Preliminary results on the Modeling of System Level Diagnosis problems with Abstract Argumentation

Stefano Bistarelli[1] and Antonio Caruso[2]

[1] University of Perugia. stefano.bistarelli@unipg.it
[2] University of Salento. antonio.caruso@unisalento.it

Abstract. System Level Diagnosis is a large subfield of fault-tolerance in parallel and distributed systems. Each node in the system can be '`faulty`' or '`fault-free`' and is able to perform a *diagnostic test* of it's neighbouring nodes. The collection of test results is called the *syndrome of the system*; the goal of the diagnosis is to provide algorithms and formal techniques to identify all faulty nodes correctly under various fault model. Abstract Argumentation is a framework used in AI to reason about belief and valid conclusion among a network of conflicting arguments, it is used by Intelligent Agents to take decisions. Subset of arguments can be collectively *accepted* or *refuted*. In this paper, we study a connection between the two models, and show that the diagnosis problem can be represented using an argumentation system: arguments can be mapped to set of units such that *refuted* arguments can be correctly diagnosed as faulty.

Keywords: Diagnosis · Argumentation

1 Introduction

Computer Science, as a research field, is so vast that it is not uncommon that ideas and concepts developed in a subarea, could be rephrased or re-discovered, in a different area, in some case without knowing the commonalities of the fundamental ideas. In this paper we start reviewing two research topic: *System Level Diagnosis*, a theory developed in the field of fault-tolerant systems to diagnose (permanent) system faults in multiprocessor and distributed systems; and *Argumentation Systems* a framework developed in the area of Artificial Intelligence to represent formal reasoning of agents about abstract topic that can be mutually conflicting. This two topics have also a different history, with System Level Diagnosis started by Preparata et. al. [20] in the '60 while the topic of argumentations has been developed quite recently by Dung [16]. In this paper, we show that a common model of system level diagnosis can be represented within an argumentation framework: arguments can be mapped to set of units, such that, their fault or fault-free status, is related to their acceptability in the argumentation system. We study this connection and see what results can be translated to the field of argumentation from system diagnosis. The rest of the paper is organized as follows. Section 2 review the field of System Level Diagnosis and the PMC model, in Section 3 we reviews the framework of Abstract Argumentation and the different semantics of acceptability of arguments. In Section 4 we present the main result of this paper: a formal proof that there exist a *mapping* of the diagnosis problem into an argumentation system. We show that under this mapping the subsets of arguments that are *refuted* by the argumentation system can be mapped back to subsets of units that can be diagnosed as faulty. We conclude the paper in Section 5.

2 System Level Diagnosis

System-level diagnosis, which was introduced by Preparata, Metze and Chien in [20] aims at diagnosing systems composed by units (usually processors) connected by point-to-point, bidirectional links. A system S is represented by the *system graph* $G = (\mathsf{V}, \mathsf{L})$, an undirected graph where nodes represent units and edges represent interconnection links. The value of $n = |\mathsf{V}|$ is called the *size* of the system.

System level diagnosis is based on a suitable set of tests between all adjacent units. The set of tests utilized for the purpose of diagnosis are represented by directed edges in the *diagnostic graph* $DG = (\mathsf{V}, \mathsf{E})$, where edge (u, v)[3] from u to v exists iff $\{u, v\} \in \mathsf{L}$, and unit u tests unit v. Edges in E are labeled with the binary test outcomes.

The system is characterized by an unknown set $\mathsf{V}_f \subseteq \mathsf{V}$ of faulty units (called the *actual fault set*). The set of all test outcomes is called a *syndrome* of the system, we denote with Σ the set of all the possible syndromes which may result from executing all tests corresponding to the edges of DG, and with σ any syndrome in Σ. Formally a syndrome is any function $\sigma : \mathsf{E} \to \{0, 1\}$. A test of v performed by unit u with outcome $\delta \in \{0, 1\}$ is denoted by $u \xrightarrow{\delta} v$. The concise notation $u \overset{\gamma\ \delta}{\longleftrightarrow} v$, with $\gamma, \delta \in \{0, 1\}$, denotes both the test of u performed by v with outcome γ and the test of v performed by u with outcome δ; if $\gamma = \delta$ we use the notation $u \overset{\gamma}{\longleftrightarrow} v$. Once a unit u has tested an adjacent unit v, the result of this test must be interpreted according to a given *fault model* trying to capture the behavior of faulty units.

2.1 The PMC Model

The *PMC Model* [20], known also as the *Symmetric invalidation model*, assumes permanent faults, perfect test coverage, and an *invalidation rule* shown in Table 1. This rule state that tests performed by fault-free units are completely reliable, while tests performed by faulty units are completely unreliable.

