

Bounded Dataflow Networks (BDNs) and Latency-Insensitive (LI) Circuits

Murali Vijayaraghavan and Arvind
Computer Science and Artificial Intelligence Laboratory
M.I.T.

WG 2.8, Chiemsee, Germany
June 12, 2009

<http://csg.csail.mit.edu>

Dataflow networks

- ◆ Kahn-Dennis networks: A network of computing stations connected by unbounded FIFOs
 - a “get” is blocking but a “put” is not
- ◆ Dataflow networks with bounded FIFOs (BDNs)
 - Hard to model as a Kahn-Dennis network
 - Varying the size of a FIFO changes the meaning (may cause a deadlock)
- ◆ Several groups are using BDNs for latency-insensitive refinements of *Synchronous Sequential Machines* (SSMs) and often encounter deadlocks

June 12, 2009

<http://csg.csail.mit.edu>

Bounded Dataflow Networks

- ◆ Can be modeled accurately in Bluespec
- ◆ Can be used a high-level structuring technique for Bluespec designs


what restrictions should be placed on BDNs such that its meaning does not change with respect to a given SSM when we vary the FIFO sizes

June 12, 2009

<http://csg.csail.mit.edu>

Examples of primitive BDNs:

A Combinational Block



f is a combinational circuit:
must accept an input value
on each input before
producing an output

Behavior


```
rule CL when ( $\neg a.empty \wedge \neg b.empty \wedge \neg c.full$ )  
   $\Rightarrow c.enq(f(a.first, b.first)); a.deq ; b.deq$ 
```

Unlike SSMs, the (red) lines only show dataflow and not all the control lines needed to make BDNs function

June 12, 2009

<http://csg.csail.mit.edu>

A fork definition



a fork that copies an input value to both its outputs simultaneously

Behavior


```
rule F when ( $\neg$ a.empty  $\wedge$   $\neg$ b.full  $\wedge$   $\neg$ c.full)
 $\Rightarrow$  b.enq(a.first); c.enq(a.first); a.deq
```

June 12, 2009

<http://csg.csail.mit.edu>

Examples of primitive BDNs:

Fork



a fork to copy an input value but the input can be dequeued only when both the outputs have accepted the input

Behavior

```
rule FO1 when ( $\neg$ a.empty  $\wedge$   $\neg$ b.full  $\wedge$   $\neg$ bDone)
 $\Rightarrow$  b.enq(a.first); bDone  $\leq$  True

rule FO2 when ( $\neg$ a.empty  $\wedge$   $\neg$ c.full  $\wedge$   $\neg$ cDone)
 $\Rightarrow$  c.enq(a.first); cDone  $\leq$  True

rule FI when ( $\neg$ a.empty  $\wedge$  bDone  $\wedge$  cDone)
 $\Rightarrow$  a.deq; bDone  $\leq$  False;
cDone  $\leq$  False
```

Initial Values

```
bDone = False
cDone = False
```


Which one do we want?

June 12, 2009

<http://csg.csail.mit.edu>

Examples of primitive BDNs:

Register



A register whose reads and writes must match

Behavior

```
rule RO when ( $\neg$ b.full  $\wedge$   $\neg$ bDone)
   $\Rightarrow$  b.enq(r); bDone  $\leftarrow$  True

rule RI when ( $\neg$ a.empty  $\wedge$  bDone)
   $\Rightarrow$  r  $\leftarrow$  a.first; a.deq; bDone  $\leftarrow$  False
```

Initial Values


```
bDone = False
r = r0
```

June 12, 2009

<http://csg.csail.mit.edu>

Examples of primitive BDNs:

Mux



A mux that accepts an input value on each input port but passes only the appropriate value to the output

Behavior

```
rule MuxO when  $\neg$ c.full  $\wedge$   $\neg$ p.empty
   $\Rightarrow$  if(p.first  $\wedge$   $\neg$ a.empty)
    then c.enq(a.first); a.deq; bCnt  $\leftarrow$  bCnt+1
    else if( $\neg$ (p.first)  $\wedge$   $\neg$ b.empty)
      then c.enq(b.first); b.deq; aCnt  $\leftarrow$  aCnt+1
rule MuxI1 when aCnt > 0  $\wedge$   $\neg$ a.empty
   $\Rightarrow$  a.deq; aCnt  $\leftarrow$  aCnt-1
rule MuxI2 when bCnt > 0  $\wedge$   $\neg$ b.empty
   $\Rightarrow$  b.deq; bCnt  $\leftarrow$  bCnt-1
```

