



ELSEVIER

Available at
www.ComputerScienceWeb.com
 POWERED BY SCIENCE @ DIRECT®

COMPUTER
 NETWORKS

Computer Networks xxx (2003) xxx–xxx

www.elsevier.com/locate/comnet

2 Overload control in QoS-aware web servers [☆]

3 Huamin Chen, Prasant Mohapatra *

4 *Department of Computer Science, University of California, One Shields Avenue, Davis, CA 95616, USA*

5 Received 14 November 2002; accepted 16 January 2003

6 Abstract

7 With the explosive use of Internet, contemporary web servers are susceptible to overloads during which their services
 8 deteriorate drastically and often lead to denial of services. Overloads are of more serious concerns for QoS-aware
 9 servers. Evaluation of performance of QoS-aware servers in terms of the number of request completion is not very
 10 meaningful. A better measure would be the number of completed sessions. In this paper, we proposed two methods to
 11 prevent and control overloads in web servers by utilizing session-based relationship among HTTP requests. We first
 12 exploited the dependence among session-based requests by analyzing and predicting the reference patterns. Using the
 13 dependency relationships, we have derived traffic conformation functions that can be used for capacity planning and
 14 overload prevention in web servers. Second, we have proposed a dynamic weighted fair sharing (DWFS) scheduling
 15 algorithm to control overloads in web servers. DWFS is distinguished from other scheduling algorithms in the sense
 16 that it aims to avoid processing of requests that belong to sessions that are likely to be aborted in the near future. The
 17 experimental results demonstrate that DWFS can improve server responsiveness by as high as 50% while providing QoS
 18 support through service differentiation for a class of application environment.

19 © 2002 Published by Elsevier Science B.V.

20 *Keywords:* Capacity planning; Dynamic weighted fair sharing; Overload control; Quality of service; Scheduling algorithm; Service
 21 differentiation; Session-based control; Web server

22 1. Introduction

23 As the widespread usage of web service grows,
 24 the number of accesses to many popular web sites
 25 is ever increasing and occasionally reaches the
 26 limit of their capacity and consequently causes the

27 servers to be overloaded. As a result, end users
 28 either receive busy signal or nothing at all before
 29 the browser indicates a time-out error or the user
 30 aborts (stops) the request. Subsequently, the server
 31 may get choked or crash causing denial of services.
 32 Such abnormality is often regarded as the servers'
 33 poor quality and compromises their long term
 34 survivability. In e-commerce applications, such
 35 server behavior could translate to sizable revenue
 36 losses.

37 Research on overload prevention and control
 38 has been limited compared to the other perfor-
 39 mance improvement issues such as web caching,
 40 and load balancing in web servers. These perfor- 40

[☆] A preliminary version of this paper appeared in the Proceedings of IEEE INFOCOM 2002. This research was supported in part by the National Science Foundation through the grants CCR-0296070 and ANI-0296034.

* Corresponding author.

E-mail addresses: chenhua@cs.ucdavis.edu (H. Chen),
prasant@cs.ucdavis.edu (P. Mohapatra).

41 mance enhancement techniques, however, are in- 89
 42 adequate in ensuring a busy web server from being 90
 43 overloaded due to the fact that the web traffic is 91
 44 highly unpredictable and bursty [10,15]. Proper 92
 45 capacity planning and forecasting methods can 93
 46 prevent servers from being overloaded under 94
 47 controlled traffic conditions. 95

48 In many web sites, especially in e-commerce, 96
 49 online brokers, and supply chain sites, majority of 97
 50 the requests in the web traffic are session-based. A 98
 51 session contains temporally and logically related 99
 52 request sequences from the same client. Sessions 100
 53 can be identified either by HTTP/1.1 persistent 101
 54 connections [12] or from the state information 102
 55 within the presence of cookies [14]. Sessions ex- 103
 56 hibit distinguishable features from individual re- 104
 57 quests. For example, session integrity requires that 105
 58 once admitted for processing, all the following 106
 59 requests within a session should be honored. 107
 60 Similarly, session affinity would require that re- 108
 61 quests belonging to the same session are handled 109
 62 by the same front-end server for security and lo- 110
 63 cality reasons. These features may complicate or 111
 64 contradict the research conclusions of the perfor- 112
 65 mance studies on web servers where the number of 113
 66 request completions have been considered as the 114
 67 primary performance measure. For example, ad- 115
 68 mission control on a per request basis may lead to 116
 69 a large number of broken or incomplete sessions 117
 70 when the system is overloaded. Incomplete ses- 118
 71 sions may be equivalent to a rejected session from 119
 72 the users viewpoint or for most e-commerce serv- 120
 73 ers. Thus, performance measure based on the 121
 74 number of request completions may not be a good 122
 75 indication of users satisfaction (the basic purpose 123
 76 of web service). Especially during overloads, the 124
 77 disparity between the two types of performance 125
 78 measures (proportion of request completion and 126
 79 proportion of session completion) is more en- 127
 80 hanced. Capacity planning schemes based on in- 128
 81 dividual requests also have the same deficiency. 129

82 Session integrity is a critical metric for com- 130
 83 mercial web service. For an online retailer, the 131
 84 more the number of sessions completed, the more 132
 85 the amount of revenue that is likely to be gener- 133
 86 ated. The same statement cannot be made about 134
 87 the individual request completions. Sessions that 135
 88 are broken or delayed at some critical stages, like

checkout and shipping, could mean loss of revenue 89
 to the web site. From the end users' perspective, 90
 this means poor service availability. Therefore, it is 91
 more useful to use session integrity to evaluate the 92
 service availability of servers, especially during 93
 high-load periods. 94

95 In this paper, we explore the session character- 96
 97 istics and their potential in overload control and 98
 99 prevention. A workload characterization study is 100
 done first to gain an insight to the load patterns in 101
 web servers. The workload characterization study 102
 was based on the server log from a popular online 103
 retailer. We found that, despite the seemingly 104
 complication of session sequences, some statistical 105
 results can be used in simplifying the session-based 106
 traffic model. Based on these results, the session 107
 logic can be utilized for capacity planning and re- 108
 quest scheduling of QoS-aware servers, which im- 109
 proves server's productivity. Server productivity 110
 quantifies the amount of useful work done by the 111
 server. Based on the session-level traffic model, we 112
 have proposed a dynamic weighted fair scheduling 113
 (DWFS) scheme that assign service weight to dif- 114
 ferent requests of a session in a dynamic manner. 115
 We have done an experimental performance anal- 116
 ysis by modifying the scheduling scheme of the 117
 Apache web server. The proposed DWFS scheme 118
 provides a performance improvement of about 50% 119
 in terms of response delay and significantly reduces 120
 the session abortion rate for the workload and 121
 system configuration used in the experimentation. 122

123 The rest of the paper is organized in the fol- 124
 125 lowing manner. Section 2 characterizes session- 126
 127 based HTTP requests. Section 3 provides capacity 128
 129 planning tools to prevent server overload. Section 130
 4 proposes a request scheduling algorithm to 131
 control server overload and improve server per- 132
 formance followed by experimental results in 133
 Section 5, which proves the feasibility and quan- 134
 tifies the performance of the proposed algorithm. 135
 The related works are discussed in Section 5, fol- 136
 lowed by the concluding remarks in Section 7. 137

2. Session-based web traffic characterization 131

A session in web accesses can be defined as a 132
 sequence of requests that form one complete 133




Fig. 1. An example of a web session represented as a state machine.

