KRR8: Diagnosis and the rest...

Yannick Pencolé
Yannick.Pencole@anu.edu.au

29 Mar 2005

Outline

Reiter Diagnosis

Definition

A Reiter Diagnosis for an observed system $(SD, COMP, OBS)$ is a minimal set $\Delta \subset COMP$ such that:

$$SD, OBS, \{\neg Ab(c), c \in COMP \setminus \Delta\}, \{Ab(c), c \in \Delta\}$$

is satisfiable.

Theorem

A Reiter Diagnosis is equivalent to a Minimal Diagnosis.

Another represenetation

An R-diagnosis is seen as a set of components and not a logical sentence. The representation are equivalent.

Reiter Diagnosis: example

Example

Davis circuit

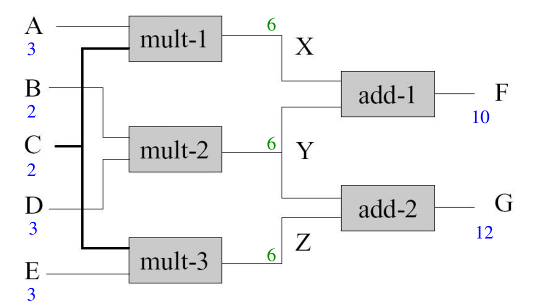

If $OBS = In1(m1, 3), In2(m1, 2), \ldots, Out(a2, 12)$ there are 4 R-diagnoses,

$$\{m1\}; \{a1\}; \{m2, m3\}; \{m2, a2\}$$

The R-diagnosis $\{m1\}$ is equivalent to the minimal diagnosis
$\neg Ab(a1) \wedge \neg Ab(a2) \wedge Ab(m1) \wedge \neg Ab(m2) \wedge \neg Ab(m3)$.

A new example: an additionner

Example

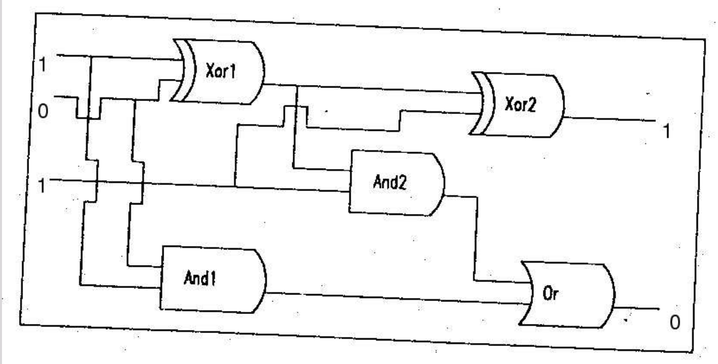

A new example: an additionner

Example

SD (behavioural model):

- $AND(x) \wedge \neg Ab(x) \Rightarrow Out(x) = and(In1(x), In2(x))$
- $OR(x) \wedge \neg Ab(x) \Rightarrow Out(x) = or(In1(x), In2(x))$
- $XOR(x) \wedge \neg Ab(x) \Rightarrow Out(x) = xor(In1(x), In2(x))$
- $AND(A1); AND(A2), OR(O1), XOR(X1); XOR(X2)$

SD (structural model):

- $Out(X1) = In2(A2)$ ...

Observations:

- $In1(X1) = 1; In2(X1) = 0; In1(A2) = 1; Out(X2) = 1; Out(O1) = 0.$

R-diagnoses:

- $\{X1\}; \{X2, O1\}; \{X2, A2\}$

Properties of R-Diagnoses

Theorem

$\emptyset$ is the only R-diagnosis for $(SD, COMP, OBS)$ iff

$$SD, OBS, \{\neg Ab(c), c \in COMP\}$$

is satisfiable.

Theorem

$\Delta \subseteq COMP$ is a R-diagnosis iff it is a minimal set such that:

$$SD, OBS, \{\neg Ab(c), c \in COMP \setminus \Delta\}$$

is satisfiable.

How to compute R-diagnoses?

Outline

R-conflicts

Definition

An R-conflict C is a set $\{c_1, c_2, \ldots, c_k\}$ with $c_i \in COMP$ such that:

$$SD, OBS, \{\neg Ab(c), c_i \in C\}$$

is not satisfiable.

An R-conflict is a set of components $C \subseteq COMP$ which cannot be together in a normal state.