Given a system, its diagnostic graph $DG = (\mathsf{V}, \mathsf{E})$, the actual fault set V_f, a syndrome σ resulting from V_f, we define:

[3] We denote with (u, v) a directed edge from unit u to v, and with $\{u, v\}$ and undirected edge.

Table 1. Invalidation rule in the PMC model

Testing unit	Tested unit	Test outcome
Faulty-Free	Faulty-Free	0
Faulty-Free	Faulty	1
Faulty	Faulty-Free	0 or 1
Faulty	Faulty	0 or 1

Definition 1. (Diagnosis): *A* diagnosis *is a partition of* V *into subsets* $(\mathsf{F}, \mathsf{K}, \mathsf{S})$. *Units in* F *are declared as faulty, units in* K *are declared as fault-free and units in* S *are declared as* suspect*, i.e., units whose status remains unidentified.*

It is immediate from Table 1 that any given fault set may yield different syndromes; conversely, any given syndrome may derive from different fault sets. Given a syndrome, the goal of a *diagnosis algorithm* is to identify a *consistent fault set* of minimum cardinality:

Definition 2. (see [18]) *Given a syndrome* σ*, any subset* $\mathsf{F} \subseteq \mathsf{V}$ *is a* consistent fault set *(CFS) of* σ *if and only if* $\forall\,(u, v) \in \mathsf{E}$

1. If $u \in \mathsf{V} - \mathsf{F}$ *and* $u \xrightarrow{0} v \implies v \in \mathsf{V} - \mathsf{F}$.
2. If $u \in \mathsf{V} - \mathsf{F}$ *and* $u \xrightarrow{1} v \implies v \in \mathsf{F}$.

That is, F is a CFS for σ if and only if the assumption that the units of F are faulty and the units in $\mathsf{V}-\mathsf{F}$ are fault-free is consistent with σ and the fault model. A system is said *one-step* t-diagnosable if it is possible to correctly diagnose all the units using only a single syndrome, under the hypothesis that $|\mathsf{V}_f| < t$. The maximum value of t, in this case, is called the *one-step diagnosability* and is limited above by minimum in-degree of nodes in DG (under the hypothesis of reciprocal tests) [20,17].

To support a number of faults higher than one-step diagnosability, a weaker notion of *sequential diagnosis* is introduced in [18]. The sequential diagnosis consists of several *diagnosis and repair* phases, the goal of each phase is the identification of at least one faulty unit. Once identified, faulty units are immediately repaired or replaced, thus reducing their number. The process is iterated until all the faulty units have been removed. A diagnosis algorithm aimed at correctly diagnosing a large fraction of units in a single phase have been proposed in [8,11,12].

2.2 Formal properties of the PMC model

The properties stated in the following lemma are immediate from the invalidation rule (Table 1) of the PMC model:

Lemma 1. (see [10]) *For any two adjacent units* u, v *in set* V *the following statements hold:*

(a) If $u \overset{1\ 0}{\longleftrightarrow} v$*, then* u *is faulty.*
(b) If $u \overset{0}{\longrightarrow} v$ *and* u *is fault-free, then* v *is also fault-free.*
(c) If $v \overset{0}{\longrightarrow} u$ *and* u *is faulty, then* v *is also faulty.*
(d) If $u \overset{1\ 1}{\longleftrightarrow} v$*, then at least one unit between* u *and* v *is faulty.*

For any given syndrome σ, we define $DG_0 = (\mathsf{V}, \mathsf{E}_0)$ as the subgraph of DG of edge set $\mathsf{E}_0 \subseteq \mathsf{E}$ where E_0 is the set of edges labeled with test outcome 0; that is, $\mathsf{E}_0 = \{(u, v) \in \mathsf{E} \mid u \overset{0}{\longleftrightarrow} v\}$. Consider the strongly connected components of DG_0 and let $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r (r \geq 1)$ be the vertex sets of such components. Since every vertex belongs to exactly one (possibly trivial) component, the collection of subsets $\{\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r\}$ is a partition of V. The following properties are immediate from Lemma 1:

Lemma 2. (see [9]) *Given syndrome* σ*, let* $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ *be the vertex sets of the strongly connected components of* $DG_0 = (\mathsf{V}, \mathsf{E}_0)$*. For every* i*,* $1 \leq i \leq r$*, in the graph* DG *we have:*

