JET PROPULSION LABORATORY INTEROFFICE MEMORANDUM

SFB: 363-89-004
1 August 1989

TO: Time Warp Folks
FROM: Steven Bellenot
RE: Jump Forward vs Lazy vs Aggressive Part 2

This is a continuation of my memo on Cancellation Policies (JPL IOM 363-89-002). Again mainly this memo is just graphs of execution time data of benchmarks under three policies:
1. Jump Forward = lazy cancellation with the jump forward optimization
2. Lazy = lazy cancellation without the jump forward optimization
3. Aggressive= aggressive cancellation
Again the TWOS's used was a special version near TWOS 2.1. However this time the configuration files were the TWOS 2.1 benchmark files. Generally the conclusions are the same. Jump forward is a loser and lazy beats aggressive. However pucks shows that aggressive can be twice as sloow as lazy.

STB88 -- 40 through 75 nodes

Figure 1 shows Jump Forward vs Lazy and Figure 2 shows Lazy vs Aggressive for STB88. There was enough variation in the execution times of STB88 that Figure 3 shows the precent difference between the fastest and slowest run times for each policy. (See the old memo for STB88 on smaller numbers of nodes.)

Warpnet

Figure 4 (Figure 6) shows Jump Forward vs Lazy for 8 through 44 nodes (40 through 75 nodes) for Warpnet. Figure 5 (Figure 7) shows Lazy vs Aggressive for 8 through 44 nodes (40 through 75 nodes) for Warpnet. Warpnet performance with Jump Forward was the best of the three benchmarks. Warpnet didn't gain with Jump Forward, it just didn't lose. Agressive cancellation is a loser by 1% to 11%.

Pucks

Figure 8 shows Jump Forward vs Lazy for 8 through 44 nodes on Pucks. Jump Forward is a constant 3% loser on Pucks in this range. The Pucks benchmark would not execute using aggressive cancellation. Actually Pucks was trapped in a infinite loop of cycles of messages for now.

A new version of pucks was created which took advantage of the new extended virtual time and had no messages for now. Figures 9 and 10 show the performance of this new pucks on 8 through 25 and 40 through 75 nodes. Aggressive is sometimes more than twice as slow as lazy. At this time it is not clear why aggressive cancellation is so much slower than lazy cancellation for Pucks, but we are working on it.

Figure 11 shows Jump Forward vs Lazy for 40 to 75 nodes using the new version of pucks. Jump Forward does a little better with this version, losing only 1% - 3% until the 6% and 10% lost at 70 and 75 nodes.

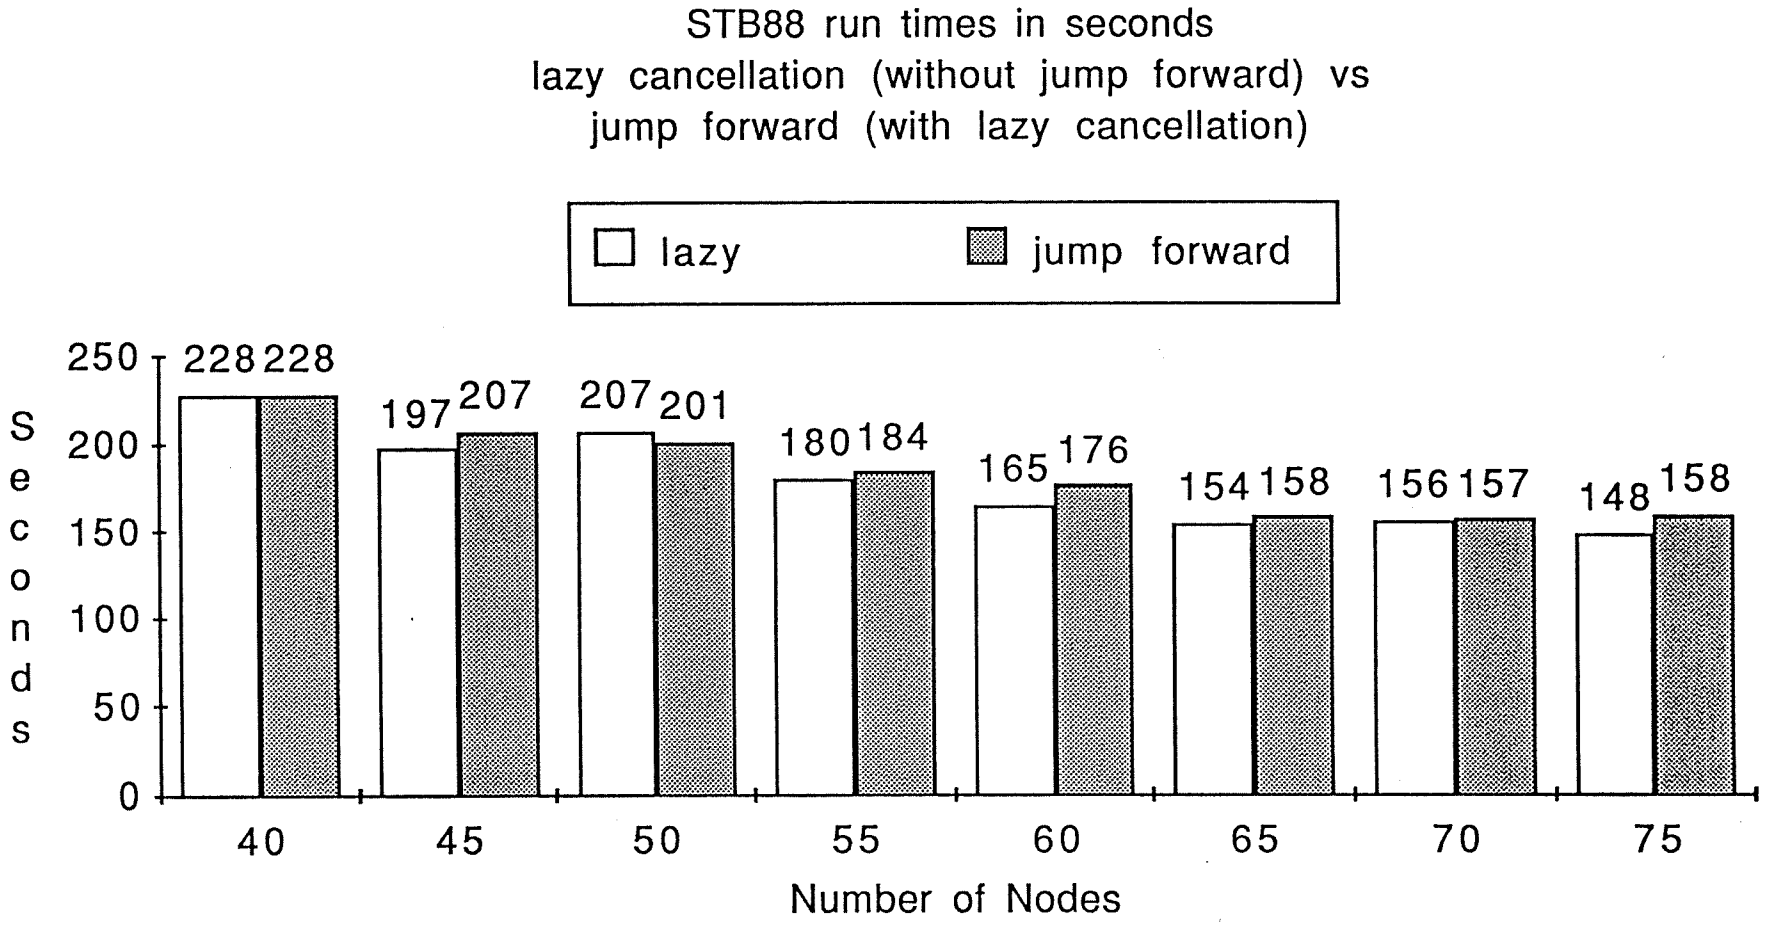
STB88 run times in seconds
lazy cancellation (without jump forward) vs
jump forward (with lazy cancellation)
lazy
jump forward
Seconds
250
200
150
100
50
0
228 228
197 207
207 201
180 184
165 176
154 158
156 157
148 158
40
45
50
55
60
65
70
75
Number of Nodes

