

Sphinx Performance Graphs

Generated on April 16, 2009. <http://notes.jschutz.net/wp-content/uploads/2009/04/sphinx-performance.pdf>

Explanation of Graphs

The graphs on the following pages have been produced by from the statistics of the Sphinx search response times as recorded in the Sphinx query.log.

Response Time Cumulative Distribution The count of all query log instances with response time less than a given value (normalised by the total number of instances).

Response Time Normalised Frequency Distribution The count of all query log instances with response time in a given bin, normalised by the total number of instances. Bin end values are 0.001s, 0.003s, 0.005s, 0.01s, 0.03s, 0.05s, 0.1s, 0.3s, 0.5s, 1s, 3s, 5s, 10s, 30s, i.e. the first bin response time range is 0 - 0.001s, second bin 0.001 - 0.003, etc.

Word Count Cumulative Distribution The count of all query log instances with query word count less than a given value (normalised by the total number of instances).

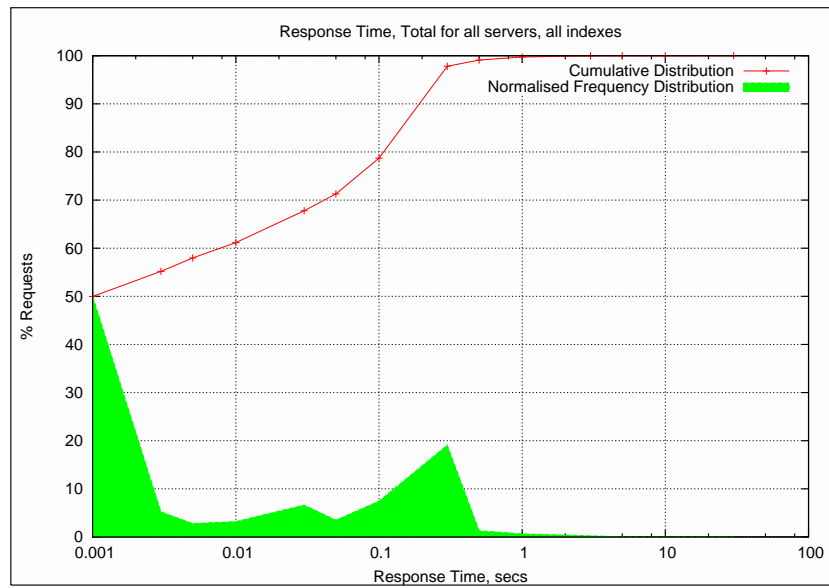
Word Count Normalised Frequency Distribution The count of all query log instances with word count in a given bin, normalised by the total number of instances. Bin end values are 0.001, 0.003, 0.005, 0.01, 0.03, 0.05, 0.1, 0.3, 0.5, 1, 3, 5, 10, 30 words.

