A unified theory of terminal-valued and edge-valued decision diagrams

Gianfranco Ciardo   Andrew S. Miner

Department of Computer Science
Iowa State University

September 20, 2014

Work supported in part by the National Science Foundation under grant CCF-1442586

Many types of decision diagrams

Extending the domain from $\mathbb{B}^L$ to $\mathcal{X} = \mathcal{X}_L \times \cdots \times \mathcal{X}_1$ is straightforward

type	domain	range
BDD	$\mathbb{B}^L$	$\mathbb{B}$
MDD	$\mathcal{X}$	$\mathbb{B}$
MTBDD	$\mathbb{B}^L$	$\mathbb{Z}$, $\mathbb{R}$
MTMDD	$\mathcal{X}$	$\mathbb{Z}$, $\mathbb{R}$
ADD	$\mathbb{B}^L$	any

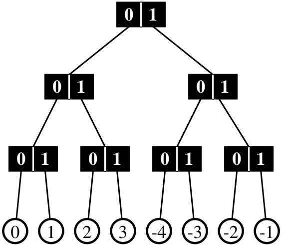

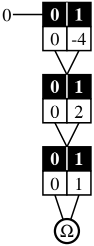

type	edge values	combinator	domain	range
EVBDD	$\mathbb{Z}$	sum	$\mathbb{B}^L$	$\mathbb{Z}$
EV$^+$MDD	$\mathbb{N} \cup \{\infty\}$	sum	$\mathcal{X}$	$\mathbb{Z} \cup \{\infty\}$
PDG	probabilities	multiply	$\mathbb{B}^L$	$[0, 1]$
EV*MDD	$[0, 1]$	multiply	$\mathcal{X}$	$\mathbb{R}^{\geq 0}$
AADD	$\mathbb{R} \times \mathbb{R}$	$(a,b) \odot (c,d) = (a+bc, bd)$	$\mathbb{B}^L$	$\mathbb{R}$

Motivating questions

Can we unify terminal and edge valued decision diagrams?
Can we achieve more elegance, simplicity, and generality?

Why can MTBDDs encode any partial function $\mathbb{B}^L \to \mathbb{Z} \cup \{\infty\}$?
Why can't EVBDDs (in their original definition) do that?
Why can EV$^+$MDDs (our canonical definition) do that?
This is why we introduced EV$^+$MDDs, but what is the key issue?

Why can EV$^+$MDDs have range $\mathbb{R}$ but EV*MDDs must have range $\mathbb{R}^{\geq 0}$?
Can EV*MDDs encode CTMC generators, not just rate matrices?

What are the advantages/disadvantages of terminal vs. edge valued DDs?
Which decision diagram encoding should I use for a particular application?

ET-monoid

Semigroup $(\mathcal{S}, \odot)$: set $\mathcal{S}$ is closed w.r.t. the associative binary operator $\odot$

Monoid $(\mathcal{S}, \odot)$: semigroup where $\mathcal{S}$ contains identity element e

Group: monoid $(\mathcal{S}, \odot)$ where every element of $\mathcal{S}$ has an inverse in $\mathcal{S}$

Total order $\succeq$: transitive, antisymmetric, and total binary relation

Definition

ET-monoid $M = (\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$: $(\mathcal{S}, \odot)$ is a monoid, $\{e\} \subseteq \mathcal{S}_E \subseteq \mathcal{S}$, $\mathcal{S}_T \subset \mathcal{S}$, $\mathcal{S}_E \cap \mathcal{S}_T = \emptyset$, $\succeq$ is a total order on $\mathcal{S}$, and

Axiom 1: $\mathcal{S}_E$ is closed over $\odot$ — $\mathcal{S}_E \odot \mathcal{S}_E \subseteq \mathcal{S}_E$

Axiom 2: $\mathcal{S}_T$ terminates $\mathcal{S}_E$ from the right — $\mathcal{S}_E \odot \mathcal{S}_T \subseteq \mathcal{S}_T$

Axiom 3: Each element in $\mathcal{S}_E$ has an inverse in $\mathcal{S}$ — $\forall a \in \mathcal{S}_E, \exists a^{-1} \in \mathcal{S}$