(A) All units in A_i *are in the same (faulty or non-faulty) state.*
(B) If there exist units $u, v \in \mathsf{A}_i$ *with* $u \overset{1}{\longrightarrow} v$*, then all units in* A_i *are faulty.*
(C) If there exist units $u \in \mathsf{A}_i$*,* $v \in \mathsf{A}_j$ *with* $i \neq j$ *and* $u \overset{0}{\longrightarrow} v$*, then all units in* A_i *are faulty.*

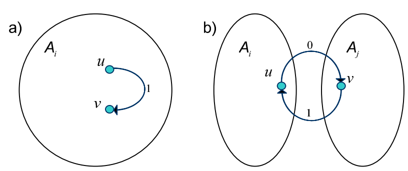

Fig. 1. Components of DG that can be declared *unconditionally faulty*

An example of an aggregate that satify Lemma 2 point (B) and of two aggregates that satify point (C) are in Fig. 1. Define F' as the union of all nodes inside the aggregates that can diagnosed as faulty using Lemma 2(B)(C). If $\mathsf{F}' \neq \emptyset$ then an incomplete, unconditionally correct diagnosis is trivially available, and the goal of sequential diagnosis is achieved since at least one faulty unit has been diagnosed. However, F' can be empty in general (and in the worst case). For this reason in the rest of the paper we assume that no faulty units are identified using rules (B) and (C) of Lemma 2, that is, $\mathsf{F}' = \emptyset$. To emphasize property (A) of Lemma 2, we call the aggregates that do not satify properties (B),(C) as *Z-aggregates*, note that for any pair of Z-aggregates A_i, A_j, they are said to be *adjacent* if there exist units $u \in \mathsf{A}_i$, $v \in \mathsf{A}_j$ such that $u \longleftrightarrow v$; it is immediate that in this case the tests outcomes must be $u \overset{1\ 1}{\longleftrightarrow} v$.

2.3 Diagnosis algorithm

Since all units in a Z-aggregates are in the same state, a diagnosis of the system consists in a (valid) labelling of the Z-aggregates as faulty, fault-free or undiagnosed. Letting α be the maximum cardinality of the Z-aggregates, that is, $|\mathsf{A}_i| \leq \alpha$ for $i = 1, \ldots, r$, we consider a diagnosis algorithm [12] that labels as fault-free all the Z-aggregates of cardinality α (at least one exists) and that labels as faulty all the Z-aggregates adjacent to some fault-free Z-aggregate.

The *fault-free core* K is defined as the union of all the Z-aggregates of cardinality α, the set of faulty units F is defined as the union of all Z-aggregates labeled as faulty, and the suspect set S is defined as the union of all the remaining Z-aggregates.

The algorithm also outputs a *syndrome-dependent bound* $t_\sigma = \alpha - 1$. In the hypothesis that the actual number of faults in the system is not above t_σ, i.e., $|\mathsf{V}_f| \leq t_\sigma$ the diagnosis is correct. In fact, under this hypothesis all Z-aggregates of cardinality α must be fault-free. The bound t_σ depends on the syndrome, it is possible to define a measure of diagnosablity that is syndrome-independent [9], such a bound can be expressed as a function $T(n)$ of the size of the system and is strictly related to the topology of the diagnostic graph.

3 Abstract Argumentation

In artificial intelligence and related fields, an argumentation framework, or argumentation system, is a way to deal with conflicting information and draw conclusions from it. In an abstract argumentation framework, entry-level information is a set of abstract arguments that, for instance, represent data or a proposition.

Conflicts between arguments are represented by a binary relation on the set of arguments. In concrete terms, you represent an argumentation framework with a directed graph such that the nodes are the arguments, and the arrows represent the attack relation.

The goal of each argumentation framework is at the end to propose a semantic for the acceptability of arguments, i.e. which subset of arguments are together maximally consistent from the point of view of an agent. We start with this basic definition:

Definition 3. *A* Dung's style *Abstract Argument System[16] is a pair* $D = \langle \mathcal{X}, \mathcal{A} \rangle$ *where* $\mathcal{X} = \{a_1, \ldots, a_k\}$ *is a finite set of* arguments*, and* $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ *is a binary* attack relation *on* $\mathcal{X}$*.*

To *solve* an argumentation system, we needs a notion of *acceptability* of arguments, and in particular which arguments are mutually compatible in the system. We start to define the following subset of arguments:

Definition 4. *initial, acceptable, conflict-free, admissible arguments:*

Initial arguments: *A subset of arguments* $\mathsf{I} \subseteq \mathcal{X}$ *is called* initial *if* $\nexists u \in \mathcal{X}$ *such that* $\exists v \in \mathsf{I}$ *with* $(u, v) \in \mathcal{A}$*. I.e. initial arguments do not have any other arguments that are in conflict with them. We denote with initial*$(\langle \mathcal{X}, \mathcal{A} \rangle)$ *the (unique) maximal such set.*

