}$	435	208	256	330	6	394	419




(a)



(b)



(c)



(d)

**Figure 2. Gamma probability density function for (a)  $\alpha = 2, \beta = 8$ , (b)  $\alpha = 8, \beta = 8$ , (c)  $\alpha = 2, \beta = 16$ , and (d)  $\alpha = 2, \beta = 32$ .**

```

(1)  $\alpha_{task} = 1/V_{task}^2$ ;  $\alpha_{mach} = 1/V_{mach}^2$ ;
    $\beta_{task} = \mu_{task}/\alpha_{task}$ 
(2) for  $i$  from 0 to  $(t - 1)$ 
(3)    $q[i] = G(\alpha_{task}, \beta_{task})$ 
   /*  $q[i]$  will be used as mean
   of  $i$ -th row of ETC matrix */
(4)    $\beta_{mach}[i] = q[i]/\alpha_{mach}$ 
   /* scale parameter for  $i$ -th row */
(5)   for  $j$  from 0 to  $(m - 1)$ 
(6)      $e[i, j] = G(\alpha_{mach}, \beta_{mach}[i])$ 
(7)   endfor
(8) endfor

```

**Figure 3. The general CVB method for generating ETC matrices.**

the effect on the distribution caused by an increase in the scale parameter from 8 to 16 to 32. The two-fold increase in the scale parameter does not change the shape of the graph (the curve is still biased to the left); however the curve now has twice as large a domain (i.e., range on x-axis).

The gamma distribution's characteristic parameters,  $\alpha$  and  $\beta$ , can be easily interpreted in terms of  $\mu_{task}$ ,  $V_{task}$ , and  $V_{mach}$ . For a gamma distribution,  $\sigma = \beta\sqrt{\alpha}$ , and  $\mu = \beta\alpha$ , so that  $V = \sigma/\mu = 1/\sqrt{\alpha}$  (and  $\alpha = 1/V^2$ ). Then  $\alpha_{task} = 1/V_{task}^2$  and  $\alpha_{mach} = 1/V_{mach}^2$ . Further, because  $\mu = \beta\alpha$ ,  $\beta = \mu/\alpha$ , and  $\beta_{task} = \mu_{task}/\alpha_{task}$ . Also, for task  $i$ ,  $\beta_{mach}[i] = q[i]/\alpha_{mach}$ .

Let  $G(\alpha, \beta)$  be a number sampled from a gamma distribution with the given parameters. (Each invocation of  $G(\alpha, \beta)$  returns a new sample.) Figure 3 shows the general procedure for the CVB ETC generation.

Given the three input parameters,  $V_{task}$ ,  $V_{mach}$ , and  $\mu_{task}$ , Line (1) of Figure 3 determines the shape parameter  $\alpha_{task}$  and scale parameter  $\beta_{task}$  of the gamma distribution that will be later sampled to build the task vector  $q$ . Line (1) also calculates the shape parameter  $\alpha_{mach}$  to use later in Line (6). In the  $i$ -th iteration of the outer **for** loop (Line 2) in Figure 3, a gamma distribution with parameters  $\alpha_{task}$  and  $\beta_{task}$  is sampled to obtain  $q[i]$ . Then  $q[i]$  is used to determine the scale parameter  $\beta_{mach}[i]$  (to be used later in Line (6)). For the  $i$ -th iteration of the outer **for** loop (Line 2), each iteration of the inner **for** loop (Line 5) produces one element of the  $i$ -th row of the ETC matrix by sampling a gamma distribution with parameters  $\alpha_{mach}$  and  $\beta_{mach}[i]$ . One complete row of the ETC matrix is produced by  $m$  iterations of the inner **for** loop (Line 5). Note that while each row in the ETC matrix has gamma distributed execution times, the execution times in columns are not gamma distributed.

