A Temporal Extension to the Parsimonious Covering Theory

Alexandre de Melo Rezende and Jacques Wainer
{arezende,wainer}@dcc.unicamp.br

Instituto de Computação, UNICAMP

Abstract. In this work we propose a temporal extension to the Parsimonious Covering Theory (PCT). PCT provides a theoretical foundation for the diagnostic reasoning process as an abductive reasoning based associations between causes with their consequences. Our temporal extension of PCT allows one to associate to a disease a temporal evolution of its symptoms.
The elimination of temporally inconsistent hypotheses minimizes one of the greatest problems of PCT: the solution for a particular diagnsotic problem may include a large number of alternative hypotheses. Furthermore, the inclusion of temporal aspects to an extension of PCT that includes probabilistic information also eliminates the problems of incorrectly rejecting hypotheses if a necessary symptom has not yet occurred.

1 Introduction

Diagnostic reasoning, that is, finding causes that explain observed symptoms, is one of the major application areas for knowledge-based systems. In some domains, time is an important aspect of the diagnostic reasoning itself: knowing when a symptom occurred may be as an important information as knowing just that the symptom did occur. In some medical diagnostic applications, temporal information about the occurrence of the symptoms is vital for a correct diagnostic and some "second generation medical expert systems" [Con89, Ham87] tried to deal with this aspect.

On the other hand, the diagnostic reasoning itself, independently of temporal considerations, does not have an agree upon theoretical foundation. [Pen90] proposed such a theoretical foundation under the name of Parsimonious Covering Theory (PCT). But PCT is still a theory of reasoning about static symptoms, in the following sense. PCT is based on a model that associates to each disease a set of symptoms it may cause. Thus, PCT assumes that, at the moment of diagnostic, all symptoms are observable and that the order of occurrence of these symptoms is irrelevant for the diagnostic.

This work extends the PCT model in such a way that to each disease one associates evolutions of symptoms (or sets of possible histories of symptoms). Thus, at diagnostic time, one will not just describe the symptoms present, as one would in a static diagnostic system, but describe the whole evolution of the symptoms. Even symptoms that are no longer present may be relevant for the diagnostic process.

This paper is organized as follows: section 2 briefly defines the simplest version of PCT. Section 3 extends this theory to incorporate temporal knowledge. Section 4 presents the solution of the problems that may arise when probabilistic knowledge is incorporated into PCT. Finally, section 5 presents the conclusions and the limitations of proposed temporal reasoning.

2 Basics of Parsimonious Covering Theory

First we will briefly introduce the PCT. In the basic version of PCT [Pen90], one uses two finite sets to define the scope of diagnostic problems (see Figure 1). They are the set D, representing all possible **disorders** d_l that can occur, and the set M, representing all possible **manifestations** m_j that may occur when one or more disorders are present.

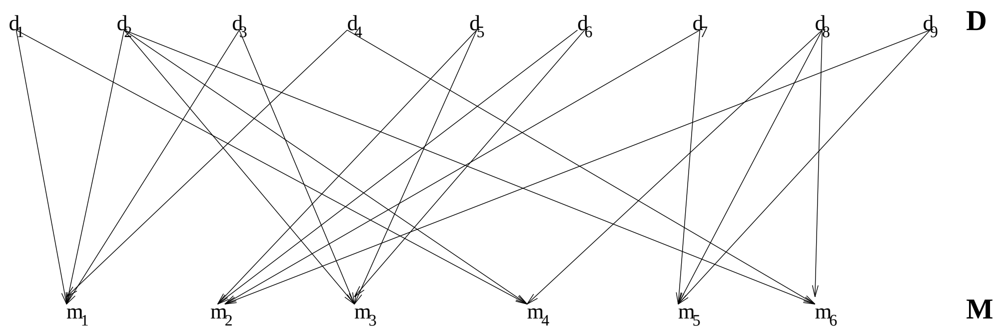

Fig. 1. Causal network of a diagnostic knowledge base $KB = \langle D, M, C\rangle$.

To capture the intuitive notion of causation, one uses the relation C, from D to M, that associates to each individual disorder its manifestations. An association $\langle d_l, m_j\rangle$ in C means that d_l may directly cause m_j; it does *not* mean that d_l *necessarily* causes m_j. The sets D, M, and C together are the knowledge base (KB) of a diagnostic problem.