Figure 1

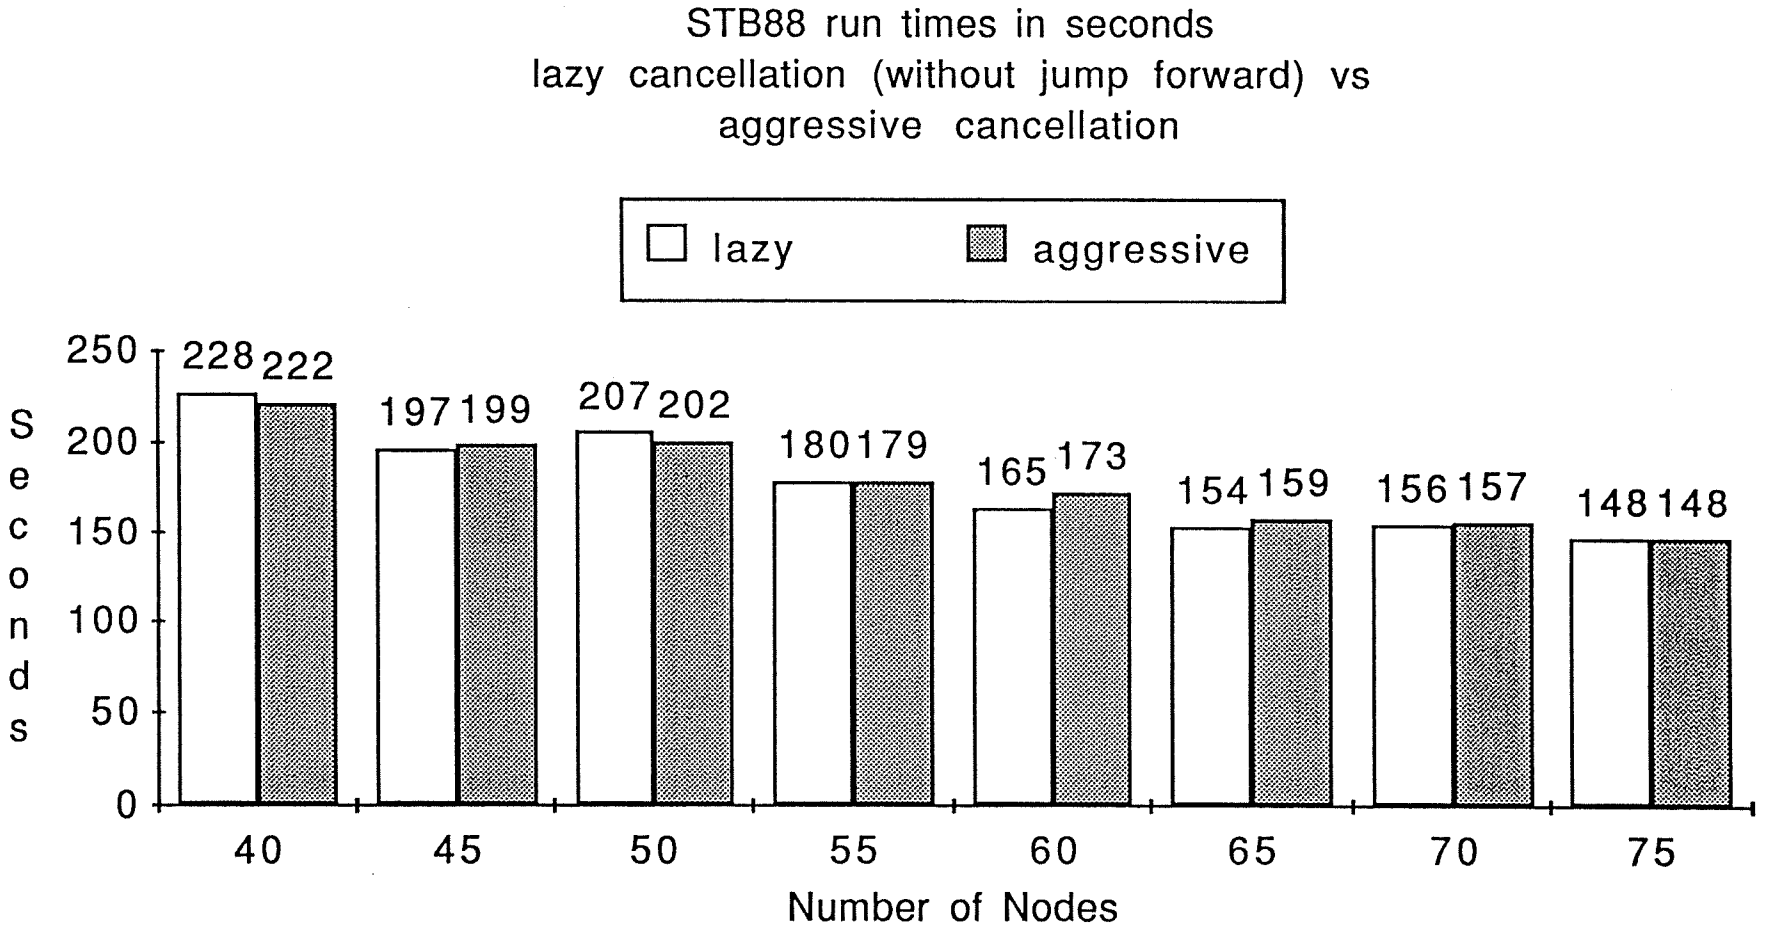
STB88 run times in seconds
lazy cancellation (without jump forward) vs
aggressive cancellation
lazy
aggressive
Seconds
250
200
150
100
50
0
228 222
197 199
207 202
180 179
165 173
154 159
156 157
148 148
40
45
50
55
60
65
70
75
Number of Nodes

