Supplementary Material

Section A: Waiting Time Approximation for Multiple Services

For the single service scenario, the waiting times are derived according to equation 9. Since such expression is analytically intractable for the multiple-service setting, we use a proportionate value of the number of resources used by each service. That is, if n is the total number of resources, the resource requirement for service type j is given by $n_j = \frac{\lambda_j / \mu_j}{\sum_j \lambda_j / \mu_j}$. This proportional value represents the share of the total number of resources dedicated to a particular service. We use this value of n_j in the waiting time expression for each service separately using the expression derived for the single service setting.

Section B: Setting Parameters for the Simulation Experiments

<u>Delay costs</u>

We have used a truncated normal distribution (so that only users with positive delay-cost/time are generated) to represent the various user populations. A normal distribution is used in accordance with previous literature (Konana, et al. 2000), also since users gain nothing from waiting for their requests to be served, it can have only positive values. We assume the medium-delay cost population in the organization is closely aligned with the prices used in the fixed-price scheme, while the high (low)-delay cost population have a higher (lower) valuation for the services. For example, for Service A with an estimate of 12.0/hr arrivals, this scheme is equivalent to charging $20 for each request. The medium-delay cost user's delay cost will be slightly lower than this value (On average, the users in this population *value* the service at $20/hr, the valuation for a service must be greater than their delay costs in order for them to enter the system and gain a positive utility (Konana, et al. 2000)). We assume this average delay cost to be 75% of the valuation, and this is assigned to the user profile with medium delay costs. The low-(high) delay cost users are assumed to have delay cost of 50% (125%) of the valuation ($20 in this case). The standard deviation in each population is kept at 25% of the mean. This way, we have three distinct user populations, with minimal amount of overlap between them.

The multiple service setting employs the same technique for each service separately. We also have tested other combinations of delay cost distributions and find the results to be robust to a range of values of delay cost. If delay cost distributions are very close to each other, the requests converge to a single priority queue – substantive differences in delay costs are necessary for the results to hold. For the results reported in this study, the parameters for the user delay-cost distributions for the three services are as follows:

Customer Profile		Low Delay-cost	Medium Delay-cost	High Delay-cost
Delay-Cost/Hr	Service A	10,2.5	15,3.75	25,6.25
	Service B	20,5	30,7.5	50,13
	Service C	80,20	120,30	200,50

Penalties

Expression 8 provides the basis of translating the fixed-prices into a dynamic price-penalty scheme. In our simulations, we used three priority classes. The penalty thus obtained, is assigned to the penalty of priority class 2, the implication being that users coming in this class obtain the same service level as in the fixed-price scheme. Since each higher priority class imposes a penalty-cost on all lower priority class requests, the difference among the penalties in these classes should not be the same, rather the difference between penalties in priority class 1 and 2 should be greater than the difference between priority classes 2 and 3. Our initial runs do confirm this conjecture. Based on this rationale and initial findings, we make the penalties in priority class 1 (3) 50% more (less) than that of priority 2.

In the multiple service scenario note that the penalties in each priority class is the same for this particular setting. This is so because the ratio of the estimated arrival rate to service rate is same (=2) for each service. If another setting was used, the penalty values would change depending on the relative value of this ratio.

Section C: Arrival Rates

Single Service Setting

Arrival Pattern: *Stationary*

Customer Profile	Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	3	6	3

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 75%

Delay Cost Distribution		Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	Month 1	3	6	3
	Month 2	3.439705564	6.879411129	3.439705564
	Month 3	2.641586385	5.28317277	2.641586385
	Month 4	3.424882055	6.84976411	3.424882055
	Month 5	3.555924507	7.111849014	3.555924507
	Month 6	2.30767455	4.6153491	2.30767455

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 80%

Customer Profile		Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	Month 1	3.083680648	6.167361296	3.083680648
	Month 2	3.59383948	7.187678959	3.59383948
	Month 3	3.587723706	7.175447413	3.587723706
	Month 4	2.924714407	5.849428815	2.924714407
	Month 5	3.01038906	6.020778119	3.01038906
	Month 6	2.921147896	5.842295792	2.921147896

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 85%

Customer Profile		Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	Month 1	3.63151	7.263019	3.63151
	Month 2	3.710258	7.420517	3.710258
	Month 3	3.10018	6.20036	3.10018
	Month 4	3.121483	6.242966	3.121483
	Month 5	3.538387	7.076775	3.538387
	Month 6	3.556418	7.112837	3.556418

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 90%

Customer Profile		Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	Month 1	3.750723	7.501445	3.750723
	Month 2	3.249004	6.498009	3.249004
	Month 3	3.124427	6.248855	3.124427
	Month 4	3.654246	7.308492	3.654246
	Month 5	3.803782	7.607563	3.803782
	Month 6	3.593391	7.186783	3.593391

