Local Consistency in Junction Graphs for Constraint-Based Inference

Le Chang and Alan K. Mackworth

Department of Computer Science, University of British Columbia
2366 Main Mall, Vancouver, B.C. Canada V6T 1Z4
{lechang, mack}@cs.ubc.ca

Abstract. The concept of local consistency plays a central role in constraint satisfaction and has been extended to handle general constraint-based inference (CBI) problems. We propose a family of novel generalized local consistency concepts for the junction graph representation of CBI problems. These concepts are based on a general condition that depends only on the existence and property of the multiplicative absorbing element and does not depend on the other semiring properties of CBI problems. We present several local consistency enforcing algorithms and their approximation variants. Theoretical complexity analyses and empirical experimental results for the application of these algorithms to both MaxCSP and probability inference are given. We also discuss the relationship between these local consistency concepts and message passing schemes such as junction tree algorithms and loopy message propagation.

1 Introduction

The concept of local consistency plays a central role in constraint satisfaction. Given a constraint satisfaction problem (CSP), local consistency can be characterized as deriving new, possibly tighter, constraints based on local information. The derived constraints simplify the representation of the original CSP without the loss of solutions. This can be seen as a preprocessing procedure. For example, a value may be removed from a variable domain by the preprocessing because it violates these derived constraints. Both systematic approaches such as inference or propagation algorithms and stochastic approaches such as local searches benefit from these simplifications. Among the family of local consistency enforcing algorithms or filtering algorithms, arc consistency [1] is one of the most important techniques for binary classic CSPs. It is straightforward to extend it as generalized arc consistency [2] to handle non-binary classic CSPs. Many stronger local consistencies [3–5] have been studied within the constraint programming community. Based on the Semiring CSP [6] and Valued CSP [7] frameworks, arc consistency has also been extended, as soft arc consistency [8, 9], to handle over-constrained and preference-based problems that can be modelled as soft CSPs. Recently, we presented a weaker condition [10], based on a commutative semiring structure, for applying the generalized arc consistency approach to constraint-based inference (CBI) problems beyond classic and soft CSPs. There problems include probability inference, possibility inference, and maximal likelihood decoding. The weaker condition proposed in [10] has also been relaxed to fit generalized approximate preprocessing schemes.

We propose in this paper a new family of generalized local consistency concepts for the junction graph representation of CBI problems. These concepts are based on a general condition that depends only on the existence and property of the multiplicative absorbing element and does not depend on other semiring properties of CBI problems. We present several local consistency enforcing algorithms with various levels of enforcement and corresponding theoretic and empirical complexity analyses. We show in this paper that some of these algorithms can be seen as generalized versions of well-known local consistency enforcing techniques in CSPs and can be exported to other domains. Other abstract local consistency concepts are novel to the constraint programming community and provide more efficient preprocessing results. We also discuss the relationship between these local consistency concepts and message passing schemes such as junction tree algorithms and loopy message propagation. Local consistencies can be achieved along with message propagation and improve the efficiency of message passing schemes.

2 A CBI Framework and Junction Graph

Constraint-Based Inference (CBI) is an umbrella term for a class of various superficially different problems including probabilistic inference, decision-making under uncertainty, constraint satisfaction problems, propositional satisfiability problems, decoding problems, and possibility inference. we abstracts these problems into a single formal framework [11] using the algebraic semiring structure $\mathbf{S} = \langle \mathbf{A}, \oplus, \otimes \rangle$ where constraint combination is represented by the abstract multiplicative operator $\otimes$ and constraint marginalization is represented by the abstract additive operator $\oplus$. A CBI problem $\mathbf{P}$ in this framework is a tuple $(\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$, where $\mathbf{X}$ is a set of variables, $\mathbf{D}$ is a set of finite domains for each variable, $\mathbf{S} = \langle \mathbf{A}, \oplus, \otimes \rangle$ is a commutative semiring, and $\mathbf{F}$ is a set of constraints. Each constraint is a function that maps value assignments of a subset of variables to values in $\mathbf{A}$. Given a CBI problem, the inference task is defined as computing $g_{CBI}(\mathbf{Z}) = \bigoplus_{\mathbf{Y}} \bigotimes_{f \in \mathbf{F}} f$. If $\oplus$ of the semiring $\mathbf{S}$ is idempotent, the allocation task is defined as computing $\mathbf{y} = \arg \bigoplus_{\mathbf{Y}} \bigotimes_{f \in \mathbf{F}} f$, where $\arg$ is a prefix of operator $\oplus$. In other words, $\arg \oplus$ is an operator that returns the arguments of the $\oplus$ operator.