Figure 2

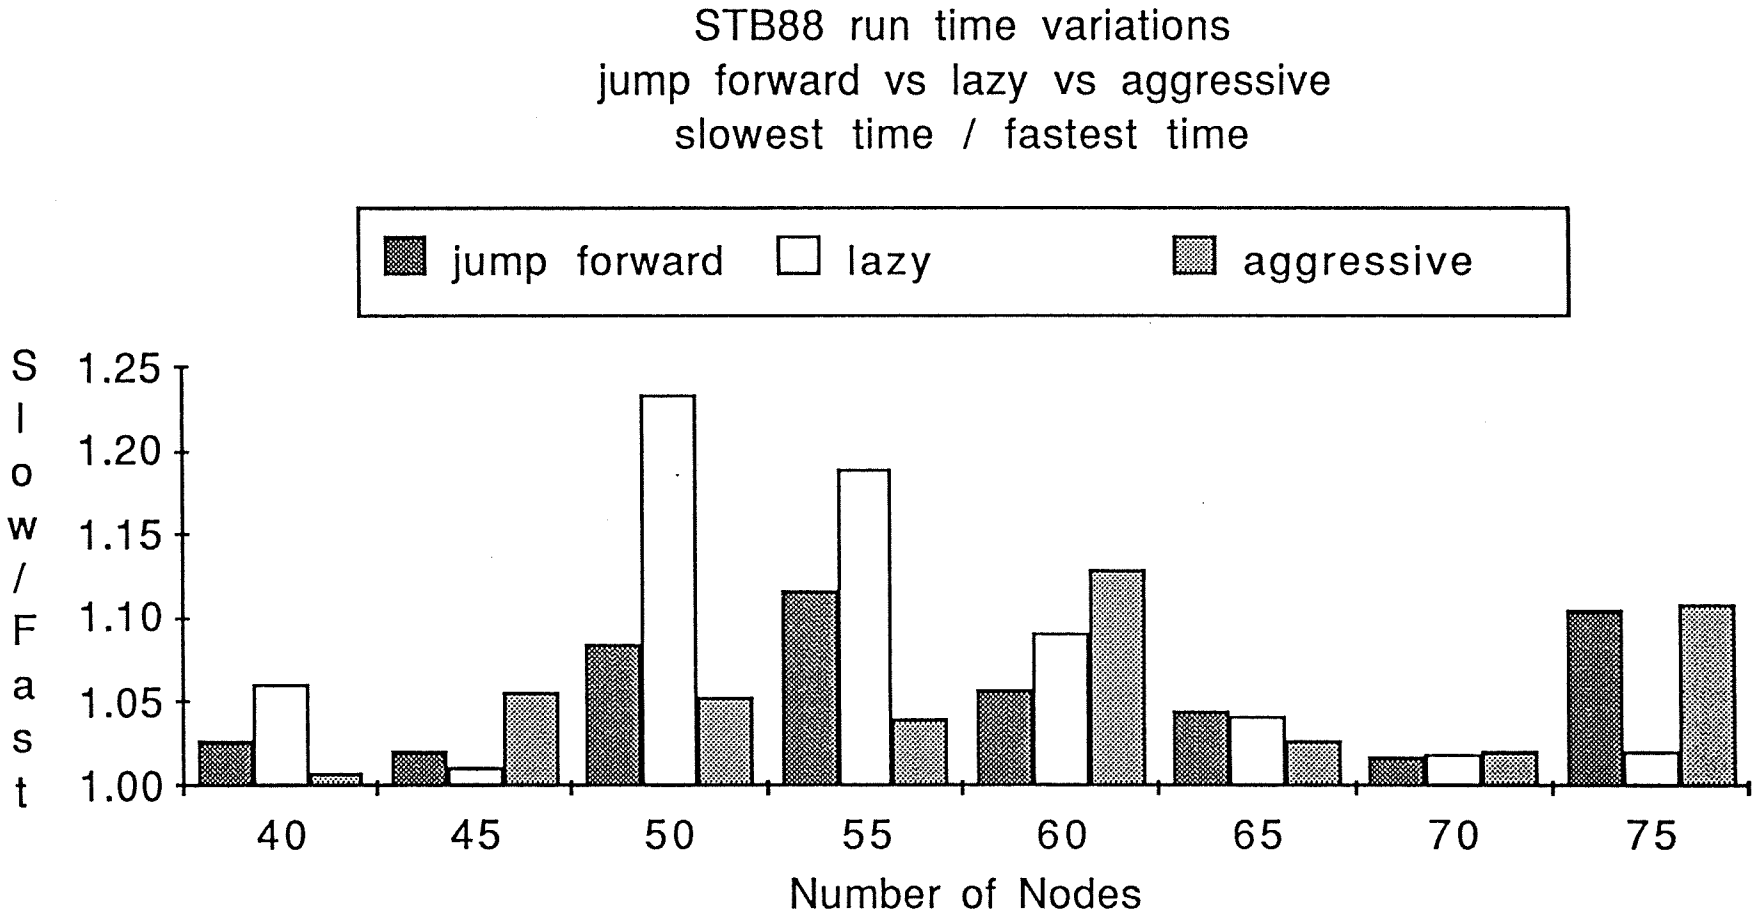
STB88 run time variations
jump forward vs lazy vs aggressive
slowest time / fastest time
jump forward
lazy
aggressive
Slow/Fast
1.25
1.20
1.15
1.10
1.05
1.00
40
45
50
55
60
65
70
75
Number of Nodes

Figure 3

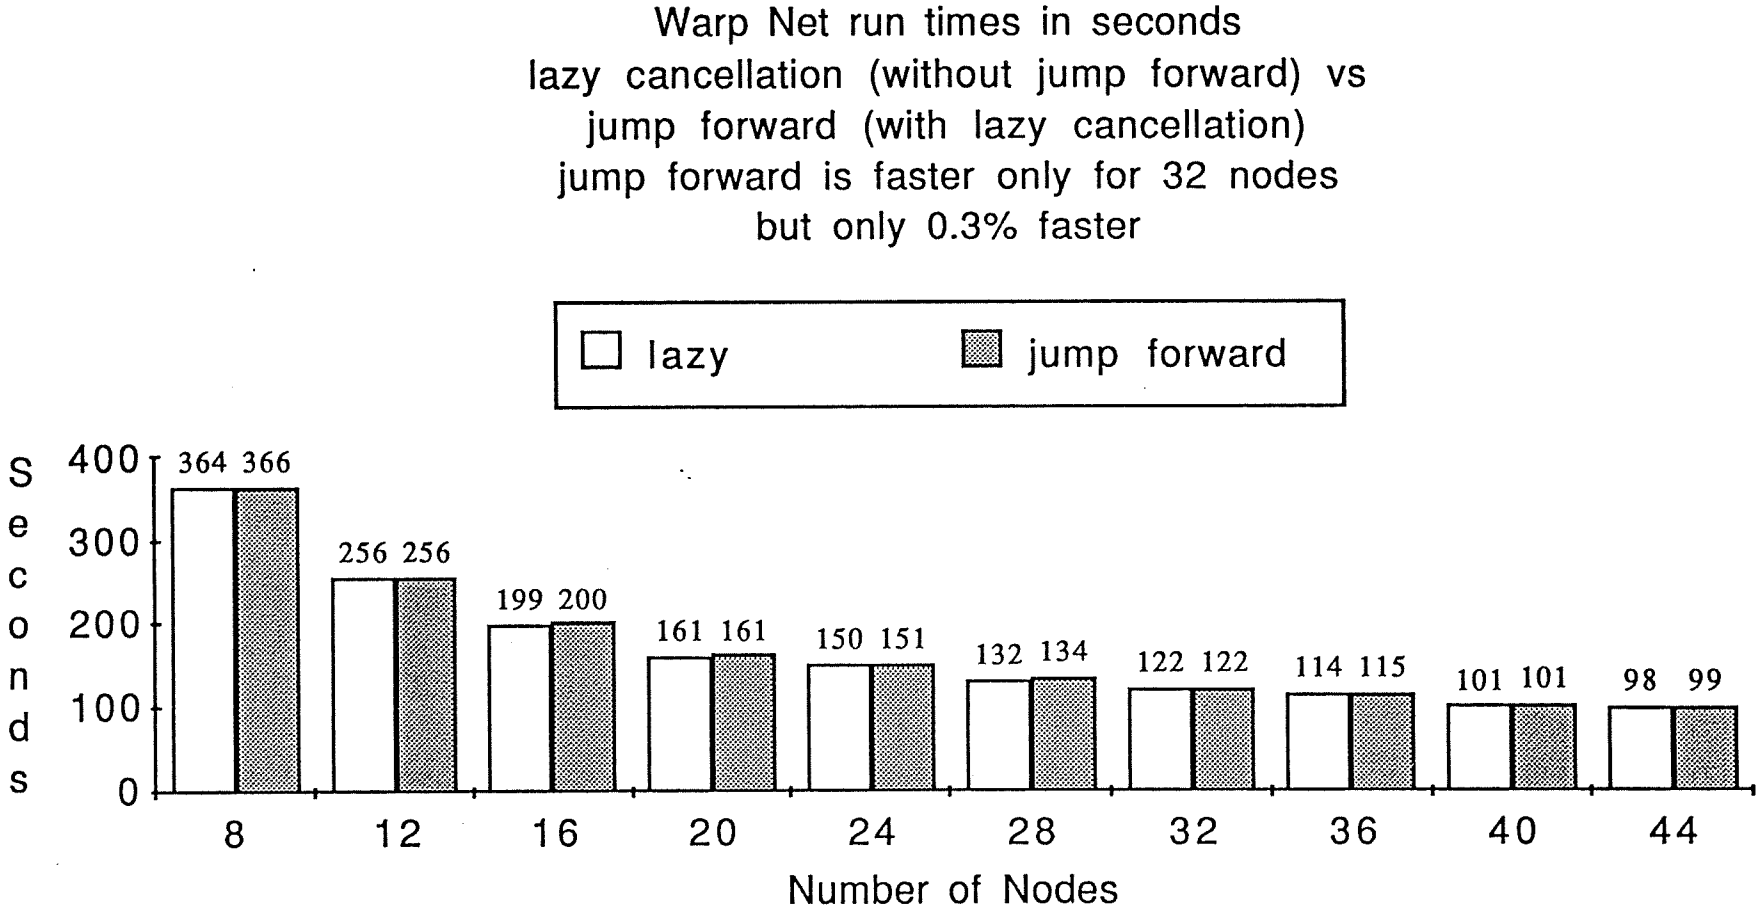
Warp Net run times in seconds
lazy cancellation (without jump forward) vs
jump forward (with lazy cancellation)
jump forward is faster only for 32 nodes
but only 0.3% faster
lazy
jump forward
Seconds
400
300
200
100
0
364 366
256 256
199 200
161 161
150 151
132 134
122 122
114 115
101 101
98 99
8
12
16
20
24
28
32
36
40
44
Number of Nodes