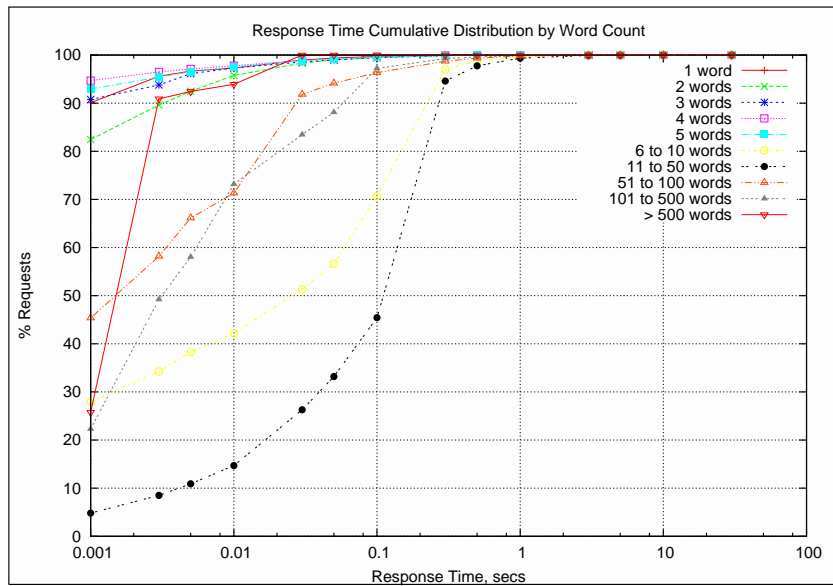
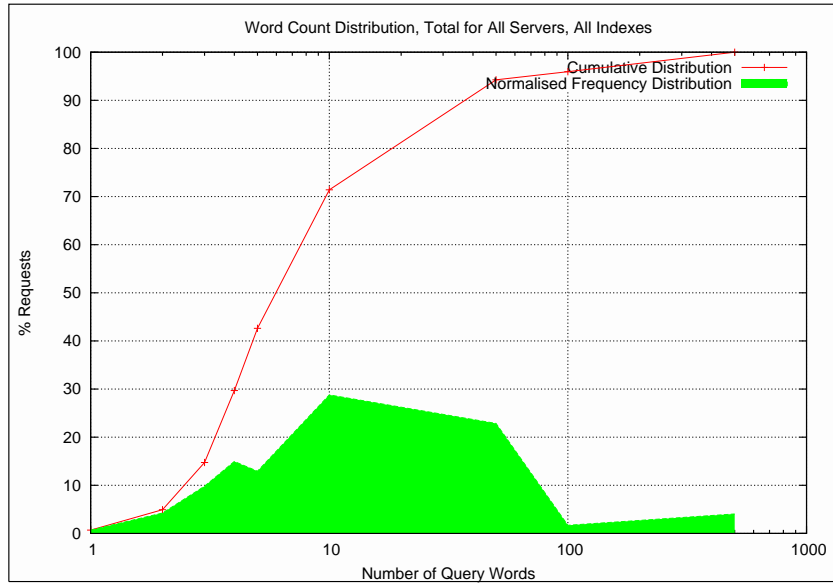
Summary Graphs

The following graphs show the sum of responses over all servers and all indexes

Dataset Sizes

This dataset contains 453495859 query log instances.