Our local consistency concepts are based on this CBI framework and apply to CBI problems with commutative semirings that are eliminative. A commutative semiring $\mathbf{S} = \langle \mathbf{A}, \oplus, \otimes \rangle$ is *eliminative* [10] if there exists a unique *multiplicative absorbing element* $\alpha_\otimes \in \mathbf{A}$ ($\alpha_\otimes \otimes a = \alpha_\otimes, \forall a \in \mathbf{A}$) and $\alpha_\otimes$ is equal to the *additive identity element* $\mathbf{0}$ ($\mathbf{0} \oplus a = a, \forall a \in \mathbf{A}$). Furthermore, our approximate local consistency concepts apply to CBI problems with commutative semirings that are eliminative and monotonic. A commutative semiring $\mathbf{S} = \langle \mathbf{A}, \oplus, \otimes \rangle$ is *monotonic* [10] if there exists a total order $\leq_{\mathbf{S}}$ on $\mathbf{A}$, the additive identity element $\mathbf{0}$ is the minimum element w.r.t. $\leq_{\mathbf{S}}$, $a \leq_{\mathbf{S}} b$ implies $a \oplus c \leq_{\mathbf{S}} b \oplus c$ and $a \otimes c \leq_{\mathbf{S}} b \otimes c$, $\forall a, b, c \in \mathbf{A}$. Examples of eliminative and monotonic semirings can be found in [10].

A junction graph $\mathcal{J} = (\mathcal{C}, \mathcal{S})$ of a CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$ is defined as following: $\mathcal{C} = \{C_1, \cdots, C_n\}$ is a set of clusters, each cluster C_i corresponds to an aggregation of variables that is a subset of $\mathbf{X}$ and has attached initially a local constraint $\phi_{C_i} = \mathbf{1}$ ($\mathbf{1}$ is the multiplicative identity element s.t. $\mathbf{1} \otimes a = a, \forall a \in \mathbf{A}$); $\mathcal{S} =$

$\{S_{ij} | C_i, C_j \in \mathcal{C}\}$ is a set of separators between C_i and C_j if $C_i \cap C_j \neq \emptyset$ and S_{ij} corresponds to an aggregation of variables that consists of the intersection variables between C_i and C_j. A junction graph satisfies the condition that for any constraint $f \in \mathbf{F}$, there exists a cluster $C_i \in \mathcal{C}$ s.t. $Scope(f) \subseteq C_i$. The definition of junction graph ensures that the subgraph induced by any variable is connected. We say a junction graph is *initialized* if for each constraint $f \in \mathbf{F}$, we choose a cluster C_i s.t. $Scope(f) \subseteq C_i$ and update ϕ_{C_i} by $\phi_{C_i} \otimes f$.

3 Local Consistency for CBI Problems

We present here novel local consistency concepts for initialized junction graphs of a CBI problem with an eliminative semiring. If the semiring used to represent a CBI problem is both eliminative and monotonic, it is straightforward to modify these concepts as approximate local consistencies using an element $\epsilon \in \mathbf{A}$ to approximate the multiplicative absorbing element $\alpha_{\otimes}$ that is equal to the additive identity element $\mathbf{0}$ for an eliminative commutative semiring, and using $\leq_{\mathbf{S}}$ to replace $\neq$ in the following definitions.