Figure 4

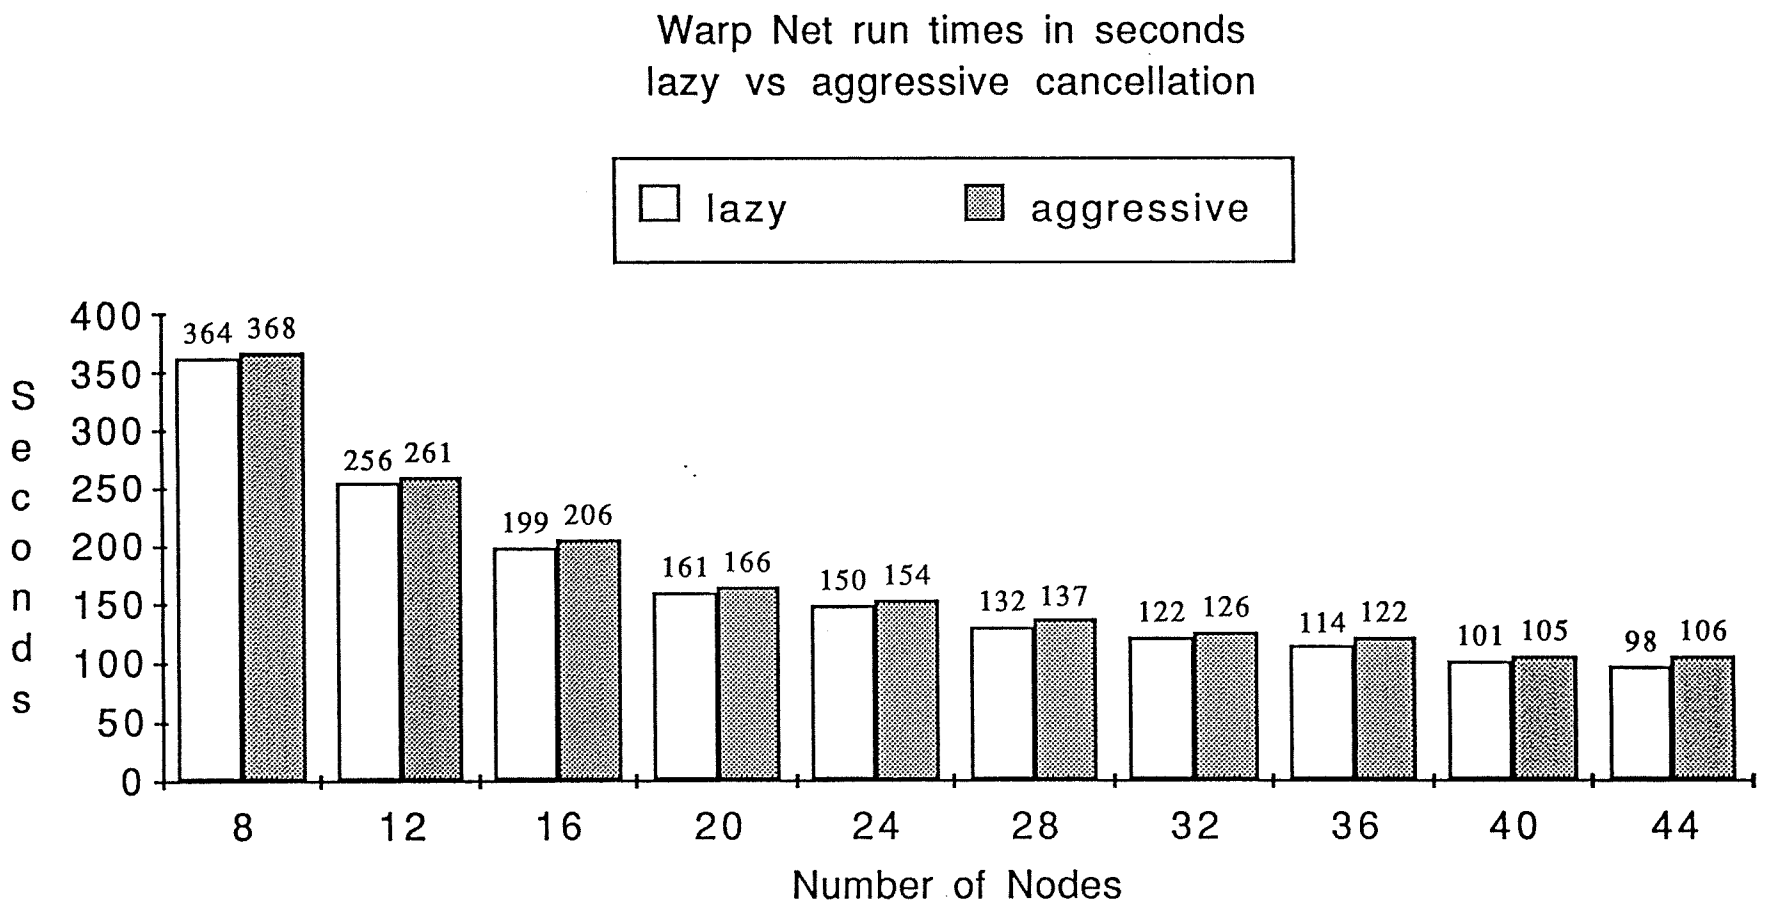
Warp Net run times in seconds
lazy vs aggressive cancellation
lazy
aggressive
Seconds
400
350
300
250
200
150
100
50
0
364 368
256 261
199 206
161 166
150 154
132 137
122 126
114 122
101 105
98 106
8
12
16
20
24
28
32
36
40
44
Number of Nodes

