This paper was presented at
The First Australia-Japan Joint Symposium on Natural Language Processing
Melbourne University, Vic, Australia, November 27-29, 1989

COSIMO: A Massively Parallel Integrated Parser

Jacob L. Cybulski and Andrew Jennings

Artificial Intelligence Systems
Telecom Research Laboratories
770 Blackburn Road
Clayton, Vic 3168, Australia
Phone: +613 541 6273, Fax: +613 543 8863

Abstract

This paper describes a method of integrating syntactic and semantic interpretation of natural language utterances via massively parallel constraint satisfaction implemented in a connectionist chart parser - COSIMO.

Preference Semantics constitutes the conceptual basis for the parser operation. During the input text analysis, individual words are assigned their semantic senses, the concept sequences trigger grammatical rules to form a chart structure, and then the rules act as pattern matching constraints driven by a form of interactive relaxation with decay algorithm.

As a result of activation spreading, inhibition, decay and pattern matching, the semantic representation of input text is constructed. The resulting structures are subsequently submitted for further processing to a planning system. Due to the massively parallel nature of the parser, syntactic and semantic interpretations of input are achieved concurrently by network relaxation. The preference for the resulting semantic structures is ranked by their level of activation, and the unlikely interpretations are discarded by decay processes.

1.0 Introduction

MEDICI is a plan-based, intelligent information management system with graphical and natural language interfaces (Jennings, Rowles and Kowalczyk 1989, Leckie and Zukerman 1989). The main application areas for the system are those in manufacturing and service where goods and services may be requested by means of electronic transaction interchange, as in electronic travel arrangements - the primary example explored in this paper. A selection of typical phrases and sentences directed at COSIMO, MEDICI's natural language interface, are collected in Figure 1. In response to each of these sentences, COSIMO analyses them into a set of semantic structures used subsequently as goals to the MEDICI planning subsystem. Finally, the planner interaction with the intension analysis module results in a set of executable plans satisfying user requirements. (Note that all examples discussed in this paper are of illustrative nature only. COSIMO's dictionary, grammar and knowledge base have, thus, been drastically simplified.)

The choice of methods and techniques used in COSIMO was motivated by two important issues in recent developments in Natural Language Understanding (NLU). The first was the failure of symbolic processing of natural language on conventional sequential machines (Dreyfus and Dreyfus 1986). The second was a rapid ascent of parallel distributed processing methods (Hillis and Steele 1986, Lippman 1987), termed as *connectionism*, which laid the foundation for the conceptual machinery in effective emulation of cognitive processes, including NLU (cf. Arbib and Caplan 1979, McClelland, Rumelhart and Hinton 1986). Unfortunately, the promise of connectionist language applications faded somewhat when some of the linguistic tasks, which seemed so trivial in the symbolic processing paradigm, were found difficult in the parallel distributed realm (e.g. variable binding, task timing, judging process completion, debugging of large systems, etc. - Waltz 1986). Connectionist models are certainly a great leap forward but definitely not a panacea.

<u>Mary</u> wants to fly to <u>Tokyo</u> with <u>Qantas</u>.	(1, 2, 3)
I wish <u>*to drive*</u> to Sydney.	(4)
I ought <u>*to go*</u> <u>*to Sydney*</u>.	(5)
Sydney wants <u>*to travel*</u> <u>*by*</u> <u>*bus*</u>.	(6)
I prefer <u>*to fly with*</u> <u>*JAL*</u>.	(7)
<u>*I want*</u> <u>*to travel overseas*</u>.	(8)
<u>*My wife*</u> <u>*wants*</u> <u>*me*</u> <u>*to come to Melbourne*</u>.	(9)
<u>*I 'd like*</u> <u>*my luggage*</u> <u>*to be sent to my hotel*</u>.	(10)
<u>*My children*</u> <u>*need*</u> <u>*a holiday*</u>.	(11)

Figure 1 - Sample travel-oriented phrases and sentences

Despite difficulties, connectionism has found its way into a number of small parsing systems. The prevalent characteristic of these systems is their ability to propagate weak constraints (lexical, syntactical, semantic and pragmatic) between input-derived terms to form an utterance interpretation. The main mechanism employed in these systems is a neural-like activation spreading between linguistic and cognitive representation elements, activation inhibition and decay (cf. Collins and Loftus 1975, Anderson 1983). Some of the systems focus on the syntactic processing of texts (Slack 1986, Jones 1987, Parisi and Nolfi 1987, Howells 1988), some look at the micro-levels of semantic analysis to resolve lexical or syntactic ambiguity (Cottrell and Small 1983, Pollack and Waltz 1984, Bookman 1987), others concentrate on language generation (Kukich 1987, Gasser 1988) or speech understanding (Bengio et al. 1989, Stanfil and Waltz 1986), few address connectionist NLU in an integrated, multi-facet way (Small 1987, also cf. Minsky 1986, Lytinen 1986). The experimental nature of the connectionist language systems and their limited perception of the linguistic task is not surprising as the researchers are confronted by a number of serious problem areas, e.g. :-

- definition of linguistic universals in terms of subsymbolic microfeatures;
- distributed grammar representation and its automatic acquisition;
- fuzzification of linguistic representations and tasks to fit neural architectures;
- parsing by neural-like processes, i.e. excitation, inhibition and decay;

- case frame instantiation by activation propagation and network relaxation;
- focus maintenance and context setting by activation-based priming; etc.