Conflict-free arguments: *a set of arguments* $\mathsf{E} \subseteq \mathcal{X}$ *is conflict-free if there is no attack between its arguments, i.e.:* $\forall a, b \in \mathsf{E}, (a, b) \notin \mathcal{A}$*.*

Acceptable arguments: *an argument* $a \in \mathcal{X}$ *is acceptable with respect to* $\mathsf{E} \subseteq \mathcal{A}$ *if and only if* E *defends* a*, that is* $\forall b \in \mathcal{X}$ *such that* $(b, a) \in \mathcal{A}, \exists c \in \mathsf{E}$ *such that* $(c, b) \in \mathcal{A}$*.*

Admissible arguments: *a set of arguments* E *is admissible if and only if it is conflict-free and all its arguments are acceptable with respect to* E*.*

Given an instance of an argumentation framework, the main problem is clearly to determine the justification state (also called the *defeat status*) of arguments, in particular: what arguments emerge undefeated from the various conflict, i.e. are acceptable? To decide if an argument can be accepted or not, or if several arguments can be accepted together, Dung defines several semantics of acceptance that allow, given an argumentation system, to compute sets of arguments, called extensions. For instance, given $S = \langle \mathcal{X}, \mathcal{A} \rangle$ we define a set E of arguments as:

Definition 5. *complete, preferred, stable and grounded extensions:*

complete extension *of* S *only if it is an admissible set and every acceptable argument with respect to* E *belongs to* E*,*

preferred extension *of* S *only if it is a maximal element (with respect to* $\subseteq$*) among the admissible sets with respect to* S*,*

stable extension *of* S *only if it is a conflict-free set that attacks every argument that does not belong in* E *(formally,* $\forall a \in \mathcal{X} \backslash \mathsf{E}, \exists b \in S$ *such that* $(b, a) \in \mathsf{E}$*)*

(unique) grounded extension *of* S *only if it is the smallest element (with respect to* $\subseteq$*) among the complete extensions of* S*.*

There exists some inclusions between the sets of extensions built with these semantics: Every stable extension is preferred, every preferred extension is complete, the grounded extension is complete, and if the system is well-founded all these semantics coincide, i.e. only one extension is grounded, stable, preferred, and complete. Given an acceptability semantics $\theta : 2^{\mathcal{X}} \rightarrow \{\top, \bot\}$, we denote with $\mathcal{E}_\theta(\langle \mathcal{X}, \mathcal{A} \rangle) = \{S \subseteq \mathcal{X} \mid \theta(S)\}$ the set of subsets accepted by θ.

Arguments can be also accepted or refuted with respect to the a given semantic, i.e. for example we consider the two notions of *credulous* or *sceptical* acceptability: $\alpha \in \mathcal{X}$ is credulous accepted if it is contained in at least an extension: $\exists\ S \in \mathcal{E}_\theta(\langle \mathcal{X}, \mathcal{A} \rangle),\ \alpha \in S$; and sceptical if it is contained in all extensions, i.e. $\forall\ S \in \mathcal{E}_\theta(\langle \mathcal{X}, \mathcal{A} \rangle),\ \alpha \in S$.

3.1 Labellings

Labellings [7] are a more expressive way than extensions to express the acceptance of the arguments. Concretely, a labelling is a mapping that associates every argument with a label `in,out,undec` respectively for arguments accepted, rejected and undefined. One can also note a labelling as a set of pairs (*argument,label*).

Such a mapping does not make sense without additional constraint. The notion of reinstatement labelling guarantees the sense of the mapping. L is a reinstatement labelling on the system $S = \langle \mathcal{X}, \mathcal{A} \rangle$ if and only if:

Definition 6. *Reinstatement labelling:*

1. $\forall u \in \mathcal{X}, L(u) = \texttt{in}$ *if and only if* $\forall v \in \mathcal{X}$ *such that* $(v, u) \in \mathcal{A}, L(v) = \texttt{out}$
2. $\forall u \in \mathcal{X}, L(u) = \texttt{out}$ *if and only if* $\exists v \in \mathcal{X}$ *such that* $(v, u) \in \mathcal{A}$ *and* $L(v) = \texttt{in}$
3. $\forall u \in \mathcal{X}, L(u) = \texttt{undec}$ *if and only if* $L(u) \neq \texttt{in}$ *and* $L(u) \neq \texttt{out}$

One can convert every extension into a reinstatement labelling: the arguments of the extension are `in`, those attacked by an argument of the extension are `out`, and the others are `undec`. Conversely, one can build an extension from a reinstatement labelling just by keeping the arguments `in`. Caminada [7] proved that the reinstatement labellings and the complete extensions can be mapped in a bijective way.