Figure 5

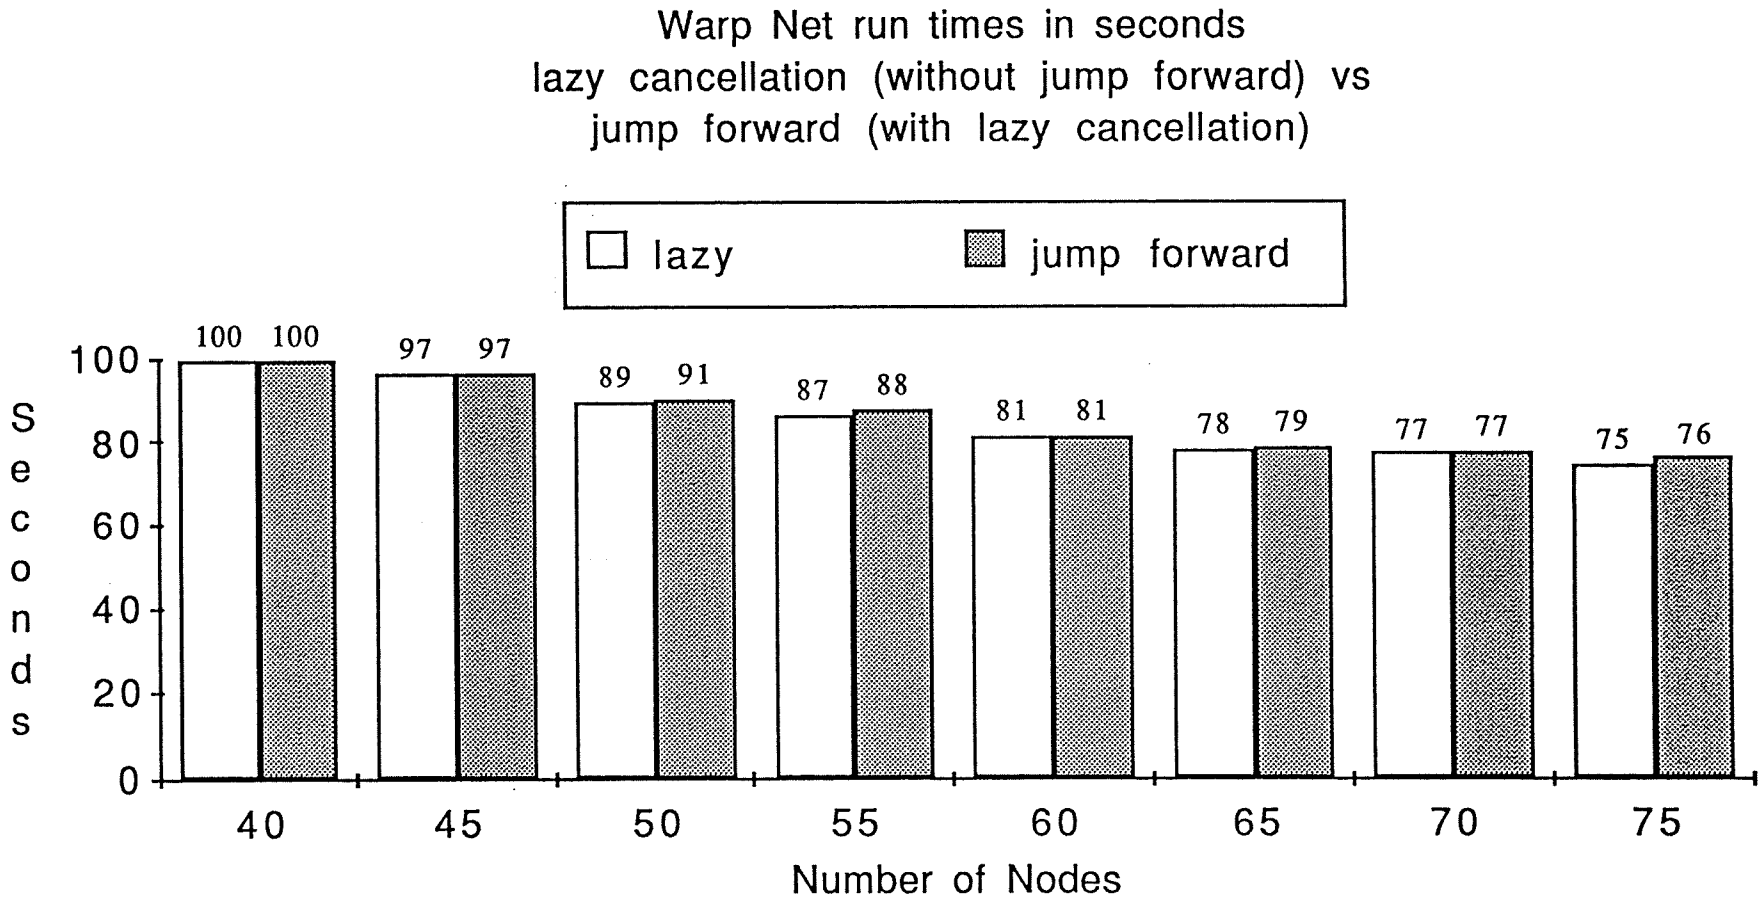
Warp Net run times in seconds
lazy cancellation (without jump forward) vs
jump forward (with lazy cancellation)
lazy
jump forward
Seconds
100
80
60
40
20
0
100 100
97 97
89 91
87 88
81 81
78 79
77 77
75 76
40
45
50
55
60
65
70
75
Number of Nodes

Figure 6

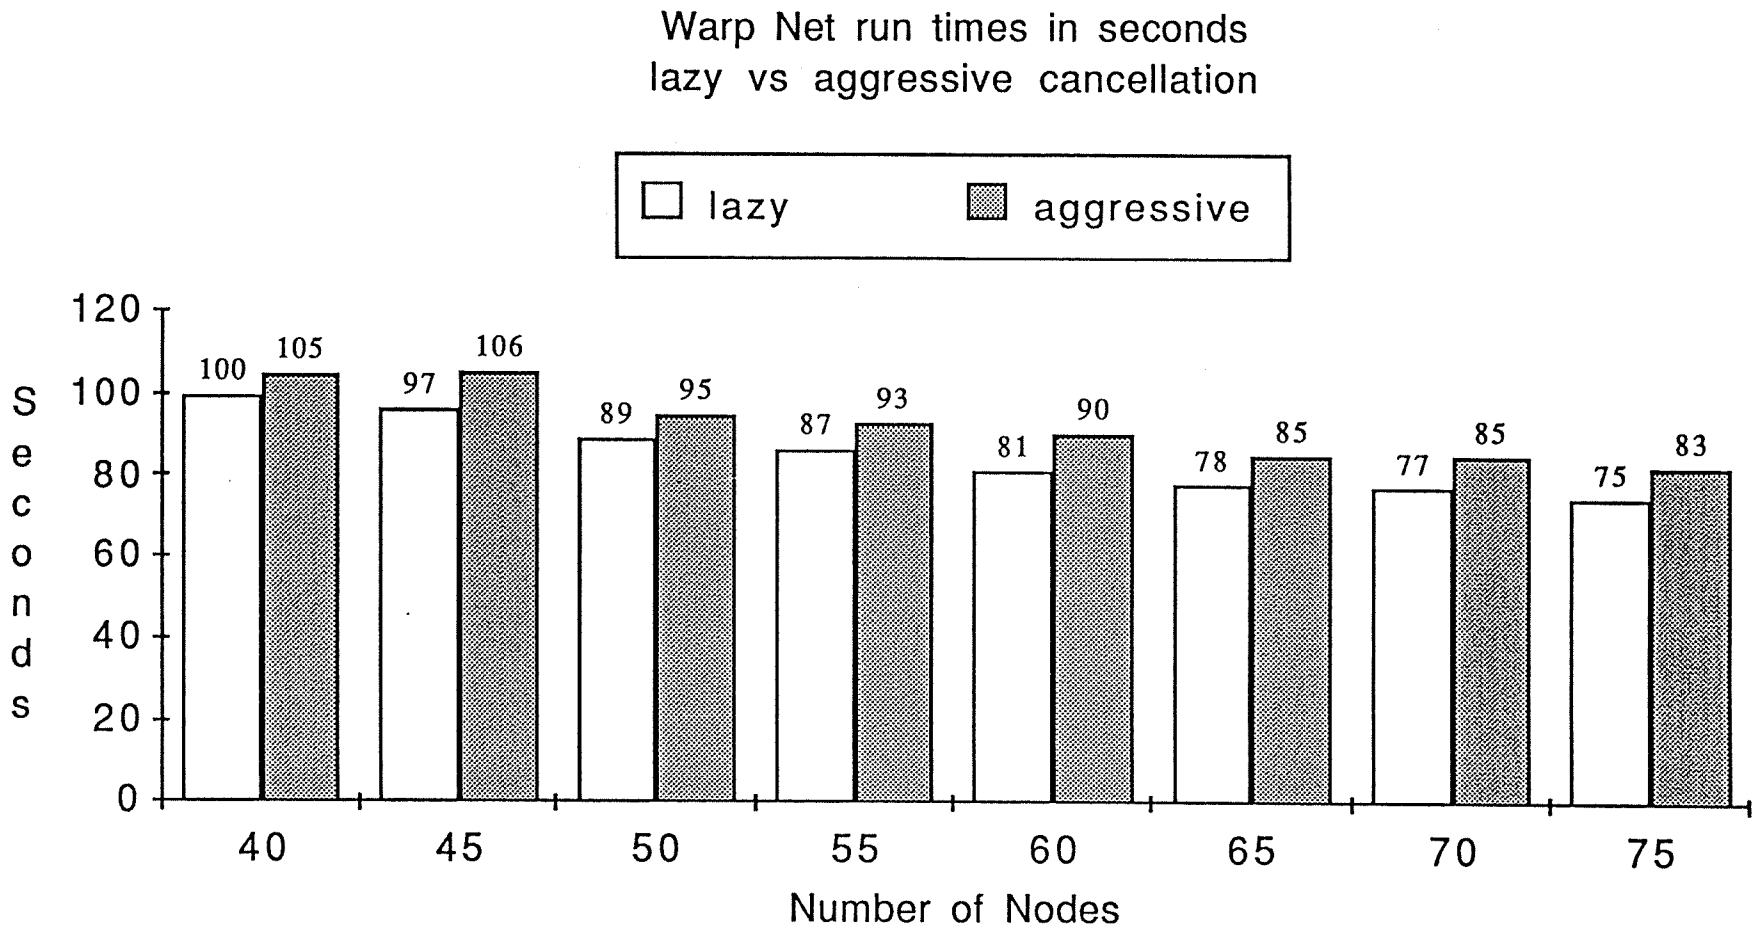
Warp Net run times in seconds
lazy vs aggressive cancellation
lazy
aggressive
Seconds
120
100
80
60
40
20
0
100
105
97
106
89
95
87
93
81
90
78
85
77
85
75
83
40
45
50
55
60
65
70
75
Number of Nodes