COSIMO was conceived to be a practical system limited to the processing of isolated sentences and definite phrases (note that a more sophisticated language processing is achieved in MEDICI via planning and intension analysis). Thus, we decided to abandon a pure connectionist approach to NLU and to enhance a mature, symbolic framework of language processing including a number of connectionist features. Several models were reviewed for the amount of parallelism, distribution and fuzziness they could accommodate, finally an *Interactive Relaxation with Decay* algorithm (McClelland and Rumelhart 1986) was decided to be a good implementation vehicle for a *chart* parser (Kay 1980) firmly set into the *Preference Semantics* paradigm (Wilks 1975). Early experiments with COSIMO showed that the decision to combine connectionist and symbolic representations opened entirely new avenues for the natural language processing leading to effective and efficient parsing techniques.

Some other work, although within symbolic processing paradigm, still had a considerable impact on the ideas presented in this paper, e.g. work on inference and interestingness in Conceptual Dependencies (Schank 1979, Schank, Lebowitz and Birnbaum 1980), possibilistic representations and fuzzy sets (Zadeh 1981), validity ratings in Blackboard Systems (Lesser and Erman 1977), finally the implications of Cascaded ATNs for parallel language parsing (Woods 1980).

2.0 Conceptual Elements

Preference Semantics sets a number of organisational and processing goals for the language understander. Paraphrasing Fass and Wilks (1983), the goals are :-

- the need for parsing with no explicit syntax processing;
- procedural semantics based on the principle of a Least Effort;
- representation reducible to domain primitives of different types;
- surface representation of linear rather than recursive nature;
- non-canonical text representation, and its "best-fit" selection;
- context and constraint oriented interpretation.

COSIMO's operation requires three different organisational components :- a *dictionary* assigning semantic word senses to character strings, *formulae* giving semantically reducible definitions of domain and universal concepts, and *templates* specifying linear language grammar. The *connectionist chart parser*, combining representational elements in preferential and constrained way, completes the requirements, first, however, let us discuss COSIMO's memory organisation.

Formulae. COSIMO's knowledge base forms a case frame system defined with a collection of *semantic formulae* reducing high level concepts into a set of well defined semantic primitives (review in Bruce and Moser 1987).

Semantic formulae define frames in clausal form (Figure 2; "::=" is a definition, whereas "$\wedge$" a conjunction operator). The formula head identifies a frame name and declares slot restriction attributes as either compulsory (indicated by "+", as in "travel"), optional (e.g. all of the "act" slots), or forbidden (indicated by "-", as in "ptrans"). The formula tail defines a set of conjunctive forms specifying the class of defined concepts (e.g. "air-travel" being a sub-class of "travel", "ptrans" and "act") and constraining the type and relationships amongst the slot fillers (e.g. an instrument of "air-travel" is restricted to "planes"). Frame class/subclass organization allows for slot and constraint inheritance (e.g. "air-travel" inherits the following slots :- "inst", "agent", "rcpt", "src", "dest", "time", "loc"). (Note that non-inheritable slots are assumed to be forbidden.)

air-travel(A) ::=
    travel(A, inst=B) $\wedge$
    plane(B);
	(1)

travel(A, +agent=B, +src=C, +dest=D, +time=E) ::=
    ptrans(A, inst=F) $\wedge$
    vehicle(F);
	(2)

*ptrans(A, dest=B, src=C, -loc=D) ::=
    act(A) $\wedge$
    place(B) $\wedge$
    place(C);
	(3)

*desire(A, agent=B, object=C) ::=
    state(A) $\wedge$
    animate(B) $\wedge$
    entity(C);
	(4)

act(A, agent=B, rcpt=C, inst=D, time=E, loc=F) ::=
    animate(B) $\wedge$
    animate(C) $\wedge$
    entity(D) $\wedge$
    time(E) $\wedge$
    place(F);
	(5)

state(A, agent=B, rcpt=B, time=C, loc=D) ::=
    animate(B) $\wedge$
    time(C) $\wedge$
    place(D);
	(6)

*user(A) ::=
    person(A) $\wedge$
    attend(agent=cosimo, rcpt=A);
	(7)

*plane(A) ::= vehicle(A);	(8)

Figure 2 - A set of sample formulae

The parser aims at mapping an arbitrary sentence or a phrase into a corresponding semantic structure readily understood by an application program (e.g. MEDICI planner). As COSIMO applications may be able to process a limited range of frame structures only, two types of formulae are distinguished :- *universal* and *domain* (domain formulae are indicated by "*", as in "ptrans", "desire", "user" and "plane"). Upon the completion of the text interpretation, it is the domain-dependent frames that are passed into an application; the universal structures, on the other hand, are used internally as mediators between the domain and lexical elements.