To complete the problem formulation we need a particular diagnostic **case**. We use M^+, a subset of M, to denote the set of **observations**, that is, manifestations that are present in the case. The set M^+ does not necessarily have to be specified all at once at the beginning of problem-solving; it can be gradually obtained from the answers to questions asked by the diagnostic system.

Definition 1. A **diagnostic problem** P is a pair $\langle KB, Ca\rangle$ where:

- $KB = \langle D, M, C\rangle$ is the **knowledge base**, composed of
  - $D = \{d_1, d_2, \ldots, d_n\}$ is a finite, non-empty set of objects, called *disorders*;
  - $M = \{m_1, m_2, \ldots, m_k\}$ is a finite, non-empty set of objects, called *manifestations*;

- $C \subseteq D \times M$ is a relation called *causation*

– $Ca = \langle M^+ \rangle$ is the case, and $M^+ \subseteq M$ is the set of observations.

For a diagnostic problem P, it is convenient and useful to define the following sets of functions based on relation C:

Definition 2. For any $d_l \in D$ and $m_j \in M$ in a diagnostic problem P

– $effects(d_l) = \{m_j | \langle d_l, m_j \rangle \in C\}$, the set of manifestation directly caused by d_l;
– $causes(m_j) = \{d_l | \langle d_l, m_j \rangle \in C\}$, the set of diseases which can directly cause m_j.

The set $effects(d_l)$ represents all manifestations that may be caused by disorder d_l, and $causes(m_j)$ represent all disorders that may cause manifestation m_j. These functions can be easily generalized to have sets as their arguments.

2.1 Solution for Diagnostic Problems

In order to formally characterize the solution of a diagnostic problem one needs to define the notion of "cover", based on the causal relation C, to define the criterion for parsimony, and to define the concept of an explanation (explanatory hypothesis).

Definition 3. The set $D_L \subseteq D$ is said to be a **cover** of $M_J \subseteq M$ if $M_J \subseteq effects(D_L)$.

Definition 4. A set $E \subseteq D$ is said to be an **explanation** of M^+ for a diagnostic problem iff E covers M^+, and satisfies a given parsimony criterion.

In the following definition we present the possible parsimony criteria:

Definition 5.
(1) A cover D_L of M_J is said to be **minimum** if its cardinality is the smallest among all covers of M_J.
(2) A cover D_L of M_J is said to be **irredundant** if none of its proper subsets is also a cover of M_J; it is **redundant** otherwise.
(3) A cover D_L of M_J is said to be **relevant** if it is a subset of $causes(M_J)$; it is **irrelevant** otherwise.

[Pen90] uses irredundancy as the preferable choice for the parsimonious criteria and in this paper we follow that choice. Acording to the authors, minimality, which is another usual criteria of parsimony, should be seen more as a domain specific heurisic than a general criteria.

In many diagnostic problems, one is generally interested in knowing all plausible explanations for a case rather than just a single explanation because they, as alternatives, can somehow affect the course of actions taken by the diagnostician. This leads to the following definition of the problem solution:

Definition 6. The **solution** of a diagnostic problem $P = \langle KB, Ca \rangle$, designated $Sol(P)$, is the set of all explanations of M^+.

Example 1. In the Figure 1, $\{d_1\}$, $\{d_2\}$, $\{d_3, d_8\}$, and $\{d_4, d_8\}$ are the only plausible explanations (i.e. irredundant covers) for $M^+ = \{m_1, m_4\}$, and therefore they are the solution of the problem.

2.2 Algorithms and Problems of PCT

[Pen90] presents the algorithm `bipartite` which incrementally and constructively compute the set of solutions of a diagnostic problem P. The algorithm processes one observation (m_j) from M^+ at a time, and incorporates $causes(m_j)$ to the set of explanations it has computed so far.

The main problem with the basic version of PCT is that the solution of a problem tends to have many alternative explanations. The reason is that irredundancy is too weak a criteria to significantly reduce the number of alternative explanations. For most practical applications a further processing to filter out some of the explanations based on domain specific heuristics or at least to order the set of explanations so that more "plausible" explanations are presented before less "plausible" ones.

