Beam search strategy for IPS
Evaluation of structure-based heuristics

Cathy Berthouzoz and Paola Merlo
Département de linguistique - LATL
Université de Genève

1 Introduction

Satisfactory resolution of ambiguity is one of the main tests of the appropriateness of the design of a system for natural language processing, as ambiguity is one of the pervasive and defining features of natural languages, at all levels.

The goal one wants to obtain in a practical system when resolving ambiguities is to produce natural preferred analyses and therefore the preferred interpretation of a given input, using limited computational resources.

One needs to determine first what is a natural preferred analysis. On the one hand, a selective ranking of alternatives seems to be the correct choice. Observationally, we remark that people impose rather strong restrictions on the preferred analyses they assign to an ambiguous sentence, which is often not even perceived as ambiguous. In some cases, speakers do not realize that a sentence is ambiguous and give it an incorrect interpretation. In some cases, speakers realize they have made a mistake only later on in the sentence or in the conversation. In these cases – the so called *garden paths* – it has been shown to be very difficult to recover the correct interpretation for the sentence (Bever 1970; Frazier 1978; Frazier and Rayner 1982; Mitchell 1987). Computationally, a program that returns all analyses, although certainly complete, is not satisfactory in practice as it would produce a considerable amount of irrelevant analyses.

On the other hand, simultaneous availability of alternatives seems to better reflect findings in experimental psycholinguistics and to be computationally more practical. Experiments on word recognition and on-line sentences processing have recently shown that all senses of a word and all syntactic alternatives are available when the input is ambiguous, even if speakers do not have a conscious perception of it (MacDonald 1994; Trueswell, Tanenhaus, and Kello 1993). Effortless and computationally efficient reanalysis is possible only to those systems that compute partial parses in parallel.

Parallel algorithms that require a bounded amount of resources – and therefore use pruning devices in the course of the analysis – reconcile the apparently contradictory requirements of efficiency and robustness (Gibson 1991; Jurafski 1996).

Psycholinguistics and computational linguistics have very different approaches to the determination of the pruning devices. By and large, there are two approaches to the resolution of ambiguity and the account for preferences in the human sentence processing literature. In one view, preferences are an effect of structural complexity, or unsatisfied grammatical functions, as parsing is a structure-building process (Frazier and Fodor 1978; Berwick and Weinberg 1984; Gibson 1991; Pritchett 1992; Stevenson 1994). In the opposite view, preferences are the result of frequency-based activation of prestored representations. Parsing consists in building structure only to a very limited degree (MacDonald 1994; MacDonald and Seidenberg 1994; Spivey-Knowlton and Sedivy 1995; Trueswell, Tanenhaus, and Kello 1993). Computational linguists simply assume that counts from a corpus are going to provide preferences by estimating probabilities of different analyses, and that the most probable analysis is the most preferred by humans according to the corpus counts (Magerman and Marcus 1991; Briscoe and Carroll 1993; Bod 1995).

Our approach to the resolution of ambiguity is the following: we adopt a chart parser, developed at LATL and described in other articles and reports (Wehrli 1992; Wehrli 1993). In this way, we are able to satisfy the requirement of full parallelism, implemented in an efficient way. In the present implementation, the space of parsing hypotheses proposed by the parser is pruned by structure-based mechanisms, which take into account the preferences studied in experimental psycholinguistics, and the way these preferences are related to different tree topologies (Walther 1993). The goal of our work is to substitute structure-based pruning mechanisms with statistical-based ones, in order to increase performance and automatic acquisition of data. In this report we outline the procedure that has been used to determine the best structure-based approach possible with the current parsing system. This is necessary to generate fair and useful evaluations of the statistical system described in other reports. All the tests, simulations and implementations which will be described below are integrated into the IPS parser.

2 Structure-based ranking strategies and pruning

In order to test a corpus-based approach to the resolution of syntactic ambiguity in principle-based parsing, we need to establish a well-implemented and fair structural alternative. This alternative configuration of the parser must be developed after careful evaluation of the performance of several filtering strategies.