Dictionary. At the first stage of input processing, all input strings, or signs, are mapped into their corresponding senses represented by semantic frame instances. After detection of a specific sign, a frame instance is created and its slots are partially filled with values, according to the rules inscribed into the COSIMO *dictionary.*

user(num={1s}, mode=hear) ::=	sign("I");	(1)
cosimo(num={1s}, mode=speak) ::=	sign("I");	(2)
desire(num={1s 2s 1p 2p 3p}, form={inf pres}) ::=	sign("want");	(3)
desire(num={3s}, form={pres}) ::=	sign("wants");	(4)
desire(num={1s 2s 3s 1p 2p 3p}, form={past past-part}) ::=	sign("wanted");	(5)
desire(num={1s 2s 3s 1p 2p 3p}, form={pres-part}) ::=	sign("wanting");	(6)
air-travel(num={1s 2s 1p 2p 3p}, form={inf pres}) ::=	sign("fly");	(7)
air-travel(num={3s}, form={pres}) ::=	sign("flies");	(8)
air-travel(num={1s 2s 3s 1p 2p 3p}, form={past}) ::=	sign("flew");	(9)
air-travel(num={1s 2s 3s 1p 2p 3p}, form={pres-part}) ::=	sign("flying");	(10)
to-loc ::=	sign("to");	(11)
to-state ::=	sign("to");	(12)
to-inf ::=	sign("to");	(13)
place-name ::=	sign("Tokyo");	(14)
place-name ::=	sign("Sydney");	(15)
person-name ::=	sign("Mary");	(16)
person-name ::=	sign("Sydney");	(17)
airline-name ::=	sign("Qantas");	(18)
airline-name ::=	sign("JAL");	(19)

Figure 3 - A set of sample senses

Each dictionary rule (Figure 3; "::=" stands for sensing) signals the type of a semantic sense (indicated by the associated frame type), together with its interpretation attributes (compulsory slot fillers, rare as the semantic category should suffice), and the word form (usually defined with optional slots). The slot filler selection may be quite arbitrary, it must however comply with the formulae-described frame constraints.

Henceforth, a dictionary entry may be a simple sense association (as in "to", "Tokyo" and "Sydney"), or it may specify a number of frame slots depending on the sense category (as in "I", "want" and "fly"). A number of standard slots are customarily used by COSIMO, to include :- "mode" determining the parsing and generation modes of the system processing (used in "I" definition), "form" specifying verb forms (as in "want" and "fly"), "num" indicating the verb conjugation information (see also "want" and "fly"), "trans" verb transitivity, "gend" word gender, etc.

Templates. *Semantic templates*, similarly to dictionary items, sense input strings and assign to them partially instantiated frames. Their role, however, extends beyond the processing of individual signs, they detect multi-pattern semantic sequences and subsequently unify them into semantic formulae. Each template, thus, acts as means of acquiring knowledge from natural language input into a system of frame instances and propagating them for further processing.

place(name=A) ::= place-name(A);	(1)
person(name=A) ::= person-name(A);	(2)
plane(owner=airline(name=A)) ::= airline-name(A);	(3)
A(form=inf) ::= to-inf⇒act(A, form={inf}, words=1);	(4)
A(dest=B) ::= ptrans(A)⇒to-loc⇒place(B);	(5)
A(inst=B) ::= ptrans(A)⇒by-inst⇒vehicle(B);	(6)
A(inst=B) ::= ptrans(A)⇒with-inst⇒vehicle(B);	(7)
B(agent=A, object=C(agent=A)) ::= person(A)⇒desire(B)⇒act(C, form=inf);	(8)
B(agent=A, object=D(agent=C)) ::= person(A)⇒desire(B)⇒person(C)⇒act(D, form=inf);	(9)
B(object=D(object=C)) ::= person(A)⇒desire(B)⇒entity(C)⇒state(D, form=inf);	(10)
B(agent=A, object=possess(agent=A, object=C)) ::= person(A)⇒desire(B)⇒entity(C);	(11)

Figure 4 - A set of sample templates

A travel example demonstrates the acquisitive character of templates. Let's have a closer look at the set of semantic templates (Figure 4 - "::=" stands for sensing, and "⇒" is a sequence operator) enabling parsing partial language phrases related to the customer desires to reach named destination locations by various means of transportation (Figure 1, note that sentences 1-11 correspond to templates 1-11 from Fig. 4).