3.1 Single, Directional and Neighborhood Cluster Consistencies

The fundamental concept of local consistency for an initialized junction graph of a CBI problem with an eliminative commutative semiring is *single cluster consistency*. Here we consider only the local constraints attached to a single cluster and do not consider the effects of other clusters. Formally:

Definition 1 (Single Cluster Consistency (SCC)). *A cluster C_i of an initialized junction graph is locally consistent if $\forall X \in Scope(\phi_{C_i})$, $\forall x \in \mathbf{D}_X$, $\exists \mathbf{w}$, a value assignment of variables $Scope(\phi_{C_i})_{-X}$, s.t. $\phi_{C_i}(x, \mathbf{w}) \neq \alpha_{\otimes}$. An initialized junction graph of a CBI problem is Single Cluster Consistent if all the clusters are consistent.*

Single cluster consistency covers the definition of Generalization of Generalized Arc Consistency (GGAC) [10], which abstracts Generalized Arc Consistency (or Hyper-Arc Consistency) in constraint programming. If the junction graph of a CBI problem is primal, in other words, there is one cluster that corresponds to exactly one constraint and separators represent common variables shared between any two constraints, SCC is identical to GGAC. If the junction graph is constructed without satisfying this special structural requirement, SCC is stronger than GGAC in general.

Figure 1 shows a generalized routine of enforcing single cluster consistency for an initialized junction graph of a CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$ with an eliminative commutative semiring $\mathbf{S}$. The procedure *REVISE* of *SCC-Enforcing* is shown in Figure 2.

We also introduce two other local consistencies for an initialized junction graph of a CBI problem that are stronger than Single Cluster Consistency. They are Directional Cluster Consistency and Neighborhood Cluster Consistency. Effects of other clusters in the junction graph are taken into account. The distinction between these two local consistencies is based on which clusters are selected for consideration.

Input: A CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$ and its initialized junction graph representation $\mathcal{J} = (\mathcal{C}, \mathcal{S})$
Output: A Single Cluster Consistent CBI problem $\mathbf{P}' = (\mathbf{X}, \mathbf{D}', \mathbf{S}, \mathbf{F}')$
for each $C_i \in \mathcal{C}$ **do**
  for each $X \in Domain(\phi_{C_i})$ **do**
    if *REVISE*(X, ϕ_{C_i}) **then**
      for each $C_j \in \mathcal{C}$ **do**
        if $X \in Scope(\phi_{C_j})$ **then**
          Remove all tuples in ϕ_{C_j} with the value that is removed from X
        end if
      end for
    end if
  end for
end for
Return $\mathbf{P}' := \mathbf{P}$

Fig. 1. Single cluster consistency enforcing algorithm (*SCC-Enforcing*).

Definition 2 (Directional Cluster Consistency (DCC)). *Given a cluster C_i of an initialized junction graph and a separator S_{ij} that connects another cluster C_j to C_i, let $g_i = \phi_{C_i} \otimes (\oplus_{C_j - S_{ij}} \phi_{C_j})$. We say C_i is pair consistent w.r.t. C_j if $\forall X \in Scope(g_i)$, $\forall x \in \mathbf{D}_X$, $\exists \mathbf{w}$, a value assignment of variables $Scope(g_i)_{-X}$, s.t. $g_i(x, \mathbf{w}) \neq \alpha_{\otimes}$. An initialized junction graph of a CBI problem is Directional Cluster Consistent given a total ordering of the clusters if all clusters of the graph are pair consistent w.r.t all its lower-order neighbor clusters.*

Definition 3 (Neighborhood Cluster Consistency (NCC)). *Given a cluster C_i of an initialized junction graph, let $g_i = \phi_{C_i} \otimes \bigotimes_{C_j \in Neighbors(C_i)} (\oplus_{C_j - S_{ij}} \phi_{C_j})$. We say C_i is neighborhood consistent if $\forall X \in Scope(g_i)$, $\forall x \in \mathbf{D}_X$, $\exists \mathbf{w}$, a value assignment of variables $Scope(g_i)_{-X}$, s.t. $g_i(x, \mathbf{w}) \neq \alpha_{\otimes}$. An initialized junction graph of a CBI problem is Neighborhood Cluster Consistent if all clusters are neighborhood consistent.*

We revise the single cluster consistency enforcing algorithm in Figure 1 to directional cluster consistency and neighborhood cluster consistency enforcing algorithms by updating the local potential ϕ_{C_i} according to the definition, as shown in Figure 3 and 4, respectively.