134 transaction. A session during web accesses can be
 135 represented as a finite state machine with each
 136 state representing a stage that a request is under-
 137 going. Fig. 1 depicts an example of such a repre-
 138 sentation. The directed arc (A,B) represents a
 139 transition from state A to state B with a proba-
 140 bility $P(A,B)$. The four states A, B, C, and D
 141 could be representing states corresponding to main
 142 menu, checkout, browsing and search in an e-
 143 commerce site.

144 In web services, a stage can be a single URL or a
 145 group of URL's that have the same reference
 146 pattern and resource claim profile. A session can
 147 be mandatory or voluntary. A mandatory session
 148 refers to the situation in which the descendant re-
 149 quests are generated by the browsers instead of
 150 clients. The requests for embedded image files
 151 within an HTML page is an example of this case.
 152 We call the page that has embedded files as the
 153 *main page*. In voluntary sessions, the descendant
 154 requests are generated by the clients explicitly. For
 155 example, the client clicks a link within the current
 156 main page to browse another page. In current web
 157 server architecture, most of the image files are
 158 served by edge servers [2] or dedicated image
 159 servers which are physically separated from those
 160 serving the main pages. Therefore, the perfor-
 161 mance of these servers is largely affected by the
 162 service of main pages. Thus, the following discus-
 163 sion is focused on voluntary sessions.

164 During our research, we obtained the trace of
 165 accesses to an e-commerce web server from a
 166 popular online retailer.¹ Based on the reference

¹ We refrain from mentioning the name of the retailer honoring a non-disclosure agreement. Without the non-disclosure agreement, we would have been able to obtain the data.

167 traces, we characterized the basic behavior of ses-
 168 sion-based traffic. The characterizations are de-
 169 rived from a typical daily traffic trace. Previous
 170 studies [10] have characterized the web traffic as
 171 very bursty, which is also observed in our result as
 172 shown in Fig. 2. From Fig. 2, it is observed that
 173 the traffic load is highest during the period of
 174 17:00–23:00 h, which accounts for over 50% of the
 175 daily traffic. The traffic volume peaked at 20:00–
 176 21:00, where nearly 10% of overall requests were
 177 initiated. So the server is presumably more heavily
 178 loaded during this period (during evening hours),
 179 which was confirmed from the server side perfor-
 180 mance data recorded by the MS Performance
 181 Monitor.

182 We further investigate the relationship between
 183 request queue length (the waiting requests and
 184 those being served) and the request processing
 185 time under heavy server load. Since the processing
 186 time for individual URL's varies, we adopt the
 187 measure called *stretch factor* from [23], which re-
 188 fers to the quotient between the current processing
 189 time and the processing time of the same URL
 190 under normal load. The stretch factor reflects the
 191 current server load. Fig. 3 depicts the queue length
 192 and the stretch factor during the 20:00–21:00 pe-
 193 riod. It is observed that the two curves show
 194 similar pattern, indicating that the server load is
 195 proportional to the queue length. While the queue
 196 length is a good indication of server load, the na-
 197 ture of the jobs in the queue also has a great im-
 198 pact on server performance. Servers whose




Fig. 2. Traffic histogram of the server trace for a day.

4

H. Chen, P. Mohapatra / Computer Networks xxx (2003) xxx-xxx

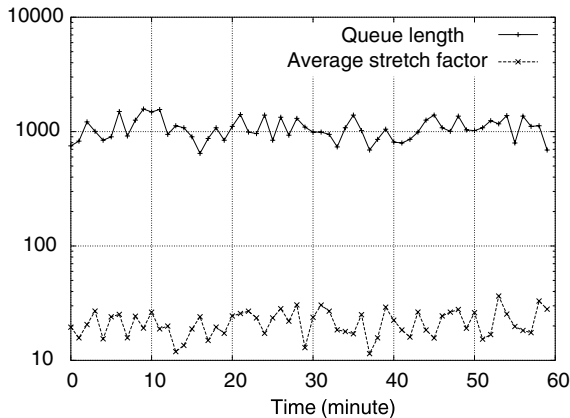


Fig. 3. Queue length and stretch factor.




Fig. 4. Queue composition for the 1 h duration.

199 workload is dominated by static files (like HTML
200 pages, images) generally perform better than the
201 ones with high proportion of dynamic files (like
202 CGI, ASP, etc). Therefore, unless the composition
203 of workload is comparable, using the queue length
204 as an indicator of server load is incorrect.

205 We sorted the requests into different queues ac-
206 cording to their nature (main menu, checkout,
207 search, browsing) and analyzed the composition of
208 workload at each time period. These queues usually
209 have different resource consumption profiles. For
210 main menu queues, the major task is static HTML
211 and image file rendering, and the update fre-
212 quency is relatively low. For checkout queues, the
213 process is SSL-secured and thus the workload is
214 very CPU intensive. For search queues, the back-
215 end database is queried and the server only receives
216 the query results and assembles them in the HTML
217 format. Characteristics of the browsing queue is
218 like main menu queue except that more image files
219 are rendered. Fig. 4 presents the nature of the
220 workload composition during the period 20:00–
221 21:00, which is observed to be relatively stable.

222 It has been realized in prior studies that the re-
223 quests within a session reveal statistically depen-
224 dent relationship [7,17]. Conclusions from these
225 studies show that historic reference patterns can be
226 exploited to predict the subsequent requests. Pre-
227 diction method of subsequent requests within a
228 session are different. In this study, we use the
229 transition probability of the state machine for
230 prediction, which derives the subsequent URL

231 from the current one. This method requires no
232 sophisticated mathematical modeling and uses less
233 computation power in practice. The probability is
234 obtained either from offline historic records like
235 server logs or from online statistics. The work in
236 this paper is based on offline record of server logs,
237 however other methods can also be applied in a
238 similar way. We will later show that even with this
239 simple prediction method, the performance im-
240 provement is significant.

241 Fig. 5 shows the transition probability (in per-
242 centage) matrix among the stage vector compo-
243 nents: main menu (MM), checkout (CO), search
244 (SR), browsing (BR) from the online retailer's
245 server log. The row vectors show the transition
246 from one stage to another. For example, the sec-
247 ond row indicates that the transition possibility
248 from CO to MM, CO, SR, and BR is 13.5%,
249 44.9%, 40.2%, and 1.4%, respectively. The server
250 log file format follows W3C extended logging [18].
251 Each request entry contains a user ID for the login
252 user which facilitates the identification of session
253 owners. Session integrity is maintained by the IIS
254 server (Microsoft Internet Information Server).