Axiom 4: For any sequence $\Sigma \in \mathcal{C}_M^+$, there exists some $\sigma \in \Sigma$ and $m \in \mathcal{D}(\sigma)$ such that σ is an optimal sequence for m in Σ

the total order $\succeq$ defines a "desirability" on (sequences of) edge values

Axiom 5: For any sequence $\Sigma \in \mathcal{C}_M^+$ and any $a \in \mathcal{S}_E$, if σ is an optimal sequence for m in Σ, then $a \odot \sigma$ is an optimal sequence for $a \odot m$ in $a \odot \Sigma$

this is quite complex, but necessary for canonicity in our general setting

A few ET-monoids $(\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$

- $(\mathbb{B}, \{0\}, \{1\}, \vee, any)$
- $(\mathbb{B}, \{1\}, \{0\}, \wedge, any)$
- $(\mathbb{B}, \mathbb{B}, \emptyset, \oplus, any)$ — $\oplus$ is exclusive–or, 0 is the identity
- $(\mathbb{Z}, \mathbb{Z}, \emptyset, +, \succeq)$ — $a \succeq b$ if $|a| < |b|$ or $|a| = |b|$ and $a \geq 0 \geq b$
- $(\mathbb{Z}, \{0\}, \mathbb{Z} \setminus \{0\}, +, \succeq)$ — 0 is the identity and the most desirable
- $(\mathbb{Z} \cup \{\infty^+\}, \mathbb{Z}, \{\infty^+\}, +, \succeq)$
- $(\mathbb{Z} \cup \{\infty^-\}, \mathbb{Z}, \{\infty^-\}, +, \succeq)$
- $(\mathbb{Z} \cup \{\infty^+, \infty^-, \mu\}, \mathbb{Z}, \{\infty^+, \infty^-\}, +, \succeq)$ — μ means "undefined"
- substitute $\mathbb{Z}$ with $\mathbb{Q}$ or $\mathbb{R}$ in the above four
- $(\mathbb{R}, \mathbb{Z}, \{n + \sqrt{2} : n \in \mathbb{Z}\}, +, \succeq)$
- $(\mathbb{Z}, \mathbb{N}, \emptyset, +, \leq)$ — 0 is the identity and the most desirable
- $(\mathbb{Z} \cup \{\infty^+\}, \mathbb{N}, \{\infty^+\}, +, \leq)$
- substitute $\mathbb{Z}$ with $\mathbb{Q}$ or $\mathbb{R}$ and $\mathbb{N}$ with $\mathbb{Q}^{\geq 0}$ or $\mathbb{R}^{\geq 0}$ in the above two

A few more ET-monoids $(\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$

- $(\mathbb{Q}, \mathbb{Q}^{>0}, \emptyset, \cdot, \succeq)$ $\quad a \succeq b$ if $|\ln a| < |\ln b|$ or $|\ln a| = |\ln b|$ and $a \geq 1 \geq b$
- $(\mathbb{Q}, \mathbb{Q}^{>0}, \{0\}, \cdot, \succeq)$ $\quad$ 1 is the identity and the most desirable
- $(\mathbb{Q} \cup \{\infty^+\}, \mathbb{Q}^{>0}, \{\infty^+\}, \cdot, \succeq)$
- $(\mathbb{Q} \cup \{\infty^+, \mu\}, \mathbb{Q}^{>0}, \{0, \infty^+\}, \cdot, \succeq)$
- substitute $\mathbb{Q}$ with $\mathbb{R}$ and $\mathbb{Q}^{>0}$ with $\mathbb{R}^{>0}$ in the above four

- $(\mathbb{Q}, \mathbb{N} \setminus \{0\}, \emptyset, \cdot, \leq)$ $\quad$ 1 is the identity and the most desirable
- $(\mathbb{Q}, \mathbb{N} \setminus \{0\}, \{0\}, \cdot, \leq)$
- $(\mathbb{Q} \cup \{\infty^+, \mu^+\}, \mathbb{N} \setminus \{0\}, \{\infty^+\}, \cdot, \leq)$
- $(\mathbb{Q} \cup \{\infty^+, \mu^+\}, \mathbb{N} \setminus \{0\}, \{0, \infty^+\}, \cdot, \leq)$