Figure 7

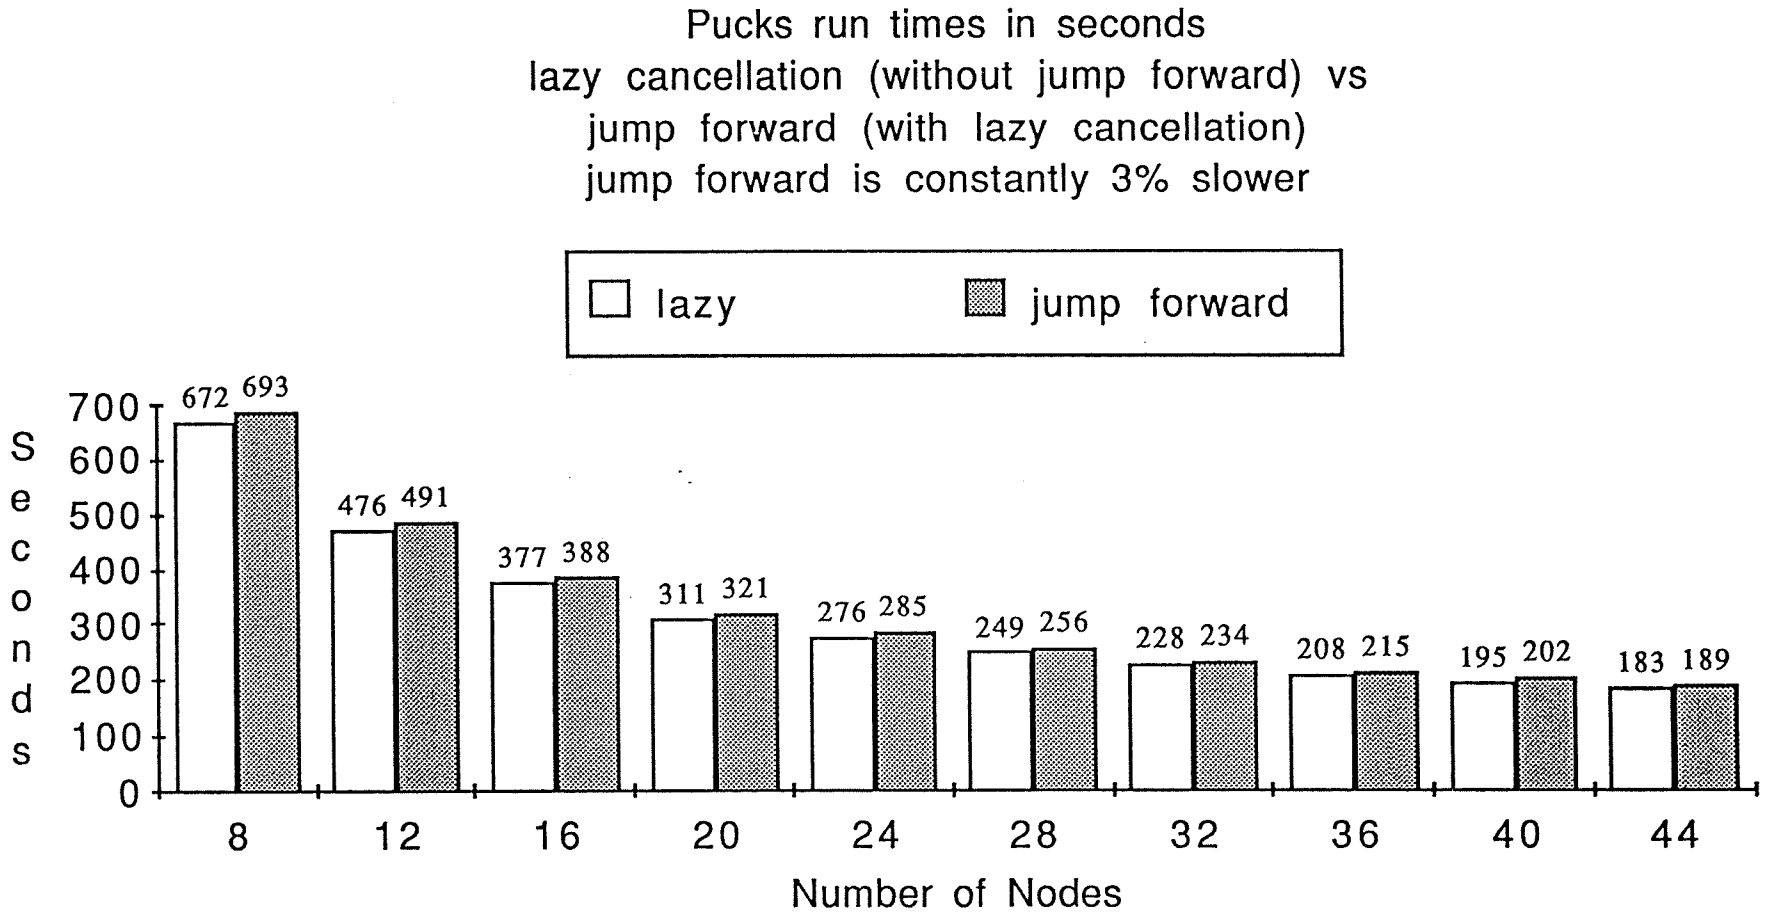
Pucks run times in seconds
lazy cancellation (without jump forward) vs
jump forward (with lazy cancellation)
jump forward is constantly 3% slower
lazy
jump forward
Seconds
700
600
500
400
300
200
100
0
672 693
476 491
377 388
311 321
276 285
249 256
228 234
208 215
195 202
183 189
8
12
16
20
24
28
32
36
40
44
Number of Nodes