The first group of templates (1-3) specifies interpretation rules for entity names and allows identification of names *Mary*, *Tokyo* and *Qantas* as being references to a person, place and a plane respectively (cf. Fig 3, senses 14, 16, 19). Group two of templates (4) consists of a single rule identifying tenseless verb forms, e.g. *to drive* (note that a "words" slot is a meta-construct allowing the rule application over one word rather than the whole phrase). The third set of rules (5-7) enables the recognition of directional and instrumental cases of "ptrans" groups of acts, e.g. *to go to Sydney*, *to travel by bus*, or *to fly with JAL*. Finally, templates (8-11) explain the interpretation of "desire" group of phrases, e.g. *I want to travel overseas*, *My wife wants me to come to Melbourne*, *I 'd like my luggage to be sent to my hotel*, or *My children need a holiday*. (Note that our definition of "desire" is greatly simplified.)

sentence(np=A, vp=C) ::= animate(A, num=B)⇒act(C, num=B);	(1)
time-passes-quickly ::= sign("time")⇒sign("flies")⇒ sign("like")⇒sign("an")⇒sign("arrow");	(2)

Figure 5 - Syntactic and idiomatic templates

COSIMO templates are capable of defining language semantic grammars. The grammar may, however, be substantially weakened, even to the level of syntax, with the use of vague semantic categories (e.g. Figure 5, example 1). On another hand, it is also possible to write templates reflecting very narrow semantic senses, colloquialisms, and idiomatic expressions consisting of specialised semantic classes, or even in entirety of low level signs (e.g. Figure 5, example 2). Thus, COSIMO templates, formulae and dictionaries integrate the whole range of natural language phenomena, its syntax, semantics and pragmatics.

3.0 Connectionist Chart Parsing

The parsing methodology used in COSIMO is based on a chart-parsing scheme (Kay 1980, Earley 1970) accounting for its short-term memory and immediate language expression interpretation. The main characteristics of COSIMO *chart* is the fact that it integrates both syntactic and semantic interpretations of the parsed sentence, another important feature is its parallel implementation (cf. Ferguson and Thau 1986), and connectionist nature (cf. Slack 1986).

Chart. First, let us illustrate the chart architecture. Figure 6 shows a chart structure resulting from parsing a simple sentence :- "I want to fly to Sydney" (refer to Figures 2-4 for definitions of a dictionary, formulae and templates). The bottom layer of the chart is constructed with instances of dictionary-extracted formulae (e.g. "user", "desire", "air-travel", etc.), the higher levels consist of instances of either formula definitions (e.g. "travel" as "ptrans" with "plane") or