Reinstatement labellings distinguish arguments not accepted because they are attacked by accepted arguments from undefined arguments — that is, those that are not defended cannot defend themselves. An argument is `undec` if it is attacked by at least another `undec`. If it is attacked only by arguments `out`, it must be `in`, and if it is attacked some `in` argument, then it is `out`.

4 Diagnosis and Argumentations

In the previous Sections we have presented the formal models of System Level Diagnosis and Abstract Argumentation, at this point, the reader could already see wide similarities between the two, and in particular the common use of a graph based theoretical model.

In this Section we formally study these similarities, starting from a mapping of System Diagnosis into a proper Argumentation Framework. The idea is to start with a diagnostic graph $DG = (\mathsf{V}, \mathsf{E})$, an unknow fault-set $\mathsf{V}_f \subseteq \mathsf{V}$ and a syndrome σ, and construct an argumentation system $D = \langle \mathcal{X}, \mathcal{A} \rangle$ such that, its solutions (acceptable arguments with respect to a give semantic) can be used to produce a diagnosis that is correct with respect to σ, and maybe also complete.

In other words, given a reinstatement labelling of D the arguments labelled as `in` can be mapped to units diagnosed as fault-free, the arguments labelled as `out` to units diagnosed as faulty, and finally arguments labelled as `undec` to undiagnosed units. If the resulting diagnosis is correct (and/or complete), we denote the labelling with the same name, i.e. a *correct* (and/or complete) *reinstatement labelling* wrt the syndrome.

Note that the simplest (most intuitive) mapping consider units of S as arguments (so $\mathcal{X} = \mathsf{V}$) and given a syndrome σ maps all diagnostic test with output 1 as conflicts, i.e. $\forall u, v \in \mathsf{V}$ such that $u \xrightarrow{1} v$ create a conflict $(u, v) \in \mathcal{A}$ between the arguments u, v. It is easy to see that such a mapping do not respect the above requirements since by definition 4 all *initial arguments* are labelled as `in` in all standard semantic for argumentation, but a diagnosis that consider them as fault-free is not correct.

Theorem 1. *There exist an actual fault set* V_f*, and a syndrome* $\sigma \in \Sigma$*, such that, if we consider the above mapping to an argumentation system* $D{=}\langle \mathcal{X}, \mathcal{A} \rangle$*, any reinstatement labelling where* $L(initial(\langle \mathcal{X}, \mathcal{A} \rangle) = \texttt{in}$ *are not correct.*

Proof. A simple counter-example is provided by a faulty-unit surrounded by other faulty-units, that collectivelly declare each other as fault-free.

Formally, consider a syndrome $\sigma \in \Sigma$ such that $\exists u \in \mathsf{V}_f$ that tests and is tested only with outcome 0, i.e. if $N(u)$ are the neighbours of u, $\forall v \in N(u), v \overset{0\ 0}{\longleftrightarrow} u$. Note that by Lemma 1 units that tests each other with 0 must be in the same state, both faulty or fault-free, and since $u \in \mathsf{V}_f$ we are considering the first case.

Such a unit is mapped to an isolated argument in the argumentation graph, and by definition it is initial. Any semantic of acceptance includes the initial arguments in the set of admissible arguments, and accordingly, in any reinstatement labelling such arguments will be labelled as `in`. Therefore, unit u is uncorrectly diagnosed as fault-free. □

The above problem can be solved using two different approaches. The first approach consider a preprocessing that transform the input of the diagnosis problem into a *contracted graph* that simplify the problem, removing all tests with outcome 0 from the syndrome. The second approach, use directly the syndrome but maps it to an extension of the Dung's Framework with *bipolar relations* [13]. This extension offer the opportunity to include the information provided in the syndrome by tests with outcome 0 into a *support* relation aside to the conflict ones. In this paper we explore the first approach, and leave the latter as future work.

4.1 The Contracted Graph

We define the *contracted graph* of a diagnostic graph DG under syndrome σ, denoted $CG_\sigma = (\mathsf{V}', \mathsf{L}')$ as the undirected, vertex-weighted graph obtained by contracting all the units of a Z-aggregate A_i (defined after Lemma 2) into a single vertex $a_i \in \mathsf{V}'$. Vertices in V' are in a one-to-one correspondence with Z-aggregates. Edge $\{a_i, a_j\} \in \mathsf{L}'$ iff Z-aggregates A_i and A_j are adjacent. Remember that we assume that the set F' of units belonging to aggregates that satisfy Lemma 2(B)(C) is empty (or we can simply remove them from the system) (see page 5) this implies that all edges in L' can be considered as labelled with 1. Each vertex a_i is weighted with $\alpha_i = |\mathsf{A}_i|$.