A more complex version of PCT (called probabilistic causal model) is also presented in [Pen90] which incorporates probabilities to the links between a disease and its manifestations, that is, the probability that the manifestation will occur provided that the disease is present. This probabilistic information allow one to rank the explanations. Furthermore, this probabilistic information allows one to filter from the set of all explanation those that contain a disease for which a necessary manifestation was not present in the case. If a disease d_i necessarily causes a manifestation m_j, that is, the probability that m_j is present given d_i is 1, then if m_j is known not to be among the observations of the case, then one can remove explanations that contain that disease. This is called categorical rejection.

3 Parsimonious Covering Theory and Time

The aim of this research is to extend PCT so that instead of associating to each disease a set of manifestation, one could associate an evolution of manifestations. Thus, the database could state that disease d_i causes first m_1 which will last between 2 and 5 days, followed in 2 to 3 days by m_2 which may last an undetermined amount of time, and so on. We accomplish this temporal representation using a graph, where vertices are manifestations and directed arcs between vertices represent temporal precedence. If there is quantitative information about the duration of the manifestation, it is associated with the corresponding node; if there is quantitative information about the elapsed time between the start of two manifestations, it is associated with the corresponding arc. Furthermore, quantitative information are not represented as a single number, but as an interval. Therefore one can state that a manifestation will follow another in 2 to 3 days. To each disease one associates one such temporal graph.

3.1 Dynamic Diagnostic Problem Formulation

Time points will be primitive objects to represent temporal information. *Intervals* are defined as non-empty convex sets of time points (points on the time line), represented by $I = [I^-, I^+]$ such that I^- and I^+ are the extreme points of interval I, respectively (obviously $I^- \leq I^+$; $I^- > I^+$ indicates an empty interval I). We use the following notations for operations on intervals:

- $I + J = [I^- + J^-, I^+ + J^+]$;
- $I \cap J = [max(I^-, J^-), min(I^+, J^+)]$;
- $I \leq p \Rightarrow I^+ \leq p$, where p is a time point.

A temporal graph is a *direct*, *acyclic*, *transitive*, and not necessarily connected graph. The existence of an arc from m_i and m_j in a temporal graph denotes the fact that the beginning of the occurrence of manifestation m_i must precede the beginning of the occurrence of m_j.

The *temporal distance* between manifestations and the *duration* of a manifestation are represented by functions on the graph, denoted by $DIST$ and DUR, respectively. The temporal distance function $DIST$ associates an interval $R = [R^-, R^+]$ to each arc of a temporal graph G_l. $DIST(G_l, (m_i, m_j)) = R$ for $(m_i, m_j) \in A_l$, which we will abbreviate as $DIST_l(m_i, m_j) = R$, states that the difference between the time of the beginning of m_j and the beginning of m_i in the temporal graph G_l of d_l must be within the interval R. The duration function DUR associates to each vertex m_i of a temporal graph G_l em G an interval J, that specifies that the duration of m_i must be within the interval J.

The transitivity of the temporal graph must be consistently carried over to the $DIST$ function: if $DIST_l(m_i, m_j) = R_1$ and $DIST_l(m_j, m_k) = R_2$ then $DIST_l(m_i, m_k) = R_1 + R_2$.

Figure 2 illustrates the temporal information about the disorders d_8 and d_9 of the diagnostic problem shown in Figure 1.

Definition 7. The **temporal graph** of a disorder $d_l \in D$, $G_l = (V_l, A_l)$, is a direct, transitive and acyclic graph defined as:

- $V_l \subseteq M \equiv$ set of objects directly caused by d_l, and
- $A_l = \{(m_i, m_j) |$ the beginning of m_i occurs *before* the beginning of m_j when the disorder d_l is said to be present$\}$.

Definition 8. The **knowledge base** of a dynamic diagnostic problem is the tuple $KB = \langle D, M, G, DIST, DUR \rangle$ where D and M are defined as before, G is a set of temporal graph, each one associated with one disease of D, $DIST$ and DUR are the temporal information functions defined above.

$DIST_8(m_4, m_5) = [1, 3]$
$DIST_8(m_5, m_6) = [2, 5]$
$DIST_8(m_4, m_6) = [3, 8]$
$DUR_8(m_6) = [1, 2]$

$DIST_9(m_2, m_5) = [2, 5]$
$DUR_9(m_2) = [10, 15]$
$DUR_9(m_5) = [7, 13]$

Fig. 2. Temporal graphs of the disorders d_8 and d_9 with their temporal distance functions and duration functions.