3.2 Approximate Local Consistencies

Given a CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$, if the commutative semiring $\mathbf{S} = \langle \mathbf{A}, \oplus, \otimes \rangle$ is both eliminative and monotonic, we propose an approximation scheme to enforce local consistency for its initialized junction graph representation with a user-controlled threshold. More specifically, we use an element $\epsilon \in \mathbf{A}$ to approximate the multiplicative absorbing element $\alpha_{\otimes}$ that is equal to the additive identity element $\mathbf{0}$ for an eliminative commutative semiring, and use $\leq_{\mathbf{S}}$ to replace $\neq$ in previous local consistency definitions. The monotonic properties for both multiplicative and additive operators in a

Input: A variable $X \in \mathbf{X}$ and a constraint f with $X \in Scope(f)$
Output: TRUE if a value is removed from the domain of X, FALSE for else

$flag := TRUE$
for each $x \in \mathbf{D}_X$ **do**
  for each value assignment $\mathbf{w}$ of $Scope(f)_{-X}$ **do**
    if $f(x, \mathbf{w}) \neq \alpha_{\otimes}$ **then**
      $flag := FALSE$
      Break loop
    end if
  end for
  if $flag$ **then**
    Remove x from $\mathbf{D}_X$
    Return TRUE
  end if
end for
Return FALSE

Fig. 2. Procedure *REVISE*(X, f) for eliminating a domain value from a variable X according to the local constraint f.

	SCC	DCC	NCC
Time	$\lvert\mathcal{C}\rvert d^{k+1}$	$(\lvert\mathcal{S}\rvert + \lvert\mathcal{C}\rvert)d^{k+1}$	$(2\lvert\mathcal{S}\rvert + \lvert\mathcal{C}\rvert)d^{k+1}$
Space	$\lvert\mathcal{C}\rvert d^{k+1}$	$\lvert\mathcal{C}\rvert d^{k+1}$	$\lvert\mathcal{C}\rvert d^{k+1}$

Table 1. Time and space upper bound comparison among various local consistency enforcing algorithms for a junction graph $\mathcal{J} = (\mathcal{C}, \mathcal{S})$ of a given CBI problem, where $d = \max_{D_i \in \mathbf{D}} |D_i|$ and $k = \max_{C_i \in \mathcal{C}} |C_i|$.

monotonic semiring ensure that this approximation always returns a lower bound estimate of the inference task for a given CBI problem [10]. Correspondingly, the procedure *REVISE* in Figure 2 is modified to handle approximate local consistency enforcing tasks, as shown in Figure 5. All the local consistency enforcing algorithms discussed in the previous section then can be modified respectively.

4 Complexities and Relation with Message Propagation in Probability Inference

The worst case space complexities of all three local consistency enforcing algorithms are the same: linear in the number of clusters in the junction graph and exponential in the maximal cluster size. The worst case time complexities are linear in the size of the junction graph and exponential in maximal cluster size too. We compare the upper bounds of time and space of local consistency enforcing algorithms for initialized junction graphs in Table 1. It shows that all of them use the same space, though achieving Single Cluster Consistency uses the least time, followed by Directional Cluster Consistency, and then Neighborhood Cluster Consistency.

As shown in Table 1, the upper bounds of both time and space for achieving local consistencies using cluster consistency enforcing algorithms proposed in this paper are

Input: A CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$, its initialized junction graph representation $\mathcal{J} = (\mathcal{C}, \mathcal{S})$, and a total ordering $\mathcal{OC}$ of clusters in $\mathcal{C}$

Output: A Directional Cluster Consistent CBI problem $\mathbf{P}' = (\mathbf{X}, \mathbf{D}', \mathbf{S}, \mathbf{F}')$

for $i = |\mathcal{C}| - 1$ to 1 **do**
    Let $C_i = \mathcal{OC}[i]$
    for $j = |\mathcal{C}|$ to $i + 1$ **do**
        Let $C_j = \mathcal{OC}[j]$
        $\phi_{C_i} := \phi_{C_i} \otimes (\oplus_{C_j - S_{ij}} \phi_{C_j})$
    end for
    if *REVISE*(X, ϕ_{C_i}) **then**
        for each $C_k \in \mathcal{C}$ **do**
            if $X \in Scope(\phi_{C_k})$ **then**
                Remove all tuples in ϕ_{C_j} with the value that is removed from X
            end if
        end for
    end if
end for
Return $\mathbf{P}' := \mathbf{P}$