The ETC generation method of Figure 3 can be used to generate high task heterogeneity high machine heterogeneity

```

(1)  $\alpha_{task} = 1/V_{task}^2$ ;  $\alpha_{mach} = 1/V_{mach}^2$ ;
    $\beta_{mach} = \mu_{mach}/\alpha_{mach}$ 
(2) for  $j$  from 0 to  $(m - 1)$ 
(3)    $p[j] = G(\alpha_{mach}, \beta_{mach})$ 
   /*  $p[j]$  will be used as mean
   of  $j$ -th column of ETC matrix */
(4)    $\beta_{task}[j] = p[j]/\alpha_{task}$ 
   /* scale parameter for  $j$ -th column */
(5)   for  $i$  from 0 to  $(t - 1)$ 
(6)      $e[i, j] = G(\alpha_{task}, \beta_{task}[j])$ 
(7)   endfor
(8) endfor

```

**Figure 4. The CVB method for generating low task heterogeneity high machine heterogeneity ETC matrices.**

ity ETC matrices, high task heterogeneity low machine heterogeneity ETC matrices, and low task heterogeneity low machine heterogeneity ETC matrices, but cannot generate low task heterogeneity high machine heterogeneity ETC matrices. To satisfy the heterogeneity quadrants of Section 2, each column in the final low task heterogeneity high machine heterogeneity ETC matrix should reflect the low task heterogeneity of the "parent" task vector  $q$ . This condition would not necessarily hold if rows of the ETC matrix were produced with a high machine heterogeneity from a task vector of low heterogeneity. This is because a given column may be formed from widely different execution time values from different rows because of the high machine heterogeneity. That is, any two entries in a given column are based on different values of  $q[i]$  and  $\alpha_{mach}$ , and may therefore show high task heterogeneity as opposed to the intended low task heterogeneity. In contrast, in a high task heterogeneity low machine heterogeneity ETC matrix the low heterogeneity among the machines for a given task (across a row) is based on the same  $q[i]$  value.

One solution is to generate what is in effect a transpose of a high task heterogeneity low machine heterogeneity matrix to produce a low task heterogeneity high machine heterogeneity one. The transposition can be built into the procedure as shown in Figure 4. The procedure in Figure 4 is very similar to the one in Figure 3. The input parameter  $\mu_{task}$  is replaced with  $\mu_{mach}$ . Here, first a machine vector,  $p$ , (with an average value of  $\mu_{mach}$ ) is produced. Each element of this "parent" machine vector is then used to generate one low task heterogeneity column of the ETC matrix, such that the high machine heterogeneity present in  $p$  is reflected in all rows. This approach for generating low task heterogeneity high machine heterogeneity ETC matrices can also be used with the range-based method.



Tables 6 through 11 show some sample ETC matrices generated using the CVB ETC generation method. Tables 6 and 7 both show high task heterogeneity low machine heterogeneity ETC matrices. In both tables, the spread of the execution time values in columns is higher than that in rows. The ETC matrix in Table 7 has a higher task heterogeneity (higher  $V_{task}$ ) than the ETC matrix in Table 6. This can be seen in a higher spread in the columns of matrix in Table 7 than that in Table 6.

Tables 8 and 9 show high task heterogeneity high machine heterogeneity and low task heterogeneity low machine heterogeneity ETC matrices, respectively. The execution times in Table 8 are widely spaced along both rows and columns. The spread of execution times in Table 9 is smaller along both columns and rows, because both  $V_{task}$  and  $V_{mach}$  are smaller.

Tables 10 and 11 show low task heterogeneity high machine heterogeneity ETC matrices. In both tables, the spread of the execution time values in rows is higher than that in columns. ETC matrix in Table 11 has a higher machine heterogeneity (higher  $V_{mach}$ ) than the ETC matrix in Table 10. This can be seen in a higher spread in the rows of matrix in Table 11 than that in Table 10.

### 3.3. Uniform Distribution in the CVB Method

The uniform distribution could also be used for the CVB ETC generation method. The uniform distribution's characteristic parameters  $a$  (lower bound for the range of values) and  $b$  (upper bound for the range of values), can be easily interpreted in terms of  $\mu_{task}$ ,  $V_{task}$ , and  $V_{mach}$ . (Recall that  $V_{task} = \sigma_{task}/\mu_{task}$  and  $V_{mach} = \sigma_{mach}/\mu_{mach}$ ). For a uniform distribution,  $\sigma = (b - a)/\sqrt{12}$  and  $\mu = (b + a)/2$  [15]. So that