Figure 8

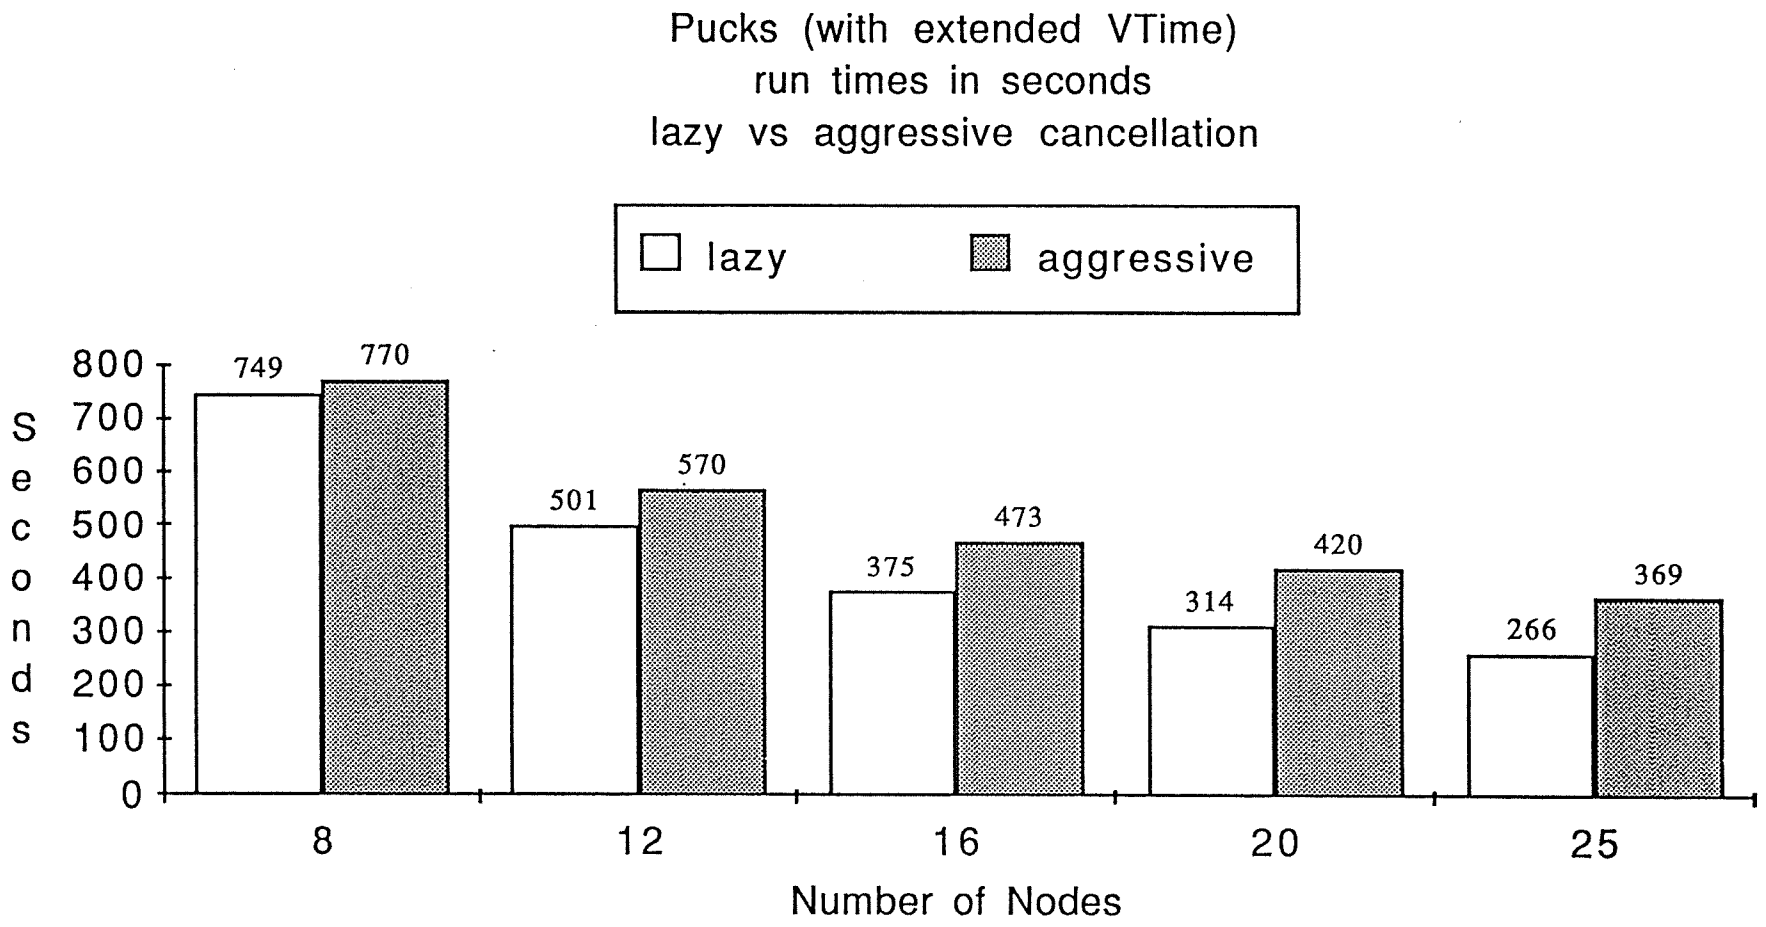
Pucks (with extended VTime)
run times in seconds
lazy vs aggressive cancellation
lazy
aggressive
Seconds
800
700
600
500
400
300
200
100
0
749
770
501
570
375
473
314
420
266
369
8
12
16
20
25
Number of Nodes

Figure 9

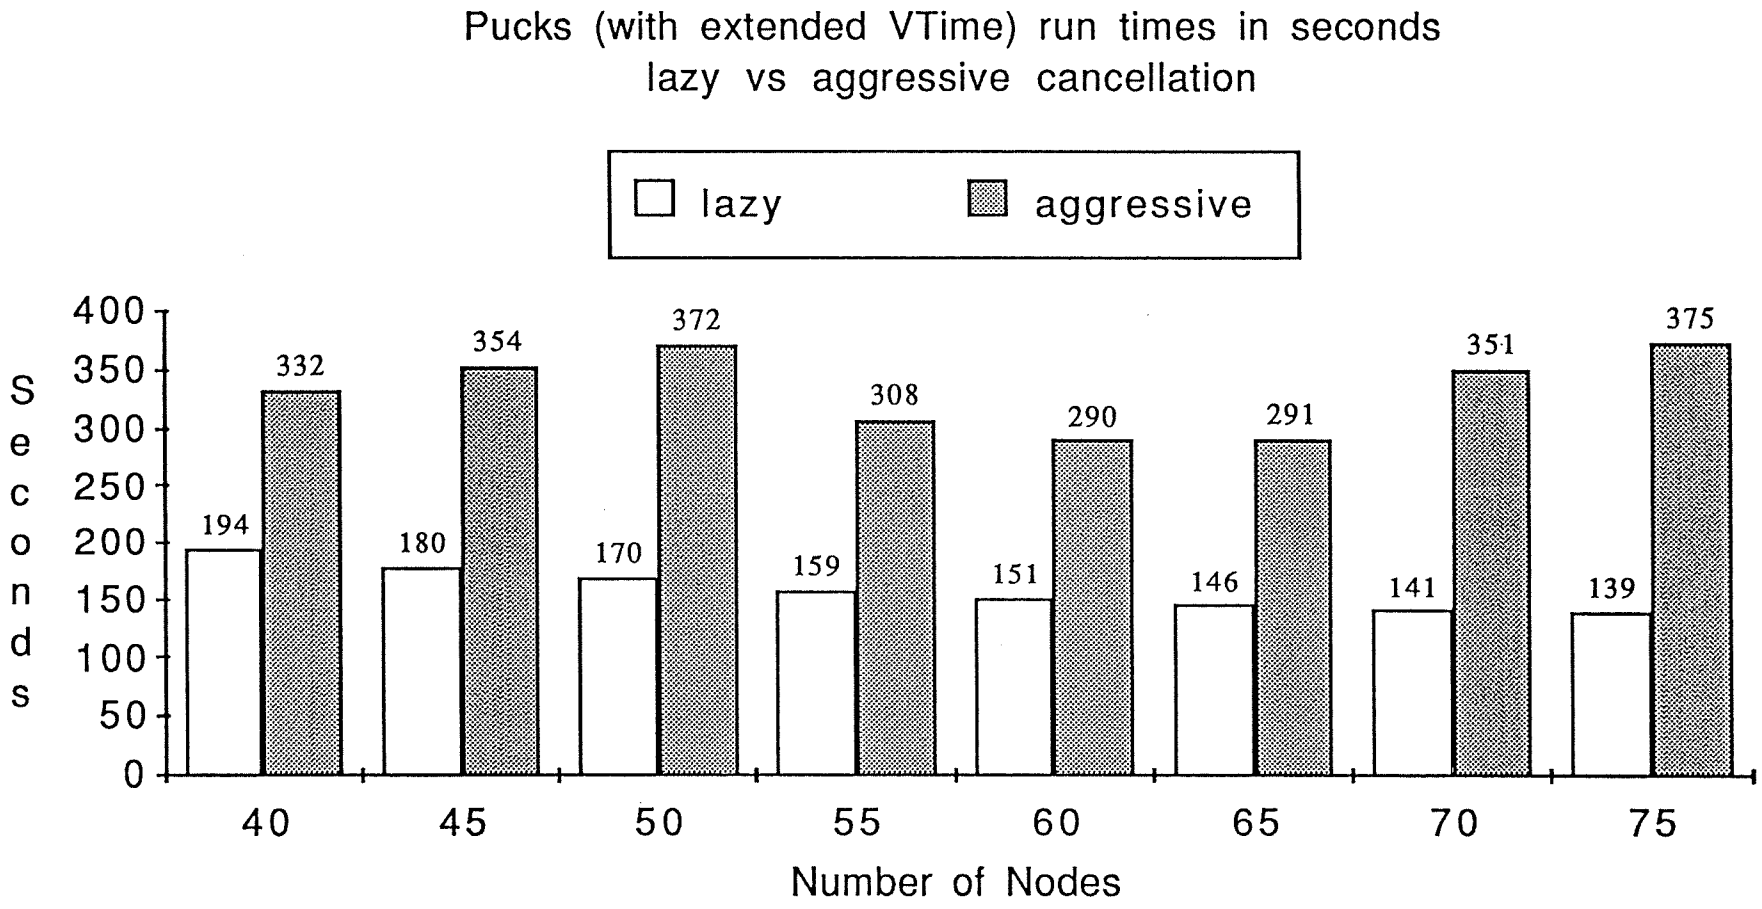
Pucks (with extended VTime) run times in seconds
lazy vs aggressive cancellation
lazy
aggressive
Seconds
400
350
300
250
200
150
100
50
0
194
332
180
354
170
372
159
308
151
290
146
291
141
351
139
375
40
45
50
55
60
65
70
75
Number of Nodes