2.1 The ranking strategies

In Walther (1993), several strategies are proposed to resolve ambiguities of attachment. Two of them are based on structural properties of the tree that is being constructed during the parse, namely the *attachment-type hierarchy* and the *locality constraint.* A third constraint is developed for prepositional phrases. This constraint does not require any syntactic information, but is based on the linear precedence of the constituents in the sentence. The attachment-type hierarchy is stated in 1 :

Definition 1 *Attachment Hierarchy functional complements* $\prec$ *subcategorized complements* $\prec$ $\{$*adjunct, non-attachment* $\}$

Evidence for this hierarchy comes from ambiguities of type transitive/intransitive showed in example (1) where the NP following the verb can be attached as subcategorized object of the verb (1b) or as specifier of the following predicate *fell* (1c). In this case, it is proposed that subcategorized complements are preferred to non attachment or to attachment as functional complement, resulting in a garden path effect for the sentence (1a) because the correct analysis is not the preferred one. This kind of ambiguity requires a reanalysis of the sentence.[1]

[1] Walther (1993) gives several other examples and relevant references.

(1)a. While John was mending the sock fell off his lap.

b. ... [$_{VP}$ mending [$_{DP}$ **the sock**]]...

c. ... [$_{VP}$ mending] [$_{TP}$ [$_{DP}$ **the sock**]]...

The locality constraint states that attachment of an argument to a local head (known as *low attachment*) is preferred to attachment to a higher one.

Definition 2 *Locality Constraint*
low attachment $\prec$ *high attachment*

In the case of modifier attachment, such as reduced relative clauses, the locality constraint would entail that the reduced relative clause *who was shot on the balcony* in the sentence below modifies the *actress*, the lowest (closes) head, as shown in (2b) instead of attachment as modifier of *the servant*, the higher head, as in (2c).

(2)a. The servant of the actress who was shot on the balcony was very rich

b. ... [$_{DP}$ the actress [$_{relP}$ who was shot ...]]

c. ... [$_{DP}$ the servant [$_{PP}$ of the actress] [$_{Relp}$ was very rich]]

For PP attachment, the locality constraint just discussed is not sufficient to select the correct analysis and a different strategy is used, based on linear precedence.

Definition 3 *PP Attachment Strategy*

1. *the first PP phrase is preferably attached to the VP (rather than attached low), as shown in (3c),*
2. *if the PP phrase is not the first one, low attachment (3d) is preferred to attachment-to-VP, as shown in (3e).*

(3)a. John gave the candies to the boy in the box.

b. ... [$_{VP}$ gave [$_{DP}$ the candies] [$_{PP}$ **to the boy**]]

c. ... [$_{VP}$ gave [$_{DP}$ the candies [$_{PP}$ **to the boy**]]]

d. ... [$_{VP}$ gave [$_{DP}$ the candies] [$_{PP}$ to the boy [$_{PP}$ **in the box**]]]

e. ... [$_{VP}$ gave [$_{DP}$ the candies] [$_{PP}$ to the boy] [$_{PP}$ **in the box**]]

2.2 Implementation in IPS

The attachment-type hierarchy, the locality constraint, and the PP-attachment strategy are part of a filtering and ranking device used to select the most probable alternative out of a set of *complete analyses* (Wehrli 1993). A sentence receives a complete analysis when all its words have been structured into a single tree according to the grammar. A complete analysis is not necessarily correct.

Filtering device

The attachment-type hierarchy is implemented as a filter on the set of candidate alternatives hypothesized by the chart parser used in IPS. Only pertinent attachment types are taken into account, *i.e.* attachments which can compete against others – for example subcategorized complements and adjuncts.