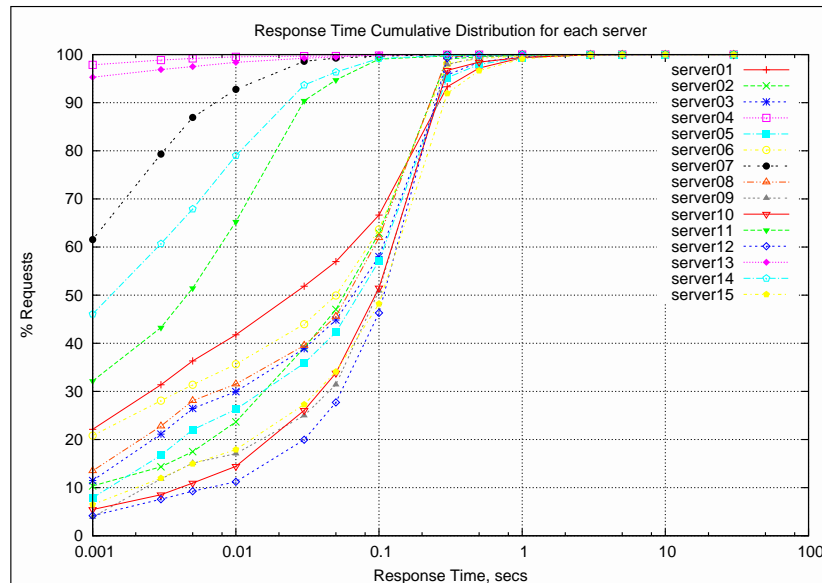
Per-server Graphs

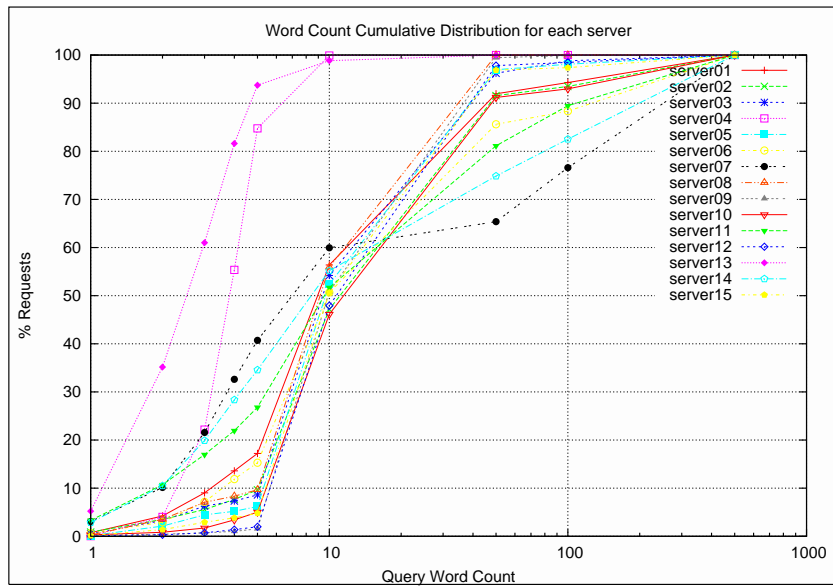
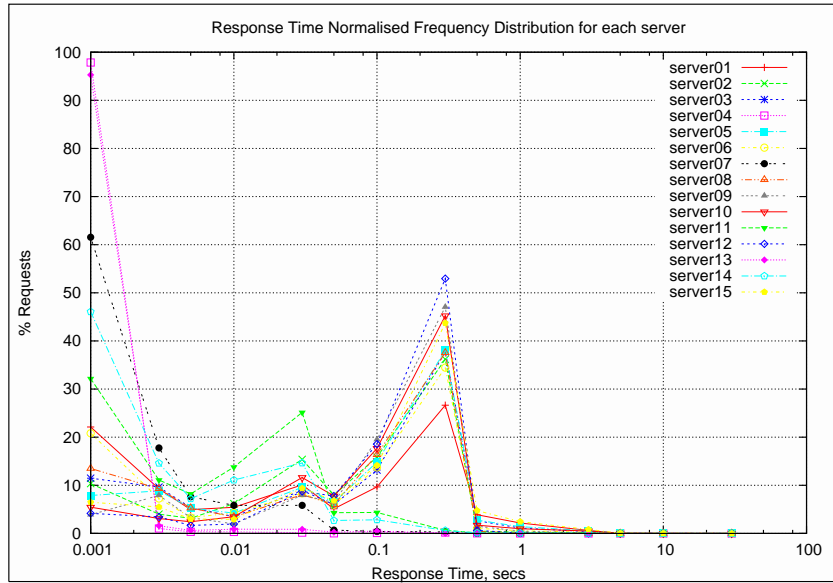
The following graphs show total responses for each server, summed over all indexes. Significant differences between servers may indicate a configuration problem.