$$a + b = 2\mu \quad (1)$$

$$a - b = -\sigma\sqrt{12} \quad (2)$$

Adding Equations (1) and (2),

$$a = \mu - \sigma\sqrt{3} \quad (3)$$

$$a = \mu(1 - (\sigma/\mu)\sqrt{3}) \quad (4)$$

$$a = \mu(1 - V\sqrt{3}) \quad (5)$$

Also,

$$b = 2\mu - a \quad (6)$$

The Equations (5) and (6) can be used to generate the task vector  $q$  from the uniform distribution with the following parameters:

$$a_{task} = \mu_{task}(1 - V_{task}\sqrt{3}) \quad (7)$$

$$b_{task} = 2\mu_{task} - a_{task} \quad (8)$$

Once the task vector  $q$  has been generated, the  $i$ -th row of the ETC matrix can be generated by sampling ( $m$  times) a uniform distribution with the following parameters:

$$a_{mach} = q[i](1 - V_{mach}\sqrt{3}) \quad (9)$$

$$b_{mach} = 2q[i] - a_{mach} \quad (10)$$

The CVB ETC generation using the uniform distribution, however, places a restriction on the values of  $V_{task}$  and  $V_{mach}$ . Because both  $a_{task}$  and  $a_{mach}$  have to be positive, it follows from Equations (7) and (9) that the maximum value for  $V_{mach}$  or  $V_{task}$  is  $1/\sqrt{3}$ . Thus, for the CVB ETC generation, the gamma distribution is better than the uniform distribution because it does not restrict the values of task or machine heterogeneities.

### 3.4. Producing Consistent ETC Matrices

The procedures given in Figures 1, 3, and 4 produce inconsistent ETC matrices. Consistent ETC matrices can be obtained from the inconsistent ETC matrices generated above by sorting the execution times for each task on all machines (i.e., sorting the values within each row and doing this for all rows independently). From the inconsistent ETC matrices generated above, partially-consistent matrices consisting of an  $i \times k$  sub-matrix could be generated by sorting the execution times across a random subset of  $k$  machines for each task in a random subset of  $i$  tasks.

It should be noted from Tables 10 and 11 that the greater the difference in machine and task heterogeneities, the higher the degree of consistency in the inconsistent low task heterogeneity high machine heterogeneity ETC matrices. For example, in Table 11 all tasks show consistent execution times on all machines except on the machines that correspond to columns 3 and 4. As mentioned in Section 1, these degrees and classes of mixed-machine heterogeneity can be used to characterize many different HC environments.

## 4. Analysis and Synthesis

Once the actual ETC matrices from a real life scenario are obtained, they can be analyzed to estimate the probability distribution of the execution times, and the values of the model parameters (i.e.,  $V_{task}$ ,  $V_{mach}$ , and  $\mu_{task}$  (or  $\mu_{mach}$ , if a low task heterogeneity high machine heterogeneity ETC matrix is desired)) appropriate for the given real life scenario. The above analysis could be carried out using common statistical procedures [9]. Once a model of a particular HC system is available, the effect of changes in the workload (i.e., the tasks arriving for service in the system) and the system (i.e., the machines present in the HC system) can be studied in a controlled manner by simply changing the parameters of the ETC model.

**Table 6. A high task heterogeneity low machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.3, V_{mach} = 0.1.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	628	633	748	558	743	684	740	692	593	554
$t_2$	688	712	874	743	854	851	701	701	811	864
$t_3$	965	1029	1087	1020	921	825	1238	934	928	1042
$t_4$	891	866	912	896	776	993	875	999	919	860
$t_5$	1844	1507	1353	1436	1677	1691	1508	1646	1789	1251
$t_6$	1261	1157	1193	1297	1261	1251	1156	1317	1189	1306
$t_7$	850	928	780	1017	761	900	998	838	797	824
$t_8$	1042	1291	1169	1562	1277	1431	1236	1092	1274	1305
$t_9$	1309	1305	1641	1225	1425	1280	1388	1268	1290	1549
$t_{10}$	881	865	752	893	883	813	892	805	873	915