Ranking device

When the sentence is completely parsed, the number of alternatives can be very high. Some of the alternatives are eliminated because they contain rather obvious mistakes, such as interrogatives or exclamatives which do not end with the correct punctuation sign, or they establish unsatisfied long distance dependencies, for instance they contain $\overline{A}$ chains which are not related to a trace in argument position. All the remaining alternatives are grammatically correct. The ranking device is used to select the most probable alternative out of a list of complete analyses. The cost of each alternative is obtained by summing up the score of each constituent of the tree structure. The score is the sum of two values : the `notes`, which encode violations of the grammar, and the `score`, which encodes preferences – based on attachment type and locality. The alternatives are then inserted into a binary tree. The root of the tree is the preferred analysis, since it is assigned the lowest score.

3 Results

Different variations of the filtering strategies were tested on the set of sentences given in Appendix A.

3.1 Materials and Method

3.1.1 The Corpus

A test suite of 61 sentences was created by collating three different sorts of sentences: 20 sentences are instances of different types of PP-attachment ambiguity, ranging from ambiguity between an NP/VP attachment, to different NP attachment sites. Testing performance on this kind of ambiguity is important, as different PP can attach to different attachment sites, generating a large number of alternative analyses (Church and Patil 1982).

Moreover, 20 sentences reflecting the so called NP/S ambiguity were selected from the annotated version of the Wall Street Journal. 10 verbs were selected from unpublished data resulting from Susan Garnsey'sGarnsey 1994 experiment. This experiment collected data on the relative frequency of completing a verb with an NP complement, or an S complement (introduced by that or not). We selected the verbs by stratified sampling, in such a way that 5 verbs were biased for an NP complement and 5 verbs were biased for an S complement, and that the differential in preference for the two continuations was strong (more than 4 on a 1 to 10 scale). The NP-biased group contains appreciate (10 - 0), assert (9 - 1), find(9 - 1), acknowledge (8 - 2) and notice (7 - 3), while the S-biased group includess hope (10 - 0), wish (9 - 1), admit (9 - 1), assume (8 - 2) and realize (7 - 3).

Moreover, 2 sentences were adapted from the text, one containing an NP resolution and one containing an S resolution. This ambiguity is going to provide us with the data necessary to compare our results with the various experimental results proposed in the psycholinguistic literature, where this ambiguity has been extensively studied (Trueswell, Tanenhaus, and Kello 1993).

Finally, 20 miscellaneous sentences from a previous test suite were added, and the classical landmark *While she was mending the sock fell off her lap* was added too, in order to test that filtering to handle specific kinds of ambiguity would not decrease performance on other constructions.

3.1.2 Metrics

The score used for evaluation is a measure of cost of an analysis, as it takes into account the distance of each filtering strategy from a non-filtering baseline. When all sentences in the test suite are taken into account, this score is going to indicate the cost of adding a filtering strategy to a basic chart parser without filtering.

Two factors are taken into account: the number of generated constituents and the number of complete analyses. The former is an index of the power of the filtering strategy, and thus the best score must maximize this distance. The latter is a measure of missing analyses that a filtering strategy can cause, thus it is a distance that must be minimized. The two factors of the measure are combined in a simple ratio. However, we want to penalize irrecoverably all those cases for which no analysis was found.

The final score for each filter is the mean of the score for each sentence. The score for each sentence is calculated as follows.

(1)

$$If\ no\ analysis\ then\ score^k = 0$$

$$else\quad score^k = \frac{B_c^k - F_c^k}{(B_a^k - F_a^k) + 1}$$

B_c and B_a are the number of constituents and the number of analyses generated by the basic, non-filtering version of the parser, respectively;

F_c and F_a are the number of constituents and the number of analyses generated by any of the filtering versions of the parser, respectively;

k ranges over all the filtering versions.

This measure ranges over the rational numbers greater or equal to 1 and it is always defined. Larger values indicate better performance.