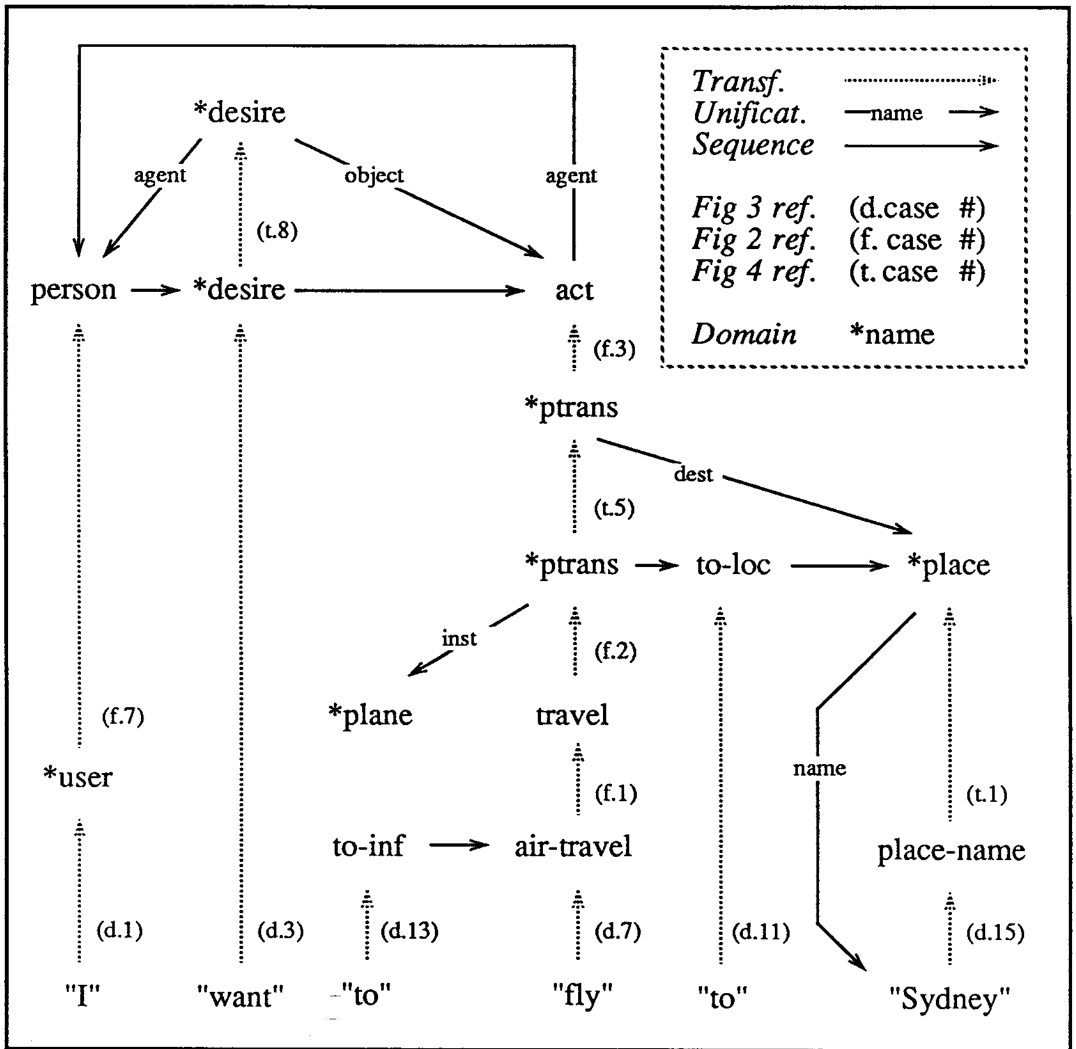

Figure 6 - A sample chart

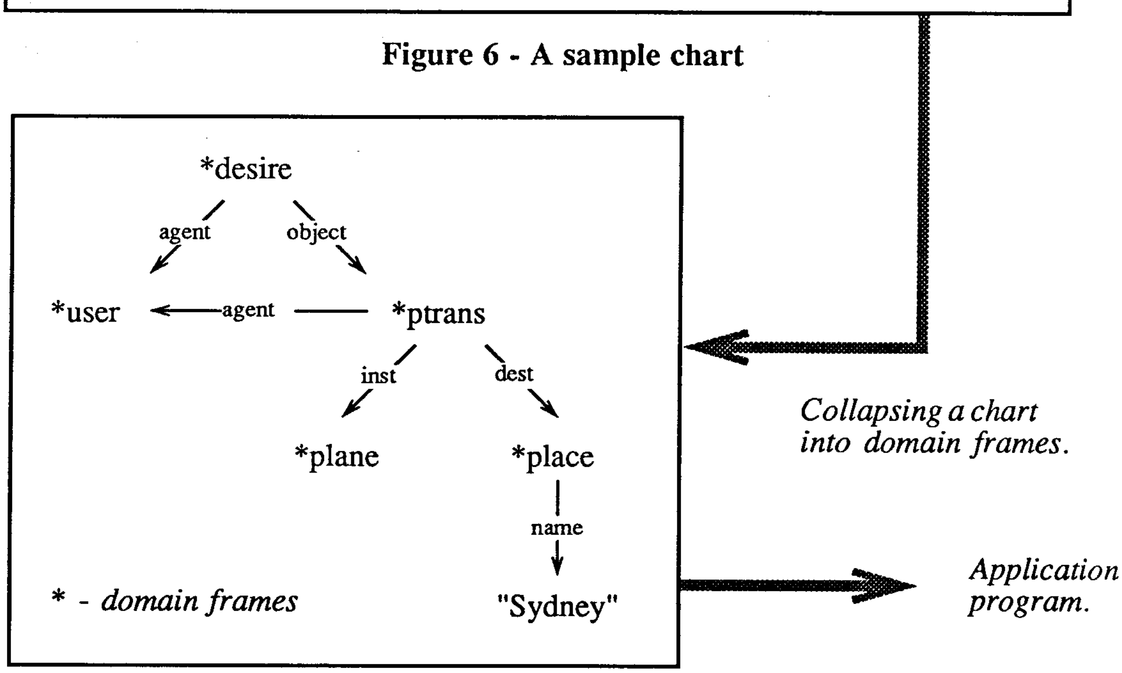

Figure 7 - Chart interpretation

templates combining their sequences (e.g. "person⇒desire⇒act" or "ptrans⇒to-loc⇒place"). The template and formula matching procedure takes into consideration type-hierarchy generalisations and specialisations (e.g. "to-inf⇒air-travel" specialised from "to-inf⇒act"). The ultimate objective of the parsing process is to construct a chart structure and then collapse it along the inheritance paths into a number of inter-related domain frames which could be passed into an application program (e.g. Figure 7).

Activation. A template and formula based chart parser could certainly account for a limited language processing, resulting in multiple interpretations of input. Such a parser, however, would have tremendous difficulties not only with exponential explosion of chart constituents, but also in dealing with certain types of linguistic phenomena requiring the use of inexact structure matching, unification and integration, e.g. understanding of anaphora, metaphor, and ill-formedness of language expressions. Fuzzyfication of representation and minimisation of structure multiplicity is achieved in COSIMO not through the deductive processes but via chart constituent activation mechanism based on an *Interactive Relaxation with Decay* algorithm (McClelland and Rumelhart 1981, Regia 1987) and inspired in particular by the VITAL parser (Howells 1988).

The model assumes all chart constituents to have associated with it a momentary energy level, represented by a real number and known as *activation*. Activations are propagated between the chart elements (or nodes) in the network of excitory and inhibitory connections, from signs up to the phrase semantic interpretations. Excitory connections strengthen the activation of neighbouring nodes and correspond to the representational associations between concepts, i.e. instance-type, frame-slot, formula-constraint, formula-template, sign-sense, or chart constituent concatenations (e.g. see all of the connections shown in Fig. 6). Inhibitory connections, on the other hand, weaken the activation of neighbouring nodes, and are related to the potentially ambiguous chart clusters, i.e. multiple sign senses, multiple formulae interpretations, multiple template associations, chart constituent overlaps, or near misses. The strength of each excitory and inhibitory link is represented with either positive or negative connection weights respectively. A chart component is said to be active when its activation level is positive. In the absence of input from its neighbours the value of a node activation slowly decays to its resting level, i.e. either zero or negative determined by the relative frequency of the constituent use (recorded with dictionary entries, templates and formulae). Over time, activation of a chart constituent a_i changes accordingly to the formula (McClelland and Rumelhart 1981) :-