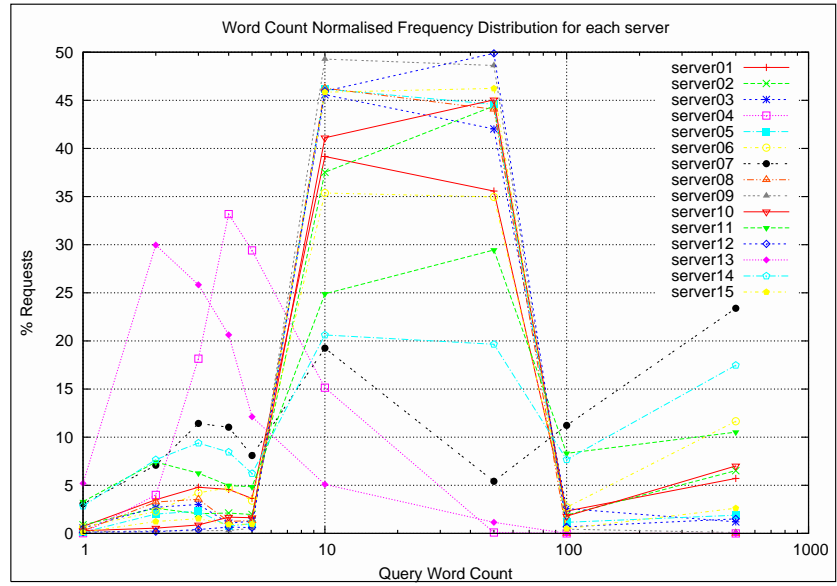
Dataset Sizes

The following table shows the number of query log instances used for each dataset.

Dataset	Size
server01	31012218
server02	16046879
server03	28153937
server04	169185140
server05	25425522
server06	37992442
server07	7218434
server08	2547959
server09	9929101
server10	29050868
server11	14580680
server12	4363270
server13	15656563
server14	22461881
server15	39870965







Per-index Graphs

The following graphs show total responses for each index across all servers.

Dataset Sizes

The following table shows the number of query log instances used for each dataset.

Dataset	Size
index001	9476
index002	9165706
index003	1992962
index004	3885037
index005	431
index006	1455107
index007	3562
index008	1711630
index009	174306
index010	1868598
index011	48
index012	6883
index013	201
index014	27203284
index015	579
index016	30
index017	34
index018	56768
index019	4552
index020	3877
index021	5536637
index022	223189
index023	359
index024	14719430
index025	11636
index026	236560
index027	259
index028	11447
index029	279325
index030	30
index031	45455
index032	168448
index033	43
index034	5946
index035	1589004
index036	11241
index037	5240652
index038	1748219
index039	886348
index040	16737476
index041	566
index042	83643

continued on next page

continued from previous page

Dataset	Size
index043	95028866
index044	825711
index045	123947
index046	7939364
index047	1676
index048	1810
index049	90206
index050	34818704
index051	3206640
index052	1371
index053	66494
index054	11
index055	259759
index056	1097
index057	10374
index058	643097
index059	2064979
index060	5687415
index061	921
index062	547415
index063	3052820
index064	227359
index065	42316
index066	98079
index067	227
index068	27999
index069	39
index070	150285
index071	310840
index072	302016
index073	44173
index074	45
index075	44977
index076	1008569
index077	330151
index078	88909
index079	501381
index080	10
index081	446947
index082	748
index083	1574235
index084	5506961
index085	68
index086	16965534
index087	132808
index088	2708598
index089	890987
index090	15
index091	185740
index092	1247226
index093	4126371
index094	1833909

continued on next page

continued from previous page

Dataset	Size
index095	7499358
index096	27
index097	60
index098	1616152
index099	604685
index100	1619028
index101	14882
index102	527
index103	851900
index104	539283
index105	106761
index106	2236365
index107	15
index108	1773844
index109	36
index110	15
index111	64117
index112	847447
index113	1049529
index114	219
index115	38687
index116	92731
index117	12
index118	968582
index119	11
index120	4688204
index121	24
index122	238
index123	14
index124	24239
index125	1917636
index126	3046176
index127	222
index128	30757
index129	1032
index130	122543
index131	53798
index132	2141504
index133	1106702
index134	4590445
index135	303881
index136	1280526
index137	294
index138	2432
index139	1106461
index140	28
index141	75
index142	271
index143	9165538
index144	81740
index145	142853
index146	18669272

continued on next page

continued from previous page

Dataset	Size
index147	1782542
index148	28
index149	184
index150	284817
index151	115939
index152	1709548
index153	1037
index154	41
index155	55893
index156	220148
index157	817854
index158	50448
index159	102108
index160	2214377
index161	5
index162	443743
index163	867103
index164	118
index165	50
index166	491992
index167	11
index168	145767
index169	1297820
index170	696948
index171	54817
index172	856627
index173	227
index174	2187969
index175	792
index176	1101279
index177	99805
index178	72771814
index179	7710430
index180	763212
index181	5