3.2 The Filters and Variations

3.2.1 Basic filter implementation

Filtering implements the idea that parsing preferences depend on an attachment type hierarchy. Functional complements, subcategorized complements, adjuncts and specifiers, and noun modifiers are partitioned into four different sets.[2][3]

The filtering procedure takes into account whether constituents are maximal projections or not. Functional complements, subcategorized complements and adjuncts, undergo filtering only if they are maximal projections. Specifiers and noun-modifiers – with the preference given to modifiers – undergo filtering in all cases. At each time unit, only the constituents belonging to the set which has the highest priority – relatively to the others – are kept.

3.2.2 Some variations

Multiple analyses are generated at two different points in the parser, and thus filters too can be applied at these points. It follows that we have three different filters:

1. Filter 1 is applied after each word sense, in which case the list to be filtered is the list of candidates waiting to be put into the chart. The elements that are marked for deletion are not put into the chart.
2. Filter 2 is applied after all the senses of a word are treated, in which case the list to be filtered is the list of constituents ending at the vertex, which is current at that point in the parse. The elements that are marked for deletion are removed from the chart.
3. Filter 12 is a combination of the two, *i.e.* the filter is first applied after each word sense and then after the word is completely processed.

Filters can be applied in their general form as described above, or in one of the three following variations.

1. Variation 1 concerns the length of the constituent to be filtered: it filters all kinds of constituents and not only maximal projections. We expect this variation to be more filtering than the basic filter.
2. Variation 2 concerns the sets of constituents to be kept: instead of keeping only one set, *i.e.* one kind of attachment, we keep pairs of sets corresponding to the different ambiguities, for example functional complements and subcategorized complements. We expect this variation to be less filtering than the basic filter.
3. Variation 3 is a combination of both variations. All level of projections are filtered, and the two top winning sets are kept.

[2]This is a specificity of English where a lot of nouns and adjectives have the same graphical form and where nouns can modify nouns. In this case adjectives, *i.e.* specifiers, are preferred.

[3]The other kinds of attachments, *i.e.* attachment of a conjunct, a trace in specifier position, a trace as subcategorized complement, a trace in adjunct position, and the projection of a new constituent on the basis of internal properties, are not filtered because they do not lead to any ambiguities.

	PP attach	NP S	others	All
Filter 1				
none	2.435	3.613	3.162	3.079
1	2.492	4.251	3.525	3.436
2	1.253	2.713	4.473	2.811
3	1.253	2.691	4.473	2.804
Filter 2				
none	4.393	3.951	4.869	4.397
1	2.453	2.387	4.505	3.103
2	13.276	11.753	5.725	10.276
3	10.801	5.379	7.346	7.802
Filter 12				
none	6.661	3.927	6.283	5.596
1	6.868	2.715	6.559	5.337
2	9.587	8.980	5.309	7.975
3	10.018	8.480	7.077	8.524

Table 1: Mean score for all filter types, without variations, (none) and with variations (1,2,3)

3.3 Discussion

In this section, we illustrate and discuss the results of the parser for the test suite, in all the different combinations of filters and variations. The scores assigned to each sentence in the suite under the 10 different conditions are shown in Appendix B. The means over the 61 observations in the 10 different conditions are shown in Table 1. Table 2 shows a qualitative intepretation of the scores in Appendix B. The number of null analyses and the number of hits were counted. A hit is the highest score assigned to a given sentence in the test suite — recall that the higher the value the better the performance of the filter.

First, we observe that filter 1 is always lower ranked than filter 2 and filter 12. Filter 1 is too selective, as is confirmed by looking at the number of zero scores in Table 1 or the number of null analyses in Table 2. This observation suggests that trying to filter word senses on a very local basis is too error-prone. Since this result is constant for all types of sentences, we no longer consider filter 1.

Version 2-2 and version 12-3 of the filter appear to strike a successful balance between overgeneration and filtering.

The direct comparison of filter 2 and filter 12 raises some substantive issues about the evaluation metrics used. Reasons must be given for choosing one over the other (as the results are not equivalent). In comparing the results in Table 1 and in Table 2, we notice the following:

1. The best filter-variation combination in the numerical metrics reported in Table 1 is not simply the one that has the highest number of hits in Table 2.