**Table 7. A high task heterogeneity low machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.5, V_{mach} = 0.1.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	377	476	434	486	457	486	431	417	429	428
$t_2$	493	370	400	420	502	472	475	440	483	576
$t_3$	745	646	922	650	791	878	853	791	756	788
$t_4$	542	490	469	559	488	498	509	431	547	542
$t_5$	625	666	618	710	624	615	618	599	522	540
$t_6$	921	785	759	979	865	843	853	870	939	801
$t_7$	677	767	750	720	797	728	941	717	686	870
$t_8$	428	418	394	460	434	427	378	427	447	466
$t_9$	263	289	267	231	243	222	283	257	240	247
$t_{10}$	1182	1518	1272	1237	1349	1218	1344	1117	1122	1260
$t_{11}$	1455	1384	1694	1644	1562	1639	1776	1813	1488	1709
$t_{12}$	3255	2753	3289	3526	2391	2588	3849	3075	3664	3312

**Table 8. A high task heterogeneity high machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.6, V_{mach} = 0.6.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	1446	1110	666	883	1663	1458	653	1886	458	1265
$t_2$	1010	588	682	1255	3665	3455	1293	1747	1173	1638
$t_3$	1893	2798	1097	465	2413	1184	2119	1955	1316	2686
$t_4$	1014	1193	275	1010	1023	1282	559	1133	865	2258
$t_5$	170	444	500	408	790	528	232	303	301	480
$t_6$	1454	1106	901	793	1346	703	1215	490	537	1592
$t_7$	579	1041	852	1560	1983	1648	859	683	945	1713
$t_8$	2980	2114	417	3005	2900	3216	421	2854	1425	1631
$t_9$	252	519	196	352	958	355	720	168	668	1017
$t_{10}$	173	235	273	176	110	127	93	276	390	103
$t_{11}$	115	74	251	71	107	479	153	138	274	189
$t_{12}$	305	226	860	554	394	344	68	86	223	120

**Table 9. A low task heterogeneity low machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.1, V_{mach} = 0.1.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	985	1043	945	835	830	1087	1009	891	1066	1075
$t_2$	963	962	910	918	1078	1091	881	980	1009	981
$t_3$	782	837	968	960	790	800	947	1007	1115	845
$t_4$	999	953	892	986	958	1006	1039	1072	1090	1030
$t_5$	971	972	913	1030	891	873	898	994	1086	1122
$t_6$	1155	1065	800	1247	980	1103	1228	1062	1011	1005
$t_7$	1007	1191	964	860	1034	896	1185	932	1035	1019
$t_8$	1088	864	972	984	736	950	944	994	970	894
$t_9$	878	967	954	917	942	978	1046	1134	985	1032
$t_{10}$	1210	1120	1043	1093	1386	1097	1202	1004	1185	1226
$t_{11}$	910	958	1046	1062	952	1054	1020	1175	850	1060
$t_{12}$	930	935	908	1155	991	997	828	1062	886	831

**Table 10. A low task heterogeneity high machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.1, V_{mach} = 0.6.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	1679	876	1332	716	1186	1860	662	833	534	804
$t_2$	1767	766	1327	711	957	2061	625	626	642	800
$t_3$	1870	861	1411	932	1065	1562	625	976	556	842
$t_4$	1861	817	1218	865	1096	1660	587	767	736	822
$t_5$	1768	850	1465	764	1066	1585	663	863	579	757
$t_6$	1951	807	1177	914	939	1483	573	961	643	712
$t_7$	1312	697	1304	921	1005	1639	562	831	633	784
$t_8$	1665	849	1414	795	1162	1593	577	791	709	774
$t_9$	1618	753	1283	794	1153	1673	639	787	563	744
$t_{10}$	1576	964	1373	752	950	1726	699	836	633	764
$t_{11}$	1693	742	1454	758	961	1781	721	988	641	793
$t_{12}$	1863	823	1317	890	1137	1812	704	800	479	848

**Table 11. A low task heterogeneity high machine heterogeneity matrix generated by the CVB method.**  
 $V_{task} = 0.1, V_{mach} = 2.0.$