	MM	CO	SR	BR
MM	74.8	0.6	19.5	5.1
CO	13.5	44.9	40.2	1.4
SR	18.1	3.4	74.3	4.3
BR	13.7	1.0	14.6	70.7

Fig. 5. State transition matrix.

255 Another issue related to the session behavior is
 256 the user thinking time between consecutive re-
 257 quests within a session. It is random in nature and
 258 varies for different web sites. In our server trace,
 259 the thinking time was usually short and less than
 260 60 s. Characterization of other traces reveals sim-
 261 ilar results. When the traffic is high, which is the
 262 case for heavily loaded servers, the long term effect
 263 of the thinking time can be ignored.

264 Finally, when evaluating the relationship be-
 265 tween the number of outstanding requests and the
 266 number of active sessions, we found that the ratio
 267 between the two is stable. Though each session can
 268 fork several requests simultaneously, the fact that
 269 some others do not send any requests offsets it,
 270 which makes the overall behavior as stable. Fig. 6
 271 shows this ratio for the trace of the period 20:00–
 272 21:00. In this figure, the average value is 0.526 with
 273 a standard deviation of 0.017, which means that
 274 the variation is small and the ratio is stable. This
 275 result is useful in estimation of the request arrival
 276 rate based on the number of active sessions.

277 Obtaining server traces from e-commerce sites
 278 has been difficult due to security and proprietary
 279 reasons. We could manage to get the traces from
 280 one corporation. Although our analysis and re-
 281 sults use only this set of trace, the proposed
 282 methodology is applicable for any other server
 283 trace. So we have laid emphasis on the method-
 284 ology, the trends, and relativeness of the results,
 285 rather than the absolute numbers (which are spe-
 286 cific for the trace).




Fig. 6. Ratio of request arrival rate to session generation rate.

3. Session-based capacity planning

287

288 We have considered QoS-aware web servers for
 289 our study in which requests are served based on
 290 their priority levels. The basis of priority assign-
 291 ment is actively discussed in [11] and is beyond the
 292 scope of this study. In QoS-aware web servers, an
 293 important and interesting performance metric is
 294 the delay bound, which is the maximum response
 295 delay a request encounters. Besides the processing
 296 time of each request, there are other latencies as-
 297 sociated with the service of a request, such as
 298 queuing delay and network transmission delay. In
 299 this paper, we only consider the delay at the server
 300 side. The study of network delay is not within the
 301 scope of this paper and is being extensively studied
 302 in the IETF architectures [11]. We assume that the
 303 service level agreement (SLA) that can be provided
 304 by a web server would consist of a bounded delay
 305 for each QoS level if the request arrival rate does
 306 not exceed an agreed amount. From the web server
 307 perspective, QoS attributes are defined in terms of
 308 the maximum rate of request arrival and the la-
 309 tency bound of each request. An SLA can thus be
 310 stated in terms of (λ, δ, s) , where λ is the maximum
 311 rate of arrival, δ is the delay bound, and s is the
 312 proportion of requests that meet the delay bound.
 313 In a QoS-aware web server, overload is said to
 314 occur when the SLA is violated for an extended
 315 period of time. Thus, we formally define overload
 316 as follows.

317 Let the predefined SLA for a specific QoS level
 318 is (λ, δ, s) . If for an extended period of
 319 time T , while the arrival rate is less than λ ,
 320 the proportion of requests that meet the delay
 321 bound stay less than s , the web server is said
 322 to be overloaded. It is implicit in this defini-
 323 tion that the SLA negotiation was done con-
 324 sidering the server capacity and workload
 325 conditions.

326 The following analysis focuses on the worst case
 327 where requests compete for CPU resources (Table
 328 1). The QoS rules are defined for each of the
 329 URL's or for URL groups. The SLA specifies the
 330 delay bound of the QoS groups when the request
 331 arrival rate is below some threshold.

Table 1
Notations used for the analysis

Notations	Description
$P_{a,b}$	State transition probability from a to b
λ_a^i	Request arrival rate of session class i to state a
μ_a^i	Request departure rate of session class i from state a
d_a^i	Delay bound of class i at stage a
T_a	Processing time at stage a
ϕ_i	Session generation rate of class i
r_{rs}	Ratio between the number of requests and the number of sessions

332 We only consider a steady-state system because
333 a transient model is mathematically intractable
334 and may be of little practical use. In a steady-state
335 system, the number of requests arriving at and
336 departing from the server must be equal. Thus, the
337 arrival rate λ_i^p is equal to the departure rate μ_i^p
338 under the steady-state assumption.

339 The turnaround time at each stage is less than
340 the lower limit, which is the sum of the delay
341 bound and the processing time. Thus,

$$\mu_i^p \geq \frac{1}{T_i + d_i^p}.$$

343 The input to each stage is the output from the
344 other stages with certain transition probability.
345 Excluding the source and sink stages, we then have

$$\lambda_j^p = \sum \mu_i^p * P_{i,j} = \mu^p \cdot P_j,$$

347 where μ^p and P_j are the vector forms of μ_i^p
348 and the second expression is the dot product.

349 With a pre-emptive priority scheduling disci-
350 pline, the high priority requests are scheduled be-
351 fore the low priority ones. Thus the departure rate
352 of a priority group is less than the service rate of
353 higher priority groups.

$$\frac{1}{T_i + d_i^p} \leq \mu_i^p \leq \frac{\lambda_i^p}{\sum_{j=1}^p \lambda_k^j * T_k}. \quad (1)$$

355 Using the steady-state assumption,

$$\frac{1}{T_i + d_i^p} \leq \lambda_i^p \leq \left[\frac{\lambda_k^p}{\sum_{j=1}^p \lambda_k^j * T_k} \right] \cdot P_i. \quad (2)$$

357 Finally, using (2) for all the stages and presenting
358 in a matrix form, we get:

Eq. (3) reveals the relationship between delay 359
and traffic volume of the QoS priority groups. To 360
infer 361

$$\left[\frac{1}{T_i + d_i^p} \right]_{m*n} \leq [\lambda_i^p]_{m*n} \leq \left[\frac{\lambda_k^p}{\sum_{j=1}^p \lambda_k^j * T_k} \right]_{m*n} * [P]_{n*n}, \quad (3)$$

where m is the number of priority groups and
 n is the number of stages. We call this expres-
sion as the traffic conformation inequation

the SLA from this function, we can analyze in a 366
stepwise manner. It is obvious that for the highest 367
priority, the request arrival rate at stage i is con- 368
strained by 369

$$\frac{1}{T_i + d_i^1} \leq \lambda_i^1.$$

In practice, the λ_i^1 can be set to its lower bound so 371
that they will not cause excessive delay for the 372
lower priority groups and maximize system utili- 373
zation by allowing more requests into the system. 374
Thus the arrival rate of the next immediate priority 375
group is bounded by 376

$$\frac{1}{T_i + d_i^2} \leq \lambda_i^2 \leq \left[\frac{\lambda_k^2}{\sum_{j=1}^2 \lambda_k^j * T_k} \right] * P_i.$$