2. The worst filter-variation combination in Table 1 is not simply the one that has the highest number of null analyses in Table 2.

3. The best filter-variation combination in Table 1 is not simply the best ratio between the one that has the highest number of hits and the one that has the highest number of null

	Score	Hits	Nulls
Filter 1			
none	3.0	1	15
1	3.4	1	12
2	2.8	1	36
3	2.8	1	36
Filter 2			
none	4.3	15	27
1	3.1	8	39
2	10.2	17	5
3	7.8	9	16
Filter 12			
none	5.5	5	10
1	5.3	7	17
2	7.9	14	8
3	8.5	19	4

Table 2: Mean score, number of hits and number of null analyses for all filter types, without (none) and with variations (1,2,3)

analyses in Table 2.

Scores in Table 1 select a best filter (2-2) which is different from the one that is selected by the measure described in 3 above, which selects 12-3 in Table 2.

The cause of this discrepancy resides both in the method of assigning the scores to each sentence and in the metrics used. First, although we have reason to think that the scores are quite finely tuned, as they were set empirically and they perform very well in selecting the appropriate analyses at the end of the sentence, still we can suppose there is room for improvement. Some sentences appear to obtain excessively high scores. As far as the metrics is concerned, we observe that in neither of the two methods to calculate a ranking is variance computed. On one hand, if we look at the numerical values of some of the scores in Appendix B we notice that the higher result of 2-2 in the scores is greatly affected by the excellent performance of a few cases. In more precise terms, the metrics in (1) does not take variance into account. A good metrics must take into account the fact that the measure we want to obtain is an average, and must privilege performances with small variances, as the scores reflect better the distribution of the population. On the other hand, the ratio between hits and null analyses treats all scores as units, artificially eliminating any effect of variance.

For the moment then, we consider filter 2-2 as the best filter, while working on the development of a more comprehensive evaluation metrics. We are going to use this version of the BIPS parser as the structural counterpart to compare to our statistical beam search method.

4 Conclusion

In this report, we have described evaluation work of the existing parser, IPS. The parser pursues all analyses in parallel in order to parse ambiguous input. This method is computationally very costly. We have tried several filtering strategies to reduce the number of active analyses at any given step in the parse, without too much loss in the number of valid alternative analyses that must be recovered by the end of the sentence. We find that the parser performs at its best if a filter is applied after each word.

References

Berwick, R. and A. Weinberg (1984). *The Grammatical Basis of Linguistic Performance.* Cambridge, MA: MIT Press.

Bever, T. (1970). The cognitive basis for linguistic structure. In J. R. Hayes (Ed.), *Cognition and the Development of Language.* New York, NY: John Wiley and Sons, Inc.

Bod, R. (1995). *Enriching Linguistics with Statistics.* Ph. D. thesis, University of Amsterdam.

Briscoe, E. and J. Carroll (1993). Generalized probabilistic lr parsing of natural language (corpora) with unification based grammars. *Computational Linguistics 19(1)*, 25–60.

Church, K. and R. Patil (1982). Coping with syntactic ambiguity or how to put the block in the box on the table. *American Journal of Computational Linguistics 8(3-4)*, 139–149.

Frazier, L. (1978). *On Comprehending Sentences: Syntactic Parsing Strategies.* Ph. D. thesis, University of Connecticut. unpublished.

Frazier, L. and J. D. Fodor (1978). The sausage machine: A new two-stage parsing model. *Cognition 6*, 291–325.

Frazier, L. and K. Rayner (1982). Making and correcting errors during sentence comprehension: Eye-movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology 14*, 178–210.

Garnsey, S. (1994). unpublished data.

Gibson, E. (1991). *A Computational Theory of Human Linguistic Processing: Memory Limitations and Processing Breakdown.* Ph. D. thesis, Carnegie Mellon University, Pittsburgh, PA.

Jurafski, D. (1996). A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science 20*(2).

MacDonald, M. (1994). Probabilistic constraints and syntactic ambiguity resolution. *Language and Cognitive Processes 9(2)*, 157–201.