	$m_1$	$m_2$	$m_3$	$m_4$	$m_5$	$m_6$	$m_7$	$m_8$	$m_9$	$m_{10}$
$t_1$	4784	326	1620	1307	3301	10	103	4449	228	40
$t_2$	4315	276	1291	1863	3712	11	91	5255	200	47
$t_3$	6278	269	1493	1181	3186	12	93	4604	235	46
$t_4$	4945	294	1629	1429	2894	14	87	4724	231	45
$t_5$	5276	321	1532	1516	2679	12	102	4621	205	46
$t_6$	4946	293	1467	1609	2661	10	96	3991	255	39
$t_7$	4802	327	1317	1668	2982	10	90	5090	252	42
$t_8$	5381	365	1698	1384	3668	12	99	5133	242	38
$t_9$	5011	255	1491	1386	3061	10	94	3739	216	42
$t_{10}$	5228	296	1489	1515	3632	12	107	4682	203	38
$t_{11}$	5367	319	1332	1363	3393	12	72	4769	221	43
$t_{12}$	4621	258	1473	1501	3124	12	96	4091	199	44

This experimental set-up can then be used to find out which mapping heuristics are best suited for a given set of model parameters (i.e.,  $V_{task}$ ,  $V_{mach}$ , and  $\mu_{task}$  (or  $\mu_{mach}$ )). This information can be stored in a “look-up table,” so as to facilitate the choice of a mapping heuristic given a set of model parameters. The look-up table can be part of the toolbox in the mapper.

The ETC model of Section 2 assumes that the machine heterogeneity is the same for all tasks, i.e., different tasks show the same general variation in their execution times over different machines. In reality this may not be true; the variation in the execution times of one task on all machines may be very different from some other task. To model the “variation in machine heterogeneity” along different rows (i.e., for different tasks), another level of heterogeneity could be introduced. For example, in the CVB ETC generation, instead of having a fixed value for  $V_{mach}$  for all the tasks, the value of  $V_{mach}$  for a given task could be variable, e.g., it could be sampled from a probability distribution. Once again, the nature of the probability distribution and its parameters will need to be decided empirically.

## 5. Related Work

To the best of the authors’ knowledge, there is currently no work presented in the open literature that addresses the problem of modeling of execution times of the tasks in an HC system (except the already discussed work [13]). However, below are presented two tangentially related works.

A detailed workload model for parallel machines has been given in [4]. However the model is not intended for HC systems in that the machine heterogeneity is not modeled. Task execution times are modeled but tasks are assumed to be running on multiple processing nodes, unlike the HC environment presented here where tasks run on single machines only.

A method for generating random task graphs is given in [17] as part of description of the simulation environment for the HC systems. The method proposed in [17] assumes that the computation cost of a task  $t_i$ , averaged over all the machines in the system, is available as  $\bar{w}_i$ . The method does provide for characterizing the differences in the execution times of a given task on different processors in the HC system (i.e., machine heterogeneity). The “range percentage” ( $\beta$ ) of computation costs on processors roughly corresponds to the notion of machine heterogeneity as presented here. The execution time,  $e_{ij}$ , of task  $t_i$  on machine  $m_j$  is randomly selected from the range,  $\bar{w}_i \times (1 - \beta/2) \leq e_{ij} \leq \bar{w}_i \times (1 + \beta/2)$ . However, the method in [17] does not provide for describing the differences in the execution times of all the tasks on an “average” machine in the HC system. The method in [17] does not tell how the differences in the values of  $\bar{w}_i$  for different machines will be modeled. That is, the method in [17] does not consider task heterogeneity.

Further, the model in [17] does not take into account the consistency of the task execution times.

## 6. Conclusions

To describe different kinds of heterogeneous environments, an existing model based on the characteristics of the ETC matrix was presented. The three parameters of this model (task heterogeneity, machine heterogeneity, and consistency) can be changed to investigate the performance of mapping heuristics for different HC systems and different sets of tasks. An existing range-based method for quantifying heterogeneity was described, and a new coefficient-of-variation-based method was proposed. Corresponding procedures for generating the ETC matrices representing various heterogeneous environments were presented. Sample ETC matrices were provided for both ETC generation procedures. The coefficient-of-variation-based ETC generation method provides a greater control over the spread of values (i.e., heterogeneity) in any given row or column of the ETC matrix than the range-based method. This characterization of HC environments will allow a researcher to simulate different HC environments, and then evaluate the behavior of the mapping heuristics under different conditions of heterogeneity.

*Acknowledgments:* The authors thank Anthony A. Maciejewski, Tracy D. Braun, and Taylor Kidd for their valuable comments and suggestions.