Figure 10

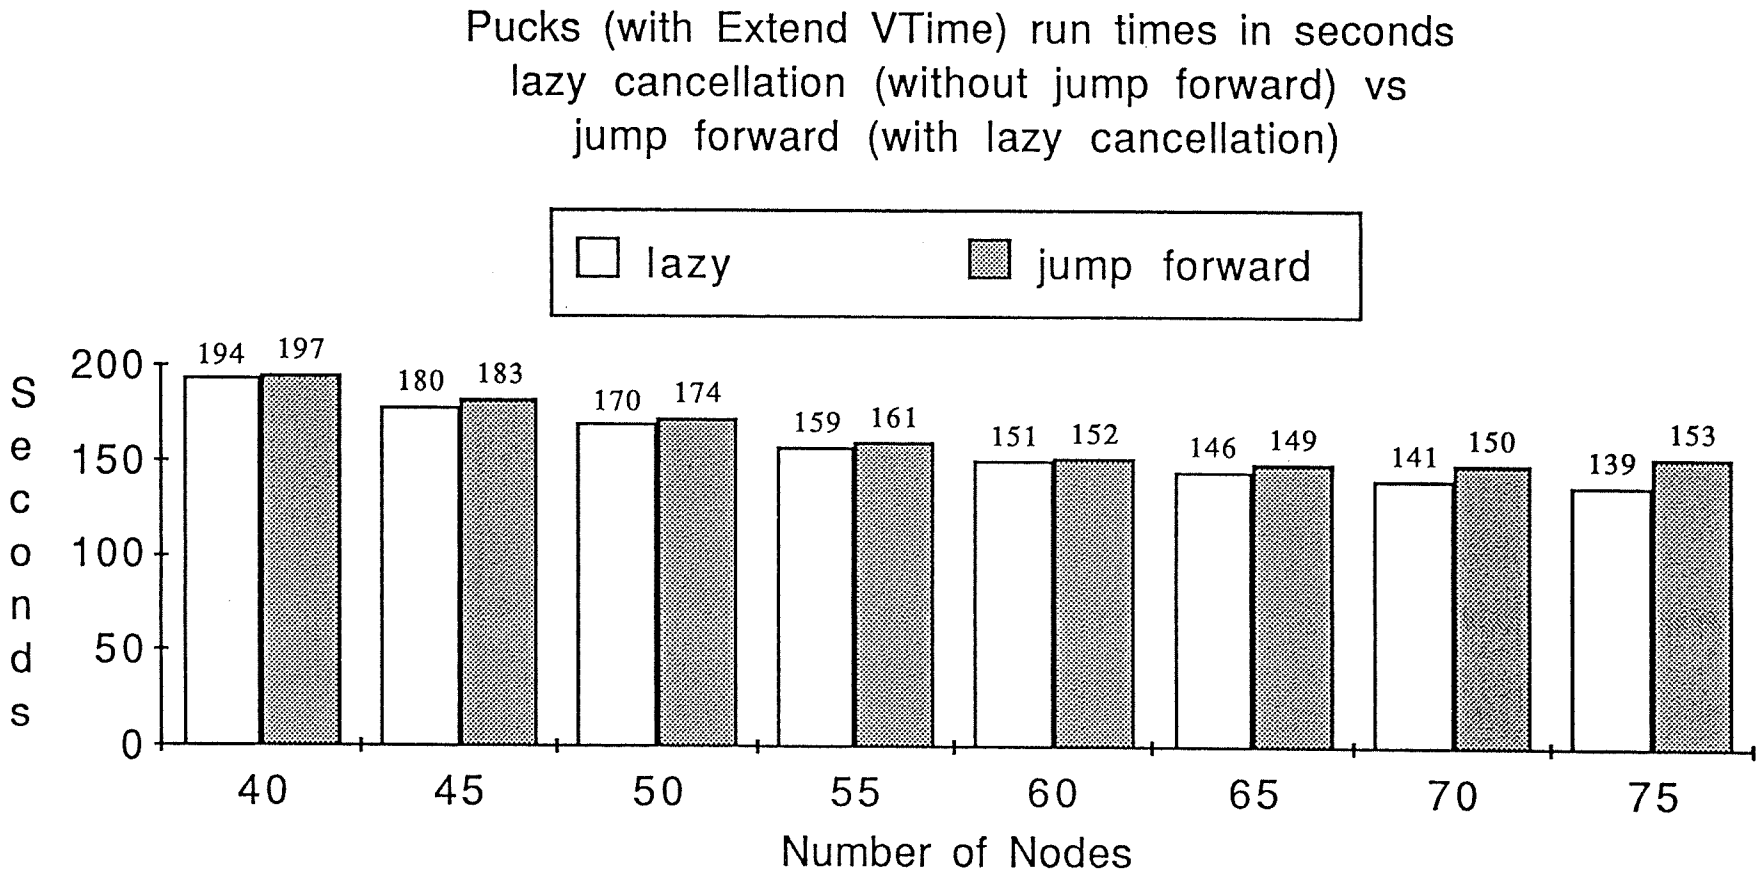
Pucks (with Extend VTime) run times in seconds
lazy cancellation (without jump forward) vs
jump forward (with lazy cancellation)
lazy
jump forward
Seconds
200
150
100
50
0
194 197
180 183
170 174
159 161
151 152
146 149
141 150
139 153
40
45
50
55
60
65
70
75
Number of Nodes

Figure 11