$$a_i(t + \Delta) = a_i(t) - \Theta_i(a_i(t) - r_i) + \varepsilon_i(t)$$

where:

$$\varepsilon_i(t) = n_i(t)(M - a_i(t)) \; \{ \text{ for } n_i > 0 \}$$

$$\varepsilon_i(t) = n_i(t)(a_i(t) - m) \; \{ \text{ for } n_i < 0 \}$$

$$n_i(t) = \sum_j \alpha_{ij} e_j(t) - \sum_k \gamma_{ik} i_k(t)$$

the symbol explanation follows.

$a_i(t + \Delta)$	new activation level	n_i	net input to the unit
$a_i(t)$	previous activation level	m, M	activation range
Θ_i	decay rate	e_j, i_k	neighbours activation
r_i	resting level	α_{ij}, γ_{ik}	exc. & inh. weights
$\varepsilon_i(t)$	neighbours effect		

The connection weights and formula constants were adapted from VITAL, thus the weights of the sign/sense connections were assumed to be proportional to the relative element frequency (in its own class), the weight of formula and template chart associations (bottom up activations) is currently a constant inversely proportional to the number of incoming links, the subcomponent associations (top down activation) is inversely proportional to the number of constituents.

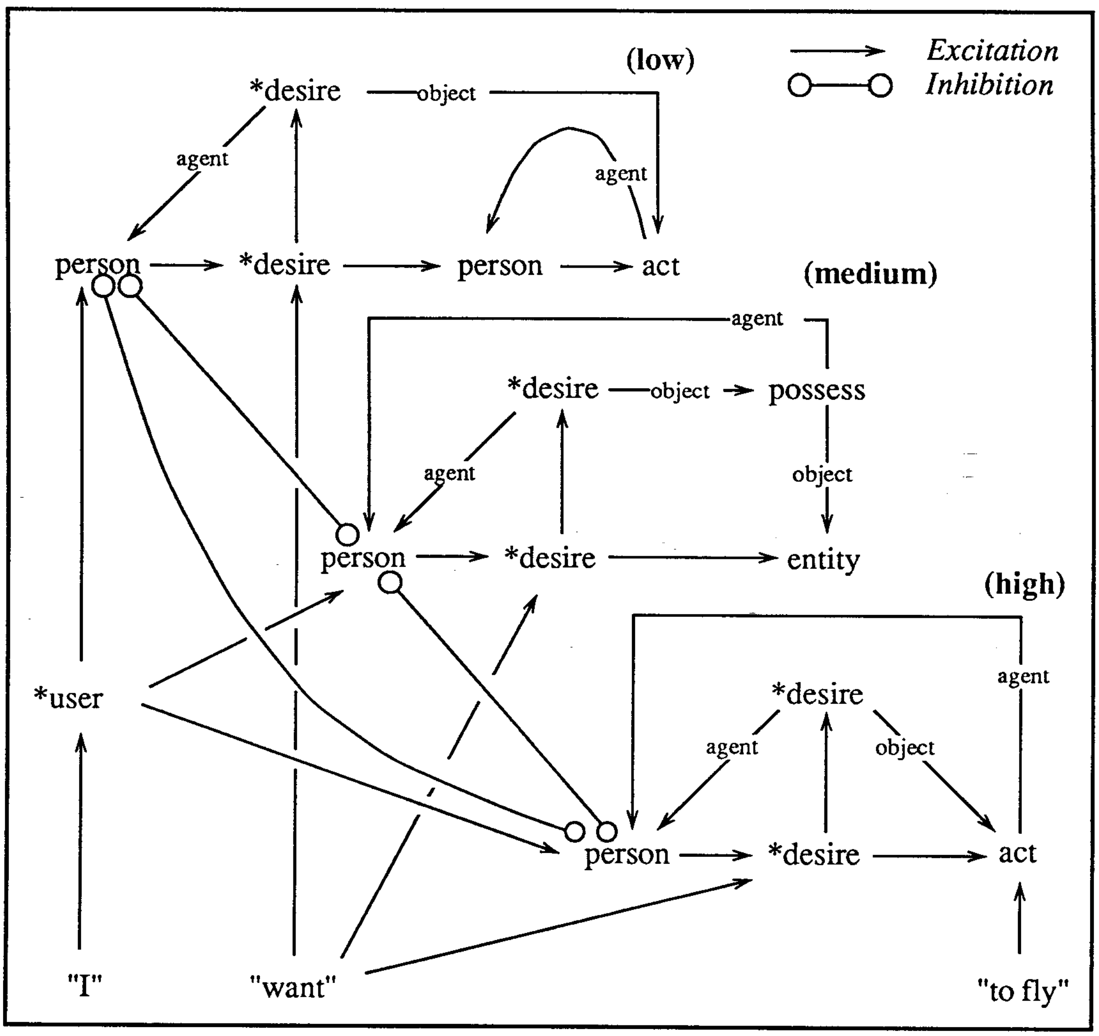

Figure 8 - Spreading activation in COSIMO chart

To allow for greater flexibility of chart parsing, we have also allowed for parsing of ill-formed input, i.e. processing of formula and template near misses. Addition of chart constituents that do not span the current input exactly, or violate the semantic constraints, causes the search for a class/subclass near miss, input token omission or token insertion. Each of the detected near misses allows for constituent insertion with the connection of reduced excitation weight. Inhibitory links are added between competing chart components, e.g. overlapping constituents, the inhibition is inversely proportional to the number of competitors.