## References

- [1] R. Armstrong, *Investigation of Effect of Different Run-Time Distributions on SmartNet Performance* Master’s thesis, Department of Computer Science, Naval Postgraduate School, 1997 (D. Hensgen, Advisor).
- [2] T. D. Braun, H. J. Siegel, N. Beck, L. Bölöni, M. Maheswaran, A. I. Reuther, J. P. Robertson, M. D. Theys, B. Yao, R. F. Freund, and D. Hensgen, “A comparison study of static mapping heuristics for a class of meta-tasks on heterogeneous computing systems,” *8th IEEE Heterogeneous Computing Workshop (HCW ’99)*, Apr. 1999, pp. 15–29.
- [3] H. G. Dietz, W. E. Cohen, and B. K. Grant, “Would you run it here... or there? (AHS: Automatic Heterogeneous Supercomputing),” *1993 International Conference on Parallel Processing (ICPP ’93)*, Vol. II, Aug. 1993, pp. 217–221.
- [4] D. G. Feitelson and L. Rudolph, “Metrics and benchmarking for parallel job scheduling,” in *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, eds., Vol. 1459, Lecture Notes in Computer Science, Vol. 1459, Springer-Verlag, New York, NY, 1998, pp. 1–15.

- [5] R. F. Freund, M. Gherrity, S. Ambrosius, M. Campbell, M. Halderman, D. Hensgen, E. Keith, T. Kidd, M. Kussow, J. D. Lima, F. Mirabile, L. Moore, B. Rust, and H. J. Siegel, "Scheduling resources in multi-user, heterogeneous, computing environments with Smart-Net," *7th IEEE Heterogeneous Computing Workshop (HCW '98)*, Mar. 1998, pp. 184–199.
- [6] A. Ghafoor and J. Yang, "Distributed heterogeneous supercomputing management system," *IEEE Computer*, Vol. 26, No. 6, June 1993, pp. 78–86.
- [7] D. A. Hensgen, T. Kidd, D. St. John, M. C. Schnaidt, H. J. Siegel, T. D. Braun, M. Maheswaran, S. Ali, J.-K. Kim, C. Irvine, T. Levin, R. F. Freund, M. Kussow, M. Godfrey, A. Duman, P. Carff, S. Kidd, V. Prasanna, P. Bhat, and A. Alhusaini, "An overview of MSHN: The Management System for Heterogeneous Networks," *8th IEEE Heterogeneous Computing Workshop (HCW '99)*, Apr. 1999, pp. 184–198.
- [8] M. A. Iverson, F. Özgüner, and G. J. Follen, "Statistical prediction of task execution times through analytic benchmarking for scheduling in a heterogeneous environment," *8th IEEE Heterogeneous Computing Workshop (HCW '99)*, Apr. 1999, pp. 99–111.
- [9] R. Jain, *The Art of Computer Systems Performance Analysis*, John Wiley & Sons, Inc., New York, NY, 1991.
- [10] M. Kafil and I. Ahmad, "Optimal task assignment in heterogeneous distributed computing systems," *IEEE Concurrency*, Vol. 6, No. 3, July-Sep. 1998, pp. 42–51.
- [11] L. L. Lapin, *Probability and Statistics for Modern Engineering*, Second Edition, Waveland Press, Inc., Prospect Heights, IL, 1998.
- [12] N. Lopez-Benitez and J.-Y. Hyon, "Simulation of task graph systems in heterogeneous computing environments," *8th IEEE Heterogeneous Computing Workshop (HCW '99)*, Apr. 1999, pp. 112–124.
- [13] M. Maheswaran, S. Ali, H. J. Siegel, D. Hensgen, and R. F. Freund, "Dynamic mapping of a class of independent tasks onto heterogeneous computing systems," *Journal of Parallel and Distributed Computing*, Special Issue on Software Support for Distributed Computing, Vol. 59, No. 2, pp. 107–131, Nov. 1999.
- [14] M. Maheswaran, T. D. Braun, and H. J. Siegel, "Heterogeneous distributed computing," in *Encyclopedia of Electrical and Electronics Engineering*, Vol. 8, J. G. Webster, ed., John Wiley, New York, NY, 1999, pp. 679–690.
- [15] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, NY, 1984.
- [16] H. Singh and A. Youssef, "Mapping and scheduling heterogeneous task graphs using genetic algorithms," *5th IEEE Heterogeneous Computing Workshop (HCW '96)*, Apr. 1996, pp. 86–97.
- [17] H. Topcuoglu, S. Hariri, and M.-Y. Wu, "Task scheduling algorithms for heterogeneous processors," *8th IEEE Heterogeneous Computing Workshop (HCW '99)*, Apr. 1999, pp. 3–14.
- [18] J. Yang, I. Ahmad, and A. Ghafoor, "Estimation of execution times on heterogeneous supercomputer architecture," *1993 International Conference on Parallel Processing (ICPP '93)*, Vol. I, Aug. 1993, pp. 219–225.