Similarly, the arrival rate of other priority 378
groups can be obtained. 379

After obtaining the request arrival threshold, it 380
is easy to define the session generation rate using 381
the ratio r_{rs} . As discussed in Section 2, r_{rs} is more 382
or less stable when the traffic volume is high. Thus 383
the session generation rate of class i is given by 384

$$\phi_i = r_{rs} \sum \lambda_j^i. \quad (4)$$

4. Productive scheduling algorithm 386

Among the several causes that are responsible 387
for the degradation of server output, scheduling of 388
requests is a critical factor. For example, queues 389
consisting of time-consuming requests have a good 390
chance of getting accumulated. As a result, they 391
would dominate in controlling CPU resources in a 392

393 round robin scheduling mode. Consequently, more
 394 time is spent on these queues and the effective
 395 overall output is degraded. Intuitively a conser-
 396 vative admission control can prevent the server
 397 from being overloaded. But such conservativeness
 398 is not easy to realize because of the burstiness of
 399 traffic and the likely-hood of leading to under-
 400 utilization of the server. We believe a relatively
 401 relaxed admission control assisted by an efficient
 402 scheduling algorithm is a better alternative. The
 403 admission control admits as many sessions as
 404 possible so long as the server is not overloaded.
 405 Our previous work on admission control based on
 406 predictable service time [8] could serve this pur-
 407 pose. In addition, the scheduling algorithm takes
 408 care of the situation when the admitted sessions
 409 are beyond the server's capacity. It seeks the best
 410 scheduling that produces as many completed ses-
 411 sions (not necessarily requests) as possible.

412 In the context of sessions, each of the waiting
 413 queue represents a particular task of the session
 414 sequence and its output serves as input to the other
 415 queues. So proper shaping of these queues by
 416 means of priority scheduling among different
 417 queues can alleviate overload conditions. This is
 418 the basic idea of our scheduling algorithm.

419 Another phenomenon that is frequently ob-
 420 served in web servers is that during overload, more
 421 number of tasks gets aborted. Before abortion,
 422 these tasks consume excessive system resources
 423 during the crunch period. To resolve this ill-effect
 424 we use a scheme in which requests of sessions that
 425 have a higher probability of getting completed are
 426 scheduled first. Such a scheduling approach helps
 427 the server in doing more useful work during
 428 overload situations, while avoiding the service of
 429 requests whose sessions are likely to get aborted.

430 4.1. Comparison of scheduling algorithms

431 The popular scheduling algorithms that are used
 432 in web servers include round robin (RR), earliest
 433 deadline first (EDF) and weighted fair sharing
 434 (WFS). RR and EDF do not consider the rela-
 435 tionship of inter-session request transition thus
 436 they cannot help in session-based overload con-
 437 trol. WFS provides higher levels of service to the
 438 tasks that have higher priority. Thus the queue

length accumulation at some stages can reach a
 steady-state by lowering the request injection rate
 and raising the service rate. However, this dis-
 crimination increases request accumulations at
 other queues, which in turn would result in request
 drops because of timeout. This domino effect dis-
 rupts the normal request transition flow into other
 queues and eventually leads to lower throughput
 in terms of the number of completed sessions.

We introduce a measure called *server's produc-*
tivity, which is defined as a function of request
 completion and error rate of requests that belong
 to ongoing sessions for a server, and can be for-
 mally expressed as follows.

Server's productivity during time interval T : If
 the number of requests completed within T is
 c , and the number of requests aborted during
 this time is e , then the *server's productivity*
 during T is $(c - e)$.

Server productivity is thus a measure of the
 amount of useful work done by a server during a
 given time period. Serving fewer requests while
 more number of requests get aborted leads to a
 negative productivity.

A request abortion occurs either because of
 some of the internal problems at the server side or
 due to the processing timeout imposed by the
 script languages like ASP and PHP. A request
 could also be aborted by the user because of im-
 patience. Ignoring the internal problems of the
 server, we use request timeout to represent all the
 server processing failures. To illustrate how the
 weight assignment could affect server productivity,
 we simulated a round robin scheduler and col-
 lected the results under different weight assign-
 ments. The simulator has four queues
 corresponding the four states (MM, CO, SR, BR)
 with the same transition probability as was listed
 in Fig. 5. The processing times of the stages are
 0.5, 1, 1.5, and 2, respectively. The timeout dura-
 tion was assumed 20 time units and the simulation
 duration was set to 10,000 time units. The more
 weight a queue is assigned, the more CPU time it
 can use. Table 2 displays the different server pro-
 ductivities. It is observed that proper weight as-
 signment can significantly improve the overall

Table 2
Server's productivity comparison

Weight assignment	Requests completed	Requests timed out	Server's productivity
(1,1,1,1)	3077	747	2330
(5,1,1,1)	3930	616	3314
(1,5,1,1)	1585	826	759
(1,1,5,1)	2520	880	1640
(1,1,1,5)	1685	879	806
(5,1,5,1)	3463	477	2986
(1,5,1,5)	1161	767	394

485 performance by increasing the number of completed requests and reducing the number of time-outs. In this set of results, the best case performance is nearly eight times better than the worst case. This inference inspires to seek productivity improvement through appropriate weight assignment.

492 Motivated by the server productivity study, we propose a DWFS scheme based on the temporal relationship in web session in such a manner that the weight distribution is not static all the time. Instead, it depends on the accumulation at the queue and the output rate with the goal to improve the server productivity. Unlike the traditional algorithms that seek short term throughput improvement, DWFS tries to smooth the domino effect of overloads in pursuit of sustained throughput.