In order to represent the case, we will need the set of observations M^+, as before, and the temporal information about these observations. The function BEG^+ associates an interval to some of the observations in M^+. $BEG^+(m_j) = I$, $m_j \in M^+$, states that m_j started at any time within interval I. Similarly, the function DUR^+ associates to some of the observations in M^+ an interval, such that the duration of the observation was anything within that interval.

Definition 9. A dynamic diagnostic problem P is a pair $\langle KB, Ca\rangle$ where KB is defined as before, and $Ca = \langle M^+, BEG^+, DUR^+\rangle$ is the case.

One can define the $effects$ and $causes$ functions in a similar way to definition 2. For example $causes(m_j) = \{d_l | m_j \in V_l$, for any temporal graph $G_l = (V_l, A_l) \in G\}$, represents the set of diseases that may cause m_j.

It is important to notice that the temporal model allows for many forms of uncertainty and incompleteness of information both at the case and at the knowledge base. For example, temporal information about the case need not to be precise but can be stated as an interval, or can be ommited althogether. In the knowledge base, not all manifestations need to be temporally related to the others: the graph need not to be connected, nor do all arcs need to have intervals associated with them. The theory uses the temporal information if it is available, otherwise it behaves as the basic PCT.

3.2 Solutions for Dynamic Diagnostic Problems

In order to define a solution for a diagnostic problem, we need to define a set of concepts about temporal inconsistency. This will eventually allow one to remove the explanations that contain diseases in which the evolution of manifestations contradicts the evolution of the observations in the case. For example, if for a certain disease m_1 precedes m_2 but in the case, the occurrence of m_1 started after the occurrence of m_2, then one can disregard all explanations that contain such disease, since it contradicts with the temporal information in the case.

Definition 10. For a dynamic diagnostic problem P let $G_l = (V_l, A_l) \in G$, $(m_i, m_j) \in A_l$, $DIST_l(m_i, m_j) = R$, $m_i, m_j \in M^+$, $BEG^+(m_i) = I_{m_i}$ and

$BEG^{+}(m_j) = I_{m_j}$. The **arc** (m_i, m_j) **is temporally inconsistent with respect to the case** iff $(I_{m_i} + R) \cap I_{m_j} = \emptyset$.

The resulting interval of operation $(I_{m_i} + R)$ corresponds to a set of valid values for the beginning of m_j. Thus if the intersection of this interval and I_{m_j} ("real" valid interval for the beginning of m_j) is empty, then the arc (m_i, m_j) is temporally inconsistent with the case. The inconsistency criterion defined above is equivalent to one described in [Con93].

Definition 11. For a dynamic diagnostic problem P let $G_l = (V_l, A_l) \in G$ the temporal graph of a disorder $d_l \in D$. The **disorder** d_l **is temporally inconsistent with the case** $Ca = \langle M^{+}, BEG^{+}, DUR^{+} \rangle$ iff

- exist at least one arc $(m_i, m_j) \in A_l$ temporally inconsistent with respect to the case , or
- exist at least a vertex $m_j \in V_l$, such that, $m_j \in M^{+}$ and $DUR_l(m_j) \cap DUR^{+}(m_j) = \emptyset$.

Finally, based on the above definitions, we formalize the notions of temporally consistent explanation and temporally consistent solution.

Definition 12. A set $E \subseteq D$ is said to be a **temporally consistent explanation of the case** for a dynamic diagnostic problem P iff

- E covers M^{+}, and
- E satisfies a given parsimony criterion, and
- for any $d_l \in E$, d_l is *not* temporally inconsistent with the case.

Definition 13. The **temporally consistent solution** of a dynamic diagnostic problem $P{=}\langle\langle D, M, G, DIST, DUR\rangle, \langle M^{+}, BEG^{+}, DUR^{+}\rangle\rangle$, designated by $Sol(P)$, is the set of all temporally consistent explanations of the case $\langle M^{+}, BEG^{+}, DUR^{+}\rangle$.

Algorithm Solution: Basic Ideas We implemented an algorithm that solves a temporal diagnostic problem. Due to space limitations, we will only present the basic ideas of the algorithm and briefly discuss an example. The full algorithm can be found in [Rez96]. The important aspect of the algorithm is that temporal consistency is not implemented as a filter, that is, it is not applied after the original `bipartite` algorithm has generated the solutions, but it is incorporated very early into the process of merging the causes on the "new" observation into the set of current explanations. Thus the algorithm has to deal with smaller sets of explanations.