MacDonald, Maryellen C., N. J. P. and M. Seidenberg (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review.*

Magerman, D. and M. Marcus (1991). Pearl: A probabilistic parser. In *Proceedings of the European Chapter of the Association for Computational Linguistics*, Berlin.

Mitchell, D. (1987). Lexical guidance in human parsing: Locus and processing characteristics. In M. Coltheart (Ed.), *Attention and Performance XII: The Psychology of Reading*, pp. 601–618. Hillsdale, NJ: Erlbaum.

Pritchett, B. (1992). *Linguistic Competence and Parsing Performance.* Chicago, IL: University of Chicago Press.

Spivey-Knowlton, M. and J. Sedivy (1995). Resolving attachment ambiguity with multiple constraints. *Cognition 55*(3), 227–267.

Stevenson, S. (1994). *A Competitive Attachment Model for Resolving Syntactic Ambiguities in Natural Language Parsing.* Ph. D. thesis, Department of Computer Science, University of Maryland. Available as technical report TR-18 from Rutgers Center for Cognitive Science, Rutgers University.

Trueswell, J., M. Tanenhaus, and C. Kello (1993). Verb specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory and Cognition 19*, 528–553.

Walther, C. (1993). Psycholinguistically-based selection of analyses in ips and fips. Notes techniques 93/3, Projet FIPS, LATL, University of Geneva.

Wehrli, E. (1992). The IPS system. In *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92)*, pp. 870–874.

Wehrli, E. (1993). Vers un système de traduction interactif. In P. Bouillon and A. Clas (Eds.), *La Traductique*, pp. 423–432. Les Presses de l'Université de Montréal.

A Test Sentences

1: He saw a man with a telescope.
2: I saw the man without clothes.
3: Paul saw Mary under the bridge.
4: She hit the man with an umbrella.
5: I followed the man with a bicycle.
6: John met his girlfriend under pressure.
7: Mary saw John and the dog over the hill.
8: He closed the windows with wooden panels.
9: John sent the books and the letter to Mary.
10: She talked to the president of the meeting.
11: John gave the candies to the boy in the box.
12: I put the chair near the table with a scratch.
13: John sent the books and the letter to Mary to Paul.
14: She warned the students of the English examination.
15: She warned the specialists of the war of the danger.
16: She placed the sofa near the armchair with a red cover.
17: She talked to the president of the company of the meeting.
18: The lightning struck the tree near the house over the hill.
19: The rain damaged the computers near the boxes with plastic tops.
20: The rain damaged the computers near the empty boxes with colour screens.
21: The company acknowledges some problems.
22: When I realized the story was over I cried.
23: I noticed the car was bouncing up and down.
24: Her sister wishes Paul had a different job.
25: While John was mending the sock fell off his lap.
26: Paul admitted his company had overstated profits.
27: We could not find the charity does anything good.
28: The biologist noticed some strange bird behaviour.
29: He asserts John can not react quickly to competition.
30: They appreciate the economic stability we have achieved.
31: I know my father would have wished a girl instead of a boy.
32: The league's promoters hope tourists will join fans like Paul.
33: While they acknowledge the president will attend several meals.
34: Paul was acknowledging one month's figures do not prove a trend.
35: They could also find a customer that accept a lower grade of oil.
36: Traders and market officials hope the damage will not be permanent.
37: There is no advantage if the president asserts a right of excision.
38: Without admitting any guilt the bureau agreed to several conditions.
39: The company said it expects to realize those gains before the end of the year.
40: The two judges conceded that they assume the quotations were deliberately altered.
41: The president said he would assume Paul's responsibilities if a successor is not found.
42: John had to go.
43: John wants to go.
44: John is happy to go.
45: Give it to this child!
46: Has he never gone there?
47: She is more patient than John.
48: Where will you go tomorrow?