503 4.2. Dynamic weighted fair sharing

504 The objective of DWFS is to improve the server's productivity through dynamic weight assignment for scheduling purpose. In an overloaded server, processing at one queue is not productive when it overwhelms other queues. For example, in an e-commerce server, the merchandise browsing and checkout are two queues. If the browsing queue is processed faster than the checkout queue, then the clients proceed to check out only to be jammed, and most of the requests get timed-out. In this case the browsing queue becomes unproductive. On the other hand, if the server distributes more weight on the check out queue such that the requests in the browsing queue will experience longer but tolerable service time, then the input to

519 the downstream checkout queue is reduced and their probability of receiving service is increased, consequently the end user can perceive a faster service for the entire session. 520 521 522

523 The other aspect of DWFS is that, if a server knows a priori that the request it is serving is unproductive, it can stop or delay processing the current request queue and transfer the weight to serve other queues to improve the server's productivity. 524 525 526 527 528

529 The modeling tool used for DWFS is a queuing network with limited waiting room. A k -waiting room queue can accommodate at most k entries waiting for service and the new arrivals are simply dropped. More specifically, as a measure to allow dynamic weighing, k is defined to be the ratio of service time and session timeout period. Each output from a queue produces a credit if the output goes to a queue that is not yet full (productive queues). The credit reflects the productiveness of the service. If service to a queue is known to be unproductive, then its weight is transferred to handle other queues. This conflicts with philosophies behind some other scheduling algorithms that seek maximal throughput from the server in all situations. In DWFS, some jobs may be dropped or delayed even though they could have been served before the deadline if more weight were assigned to them. However, in the long run, more incoming requests can meet their service expectation thus the overall throughput increases. If the drop rate is not high to overshadow the throughput, the server's productivity improves. 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544 545 546 547 548 549 550 551

Productivity function:

$$f(n) = \sum_{i=1}^n \sum_{j=1}^n P_{i,j} * 1 \left(L_j * \frac{T_j}{w_j} < \text{Timeout} \right), \quad (5)$$

where w_j is the normalized weight assigned to queue j and L_j is the queue length. $1(x) = 1$ if x is true, 0 otherwise.

552 The productivity function defined in Eq. (5) states that if the output from a queue is served 553 554 555 556 557 558

559 before the deadline, then it is considered produc- 602
 560 tive. The productiveness of the output is deter- 603
 561 mined by the destination it leaves for; if the 604
 562 destination queue is full, then no credit is earned 605
 563 but a penalty is imposed, otherwise a credit is 606
 564 added. The credit can also be inversely propor- 607
 565 tional to the queue length to avoid filling up the 608
 566 queue. Since there are more than one possible 609
 567 destinations for the output, the credit is reflected
 568 by multiplying it with the transition probability.
 569 The capacity of the waiting room depends on the
 570 service weight; the more weight assigned to the
 571 queue, the less time a request takes to complete,
 572 and more requests can be served before the time-
 573 out period expires.

574 Eq. (5) can be rewritten as

$$f(n) = \sum_{j=1}^n Q_j * 1 \left(L_j * \frac{T_j}{w_j} < \text{Timeout} \right), \quad (6)$$

576 where $Q_j = \sum_{i=1}^n P_{i,j}$. The maximization of $f(n)$
 577 can be achieved through dynamic programming
 578 method as illustrated in Algorithm 1, the com-
 579 plexity of which is $O(n^2)$. When the number of
 580 stages is small, the computational overhead is
 581 trivial.

582 **Algorithm 1.** Pseudocode of productivity function
 583 solver

584 • **INPUT:**
 585 L[1..n], T[1..n], Q[1..n], and timeout have
 586 the same meaning as in Eq. (5)
 587 w: remaining weight
 588 j: current queue to calculate
 589 • **OUTPUT:**
 590 maximum productivity
 591 • **ALGORITHM:**
 592 int max_fn(w, j) {
 593 if(j >= n)
 594 return 0;
 595 min_w = T[j] * L[j]/timeout;
 596 if(min_w > w)
 597 return max_fn(w, j + 1);
 598 f1 = max_fn(w - min_w, j + 1) + Q[j];
 599 f2 = max_fn(w, j + 1);
 600 return max(f1, f2);
 601 }

It is inferred from the productivity function that,
 when other parameters are fixed, the timeout value
 determines how many jobs can be queued in each
 stage. Bigger the timeout value, the more number
 of jobs that can be queued leading to longer mean
 queueing time. We were thus inspired to differen-
 tiate QoS based on timeout values. This impact is
 further explored in Section 5.

5. Experimental performance evaluation 610

611 From the above discussion, we can see that the 611
 612 DWFS's approach to relieve an overloaded server 612
 613 is to add weights on those requests whose de- 613
 614 scendant requests within the same session can be 614
 615 honored. Those requests whose descendant re- 615
 616 quests are predicted to miss their delay bound are 616
 617 delayed for later processing. To verify the feasi- 617
 618 bility of this scheme, we implemented the algo- 618
 619 rithm in an Apache web server and evaluated its 619
 620 performance.

5.1. Experimental setup 621

622 The test-bed contains a web server and several 622
 623 clients. The server has an Intel PIII 733MHZ CPU 623
 624 and 128 MB RAM, running Apache 1.3.19 for MS 624
 625 Windows 2000. Apache [3] is an open source, 625
 626 widely used web server. Fig. 7 illustrates the re- 626
 627 quest handling process in the Apache server. At 627
 628 runtime, the Apache server consists of worker 628
 629 processes (or threads in some systems like MS 629
 630 Windows) and one listener. The listener listens on 630
 631 HTTP port (usually TCP 80) and accepts new 631
 632 connections. It adds the connections into a job 632
 633 queue by calling add_job(). At this time, ad- 633
 634 d_job() does not parse the HTTP request lines. 634
 635 Each of the worker processes unlinks a job from 635
 636 the job queue by calling remove_job(). Only then 636
 637 the request line is parsed and the priority group 637
 638 and stage the job belongs to are known. This 638
 639 working mechanism does not fit DWFS in the 639
 640 sense that the worker processes have no control 640
 641 over the job queues. We have changed the way new 641
 642 jobs are added.

643 As shown in Fig. 8, instead of connections, in 643
 644 our modification, the jobs are HTTP requests and 644




Fig. 7. Apache implementation of request handling

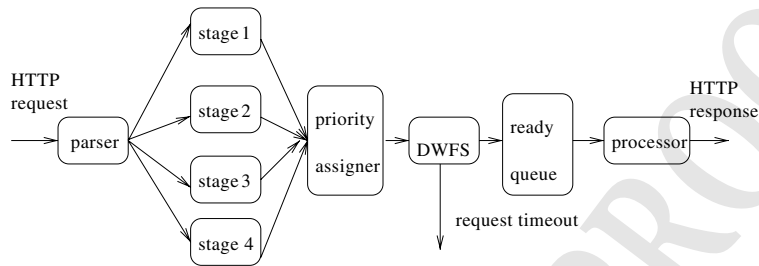


Fig. 8. Revised implementation of request handling

645 the parser sorts them into different queues based
 646 on their URL's and records the time-stamp they
 647 are added into the queue. Then the priority as-
 648 signer assigns different priority to the requests
 649 based on their IP addresses. The DWFS module is
 650 invoked when the ready queue is empty to exercise
 651 the DWFS algorithm to assign jobs to the worker
 652 processes and drops the requests that are timed
 653 out. Since Apache server itself has no control over
 654 the CPU time slot, the weights are assigned as the
 655 number of requests from the queues to be dis-
 656 patched to the ready queue. At this time, the re-
 657 quest that exceed the timeout are dropped. Finally
 658 the processor handles the requests within the ready
 659 queue.