Fig. 3. Directional cluster consistency enforcing algorithm (*DCC-Enforcing*).

bounded by the maximum cluster size as well as the structure of the junction graph for a given CBI problem. Intuitively a simple junction graph implies large cluster sizes, so there is a tradeoff between the size of the graph and the largest cluster when constructing a junction graph. Various heuristic search approaches are discussed in [11] that can be used to construct junction graphs.

The junction tree representation is a special case of junction graphs that satisfies the tree property. The junction tree algorithm [12] is a widely studied inference algorithm in probability inference that utilizes the properties of the junction tree structure. It is also generalized to handle constraint-based inference problems [11], based on the seminal work on constraint programming [13, 14] and the latest general algorithmic framework [15]. Given the identical message representation and updating scheme in the inward phase of the junction tree algorithm and our directional cluster consistency enforcing algorithm (with a cluster order given by the width-first traverse starting from the root cluster), it is straightforward to show that the directional cluster consistency can be achieved along with the inward message passing in the junction tree algorithm, if the junction graph of a given CBI problem satisfies junction tree properties. A CBI problem processed by such a *DCC-enforcing* procedure then can be solved through backtrack-free search starting from the root cluster, which is a process equivalent to outward message propagation in the junction tree algorithm. This observation ensures that we can perform the message passing of junction tree algorithms and at the same time simplify the original problem representation according to the directional cluster consistency enforcement. Performing the message propagation and the simplification together reduces both the time and space complexities of the junction tree algorithm. The nature of the message passing scheme in the junction tree algorithm ensures that the directional cluster consistency enforcing can be performed in parallel for clusters

Input: A CBI problem $\mathbf{P} = (\mathbf{X}, \mathbf{D}, \mathbf{S}, \mathbf{F})$ and its initialized junction graph representation $\mathcal{J} = (\mathcal{C}, \mathcal{S})$
Output: A Single Cluster Consistent CBI problem $\mathbf{P}' = (\mathbf{X}, \mathbf{D}', \mathbf{S}, \mathbf{F}')$

for each $C_i \in \mathcal{C}$ **do**
  for each $C_j \in Neighbors(C_i)$ **do**
    $\phi_{C_i} := \phi_{C_i} \otimes (\oplus_{C_j - S_{ij}} \phi_{C_j})$
  end for
  if *REVISE*(X, ϕ_{C_i}) **then**
    for each $C_k \in \mathcal{C}$ **do**
      if $X \in Scope(\phi_{C_k})$ **then**
        Remove all tuples in ϕ_{C_j} with the value that is removed from X
      end if
    end for
  end if
end for
Return $\mathbf{P}' := \mathbf{P}$

Fig. 4. Neighborhood cluster consistency enforcing algorithm (*NCC-Enforcing*).

in different branches of the tree. We will investigate different parallel and hybrid DCC-enforcing techniques following the results of [14] in future work.

Loopy message propagation [16] is another widely studied approximation inference approach based on the junction graph representation in probability inferences. It is also generalized to apply to other CBI problems [11] using the semiring concepts. Neighborhood cluster consistency can be achieved along with each message updating step in the generalized loopy message propagation without additional computational cost except invalid value detection at each cluster. The time and space complexities of loopy message propagation are reduced after invalid values are removed from the CBI problem following NCC enforcement. The message updating step as well as NCC enforcement of the generalized loopy message propagation can be performed in all clusters in parallel saving significant computational cost if parallel computing is feasible.