Definition

An R-conflict is minimal iff there is no strict subset which is also an R-conflict.

R-conflicts: Example 1

Example

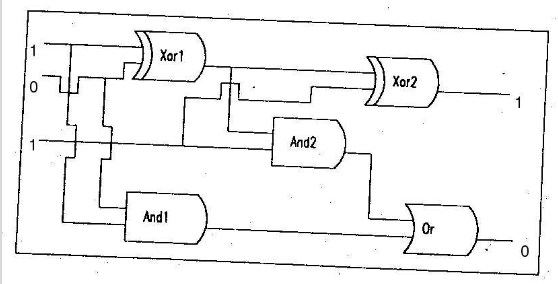

There are 2 minimal R-conflicts:

1. ?
2. ?

R-conflicts: Example 2

Example

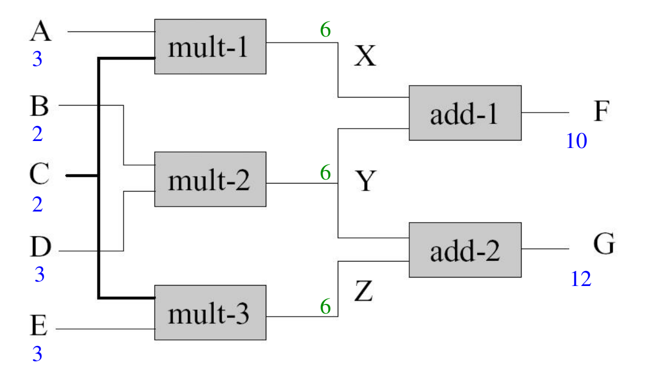

There are 2 minimal R-conflicts:

1. ?
2. ?

R-conflict and R-Diagnosis

Theorem

$\Delta \subseteq COMP$ is an R-diagnosis for $(SD, COMP, OBS)$ iff Δ is a minimal set such that $COMP \setminus \Delta$ is not an R-conflict.

This theorem is the basis of the algorithm DIAGNOSE from Reiter: it is a lattice exploration.

Definition

A lattice is (roughly) a non-empty partial order set $(S, \subseteq)$ such that every element a, b have an infimum $inf(a, b)$ (a "lower bound" element) and a supremum $sup(a, b)$ (an "upper bound" element).

Search space for R-diagnoses

Example

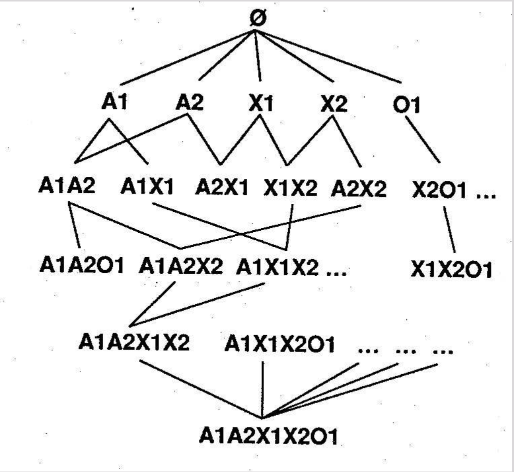

The search space is a lattice.

DIAGNOSE algorithm

Algorithm

Breadth-first search on the lattice from the empty set $\emptyset$

1. let X the current node in the search
2. Call a theorem prover and ask:

$$\text{Is } COMP \setminus X \text{ an R-conflict ?}$$

3. if **yes**, eliminate the nodes X' such that $X' \cap (COMP \setminus X) = \emptyset$
   - X' cannot be a minimal diagnosis.
4. if **no**, X is a minimal diagnosis, eliminate the descendants

DIAGNOSE algorithm: example

Example

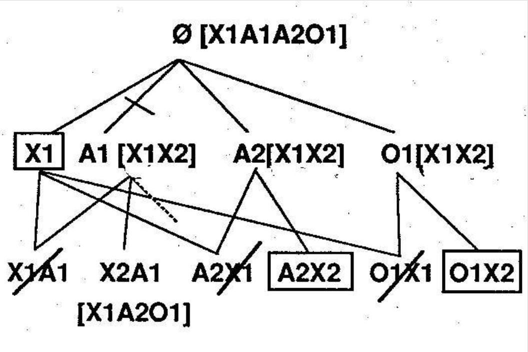

Sets in brackets are R-conflicts.
Three minimal diagnoses: $\{X1\}$; $\{X2, O1\}$; $\{X2, A2\}$

Another way to solve the problem

The intersection between a diagnosis and any R-conflicts is not empty $\Rightarrow$ Hitting set

Theorem

$\Delta \subseteq COMP$ is an R-diagnosis for $(SD, COMP, OBS)$ iff Δ is a minimal hitting set for the set of minimal conflicts of $(SD, COMP, OBS)$

General diagnosis engine (GDE) from de Kleer.