## Biographies

**Shoukat Ali** is pursuing a PhD degree from the School of Electrical and Computer Engineering at Purdue University, where he is currently a Research Assistant. His main research topic is dynamic mapping of meta-tasks in heterogeneous computing systems. He has held teaching positions at Aitchison College and Keynesian Institute of Management and Sciences, both in Lahore, Pakistan. He was also a Teaching Assistant at Purdue. Shoukat received his MS degree in electrical and electronic engineering from Purdue University in 1999, and his BS degree in electrical and electronic engineering from the University of Engineering and Technology, Lahore, Pakistan in 1996. His research interests include computer architecture, parallel computing, and heterogeneous computing.

**Howard Jay Siegel** is a Professor in the School of Electrical and Computer Engineering at Purdue University. He is a Fellow of the IEEE and a Fellow of the ACM. He received BS degrees in both electrical engineering and management from MIT, and the MA, MSE, and PhD degrees from the Department of Electrical Engineering and Computer Science at Princeton University. Prof. Siegel has coauthored over 250 technical papers, has coedited seven volumes, and wrote the book *Interconnection Networks for Large-Scale Parallel Processing*. He was a Coeditor-in-Chief of the *Journal of Parallel and Distributed Computing*, and was on the Editorial Boards of the *IEEE Transactions on Parallel and Distributed Systems* and the *IEEE Transactions on Computers*. He was Program Chair/Co-Chair of three conferences, General Chair/Co-Chair of four conferences, and Chair/Co-Chair of four workshops. He is an international keynote speaker and tutorial lecturer, and a consultant for government and industry.

**Muthucumar Maheswaran** is an Assistant Professor in the Department of Computer Science at the University

of Manitoba, Canada. In 1990, he received a BSc degree in electrical and electronic engineering from the University of Peradeniya, Sri Lanka. He received an MSEE degree in 1994 and a PhD degree in 1998, both from the School of Electrical and Computer Engineering at Purdue University. He held a Fulbright scholarship during his tenure as an MSEE student at Purdue University. His research interests include computer architecture, distributed computing, heterogeneous computing, Internet and world wide web systems, metacomputing, mobile programs, network computing, parallel computing, resource management systems for metacomputing, and scientific computing. He has authored or coauthored 15 technical papers in these and related areas. He is a member of the Eta Kappa Nu honorary society.

**Debra Hensgen** is a member of the Research and Evaluation Team at OpenTV in Mountain View, California. OpenTV produces middleware for set-top boxes in support of interactive television. She received her PhD in the area of Distributed Operating Systems from the University of Kentucky. Prior to moving to private industry, as an Associate Professor in the systems area, she worked with students and colleagues to design and develop tools and systems for resource management, network re-routing algorithms and systems that preserve quality of service guarantees, and visualization tools for performance debugging of parallel and distributed systems. She has published numerous papers concerning her contributions to the Concurra toolkit for automatically generating safe, efficient concurrent code, the Graze parallel processing performance debugger, the SAAM path information base, and the SmartNet and MSHN Resource Management Systems.

**Sahra Ali** is pursuing a PhD degree at the School of Electrical and Computer Engineering at Purdue University, where she is currently a Research Assistant. Her main research topic is the modeling of reliability in software intensive communication networks. She has also been working as a software developer for Cisco Systems since 1997. She was previously a Teaching Assistant and lab coordinator at Purdue. Sahra received her MS degree in electrical engineering from Purdue University in 1998, and her BS degree in electrical and electronic engineering from Sharif University of Technology, Tehran, Iran in 1995. Her research interests include testing, reliability and availability of communication networks, as well as multimedia.