As an example, let us consider a partial parse of the sentence from Figure 6 (see Figure 8). The initial processing of string "I want" results in addition of three competing templates (*high, medium* and *low*), only one matches the expected continuation "to fly" (*high*), the remaining two are in fact near misses to the input (*low* omits the person, *medium* is penalised for the use of incorrect semantic category). The spreading activation will prefer the structure with the highest number of shortest connections coming from input signs, thus the correctly matching template gets the highest excitation, inhibits the remaining two interpretations from receiving activation from input, then they slowly decay into their resting energy levels at which they get removed from the chart. The remaining chart constituents achieve equilibrium of up and down activation spreading after a number of chart cycles. If a network relaxes into a number of disjoint interpretations, their selection may use the level of activation as a degree of interpretation correctness.

4.0 Conclusions

We have described the working principles of the COSIMO system. The system prototype was implemented in CommonLisp and C on the Sun 3/60, and is still in its experimental stage. The tests of COSIMO were performed on a dictionary and grammar of limited size; the parser, however, exhibits enough competence to be used in its narrow application domain (i.e. travel). Early work with the COSIMO system focused on purely semantic parsing with formulae, it was shown, however, that addition of large numbers of templates greatly improves the accuracy of the text interpretation. Thus, current research concentrates on increasing the size of COSIMO's dictionary (2,000 of vague semantic entries), better coverage of the application domain with formulae (to about 200 domain specific entries), and a significant extension to the pool of available templates. As the connectionist architecture was actually simulated in Lisp, the system lacks speed and efficiency; we do, however, plan to consider the use of more efficient parallel chart algorithms (e.g. Ferguson and Thau 1986), perceptron-based dictionaries (cf. Cybulski, Ferra, Kowalczyk and Szymanski 1989 a,b), and finally parallel computer architectures and VLSI implementation of Neural Nets, all to improve COSIMO processing efficiency.

Overall, the results of even our preliminary experiments convince us that the addition of connectionist chart methods to symbolic NLU, and Preference Semantics in particular, greatly enhances the capability of natural language parsers.

5.0 Acknowledgements

We acknowledge the permission of Executive General Manager of Telecom Australia Research Laboratories to publish this paper.

6.0 References

Anderson, J.R. (1983): "A spreading activation theory of memory," *Journal of Verbal Learning and Verbal Behavior*, V22, 261-295. Also in Collins and Smith 1988, 137-154.

Arbib, M.A. and Caplan, D. (1979): "Neurolinguistics must be computational," *The Behavioral and Brain Sciences*, V2, 449-483.

Bengio, Y., Cardin, R., DeMori, R., and Merlo, E. (1989): "Programmable execution of multi-layered networks for automatic speech recognition," *CACM*, V32, N2, 195-199.

Bookman, L.A. (1987): "A microfeature based scheme for modelling semantics," *IJCAI 87*, 611-614.

Bruce, B. and Moser, M.G. (1987): "Grammar, Case," in S.C. Shapiro (ed.) *Encyclopedia of Artificial Intelligence*, New York, Wiley and Sons, 333-339.

Collins, A. and Smith, E.E. (eds 1988): *Readings in Cognitive Science*, San Mateo, California, Morgan Kaufmann Publishers.

Collins, A.M. and Loftus, E.F. (1975): "A spreading activation theory of semantic processing," *Psychological Review*, V82, 407-428. Also in Collins and Smith 1988, 126-136.

Cottrell, G.W. and Small, S.L. (1983): "A connectionist scheme for modelling word sense disambiguation," *Cog. and Brain Theory*, 6(1), 89-120.

Cybulski, J.L., Ferra, H.L., Kowalczyk, A., and Szymanski, J. (1989a): "Determining Word Lexical Categories with a Multi-Layer Binary Perceptron," *Proc. Conf. IASTED Expert Systems and Neural Networks*, Honolulu, Hawaii; also Telecom Australia Research Laboratories, CSS Branch Paper 182.

Cybulski, J.L., Ferra, H.L., Kowalczyk, A., and Szymanski, J. (1989b): "Experiments with Multi-Layer Perceptrons," to appear in *Proc. Conf. AI'89*, Melbourne, Australia; also Telecom Australia Research Laboratories, CSS Branch Paper 181.

Dreyfus, H.L. and Dreyfus, S.E. (1986): *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, New York, The Free Press, A Division of Macmillan, Inc.