R-diagnosis: a minimal hitting set problem

Definition

Let $\mathcal{S} = \{S_1, \ldots, S_n\}$ be a set of sets, H is a hitting set of $\mathcal{S}$ iff

$$H \subseteq_{S_i \in \mathcal{S}} S_i$$

and

$$\forall S_i \in \mathcal{S}, H \cap S_i \neq \emptyset$$

Example

$\mathcal{S} = \{\{a, b\}, \{c, b\}, \{e, f\}\}$ The following sets are hitting sets of $\mathcal{S}$:

- $H = \{a, b, c, e\}$
- $H = \{b, e\}$ (H is minimal)
- $H = \{a, c, f\}$ (H is minimal)

The following sets are not hitting sets of $\mathcal{S}$:

- $H = \{a, b\}$
- $H = \{b, e, g\}$

GDE algorithm

Algorithm

1. Computation of all the minimal R-conflicts.
   - Use of an ATMS (Assumption Truth Maintenance System)
   - Update of beliefs about assumptions by retractation of knowledge and declaration of new ones
2. Computation of the minimal hitting set on the obtained R-conflicts

R-conflict and R-Diagnosis: examples

Example

Additionner:
The 2 minimal R-conflicts
$\{X1, X2\}$ and $\{X1, A2, O1\}$
correspond to the 3 minimal diagnoses:
$\{X1\}$; $\{X2, O1\}$; $\{X2, A2\}$

Davis circuit:
The 2 minimal R-conflicts:
$\{a1, m1, m2\}$ and $\{a1, a2, m1, m3\}$
correspond to the 4 minimal diagnoses:
$\{m1\}$; $\{a1\}$; $\{a2, m2\}$; $\{m2, m3\}$

Outline

Incremental diagnosis

GDE or DIAGNOSE solve the diagnosis problem in a off-line way.

- The observation set is supposed to be complete

Observations/Tests

In some systems, an observation is the result of a test, an action, a measurement from the environment to the system.

Definition

The incremental diagnosis problem is to:

- compute a diagnosis based on a partial set of observations
- choose what could be the next measurement to perform in the system: prediction

Predicted observations

Definition

An R-diagnosis Δ predicts O iff

$$SD, OBS, \{\neg Ab(c), c \in COMP \setminus \Delta\}, \{Ab(c), c \in \Delta\} \models O$$

Given the system SD, the current set of observations OBS and the current diagnosis Δ, the system should produce the observation O.

Predicted observations: example

Example

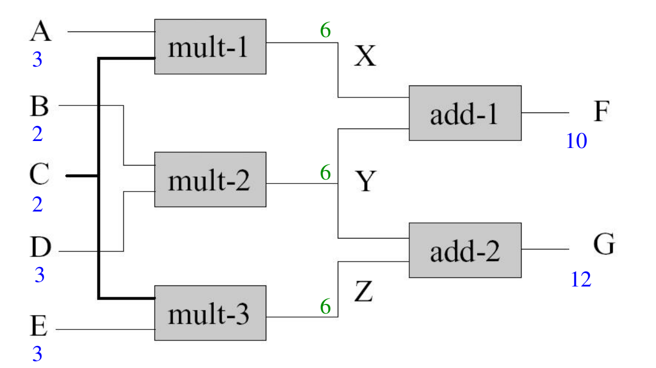

- $\Delta_1 : \{m1\}$ predicts $Out(m2) = 6$
- $\Delta_2 : \{m2, m3\}$ predicts $Out(m2) = 4$ and $Out(m3) = 6$.

Updating an R-Diagnosis

Theorem

Confirmation: A R-diagnosis for $(SD, COMP, OBS)$ which predicts O is a R-diagnosis for $(SD, COMP, OBS \wedge O)$.