At the beginning of a new cycle, after a new observation has been entered, the disorders evoked by the new observation are checked for temporal consistency with the case information so far. Then hypotheses that contain the temporally inconsistent disorders are eliminated from the set of current hypotheses and the new temporally consistent evoked disorders are used to update the set of hypothesis.

The example below illustrates the basic ideas of the algorithm. For example in Figure 1, we have that $S_1 = \{\{d_1\}, \{d_2\}, \{d_3, d_8\}, \{d_4, d_8\}\}$ is the set of all explanations (irredundant covers) of $M^+ = \{m_1, m_4\}$ temporally consistent with $BEG^+(m_4) = [10, 10]$ and $DUR^+ = \emptyset$. Each time a new observation is discovered and temporal information is available for it, the system verifies the temporal consistency of the hypotheses in S_1, and update the hypotheses in the correct way. Thus, consider m_5 new observation of M^+, and $BEG^+(m_5) = [16, 18]$ and $DUR^+(m_5) = [2, 3]$. First, we obtain the disorders evoked by m_5 (i.e. $causes(m_5) = \{d_7, d_8, d_9\}$) that are temporally inconsistent with $BEG^+(m_5)$ and $DUR^+(m_5)$. As an illustration, consider d_8 and d_9 in Figure 2. Disorder d_8 is temporally inconsistent because the arc (m_4, m_5) with label $[1, 3]$ is inconsistent with $BEG^+(m_4)$ and $BEG^+(m_5)$. On the other hand, disorder d_9 is temporally inconsistent because the duration of m_5 in d_9 is inconsistent with $DUR^+(m_5)$. In the next step, we remove all explanations in S_1 that contain these temporally inconsistent disorders. Thus, $S_2 = \{\{d_1\}, \{d_2\}\}$ is the set of all explanations that are not inconsistent with this new information. Finally, the consistent disorders (only d_7 in this case) are used to update the current explanations. $S_3 = \{\{d_1, d_7\}, \{d_2, d_7\}\}$ is thus the set of all explanations temporally consistent with the case (with m_5 added). If no other manifestation is present than S_3 represents the temporally consistent solution.

4 Categorical Dynamic Diagnostic Problems

We mentioned that the basic PCT can be extended so that probabilities can be associated to each manifestation in a disease. This allows one to eliminate a disease in the absence of an observation if that manifestation is necessary for the disease. But when time is added, this categorical rejection may pose some problems. It may happen that a necessary manifestation was not observed because there was not enough time for it to occur. Thus, some form of temporal reasoning must be performed in order to ascertain whether a disease can be categorically rejected.

4.1 Problem Formulation and its Solutions

In this paper we are not interested in a general probabilistic (numeric) information relating manifestations and disorders, but just some information whether the disorder necessarily causes the manifestation, or whether the causation is only possible. Thus, in the knowledge base KB we have to add a function $POSS$ that attributes to each vertex of each temporal graph either the label N, for necessary, or the label P, for possible. Thus, $POSS(G_l, m_j) = N$, abbreviated as $POSS_l(m_j) = N$, states that disorder d_l necessarily causes the manifestation m_j.

For categorical diagnostic problems, one is interested in manifestations known to be absent in the case, called **negative observations**. Thus we add, M^-, the set of negative observations, and I_{now}, the time point that represents the moment of diagnosis, to M^+, BEG^+, DUR^+ as the components of the case Ca.

Definition 14. Let $P{=}\langle KB, Ca\rangle$ be an open dynamic diagnostic problem and $G_l = (V_l, A_l) \in G$. The **disorder d_l is categorically inconsistent with the case** iff

- exist an arc (m_j, m_k) in A_l, such that, $POSS_l(m_j) = N$, $m_j \in M^-$ e $m_k \in M^+$, or
- exist an arc (m_i, m_j) in A_l, such that, $POSS_l(m_j) = N$, $m_j \in M^-$, $m_i \in M^+$ and $BEG^(m_i) + DIST(m_i, m_j) \leq I_{now}$.