5 Experimental Results

We discuss in this section experimental results of applying the local consistency enforcing algorithms proposed in this paper to the junction graph representation of Weighted CSP and Probability Assessment that can be modelled as CBI problems. These preprocessing or filtering algorithms simplify the original problem so that inference algorithms can then be applied with less computational complexity. A workstation with Pentium 4 3.0GHz CPU and 1 Gigabyte memory running SuSE Linux 9.1 is used to run the experiments in this section.

5.1 Weighted CSP

Weighted CSPs are direct extension of MaxCSPs where each value assignment in a constraint corresponds to a non-negative integer or weight instead of 0 for legal and

Input: A variable $X \in \mathbf{X}$, a constraint f, an element $\epsilon \in \mathbf{A}$
Output: TRUE if a value is removed from the domain of X; FALSE if else

```
flag := TRUE
for each x ∈ D_X do
    for each value assignment w of Scope(f)_{-X} do
        if ε ≤_S f(x, w) then
            flag := FALSE
            Break loop
        end if
    end for
    if flag then
        Remove x from D_X
        Return TRUE
    end if
end for
Return FALSE
```

Fig. 5. Procedure ϵ-*REVISE*(X, f, ϵ) for eliminating a domain value from a variable X according to the approximate threshold ϵ of a local constraint f.

1 for forbidden in MaxCSPs. Two constraint tuple weights are combined with arithmetic plus and the goal of the inference is to find a value assignment of all variables that minimizes the combination of all constraints in the problem. Weighted CSPs then can be easily embedded into the semiring-based CBI framework using the semiring $\mathbf{S_{WCSP}} = \langle \mathbb{Z}^+ \cup \{0\}, \min, + \rangle$. Because the multiplicative absorbing element $\alpha_\otimes$ of the semiring $\mathbf{S_{WCSP}}$ is equal to the additive identity element **0** that is equal to $+\infty$, $\mathbf{S_{WCSP}}$ is eliminative. Also we can show that $\mathbf{S_{WCSP}}$ is monotonic, so both the exact and approximate local consistency enforcing schemes in this paper apply to Weighted CSPs.

We study a random binary Weighted CSP with 100 variables and 200 constraints. The domain of each variable consists of 5 values. We choose randomly a weight from 0 to 10 for each value assignment of every constraint. We construct junction graphs through restricting the maximum cluster size from 2 to 4. Then we apply SCC, DCC and NCC enforcing algorithms to preprocess this Weighted CSP with various ϵ that approximates the multiplicative absorbing element $\alpha_\otimes = \infty$. The efficiency of the preprocessing algorithms is characterized as the average variable domain size. We show the experimental results in Table 2. For the purpose of comparison, we normalize the local constraint at each cluster before performing invalid value detection. Given these experimental results, we conclude: (1) For all of these approximate local cluster consistency enforcing algorithms, the closer ϵ is to the exact multiplicative absorbing element $\alpha_\otimes = \infty$, the fewer domain values are eliminated during the preprocessing. (2) The preprocessing time for each local cluster consistency enforcing algorithm is affected by the structure of the junction graph, but it does not change monotonically with the maximal cluster size. (3) In sequential computing schemes, SCC uses the least preprocessing time, followed by DCC and then NCC. The time used by DCC or NCC can be reduced if parallel computing is introduced. (4) DCC has the strongest preprocessing ability