If the predicted observation O is real (the measurement gives O), then the diagnosis is confirmed by the observation O.

Theorem

Invalidation: A R-diagnosis for $(SD, COMP, OBS)$ which predicts $\neg O$ is not a R-diagnosis for $(SD, COMP, OBS \wedge O)$.

If a diagnosis predicts something which is not true, it means that the diagnosis becomes a wrong hypothesis and is invalid.

Updating an R-Diagnosis

Algorithm

1. Input: $(SD, COMP, OBS)$ an observed system, O a new observation
2. Check if Δ predicts O
3. if **yes** then Δ is confirmed
   - Δ is a diagnosis of $(SD, COMP, OBS \wedge O)$
4. Check if Δ predicts $\neg O$
5. If **yes** then look at supersets of Δ

Updating an R-Diagnosis: example

Example

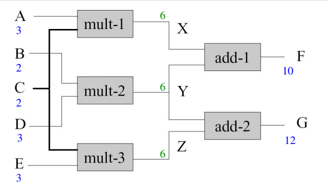

- $\Delta_1 = \{m1\}$ predicts $Out(m2) = 6$
- $\Delta_2 = \{m2, m3\}$ predicts $Out(m2) = 4$
- $\Delta_3 = \{a1\}$ predicts $Out(m2) = 6$
- $\Delta_4 = \{a2, m2\}$ predicts $Out(m2) = 4$

If O is $Out(m2) = 5$, every diagnosis is invalidated. The new ones are supersets: $\{a2, m1, m2\}$, $\{a1, m2, m3\}$, $\{a1, a2, m2\}$, $\{m1, m2, m3\}$

Discriminability/ Diagnosability

Definition

Let O be an observation which confirms Δ_1 and invalidates Δ_2, we say that O discriminates.

Definition

A system $(SD, COMP)$ is diagnosable if for any set of possible measurements (any complete set of observations) we have a unique diagnosis.

In a diagnosable system, we have enough information (observations) to discriminate between all the diagnoses and to get only one.

Using an incremental diagnosis algorithm on a diagnosable system, we have the guarantee that it converges to one diagnosis.

Diagnosability: example

Example

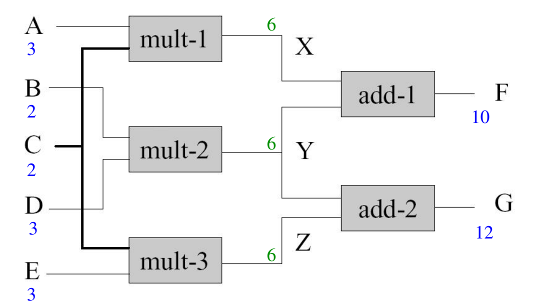

If we can can observe only A, B, C, D, E, F, G then the system is not diagnosable. If we observe $A, B, C, D, E, F, G, X, Y, Z$ then the system is diagnosable.
The observations from B, C, D, E, G do not allow to discriminate between diagnoses involving $m2$, $m3$, $a2$.

Summary

- Theory of Reiter: notions of R-Diagnosis, R-conflicts
- Logic representation ≡ set representations (minimal diagnoses)
- Algorithms:
  - DIAGNOSE: use of a theorem prover, exploration a lattice
  - GDE: computation of conflicts and hitting sets computation
- Incremental diagnosis: update the diagnoses with new measurements
- Discriminality-Diagnosability of systems
  - The more information we have, the less numerous are the diagnoses.

Outline

So many things...

- Non-monotonic reasoning
  - **Monotonicity**: if *KB* $\models \alpha$ then with a new information β, we still have *KB* $\wedge \beta \models \alpha$
  - The world is full of **exceptions**: every bird can fly, so the emu does!
  - Nonmonotonic logics: **Default logic**, **Circumscription**
- Uncertainty
  - Strong assumption: our knowledge is complete!
  - How to express and make reasoning about ignorance, incompleteness
  - Use of probability theory (Bayesian networks, Markov Decision Process, Fuzzy logic)

So many things...

- Inconsistency
  - Always reasoning with consistency! boring! and bounded! (incompleteness)
  - What about reasoning about inconsistencies: 1 + 1 = 3 for 1 big enough !
  - **Paraconsistent logics**...
- I give up, I do not have time...