Multiple-Service Setting

Arrival Pattern: *Stationary*

Customer Profile		Low Delay-cost	Medium Delay-cost	High Delay-cost
Arrival Rate (/hr)	Service A	4	8	4
	Service B	2	4	2
	Service C	0.5	1	0.5

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 75%

Delay Cost Distribution		Low Delay-cost	Medium Delay-cost	High Delay-cost
Service A:				
Arrival Rate (/hr)	Month 1	3.705932	7.411865	3.705932
	Month 2	3.251463	6.502925	3.251463
	Month 3	3.667351	7.334702	3.667351
	Month 4	3.755288	7.510575	3.755288
	Month 5	3.294362	6.588725	3.294362
	Month 6	3.567109	7.134217	3.567109
Service B:				
Arrival Rate (/hr)	Month 1	1.97146	3.94292	1.97146
	Month 2	1.569626	3.139252	1.569626
	Month 3	1.985476	3.970952	1.985476
	Month 4	1.63337	3.26674	1.63337
	Month 5	1.938882	3.877763	1.938882
	Month 6	1.849061	3.698121	1.849061
Service C:				
Arrival Rate (/hr)	Month 1	0.407641	0.815283	0.407641
	Month 2	0.475846	0.951691	0.475846
	Month 3	0.434244	0.868488	0.434244
	Month 4	0.420983	0.841967	0.420983
	Month 5	0.407246	0.814491	0.407246
	Month 6	0.442365	0.884731	0.442365

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 80%

Delay Cost Distribution		Low Delay-cost	Medium Delay-cost	High Delay-cost
Service A:				
Arrival Rate (/hr)	Month 1	4	8	4
	Month 2	4.08777635	8.1755527	4.08777635
	Month 3	3.614761241	7.229522482	3.614761241
	Month 4	3.487500908	6.975001817	3.487500908
	Month 5	3.865451605	7.73090321	3.865451605
	Month 6	4.574339381	9.148678761	4.574339381
Service B:				
Arrival Rate (/hr)	Month 1	2	4	2
	Month 2	1.577478691	3.154957382	1.577478691
	Month 3	1.850125256	3.700250512	1.850125256
	Month 4	1.685801684	3.371603369	1.685801684
	Month 5	1.864005882	3.728011764	1.864005882
	Month 6	1.507430182	3.014860364	1.507430182
Service C:				
Arrival Rate (/hr)	Month 1	0.5	1	0.5
	Month 2	0.489027956	0.978055913	0.489027956
	Month 3	0.545615595	1.09123119	0.545615595
	Month 4	0.564824394	1.129648788	0.564824394
	Month 5	0.503612719	1.007225438	0.503612719
	Month 6	0.411254843	0.822509686	0.411254843

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 85%

Delay Cost Distribution		Low Delay-cost	Medium Delay-cost	High Delay-cost
Service A:				
Arrival Rate (/hr)	Month 1	4.063347781	8.126695561	4.063347781
	Month 2	4.898065527	9.796131055	4.898065527
	Month 3	3.781042505	7.562085011	3.781042505
	Month 4	4.142380044	8.284760088	4.142380044
	Month 5	4.149891968	8.299783936	4.149891968
	Month 6	4.581821433	9.163642867	4.581821433
Service B:				
Arrival Rate (/hr)	Month 1	1.663841427	3.327682853	1.663841427
	Month 2	1.504552633	3.009105266	1.504552633
	Month 3	1.571899962	3.143799924	1.571899962
	Month 4	1.751907494	3.503814988	1.751907494
	Month 5	2.02110911	4.04221822	2.02110911
	Month 6	2.001595764	4.003191528	2.001595764
Service C:				
Arrival Rate (/hr)	Month 1	0.507918473	1.015836945	0.507918473
	Month 2	0.612258191	1.224516382	0.612258191
	Month 3	0.472630313	0.945260626	0.472630313
	Month 4	0.517797506	1.035595011	0.517797506
	Month 5	0.518736496	1.037472992	0.518736496
	Month 6	0.572727679	1.145455358	0.572727679

Arrival Pattern: *Non-stationary*: Utilization Rate ≈ 90%

Delay Cost Distribution		Low Delay-cost	Medium Delay-cost	High Delay-cost
Service A:				
Arrival Rate (/hr)	Month 1	3.969956	7.939913	3.969956
	Month 2	3.724581	7.449162	3.724581
	Month 3	3.668444	7.336888	3.668444
	Month 4	4.22811	8.45622	4.22811
	Month 5	3.659025	7.318051	3.659025
	Month 6	4.68388	9.36776	4.68388
Service B:				
Arrival Rate (/hr)	Month 1	1.889619	3.779238	1.889619
	Month 2	2.083885	4.16777	2.083885
	Month 3	2.191929	4.383859	2.191929
	Month 4	2.09303	4.18606	2.09303
	Month 5	1.862091	3.724181	1.862091
	Month 6	1.844991	3.689982	1.844991
Service C:				
Arrival Rate (/hr)	Month 1	0.506646	1.013291	0.506646
	Month 2	0.583271	1.166542	0.583271
	Month 3	0.592428	1.184855	0.592428
	Month 4	0.574226	1.148452	0.574226
	Month 5	0.593307	1.186614	0.593307
	Month 6	0.598831	1.197661	0.598831