660 The clients used in the test-bed are several Sun
 661 UltraSparc 10 workstations running Solaris 2.8.
 662 The benchmarking tools used in the experiment is
 663 a modified version of WebStone 2.5 [22]. The re-
 664 quest distribution in the original WebStone
 665 benchmark is not adequate for our purpose. The
 666 way a client picks a URL is randomized and the
 667 URL array size is too small in the original
 668 benchmark. We modified the source so that each
 669 client sequentially picks a URL from the array to
 670 simulate a session and all the clients together re-
 671 play the request trace of the online retailer that we
 672 have characterized earlier.

673 The working mechanism of the experiment is as
 674 follows. A master process instructs a configurable

number of clients to send HTTP requests to a web
 server for a specified time period. During the test,
 the clients choose a URL from an array and keep
 track of the connection time, response time, and
 other parameters. After the test finishes, the clients
 send the statistics to the webmaster process and
 the latter reports the test results after collecting all
 the clients' data. The final results contain infor-
 mation such as server connection rate (the number
 of connections the server accepts; the higher the
 better), average response time (the average time-
 span the client receives the whole response; the
 smaller the better), average throughput (the server'
 throughput during the test time; the higher the
 better). There are other HTTP benchmarking tools
 available like httpperf [16], SpecWeb99 [20],
 SURGE [6], etc. But WebStone's ability to control
 the number of clients and URL array and its ex-
 cellent result reporting tools was appropriate and
 the modified version was adequate for our exper-
 iment.

696 We created over 4,000 files in the web server and
 697 divided them into four queues as stated in the
 698 previous section. The average processing time of
 699 URL's in each queue under normal server load is
 700 0.5, 20, 10 and 5 s respectively. In the test, the
 701 WebStone clients were instructed to replay the
 702 trace of the online retailer's server collected during
 703 the period 20:00–21:00.

704 5.2. Experimental results

705 From the test, we found that some of the per-
 706 formance parameters such as server connection
 707 rate, number of completed requests and server
 708 throughput (in terms of the number of requests
 709 served) were more or less the same, while the av-
 710 erage response time varied significantly. For better
 711 comparison, we have chosen the response time as
 712 the primary performance measure for our analysis.
 713 We also analyzed parameters such as request
 714 timeout rate due to DWFS, the queueing time, and
 715 the processing time of each request from the server
 716 log at the web server side. We observed that the
 717 web server was overloaded when the number of
 718 clients reached 40 and essentially became saturated
 719 after being requested by 60 clients, so the following
 720 results and discussions are focused on the results
 721 with the number of clients as 40, 50, and 60.

722 The experiments were conducted under different
 723 DWFS session timeout settings and the results
 724 were compared to the original Apache server with
 725 the same configurations. In the DWFS tests, the
 726 clients were evenly divided into three priority
 727 classes. Each priority class has a scheduling time-
 728 out, the shorter timeout a request is assigned, the
 729 higher is its priority and thus it gets quicker ser-
 730 vice. In WebStone, the webmaster process always
 731 tries to evenly divide the number of total processes
 732 on the client workstations, making the number of
 733 processes in each of the priority class approxi-
 734 mately equal.

735 Tables 3–5 show the DWFS results with differ-
 736 ent timeout values, and Table 6 shows the results
 737 from the unmodified Apache server with the same
 738 configuration. Fig. 9 depicts the response time
 739 comparison of the four configurations. In Tables
 740 3–5, the values in *Number of Clients* is represented
 741 as an fraction because the number of processes in

Table 3
DWFS with session timeout period of 10 s

Number of clients	Average response time (s)	Request timeout rate (%)
40/3	6.902	0.16
50/3	7.152	0.28
60/3	6.968	0.44

Table 4
DWFS with session timeout period of 15 s

Number of clients	Average response time (s)	Request timeout rate (%)
40/3	8.537	0.11
50/3	8.357	0.20
60/3	8.772	0.35

Table 5
DWFS with session timeout period of 20 s

Number of clients	Average response time (s)	Request timeout rate (%)
40/3	9.512	0.01
50/3	9.864	0.16
60/3	10.584	0.24

Table 6
Response time behavior of the Apache server

Number of clients	Average response time (s)
40	9.623
50	11.885
60	14.267

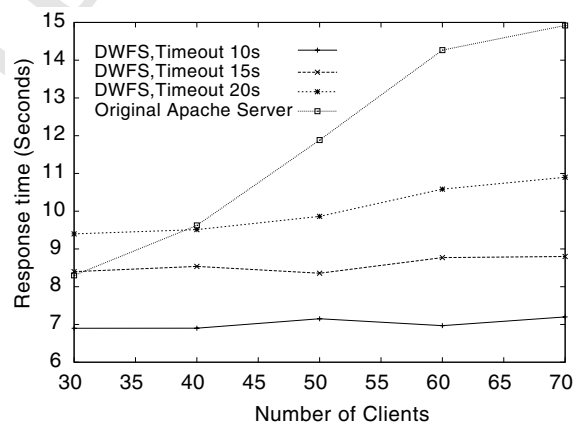


Fig. 9. Response time comparison.

742 each priority class varies in each test iteration but
 743 is approximately one-third of the total number of
 744 clients. Since the original Apache server does not
 745 drop requests because of timeout, only the average
 746 response time is listed in Table 6.

747 It is observed from the tables that, DWFS can
 748 significantly improve server performance by re-
 749 ducing the response time up to 52%. The smaller

750 the session timeout, the shorter is the response
 751 time and thus better is the service. DWFS incurs
 752 additional request timeout. The timeout rate rises
 753 significantly with respect to the number of clients
 754 and the timeout value. But such occurrence is less
 755 than 0.5%, the impact of which is insignificant on
 756 the service availability provided by an overloaded
 757 server. This observation also distinguishes DWFS
 758 from the shortest-job-first scheduling in the sense
 759 that the shortest-job-first scheduling starves long
 760 jobs which may lead to more request timeouts.
 761 Finally, the relatively steep slope of the original
 762 Apache server curve compared to DWFS curves in
 763 Fig. 9 reveals that the server performance using
 764 DWFS is more scalable when the number of the
 765 clients increases.

766 To investigate the underlying factors of the re-
 767 sponse time difference, we further analyzed the
 768 anatomy of the response time. We divided the re-
 769 sponse time into queueing time and processing
 770 time. Queueing time is the period during which a
 771 request remains queued and processing time is the
 772 interval between when the request was read and
 773 the time when the HTTP response sent. In Apache
 774 architecture, the queueing time starts when a job is
 775 accepted and ends at the point when the request
 776 starts getting processed by a worker thread; its
 777 processing time starts at that point. For a DWFS
 778 enabled Apache server, however, since every re-
 779 quest line is read immediately after it has been
 780 accepted, the queueing time spans from the time
 781 the request is read to the moment when a worker
 782 thread begins to process it. Fig. 10 presents the
 783 anatomy of the response time under different

784 configurations. The x -axis label specifies server
 785 type/timeout value/number of clients. The timeout
 786 value is zero for the original Apache server be-
 787 cause it has no scheduling restraints.