Earley, J. (1970): "An efficient context-free parsing algorithm," *CACM*, V13, N2, 415-462. Also in Grosz, Spark-Jones and Webber 1986, 25-34.

Engelmore, R. and Morgan, T. (eds 1988): *Blackboard Systems*, Reading, Massachusetts, Addison-Wesley Publishing Co.

Fass, D. and Wilks, Y. (1983): "Preference Semantics, ill-formedness, and metaphor," *AJCL*, V9, N3-4, 178-187.

Ferguson, S. and Thau, R. (1986): "Context-free parsing on the Connection Machine system," Thinking Machine Corp.

Gasser, M.E. (1988): "Connectionist model of sentence generation in a first and second language," University of California, Computer Science Department, Report CSD-880050.

Grosz, B.J., Sparck-Jones, K., and Webber, B.L. (eds 1986): *Readings in Natural Language Processing*, Los Altos, California: Morgan Kaufmann Publishers, Inc.

Hillis, W.D. and Steele, G.L. Jr. (1986): "Data parallel algorithms," *CACM*, V29, N12, 1170-1183.

Howells, T. (1988): "VITAL: A connectionist parser," *Proc. Conf. Cog. Sci.* 88, 18-25.

Jennings, A., Rowles, C.D. and Kowalczyk, A. (1989): "Natural language understanding in the MEDICI project," *Proc. Int. Conf. on Comp. and Inf.*, Toronto, Canada 1989.

Jones, M.A. (1987): "Feedback as a coindexing mechanism in connectionist architectures," *IJCAI 87*, 602-610.

Kay, M. (1980): "Algorithm schemata and data structures in syntactic processing," CSL-80-12, Xerox Corporation. Also in Grosz, Sparck-Jones and Webber 1986, 35-70.

Kukich, K. (1987): "Where do phrases come from: Some preliminary experiments in connectionist phrase generation," in Kempen, G. (ed.) *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics, Proc. of the NATO Advanced Research Workshop*, 405-421.

Leckie, C.A. and Zukerman, I. (1989): "Planning in co-operative domains: A case study," Telecom Research Laboratories, Report CSS-198.

Lesser, V.R. and Erman, L.D. (1977): "A retrospective view of Hearsay-II architecture," *Proc. IJCAI-77*, 790-800. Reproduced in Engelmore and Morgan 1988, 87-124.

Lippman, R.P. (1987): "An introduction to computing with neural nets," *IEEE ASSP Magazine*, April, 4-22.

Lytinen, S.L. (1986): "Dynamically combining syntax and semantics in natural language processing," *AAAI 86*, 574-578.

Mamdani, E.H. and Gaines, B.R. (1981): *Fuzzy Reasoning and its Applications*, London, Academic Press.

McClelland, J.L. and Rumelhart, D.E. (1981): "An interactive activation model of context effects in letter perception," *Psychological Reviews* 88, 375-407.

McClelland, J.L., Rumelhart, D.E. and Hinton, G.E. (1986): "The appeal of parallel distributed processing," in Rumelhart, D.E., McClelland, J.L., and the PDP Research Group (eds), *Parallel Distributed Processing*, V1, 3-40. Also in Collins and Smith 1988, 52-72.

Minsky, M. (1986): *The Society of Mind*, New York: Simon and Shuster.

Parisi, D. and Nolfi, S. (1987): "Connectionist modelling of syntactic constraints on sentence processing, " *Proc. IEEE First Int. Conf. on Neural Networks*, V2, 505-512.

Pollack, J.B. and Waltz, D.L. (1984): "Parallel Interpretation of Natural Language," *Proc. Int. Conf. 5th Gen. Comp. Sys.*, 686-691.

Reggia, J.A. (1987): "Properties of competition-based activation mechanism in neuromimetic network models," *Proc. IEEE First Int. Conf. on Neural Networks*, V2, 131-138.

Schank, R.C. (1979): "Interestingness: Controlling inferences," *Artificial Intelligence*, V12, 273-297.

Schank, R.C., Lebowitz, M. and Birnbaum, L. (1980): "An integrated understander," *AJCL*, V6, N1, 13-30.

Slack, J.M. (1986): "Distributed memory: A basis for chart parsing," *COLING'86*, 476-481.

Small, S.L. (1987): "A distributed word-based approach to parsing," in L. Bolc (ed.) *Natural Language Parsing Systems*, Springer-Verlag, 161-202.

Stanfill, C. and Waltz, D.L. (1986): "Toward memory-based reasoning," *CACM*, V29, N12, 1213-1228.

Waltz, D.L. (1986): "Connectionist models: Not just a notational variant, not a panacea," *TINLAP 3*, 56-62.

Wilks, Y. (1975): "An intelligent analyser and understander of English," *CACM* 18, 5,264-274. Also in Grosz, Sparck-Jones and Webber 1986, 193-204.

Woods, W.A. (1980): "Cascaded ATN Grammars," *AJCL* 6, 1, 1-12.

Zadeh, L.A. (1981): "PRUF - a meaning representation language for natural languages," in Mamdani and Gaines 1981, 1-66.