	$k=2$				$k=3$				$k=4$			
$\vert\mathcal{C}\vert$	200				107				69			
$\vert\mathcal{S}\vert$	1500				772				520			
Max Degree	26				24				26			
Algorithm	SCC	DCC	NCC-1	NCC-3	SCC	DCC	NCC-1	NCC-3	SCC	DCC	NCC-1	NCC-3
$\epsilon = 7$	4.84	2.40	3.17	1.38	4.70	2.69	2.90	1.44	4.41	2.94	3.02	1.50
$\epsilon = 10$	5.00	2.88	4.21	1.58	4.97	3.36	3.97	1.65	4.94	3.51	4.05	1.87
$\epsilon = 15$	5.00	3.35	4.85	1.87	5.00	4.05	4.79	2.02	4.98	4.22	4.76	2.28
$\epsilon = 25$	5.00	3.97	5.00	2.36	5.00	4.44	4.96	2.61	5.00	4.65	4.97	2.88
$\epsilon = 50$	5.00	4.32	5.00	3.09	5.00	4.88	5.00	3.39	5.00	4.91	5.00	3.75
$\epsilon = 75$	5.00	4.55	5.00	3.53	5.00	4.94	5.00	3.93	5.00	4.97	5.00	4.22
$\epsilon = 100$	5.00	4.66	5.00	3.88	5.00	4.97	5.00	4.22	5.00	4.98	5.00	4.46
$\epsilon = 125$	5.00	4.77	5.00	4.08	5.00	4.98	5.00	4.36	5.00	4.99	5.00	4.59
$\epsilon = 150$	5.00	4.82	5.00	4.24	5.00	4.99	5.00	4.53	5.00	5.00	5.00	4.66
$\epsilon = 200$	5.00	4.90	5.00	4.46	5.00	5.00	5.00	4.65	5.00	5.00	5.00	4.75
$\epsilon = 500$	5.00	4.94	5.00	4.89	5.00	5.00	5.00	4.99	5.00	5.00	5.00	4.92
Avg. Time (s)	0.16	16.25	42.88	180.65	2.62	12.72	30.86	137.48	3.55	16.25	43.85	234.80

Table 2. Average variable domain sizes for a Weighted CSP that is preprocessed by three approximate local cluster consistency enforcing algorithms with different junction graph representations and different approximate element ϵ. Here SCC stands for Single Cluster Consistency; DCC stands for Directional Cluster Consistency; NCC-n stands for Neighborhood-Cluster Consistency with n steps of message updating. k is the size of the maximal cluster in a junction graph $\mathcal{J} = (\mathcal{C}, \mathcal{S})$. The original average domain size is 5 and we normalize the local constraint at each cluster for comparison.

due to its "global" property. In other words, message passing from lower order clusters contains information from clusters that are lower than them. NCC with one step message updating is slightly better than SCC in that a cluster in NCC collects information from all its immediate neighbors. If we perform several steps (3 in our experiments) of message updating, more values are removed.

5.2 Probability Assessment

Probability inference problems can be seen as constraint-based inference by treating conditional probability distributions (CPDs) as soft constraints over variables. The probability distribution over all the variables is given by the combination of the CPDs using the arithmetic product. The probability assessment problem computes the posterior marginal probability of a subset of variables, given values for some variables as known evidence. The probability assessment problem can be represented as a CBI problem using the commutative semiring $\mathbf{S_{PROB}} = \langle \mathbb{R}^+ \cup \{0\}, +, \times \rangle$. It is easy to show that $\alpha_\otimes = \mathbf{0} = 0$ in $\mathbf{S_{PROB}}$ and $\mathbf{S_{PROB}}$ is monotonic that both the exact and approximate local consistency enforcing schemes in this paper apply to probability assessment problems.

The Bayesian network used here is the Insurance network from the Bayesian Network Repository [17]. The network has 27 variables and 27 non-binary constraints

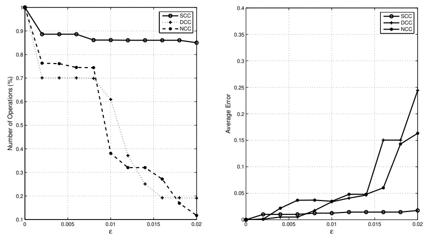

A: Inference Speed B: Average Error

Fig. 6. The number of binary operations required for probability assessment after using the local cluster consistency enforcing algorithms (shown as a fraction of the number required without preprocessing) and the resultant average error of the marginal probability for the Insurance network as a function of ϵ

(CPDs). In our experiments, we randomly choose one variable as observed. The approximate local cluster consistency enforcing algorithms are used to preprocess the problem based on a junction graph representation with the maximal cluster size of 5. The junction tree algorithm in Lauritzen-Spiegelhalter architecture [18] is used to infer the marginal probability of every unobserved variable. We compare the number of binary operations required for probability assessment after using the preprocessing algorithms (shown as a fraction of the number required without preprocessing) and the resultant error of the marginal probability for the Insurance network as a function of ϵ in Figure 6. At each value of ϵ, we collect data for 5 runs. It is clear that ϵ controls the tradeoff of the precision and the speed of the inference. DCC and NCC have stronger ability of speeding up the inference but introduce more errors than SCC, so they are more sensitive to the selection of the approximation threshold ϵ.