788 From Fig. 10, it is observed that the queueing
 789 time of the original Apache server is dependent on
 790 the number of the clients. The more the number of
 791 clients, the longer is the queueing time. We believe
 792 that this effect is a direct result of its best effort
 793 scheduling discipline, where requests are queued
 794 on a first-come-first-serve manner and short re-
 795 quests have to wait for the completion of long
 796 requests even if these long requests' turnaround
 797 time may exceed the delay bound and be aborted
 798 by the impatient clients. For the DWFS's case, the
 799 queueing time also varies for different timeout
 800 settings but is much shorter than those of the
 801 original Apache server. Under the same timeout
 802 setting (thus the same priority class), the queueing
 803 time remains stable, which in turn means that the
 804 session level serviceability is maintained. This is
 805 due to the pre-emptive scheduling and the early
 806 dropping of those requests whose pre-assigned
 807 service time cannot be guaranteed for the descen-
 808 dant requests within the same sessions.

809 In Section 4.2, we claimed that by varying the
 810 timeout value, the number of requests in each
 811 queue will change. For bigger timeout value, more
 812 requests reside in the queues, as a result of which
 813 the mean queueing time increases. This was found
 814 to be true in the experiment. Figs. 11 and 12 plot
 815 the queue length under timeout value 10 and 20 s
 816 with 40 clients. Table 7 lists the average request
 817 latency at each queue using different timeout val-
 818 ues. It is observed that queue lengths under time-
 819 out 20 s are bigger than those in 10 s and queue
 820 latencies under timeout 20 s are correspondingly
 821 longer. These results verify that varying timeout
 822 value can provide service differentiation.

823 Finally, we compared the session abortion rate
 824 under different configurations. We assume that a
 825 session is aborted if one of the requests within it
 826 does not receive service before its delay bound. In
 827 the real world, when this occurs, the end users get
 828 impatient and abort the session. In DWFS, a ses-
 829 sion gets aborted when the request has not been
 830 processed before its timeout value. While there is
 831 no timeout constraint in Apache, we assume that

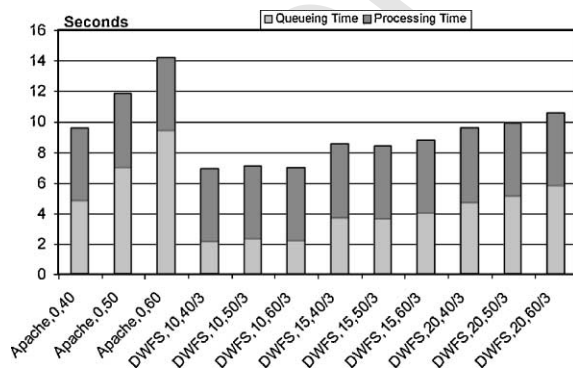


Fig. 10. Anatomy of the average response time.

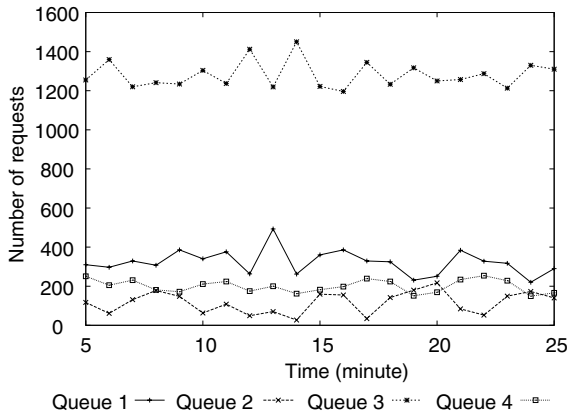


Fig. 11. Queue length at timeout 10 s.




Fig. 12. Queue length at timeout 20 s.

Table 7
Average latency

Queue	Timeout 10 s	Timeout 20 s
1	1.6	2.3
2	27.5	27.9
3	14.5	20.9
4	6.0	6.6

832 every session has a timeout value of 15 s. Fig. 13
 833 depicts the comparison. It is observed that for
 834 DWFS, the aborted session rate is relatively small
 835 and the priority class with longer timeout value
 836 usually has lower abortion rate. For Apache ser-
 837 ver, the aborted session rate is sensitive to the
 838 number of clients which means more and more
 839 sessions get aborted as the server load increases.




Fig. 13. Session abortion rate comparison.

6. Related work

840

A limited number of work has been reported on
 sessions characterization in web servers. In [4], the
 authors provide parameters based on the World
 Cup 98 server log, which include session length,
 inter-session time, and their implication on server
 performance. Our characterization work provide
 complementary results on workload composition,
 session stage transition, and the ratio of request
 arrival rate to session generation rate. These em-
 pirical results can be exploited to recognize user
 browsing behavior and capacity planning.

841
842
843
844
845
846
847
848
849
850
851

In the context of capacity planning, [13] pro-
 vides a model based on bandwidth demands for
 memory, processors data bus, NIC and I/O buses.
 It is practical for server configuration. Our ca-
 pacity planning model is targeted towards session
 level SLA specification and overload prevention.

852
853
854
855
856
857

Although a plethora of work on web servers
 have addressed performance issues in web servers,
 the studies on overload control has been limited.
 An approach for overload control by content ad-
 aptation has been proposed in [1]. Under high load
 the servers resorts to low fidelity images that
 consume less system resources, thus reducing the
 load. Content adaptation is applicable mainly to
 static web content. Overload control using oper-
 ating system support has been studied in [5,15,21].
 The server behavior under overload has been an-
 alyzed in [15] and three solutions are proposed to
 help relieve an overloaded server. These solutions

858
859
860
861
862
863
864
865
866
867
868
869
870

871 include direct control over kernel timeouts and
 872 resource limits, resource introspection, and disas-
 873 ter management. Three kernel-based mechanism
 874 that prevent server from being overloaded by ad-
 875 mission control and service differentiation are
 876 presented in [21]. Their mechanisms include TCP
 877 SYN policing that control the TCP connection
 878 rate, prioritized listen queue and HTTP header-
 879 based connection control that provides service
 880 differentiation. A new kernel facility called re-
 881 source container which can effectively audit overall
 882 resource usage by each process is presented in [5].
 883 This scheme is useful for service differentiation as
 884 well as overload control. In [19], the authors have
 885 studied web server overload control through three
 886 different schemes. The first approach is based on
 887 the network interface level request dropping. The
 888 second approach refers to a feedback mechanism
 889 from the application level to throttle the traffic
 890 volume. The third approach is a hybrid of the
 891 other two schemes. These schemes significantly
 892 improve server throughput under high load. All
 893 these solutions have not considered the session
 894 integrity and hence have limited applications for
 895 session-based web traffic.