In Figure 2 the contracted graph of a system with toroidal grid interconnections is shown. This example has been obtained by randomly distributing 40 faulty units into a 10×10 toroidal grid that determine a partition of the grid into 13 aggregates.[4]

The fault/fault-free status of a node in the contracted graph correspond directly to the fault/fault-free status of the nodes inside each Z-aggregate. Thus, any CFS of the contracted graph is in a one-to-one correspondence with a CFS

[4] Z-Aggregates $1, 2, 4$ wrap-around the edges of the toroidal grid

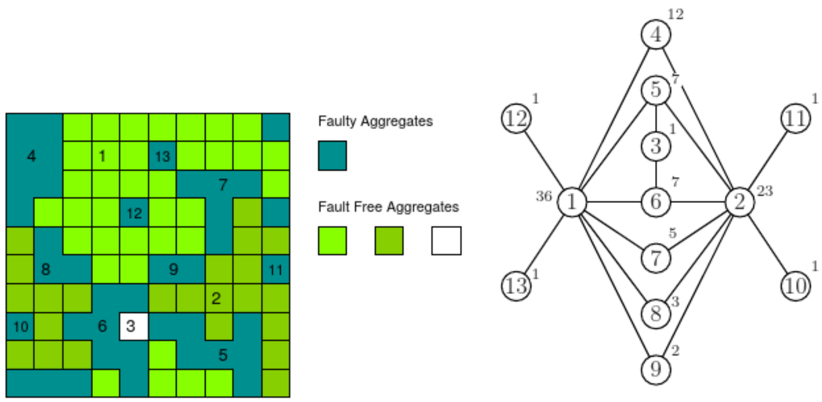

(a) The toroidal grid (b) The contracted graph

Fig. 2. Example of a toroidal grid partitioned by 40 faults into 13 Z-Aggregates and the associated contracted graph.

of the diagnostic graph. Moreover, the properties of the diagnostic model remain valid, in particular Lemma 1 can be applied to the contracted graph by considering the vertices of the contracted graph as single units. Let $\Gamma(\mathsf{A})$ is the set of nodes in a graph adjacent to some nodes in set A. We use an important property that relate, consistent fault sets of σ and independent sets[5] of the contracted graph:

Theorem 2 (Independence). *Given a system graph G, and a syndrome σ. Let $CG_\sigma = (\mathsf{V}', \mathsf{L}')$ be the contracted graph of G under syndrome σ. All subsets $\mathsf{X} \subseteq \mathsf{V}'$ of nodes declared fault free are independent, and the set $\Gamma(\mathsf{X})$ is a consistent fault set with respect to σ.*

Proof. By definition, for any pair a_i, a_j of adjacent nodes of CG_σ, the outcomes of the mutual tests between units in A_i and units in A_j are 1. If we assume that a_i is fault-free, from properties (d) of Lemma 1 every a_j adjacent to a_i in CG_σ must be faulty. This means that fault-free nodes cannot be mutually adjacent (so they are independent). For the same reason, the set of nodes $\mathsf{U} = \Gamma(\mathsf{X})$ is a consistent fault set of σ.

At this point a mapping from the contracted graph to an argumentation system is straightforward:

Definition 7. *Mapping* *Given a contracted graph $CG_\sigma = (\mathsf{V}', \mathsf{L}')$ from the syndrome σ, we map it to the argumentation system $D = \langle \mathcal{X}, \mathcal{A} \rangle$ with $\mathcal{X} = \mathsf{V}'$ (nodes of the graph represent arguments) and $(u, v), (v, u) \in \mathcal{A}$ for each $(u, v) \in \mathsf{L}'$ (test with outcome 1, represent conflicts).*

[5] A subset I of nodes is independent iff $\forall u, v \in \mathsf{I}$ u, v are not connected.

With this mapping, the terms: *argument*, *node in the contracted graph* or *Z-aggregate* (or the set of units it represent) are all equivalent, since we can map directly one concept to the others. The main result of the paper, presented in the following Theorem, shows that *refuted arguments* are directly related to consistent fault set of a syndrome, i.e. set of units that collectivelly can be diagnosed as faulty.