Section C: Experimental Set Up:

As described in section 4 we use a discrete event simulation package (Simprocess) with hierarchical process modeling capabilities in conjunction with a mathematical software (Maple) to conduct our experiments. Details of the simulation program as well as the price generator are provided below.

Simulation environment – Figure 4 shows a screen shot of the highest level of the service set up. There are three 'generate' processes (one for each service), one 'process calls' process, and a 'dispose calls' process.

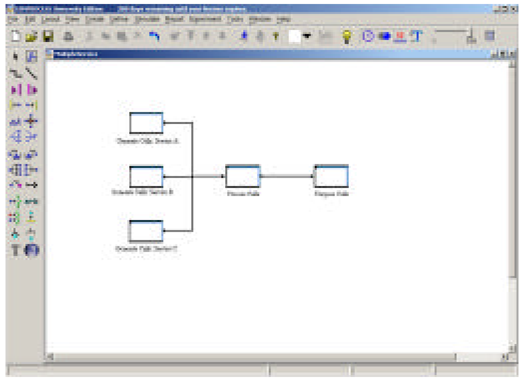

Figure 4: Main Simulation Model

Figure 5 shows the details of the 'generate calls' process for service A. Entities generated here are associated with two distributions. A random number from a Poisson distribution for the arrival rates, and a random number for the delay costs. Each generated entity is tagged so that individual delay costs, waiting time and welfare values could be captured programmatically.

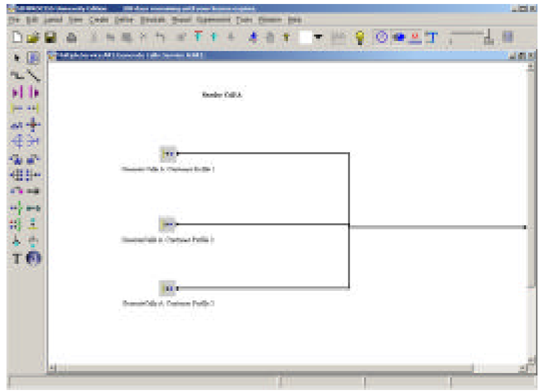

Figure 5. Details of the generate Calls (Service A) Process

Figure 6 shows the details of the process calls process. Here the branch activity implements code to separate service requests into priority classes. Since the entities are tagged, the waiting time spent for each individual service call is recorded at the 'process calls' activity. The waiting times are used in calculating individual welfare values for each user in the customer organization and the total welfare for the provider. The update welfare entity is set to be triggered every 7 days. This entity writes the past arrival rates into flat files and gets the price updates from the price generator module. Also, the sum of all welfare values for both the customer and provider are recorded and appended in output flat files. After writing, these values are refreshed every output period.

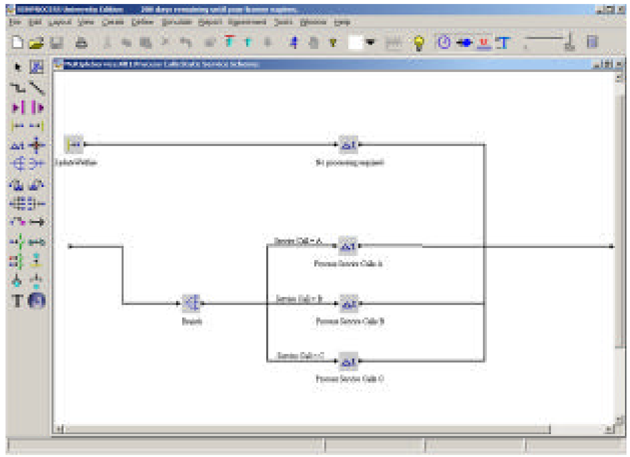

Figure 6: Details of 'Process Calls'

Price-Generator:

This is a mathematical program that runs synchronously with the simulation program to periodically update the prices. The simulation program outputs the number of resources and the arrival rates in the past interval into flat files. The Maple program reads these values and calculates the waiting times $\overline{\omega}_{jk}$, their partial derivatives $\overline{\omega'}_{jk}$, and the subsequent prices (S_{jk}). Details of the price generation process are given in Appendix A.