- $(\mathbb{R} \times \mathbb{R}, \mathbb{R} \times \mathbb{R}^{>0}, \emptyset, \odot, \succeq)$ $\quad (a, b) \odot (c, d) = (a + bc, bd)$
  $\succeq$ is such that the identity $(0, 1)$ is the most desirable

ETDDs: edge-and-terminal valued decision diagrams

First, a non-canonical version of an ETDD (forest):

Definition

Given domain $\mathcal{X} = \mathcal{X}_L \times \cdots \mathcal{X}_1$ and ET-monoid $M = (\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$, an ordered, ET-valued decision diagram (ETDD) over $(\mathcal{X}, M)$ is an acyclic, node-labeled, and edge-labeled multi-graph where:

- Each node p is at a level $p.lvl = k$, with $L \geq k \geq 0$
- Only terminal node Ω is at level 0
- Node p at level $k > 0$ has $n_k = |\mathcal{X}_k|$ edges;
  for $i_k \in \mathcal{X}_k$, $p[i_k] = \langle a, q \rangle$ means edge i_k has value $a \in \mathcal{S}_E \cup \mathcal{S}_T$ and points to node q at level $h < k$; let $p[i_k].val = a$ and $p[i_k].node = q$
- There is a non-empty set of root edges $\mathcal{R}$; for any root edge $\langle a_\star, p_\star \rangle \in \mathcal{R}$, $a_\star \in \mathcal{S}_E \cup \mathcal{S}_T$ and $p_\star$ is a node at level $k_\star$, $L \geq k_\star \geq 0$
- For any edge $\langle a, q \rangle$, including a root edge, if $a \in \mathcal{S}_T$, then $q = \Omega$
- Every node in the graph is reachable from some root edge

Toward canonical ETDDs

As usual, no duplicate or redundant nodes

Our general setting also requires constraints on edge values and nodes:

- If $(\mathcal{S}_E, \odot)$ is a group and $\mathcal{S}_T = \emptyset$, force $p[0].val = e$ (e.g., EVBDDs)
- If not, canonicity is surprisingly elusive, we need to use desirability

e.g., ET-monoid $(\mathbb{Q} \cup \{\infty^+, \infty^-, \mu^+, \mu^-\}, \mathbb{Z}\setminus\{0\}, \{0, \infty^+, \infty^-\}, \cdot, \succeq)$

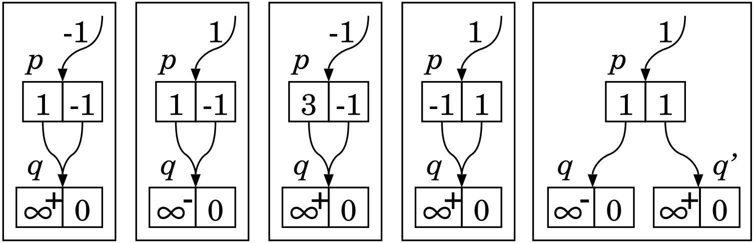

Edge normalization ⇒ use the **representative** for the equivalence class

Node normalization ⇒ divide the node by its **most desirable divisor**

Canonicity

Definition

An ETDD is (fully) reduced if its edges and non-terminal nodes are normalized and contains no duplicate nodes

Theorem

Given a nonempty, finite set of vectors $\mathcal{V} \subseteq \mathcal{X} \to \mathcal{S}_E \cup \mathcal{S}_T$, there exists a reduced ETDD with root edges $\mathcal{R}$ such that $\mathcal{V}(L, \mathcal{R}) = \mathcal{V}$

Definition

An ETDD is scalar–independent if, for any nodes p, q with
$p.lvl = q.lvl = k$, if $\mathbf{v}(p) = a \odot \mathbf{x}$ and $\mathbf{v}(q) = b \odot \mathbf{x}$
for vector $\mathbf{x} : \mathcal{X}_k \times \cdots \times \mathcal{X}_1 \to \mathcal{S}_E \cup \mathcal{S}_T$ and scalars $a, b \in \mathcal{S}_E$
then $p = q$

Theorem

Every reduced ETDD is scalar–independent

Space comparisons: ETDDs over the same ET-monoid

Theorem

Given an ET-monoid $M = (\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$, an ETDD G over $(\mathcal{X}, M)$ with root edges $\mathcal{R}_G$, and a **reduced** ETDD H over $(\mathcal{X}, M)$ with root edges $\mathcal{R}_H$, if $\mathcal{V}(L, \mathcal{R}_G) = \mathcal{V}(L, \mathcal{R}_H)$, then G is homomorphic to H...