49: This book is easy to talk of.
50: The apple is very easy to eat.
51: John thinks that Mary will come.
52: John wanted to introduce him to us.
53: Is Mary coming because she wants to?
54: John seems to have liked Mary's eyes.
55: At what was the cat looking yesterday?
56: Have they promised each other to look at themselves?
57: Mary said that John made her read it to the children.
58: Paul wanted to have his car fixed by his brother's friend.
59: The small white mouse has been looked at by my sister's cat.
60: I was reading the Sunday afternoon newspapers that my brother bought.
61: The last suicide attempt of the king has been reported in the Sunday newspapers.

B Scores

Filter 1				Filter 2				Filter 12			
1-0	1-1	1-2	1-3	2-0	2-1	2-2	2-3	12-0	12-1	12-2	12-3
3.000	3.000	3.000	3.000	0.000	0.000	11.000	5.667	4.333	4.333	11.000	11.000
1.625	1.625	0.000	0.000	3.333	3.556	6.000	5.000	3.000	3.250	6.000	9.000
0.000	0.000	0.000	0.000	0.000	0.000	8.000	8.000	0.000	0.000	8.000	8.000
1.500	1.500	0.000	0.000	0.000	0.000	3.000	6.000	3.500	3.500	3.000	3.000
2.000	2.000	1.000	1.000	0.000	0.000	9.000	13.000	7.000	6.500	10.000	9.000
1.000	1.000	0.000	0.000	0.000	0.000	1.000	1.000	2.500	2.500	1.000	1.000
2.438	2.438	0.000	0.000	0.000	0.000	4.000	3.938	0.000	0.000	4.000	4.000
1.000	1.000	0.000	0.000	0.000	0.000	4.000	7.000	4.000	4.000	4.000	4.000
0.000	0.000	0.000	0.000	27.000	27.000	15.000	18.000	21.000	21.000	15.000	15.000
0.000	0.000	0.000	0.000	9.000	9.000	5.000	6.000	6.000	7.000	5.000	6.000
2.000	2.000	0.000	0.000	9.875	0.000	7.000	9.000	6.500	8.111	7.000	7.833
4.000	4.000	1.000	1.000	0.000	0.000	27.000	13.000	10.200	10.600	27.000	29.00
1.000	1.000	0.000	0.000	0.000	0.000	17.000	36.000	17.000	17.000	17.000	14.00
5.250	5.250	5.667	5.667	0.000	0.000	21.000	37.000	9.000	9.000	10.000	10.00
3.000	3.000	0.000	0.000	0.000	0.000	24.000	16.667	9.200	9.200	24.000	24.00
7.556	7.556	8.000	8.000	0.000	0.000	54.000	19.250	11.333	11.222	11.857	11.857
1.500	1.500	0.000	0.000	9.500	9.500	14.000	11.500	10.500	10.500	14.000	14.000
4.542	5.583	0.000	0.000	8.880	0.000	6.077	0.000	0.000	0.000	6.000	10.154
0.000	0.000	0.000	0.000	9.000	0.000	23.000	0.000	0.000	0.000	0.000	0.000
7.286	7.381	6.391	6.391	11.282	0.000	6.448	0.000	8.154	9.641	7.879	9.515
1.000	1.000	0.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000
1.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	3.083	3.333	0.000	0.000
0.000	0.000	0.000	0.000	3.826	3.826	0.000	0.000	1.765	1.765	0.000	0.000
6.000	10.000	0.000	0.000	0.000	0.000	2.000	0.000	10.500	0.000	2.000	15.000
2.357	2.357	0.000	0.000	0.000	0.000	15.000	5.784	0.000	0.000	16.000	4.714