6 Conclusion

As the first contribution of this paper we propose a family of novel generalized local consistency concepts for the junction graph representation of CBI problems. These concepts apply to a broader coverage of inference problems from various fields based only on the general condition that depends on the properties of semirings that are used to abstract these problems. Second, we present several local consistency enforcing algorithms, including single, directional and neighborhood cluster consistency enforcing algorithms and their corresponding approximation variants. Third, theoretical space and time complexities of these preprocessing or consistency enforcing algorithms are discussed and experimental results of applying them to both MaxCSP and probability assessment problems are given. Finally, we discuss the relationship between these local cluster consistency concepts and message passing schemes such as junction tree algorithms and loopy message propagation. We will study efficient approaches to combining together the local cluster consistency enforcing with the message propagation for general CBI problems in future work.

References

1. Mackworth., A.K.: Consistency in networks of relations. Artificial Intelligence **8** (1977) 99–118
2. Mackworth, A.K.: On reading sketch maps. In: IJCAI77. (1977) 598–606
3. Mohr, R., Henderson, T.: Arc and path consistency revisited. Artificial Intelligence **28** (1986) 225–233
4. Wallace, R.: Directed arc consistency preprocessing as a strategy for maximal constraint satisfaction. In: Proc. of ECAI'94 Workshop on Constraint Processing, Amsterdam (1994)
5. Larrosa, J., Schiex, T.: In the quest of the best form of local consistency for weighted csp. In: Proc. of IJCAI-03, Acapulco, Mexico (2003) 239–244
6. Bistarelli, S., Montanari, U., Rossi, F.: Semiring-based constraint satisfaction and optimization. J. ACM **44** (1997) 201–236
7. Schiex, T., Fargier, H., Verfaillie, G.: Valued constraint satisfaction problems: Hard and easy problems. In: IJCAI95, Montreal (1995) 631–637
8. Cooper, M., Schiex, T.: Arc consistency for soft constraints. Artificial Intelligence **154** (2004) 199–227
9. Bistarelli, S.: Semirings for Soft Constraint Solving and Programming. Springer-Verlag (2004)
10. Chang, L., Mackworth, A.K.: A generalization of generalized arc consistency: From constraint satisfaction to constraint-based inference. In: Proc. of 5th Workshop on Modelling and Solving Problems with Constraints, Edinburgh (2005) 68–75 www.cs.ubc.ca/~lechang/publications/IJCAI05_Wshp.pdf.
11. Chang, L.: Generalized constraint-based inference. Master's thesis, Dept. of Computer Science, Univ. of British Columbia (2005)
12. Shenoy, P.P., Shafer, G.: Axioms for probability and belief-function propagation. In: Proc. of UAI90. (1990) 169–198
13. Dechter, R., Pearl, J.: Tree clustering for constraint networks. Artif. Intell. **38** (1989) 353–366
14. Zhang, Y., Mackworth, A.K.: Parallel and distributed finite constraint satisfaction: Complexity, algorithms and experiments. In: Parallel Processing for Artificial Intelligence. Elsevier Science Publishers (1994) 305–334
15. Kask, K., Dechter, R., Larrosa, J., Dechter, A.: Unifying cluster-tree decompositions for reasoning in graphical models. Artificial Intelligence **166** (2005) 165–193
16. Murphy, K.P., Weiss, Y., Jordan, M.I.: Loopy belief propagation for approximate inference: An empirical study. In: Proc. of UAI99. (1999) 467–475
17. Friedman, N., Goldszmidt, M., Heckerman, D., Russell, S.: (Bayesian network repository, http://www.cs.huji.ac.il/labs/compbio/repository/)
18. Lauritzen, S.L., Spiegelhalter, D.J.: Local computations with probabilities on graphical structures and their application to expert systems. Journal of the Royal Statistical Society, Series B **50** (1988) 157–224