The reduced ETDD encoding is minimal (for a given $\mathcal{X}$ and M)

Theorem

...if G is scalar–independent with no redundant nodes, it is isomorphic to H

Any normalization that guarantees scalar independence works equally well

Implications for ETDDs over the same ET-monoid

Applying desirability to sequences left-to-right vs. right-to-left works equally well

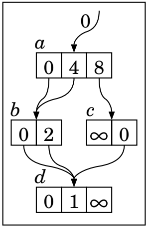

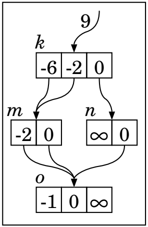

Using EVBDDs with $p[0].val = 0$ vs. $\min\{p[i_k].val : i_k \in \mathcal{X}_k\} = 0$ works equally well

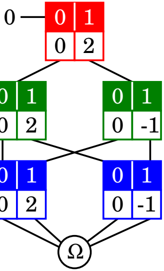

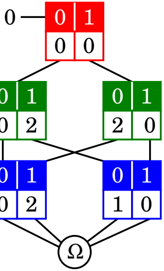

Space comparisons: ETDDs over different ET-monoids

Definition

Given $M_G = (\mathcal{G}, \mathcal{G}_E, \mathcal{G}_T, \odot, \succeq_G)$ and $M_H = (\mathcal{H}, \mathcal{H}_E, \mathcal{H}_T, \oplus, \succeq_H)$, M_G is homomorphic to M_H via function $f : \mathcal{G} \to \mathcal{H}$ if

- $a \in \mathcal{G}_E$ and $b \in \mathcal{G}_E \cup \mathcal{G}_T \Rightarrow\ f(a \odot b) = f(a) \oplus f(b)$
- $a \in \mathcal{G}_E \Rightarrow f(a) \in \mathcal{H}_E$ $\qquad$ $t \in \mathcal{G}_T \Rightarrow f(t) \in \mathcal{H}_E \cup \mathcal{H}_T$

If f is one-to-one, M_G is lossless–homomorphic to M_H via f
Otherwise, M_H is lossy–homomorphic to M_H via f

M_G is isomorphic to M_H via bijection f if M_G is homomorphic to M_H via f and M_H is homomorphic to M_G via f^{-1}

Implications for ETDDs over different ET-monoids

Theorem

If M_G is homomorphic to M_H via f, the reduced ETDD over $(\mathcal{X}, M_G)$ with root edges $\mathcal{R}_G$ s.t. $\mathcal{V}(L, \mathcal{R}_G) = \mathcal{V}$, is homomorphic to the reduced ETDD over $(\mathcal{X}, M_H)$ with root edges $\mathcal{R}_H$ s.t. $\mathcal{V}(L, \mathcal{R}_H) = f(\mathcal{V})$

Compl.-edge BDDs $(\mathbb{B},\mathbb{B},\emptyset,\oplus,\succeq)$ never worse than BDDs $(\mathbb{B},\{0\},\{1\},\vee,\succeq)$

Lemma

If G is isomorphic to H via some f, the reduced ETDD over $(\mathcal{X}, G)$ is isomorphic to the reduced ETDD over $(\mathcal{X}, H)$ encoding the same vector

Any isomorphic ET-monoid works equally well

Lemma

ET-monoid $M' = (\mathcal{S}, \{e\}, \mathcal{S}_E \cup \mathcal{S}_T \setminus \{e\}, \odot, \succeq)$ is lossless–homomorphic to any ET-monoid $M = (\mathcal{S}, \mathcal{S}_E, \mathcal{S}_T, \odot, \succeq)$ via the identity function

A reduced ETDD is never worse than the equivalent reduced MTMDD

Conclusion and ongoing work

Our ETDD framework unifies many types of decision diagrams

We found the key canonicity requirements in a general setting

Our general theorems provide results for popular decision diagram classes

We are still completing work on the time complexity of ETDD algorithms (complexity improves if ET-monoid has more structure, e.g., $\mathcal{S}_E$ is a group)

The current paper is already 53 pages so far :-(