4.000	4.000	0.000	0.000	14.000	14.000	8.000	11.000	0.000	0.000	9.000	8.000
9.000	9.000	3.200	3.200	0.000	0.000	24.000	15.500	3.357	3.357	3.700	3.700
0.000	2.000	0.000	0.000	9.333	9.333	8.000	9.000	3.667	6.333	8.000	10.000
0.000	0.000	0.000	0.000	5.000	0.000	2.000	3.000	3.000	0.000	2.000	2.000
0.000	0.000	0.000	0.000	5.333	5.667	14.000	15.000	5.333	5.667	14.000	15.000
2.000	2.000	2.000	2.000	9.000	8.250	5.000	5.000	7.000	9.000	4.000	5.000
5.736	6.180	8.103	8.103	7.542	0.000	8.567	8.061	6.617	7.376	7.881	7.992
1.000	3.250	3.800	3.800	4.045	0.000	10.000	0.000	3.909	4.000	4.200	3.571
15.143	15.143	14.286	14.286	0.000	0.000	11.143	17.714	0.000	0.000	19.571	19.571
4.000	4.000	2.000	2.000	0.000	0.000	17.000	0.000	0.000	0.000	17.000	17.000
5.250	5.906	11.062	10.600	8.091	0.000	8.778	9.136	6.718	8.179	12.062	12.050
3.297	3.486	4.000	4.000	0.000	0.000	34.000	5.025	4.289	4.489	5.087	4.815
3.543	5.102	0.000	0.000	5.581	0.000	29.000	0.000	5.347	0.000	29.000	0.000
3.571	3.714	0.000	0.000	0.000	0.000	18.000	0.000	6.357	0.000	18.000	28.000
6.059	7.588	5.833	5.833	6.925	9.054	10.333	8.741	7.357	0.000	11.083	17.583
2.923	3.538	2.692	2.692	3.303	0.000	21.000	0.000	3.172	3.517	5.000	4.077
0.000	0.000	0.000	0.000	3.000	3.000	0.000	0.000	1.000	1.000	0.000	0.000
0.000	1.000	0.000	0.000	11.000	11.000	2.000	2.000	5.000	5.000	2.000	2.000
2.000	2.000	0.000	0.000	0.000	0.000	4.000	4.000	9.000	9.000	4.000	4.000
0.000	0.000	0.000	0.000	5.000	4.000	0.000	0.000	2.000	3.000	0.000	0.000
0.000	0.000	0.000	0.000	2.250	2.250	5.000	3.500	2.667	2.667	5.000	5.000
15.000	15.000	17.000	17.000	0.000	0.000	10.000	10.000	20.000	23.000	23.000	23.000
0.000	2.000	2.333	2.333	3.750	3.750	3.000	2.333	2.500	3.000	3.000	3.667
2.235	2.294	27.000	27.000	2.629	3.600	3.333	2.778	2.611	2.917	11.000	12.333
6.400	7.000	6.333	6.333	11.750	7.727	4.500	6.889	7.333	9.333	6.000	6.429
4.000	14.000	0.000	0.000	19.000	23.000	7.000	12.000	11.000	18.000	7.000	7.000
1.000	1.000	0.000	0.000	0.000	0.000	0.000	0.000	2.500	2.500	0.000	1.000
0.000	1.500	0.000	0.000	10.000	0.000	5.000	7.000	5.500	6.500	5.000	5.000
4.000	4.000	4.000	4.000	14.000	14.000	5.000	6.000	10.000	10.000	7.000	7.000
0.000	0.000	0.000	0.000	5.000	5.000	2.000	5.000	3.000	3.000	2.000	2.000
1.500	1.500	0.000	0.000	0.000	0.000	3.000	4.000	13.000	13.000	3.000	3.000
2.000	2.000	0.000	0.000	0.000	0.000	8.000	16.000	0.000	0.000	8.000	12.000
5.500	5.500	1.000	1.000	0.000	0.000	9.000	3.091	6.545	6.545	9.000	9.000
9.889	0.000	11.000	11.000	10.000	12.778	12.667	21.333	11.154	0.000	11.182	11.636
0.000	2.000	10.400	10.400	0.000	0.000	19.000	41.000	0.000	0.000	0.000	13.667
9.714	9.714	10.400	10.400	0.000	0.000	12.000	0.000	10.857	12.714	0.000	13.800