Theorem 3. *Given a diagnostic graph DG, a syndrome σ, and the derived contracted graph CG_σ. Consider an argumentation system $D = \langle \mathcal{X}, \mathcal{A} \rangle$ produced from CG_σ with the above mapping. In any reinstatement labelling L of D (from an* preferred *extension of D) the set of* refuted *arguments can be mapped back to a Consistent Fault Set of σ.*

Proof. Given σ and CG_σ, consider $D = \langle \mathcal{X}, \mathcal{A} \rangle$ built from the above mapping. Solve D, i.e. consider a reinstatement labelling L of D produced using a *preferred extension*. Refuted arguments are labelled as `out` if they are in conflict with arguments labelled `in` by L. The latter belong to a maximal *admissible* extension of D, a set of arguments K that satisfy two conditions:

$$\forall u, v \in \mathsf{K}, (u, v) \notin \mathcal{A} \quad \textit{conflict-free} \tag{1}$$

$$\forall u \in \mathsf{K}, \forall v \in \mathcal{X} \text{ with } (v, u) \in \mathcal{A}, \exists w \in \mathsf{K}, \text{ with } (w, v) \in \mathcal{A} \quad \textit{admissible} \tag{2}$$

First we note that (2) is *redundant* if the argumentation graph is symmetric. In fact, if $\forall u \in \mathsf{K}$ and any $(v, u) \in \mathcal{A}$ (by simmetry of the conflict relation) we have $(u, v) \in \mathcal{A}$ and (2) is always true (with trivial $w = u$). So, K is a conflict-free set in D iff all its members are pairwise disjoint, but this is just the definition of an independent set. Moreover, since the obtaimned AF is symmetric, preferred extensions coincide with stables [15] and thus there are no `undec` arguments [7]. From Theorem 2 the result follow. □

Theorem 3 states that the two field of Abstract Argumentation and System Level Diagnosis are strongly related. In fact, in both cases the solution space is made of maximal independent sets of the graph used to model the system: the graph of arguments and conflicts in the case of Argumentation, and the contracted graph of Z-Aggregates in the case of Diagnosis.

5 Conclusions and Future Works

In this paper we studied how to cast the problem of Diagnosis of faulty units, studied in the area of System Level Diagnosis, that we reviewed in Section 2 in the framework of abstract argumentation, reviewed in Section 3. We proved in Theorem 1 that the simplest, most intuitive mapping do not work, and provide a different mapping based on the notion of a contracted graph. We proved in Theorem 3 that under this mapping, the complement of a consistent fault set for a given syndrome, and the set of accepted arguments are both described by independent set of a graph.

However, computing a correct diagnosis using the mapping in Definition 7 and the solutions of the argumentation system, requires a last step. Solving the argumentation system (for instance using a solver like ConArg[6] [5,3]), gives us the set of solutions, i.e. $\mathcal{E}_{\text{pref}}(\langle \mathcal{X}, \mathcal{A} \rangle)$, but we cannot distinguish among them. In argumentation this is clearly reasonable, since there is no reason to prefer a particular solution to another. However, the diagnosis must distinguish the correct consistent fault set, under some hypothesis on the diagnosability of the system.

As an example of this problem, consider Fig. 2 (b), and two solutions can be $S = \{1, 3, 2\}$, or $S' = \{12, 13, 4, 5, 6, 7, 8, 9, 10, 11\}$, both are maximal independent sets. The two sets represent two different diagnosis, both consistent with σ, but qualitatively different: S consider as fault-free a large fraction of the system (60 units), while S' declare fault-free a minority of units. Since the diagnosability of any system must be at least $n/2$ (the majority of units must be fault-free) the first diagnosis is also correct and complete, while the other is not.

A possible way to solve to the above problems would be to consider, among all solutions computed by the argumentation system, i.e. all $S \in \mathcal{E}_{\text{adm}}(\langle \mathcal{X}, \mathcal{A} \rangle)$ the ones with maximum weight. Use the weights in the contracted graph to compute a value $w(S) = \sum_{v \in S} w(v)$ for each set S, and consider the set S' that maximize $w(S')$. We diagnose all units inside the Z-aggregates accepted (as arguments) in S' as fault-free, and the others as faulty; this diagnosis is correct if the number of faults in the system is strictly below $w(S')$, a value (called the diagnosability bound) that is returned to the user together with the diagnosis. This will be part of our future work, by using extensions of the Dung's framework, that allows, for example, weighted conflict [2,4,14,19]. Also, some work on coalitions [6,1], can be used to work on the contracted graph natively in the AF framework. Other further work will be the use of both fault and fault-free diagnosis in the syndrome by using bipolar Argumentation frameworks [13].