The definition above has two conditions. For both of them, the disorder d_l is categorically inconsistent due to the combination of two factors: a necessary manifestation is not present ($POSS_l(m_j) = N$ and $m_j \in M^-$) and there has been enough time for it to happen. In the first condition the second factor is warranted because a later manifestation has already occurred ((m_j, m_k) in A_l and $m_k \in M^+$). In the second one, this factor is warranted because all values of a set of valid values (time points) for the beginning m_j are lower or equal than the actual instant ($BEG^+(m_i) + DIST_l(m_i, m_j) \leq I_{now}$). The categoric rejection problem occurs when one considers only the first factor above as a sufficient condition to classify a disorder as categorically inconsistent with the case.

Finally, we define an explanation of a categorical dynamic diagnostic problem.

Definition 15. A set $E \subseteq D$ is said to be a **consistent explanation** of the case for an open dynamic diagnostic problem $P = \langle KB, Ca\rangle$ iff

- E covers M^+, and
- E satisfies a given parsimony criterion, and
- for any $d_l \in E$, d_l is *not* temporally inconsistent, and
- for any $d_l \in E$, d_l is *not* categorically inconsistent.

As with the basic PCT plus time, we developed an algorithm that solves categorical dynamic diagnostic problems. It can be found in [Rez96].

5 Conclusion

This paper presented temporal and categorical extensions of the parsimonious cover theory which could serve as the core of a diagnostic system specially for problems where the temporal evolution of the manifestations is an important aspect.

5.1 Implementation

We developed a small example of a medical diagnostic system as a test for the theory developed herein. This diagnostic system deals with food-borne diseases [Man90] which is a domain of application where temporal information is very important. The domain included 28 diseases and 60 different symptoms. Because of the simplicity of the PCT model of diagnostic, in which no heuristic information needs to be included into the knowledge base, the whole diagnostic system was developed in two days, mainly from textbook information [Man90], with only one consultation to a specialist to resolve ambiguities in the text. The results of diagnostic cases were also verified by the specialist.

With the introduction of the temporal and categorical extension there was a significant reduction on both the number of hypotesis in the solution and the time to compute them, as compared to the basic PCT. In a particular case, the number of hypothesis in the solution was reduced from 73 to 2, and the time to compute the solution was reduced by 70 %.

5.2 Limits of the extension of PCT

The main limitation of the theory refers to multiple simultaneous disorders. The PCT assumes that multiple disorders that cause the same manifestation do not interfere with each other. That is, if both d_i and d_j cause m_k then they can both be part of an hypothesis that explains the observation m_k. Unfortunately, in the presence of temporal information is very unlikely that two disorders will not interfere with each other. As an example, let us suppose that d_i causes m_k with duration I and d_j causes m_k with duration J. Then certainly the presence of both disorders simultaneously will cause some change on the duration of m_k (the same can be true for the temporal relation of m_k with other manifestations in both d_i and d_j). This has been documented in the area of medical diagnostics [Pat81]. PCT, and therefore our extension to it, cannot represent and deal with this interference.

References

[Con89] L.Console, A. J. Rivolin, D. T. Dupré, and P. Torasso, Integration of causal and temporal reasoning in diagnostic problem solving, *in:Proc. 9th International Workshop on Expert Systems and Their Applications (Conference on 2nd Generation Expert Systems)*, (1989).

[Con93] L. Console and P. Torasso, Temporal constraint satisfaction on causal models, *Information Sciences*, 68 (1993) 1-32.

[Ham87] I. Hamlet and J. Hunter, A representation of time for medical expert systems, *Lecture Notes Med. Informatics*, 33 (1987) 112-119.

[Man90] G. L. Mandell, R. G. Douglas and J. E. Bennett, *Principles and pratice of infectious diseases* 3ed. (1990).

[Pat81] R. S. Patil, Causal representation of patient illness for eletrolyte and acid-base diagnosis, *MIT Laboratory for Computer Science*, Technical Report, MIT/LCS/TR-267, (1981).

[Pen90] Y. Peng and J. A. Reggia, *Abductive Inference Models for Diagnostic Problem-Solving*, Springer-Verlag, (1990).

[Rez96] A. de M. Rezende, Uma Extensão Temporal à Teoria das Coberturas Parcimoniosas. *Master Thesis (in portuguese)* Instituto de Computação, UNICAMP, Jul. (1996).

This article was processed using the LaTeX macro package with LLNCS style