896 Most of the prior work on overload control
 897 having examined performance on per request ba-
 898 sis, which may not be adequate for many appli-
 899 cations that require session-based overload
 900 control. A session-based admission control scheme
 901 has been reported in [9], which prevents overload
 902 by efficient admission control. They monitor the
 903 server load periodically and estimate the load in
 904 near future. If the predicted load is higher than a
 905 predefined threshold, no new requests are admit-
 906 ted. This situation may lead to denial of services.
 907 The proposed DWFS scheme is targeted for effi-
 908 cient scheduling of requests and complements the
 909 work reported in [9] in maintaining long term
 910 server availability.

911 7. Conclusion

912 Overload control ensures service availability in
 913 varying workload and is an indispensable part of
 914 network server engineering. This paper presents
 915 QoS capacity planning and scheduling algorithm

916 for overload control based on characterization of a
 917 commercial web server log. The main idea of the
 918 proposed scheme is to use session-based overload
 919 control. Performance measures of web services in
 920 terms of sessions is more meaningful than the
 921 measures based on individual requests. We have
 922 targeted QoS-aware web servers that provide
 923 guaranteed QoS based on the requirement of ses-
 924 sions. The traffic conformation function provides
 925 quantitative solution for SLA specification and
 926 can be used in commercial servers. We have pro-
 927 posed and evaluated a new scheduling algorithm
 928 called DWFS, which discriminates the scheduling
 929 of requests on the basis of the probability of
 930 completion of the session that the requests belong
 931 to. The proposed scheduling algorithm improves
 932 server productivity under heavy load by more than
 933 50% in the configuration studied in this paper.
 934 This work can be used as a framework for further
 935 development and deployment of session-based
 936 overload control techniques.

References

- 937
- [1] T. Abdelzaher, N. Bhatti, Web content adaptation to
improve server overload behavior, in: Proceedings of the
8th WWW Conference, Toronto, Canada, May 1999. 938
 - [2] Akamai Corporation, <http://www.akamai.com>. 939
 - [3] Apache Group, <http://www.apache.org>. 940
 - [4] M. Arlitt, Characterizing web user sessions, in: Proceedings
of the Performance and Architecture of Web Servers
Workshop, Santa Clara, CA, June 2000. 941
 - [5] G. Banga, P. Druschel, J. Mogul, Resource containers: a
new facility for resource management in server systems, in:
Proceedings of OSID, February 1999. 942
 - [6] P. Barford, M. Crovella, Generating representative web
workloads for network and server performance evaluation,
in: Proceedings of the ACM Sigmetrics Conference on
Measurement and Modeling of Computer Systems, Mad-
ison, WI, June, 1998. 943
 - [7] J. Borges, M. Levene, Data mining of user navigation
patterns, in: Proceedings of the Workshop on Web Usage
Analysis and User Profiling, August 1999, San Diego, CA,
USA. 944
 - [8] X. Chen, P. Mohapatra, H. Chen, An admission control
scheme for predictable server response time for web
accesses, in: Proceedings of the 10th WWW Conference,
Hong Kong, May 2001. 945
 - [9] L. Cherkasova, P. Phaal, Session-based admission control-
a mechanism for improving performance of commercial
web servers, in: Proceedings of IEEE/IFIP IWQoS 99. 946

- 965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
- [10] M.E. Crovelli, A. Bestavros, Self-similarity in world wide web traffic: evidence and possible causes, *IEEE/ACM Transactions on Networking* (December) (1997).
- [11] Internet Engineering Task Force, <http://www.ietf.org>.
- [12] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee, RFC 2616: Hypertext Transfer Protocol-HTTP/1.1, <http://www.ics.uci.edu/pub/ietf/http/rfc-2616.txt>.
- [13] K. Kant, Y. Won, Server capacity planning for web traffic workload, *IEEE Transactions on Knowledge and Data Engineering* 11 (5) (1999).
- [14] D.M. Kristol, L. Montulli, Internet draft: HTTP state management mechanism, <http://www.ics.uci.edu/pub/ietf/http/draft-ietf-http-state-man-mec-12.txt>.
- [15] J.C. Mogul, Operating system support for busy internet services, in: *Proceedings of the 5th Workshop on Hot Topics in Operating Systems*, Orcas Island, WA USA, May, 1995.
- [16] D. Mosberger, T. Jin, httpperf—A tool for measuring web server performance, in: *Workshop on Internet Server Performance*, Madison, WI, USA, June, 1998.
- [17] J. Pitkow, P. Pirolli, Mining longest repeating subsequences to predict world wide web surfing, in: *USENIX Symposium on Internet Technologies and Systems*, Boulder, CO, 1999.
- [18] Logging properties reference, <http://www.windows.com/windows2000/en/server/iis/html/core/iiintlg.html>.
- [19] V. Tewari, R. Iyer, K. Kant, Overload control mechanisms for web servers, in: *Performance and QoS of Next Generation Network*, Nagoya, Japan, November, 2000.
- [20] The Standard Performance Evaluation Corporation, <http://www.spec.org>.
- [21] T. Voigt, R. Tewari, D. Freimuth, A. Mehra, Kernel mechanisms for service differentiation in overloaded web servers, in: *Proceedings of the 2001 USENIX Annual Technical Conference*, Boston, MA, USA, June, 2001.

- [22] WebStone, <http://www.mindcraft.com/webstone>.
- [23] H. Zhu, H. Tang, T. Yang, Demand-driven service differentiation in cluster-based network Servers, in: *Proceedings of the IEEE INFOCOM 2001*, Anchorage, April, 2001.

1004
1005
1006
1007
1008



Huamin Chen received the B.S. and M.S. degrees from Huazhong University of Science and Technology, Wuhan, China in 1996 and 1999, respectively. He is working towards the Ph.D. degree in the department of Computer Science, University of California at Davis. His research interests are in Internet server performance improvement and QoS provisioning, Web Services, computer networks, and operating system optimization.



Professor Prasant Mohapatra is currently an Associate Professor in the Department of Computer Science at the University of California, Davis. In the past, he was on the faculty at Iowa State University and Michigan State University. He has also held Visiting Scientist positions at Intel Corporation and Panasonic Technologies. Dr. Mohapatra received his Ph.D. in Computer Engineering from the Pennsylvania State University in 1993. Dr. Mohapatra is currently on the editorial board of the *IEEE Transactions on Computers* and has been on the program/organizational committees of several international conferences. He was the Program Chair for the PAWS Workshop during 2000 and 2001, and the Program Vice-Chair for the International Conference on Parallel Processing, 2001. He is also the Co-Editor of the January 2003 issue of *IEEE Network*.

Dr. Mohapatra's research interests are in the areas of wireless mobile networks, Internet protocols and QoS, and Internet servers. Dr. Mohapatra's research has been funded through grants from the National Science Foundation, Intel Corporation, Panasonic Technologies, Hewlett Packard, and EMC Corporation.