References

1. Bistarelli, S., Foley, S.N., O'Sullivan, B., Santini, F.: From marriages to coalitions: A soft CSP approach. In: Oddi, A., Fages, F., Rossi, F. (eds.) Recent Advances in Constraints, 13th Annual ERCIM International Workshop on Constraint Solving and Constraint Logic Programming, CSCLP 2008, Rome, Italy, June 18-20, 2008, Revised Selected Papers. Lecture Notes in Computer Science, vol. 5655, pp. 1–15. Springer (2008)
2. Bistarelli, S., Rossi, F., Santini, F.: A collective defence against grouped attacks for weighted abstract argumentation frameworks. In: FLAIRS Conference. pp. 638–643 (2016)
3. Bistarelli, S., Rossi, F., Santini, F.: Conarg: A tool for classical and weighted argumentation. In: Baroni, P., Gordon, T.F., Scheffler, T., Stede, M. (eds.) Computational Models of Argument - Proceedings of COMMA 2016, Potsdam, Germany, 12-16 September, 2016. Frontiers in Artificial Intelligence and Applications, vol. 287, pp. 463–464. IOS Press (2016)

[6] http://www.dmi.unipg.it/conarg/

4. Bistarelli, S., Santini, F.: A common computational framework for semiring-based argumentation systems1, 2. In: ECAI 2010: 19th European Conference on Artificial Intelligence, 16-20 August 2010, Lisbon, Portugal: Including Prestigious Applications of Artificial Intelligence (PAIS-2010): Proceedings. vol. 215, p. 131. IOS Press (2010)
5. Bistarelli, S., Santini, F.: Modeling and solving afs with a constraint-based tool: Conarg. In: Modgil, S., Oren, N., Toni, F. (eds.) Theorie and Applications of Formal Argumentation - First International Workshop, TAFA 2011. Barcelona, Spain, July 16-17, 2011, Revised Selected Papers. Lecture Notes in Computer Science, vol. 7132, pp. 99–116. Springer (2011)
6. Bistarelli, S., Santini, F.: Coalitions of arguments: An approach with constraint programming. Fundam. Inform. **124**(4), 383–401 (2013)
7. Caminada, M.: On the issue of reinstatement in argumentation. In: Fisher, M., van der Hoek, W., Konev, B., Lisitsa, A. (eds.) Logics in Artificial Intelligence. pp. 111–123. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
8. Caruso, A., Chessa, S., Maestrini, P., Santi, P.: Diagnosis of regular structures. In: IEEE Proc. Dependable Systems and Networks. pp. 213–222. New York, NY (June 2000)
9. Caruso, A., Chessa, S., Maestrini, P., Santi, P.: Diagnosability of regular systems. Journal of Algorithms **45**(2), 126–143 (November 2002)
10. Caruso, A., Chessa, S., Maestrini, P., Santi, P.: Evaluation of a diagnosis algorithm for regular structures. IEEE Transactions on Computers **51**(7), 850–865 (July 2002)
11. Caruso, A., Chessa, S., Maestrini, P., Santi, P.: Fault-diagnosis of grid structures. Theoretical Computer Science **290**(2), 1149–1174 (1 2003)
12. Caruso, A., Luiz. P., A., Maestrini, P.: A New Diagnosis Algorithm for Regular Interconnected Structures. In: First Latin American Symposium on Dependable Computing (LADC 2003) - LNCS 2847 (Oct 2003)
13. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolar abstract argumentation systems, pp. 65–84. Springer US, Boston, MA (2009)
14. Coste-Marquis, S., Konieczny, S., Marquis, P., Ouali, M.: Weighted attacks in argumentation frameworks. In: 13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012. pp. 593–597 (2012)
15. Coste-Marquis, S., Devred, C., Marquis, P.: Symmetric argumentation frameworks. In: Godo, L. (ed.) Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6-8, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3571, pp. 317–328. Springer (2005)
16. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–357 (1995)
17. Hakimi, S.L., Amin, A.T.: Characterization of connection assignment of diagnosable systems. IEEE Transactions on Computers **C-23**(1), 86–88 (January 1974)
18. Huang, S., Xu, J., Chen, T.: Characterization and design of sequentially t-diagnosable systems. In: IEEE Proc. 19^{th} International Symposium on Fault-Tolerant Computing. pp. 554–559. Washington, D.C., USA (June 1989)
19. Martínez, D., García, A., Simari, G.: An abstract argumentation framework with varied-strength attacks. In: Principles of Knowledge Representation and Reasoning: Proceedings of the 11th International Conference, KR 2008. pp. 135–143 (2008)
20. Preparata, F.P., Metze, G., Chien, R.T.: On the connection assignment problem of diagnosable systems. IEEE Transactions on Computers **EC-16**(12), 848–854 (December 1967)