Initial values


```
aCnt = 0
bCnt = 0
```

June 12, 2009


<http://csg.csail.mit.edu>

Composition of BDNs

- ◆ If R_1 and R_2 are BDNs then so is the parallel composition of R_1 and R_2 ($R = R_1 \oplus R_2$)



- ◆ R_1 is a BDN then so is the (I_i, O_j) iterative composition of R_1 ($R = (i, j) \otimes R_1$) provided $I_i \notin \text{Depends-on}(O_j)$ *




* No direct combinational path

June 12, 2009

<http://csg.csail.mit.edu>

BDN as a refinement of an SSM



- ◆ There is a bijective mapping between the inputs (outputs) of S and R
- ◆ for all $n > 0$,


$I(k)$ matches for S and R ($1 \leq k \leq n$)	}	Cycle
$\Rightarrow O(j)$ matches for S and R ($1 \leq j \leq n$)		

In general it is difficult to compare an SSM and a BDN because a BDN can deadlock. We will restrict our attention to a class of BDNs with some "desirable properties"

June 12, 2009

<http://csg.csail.mit.edu>

Deadlock-free BDN



- ◆ Assuming an infinite sink, a BDN is deadlock-free if for all $n > 0$, if n values are enqueued into I then eventually n values will be dequeued from both O and I
 - we need a stronger property for deadlock-freeness to be preserved under composition

June 12, 2009

<http://csg.csail.mit.edu>

NED-BDN:


BDNs with *no extraneous dependencies*

- ◆ A BDN is said to have *no extraneous dependencies* if its output O_i is not enqueued n times, assuming it is not full and all the inputs are enqueued $n-1$ times, then it must be that one of the inputs in $\text{Depends-on}(O_i)$ is not enqueued n times
- ◆ Note that this is a property of BDN – it is different from the condition for iterative composition

June 12, 2009

<http://csg.csail.mit.edu>

NED-BDN violation 1



a fork that copies an input value to both its outputs simultaneously


Behavior

```
rule F when ( $\neg a.empty \wedge \neg b.full \wedge \neg c.full$ )
   $\Rightarrow b.enq(a.first); c.enq(a.first); a.deq$ 
```

June 12, 2009

<http://csg.csail.mit.edu>

NED-BDN violation 2



Possible Behaviors

```
rule O when ( $\neg a.empty \wedge \neg b.empty \wedge \neg c.full \wedge \neg d.full$ )
   $\Rightarrow c.enq(f(a.first, b.first)); d.enq(b.first);$ 
   $a.deq; b.deq$ 
```


```
rule O1 when ( $\neg a.empty \wedge \neg b.empty \wedge \neg c.full \wedge \neg cDone$ )
   $\Rightarrow c.enq(f(a.first, b.first)); cDone \leq True$ 
rule O2 when ( $\neg b.empty \wedge \neg d.full \wedge \neg dDone$ )
   $\Rightarrow d.enq(b.first); dDone \leq True$ 
rule In when ( $cDone \wedge dDone$ )
   $\Rightarrow a.deq; b.deq$ 
```

June 12, 2009

<http://csg.csail.mit.edu>

Latency-Insensitive (LI) BDN

- ◆ LI-BDN is an NED-BDN which is refinement of an SSM



June 12, 2009

<http://csg.csail.mit.edu>

LI refinement Theorem

- ◆ LI-BDNs are composable under parallel and sequential composition

- If R_1, R_2 are refinements of $S_1, S_2 \Rightarrow$
 - ◆ $R_1 \oplus R_2$ is the refinement of $S_1 + S_2$
 - ◆ $(i, j) \otimes R_1$ is the refinement of $(i, j) \times S_1$

- ◆ Basically this ensures that the composition of LI-BDNs are deadlock-free and cycle-accurate w.r.t. the original SSMs

June 12, 2009

<http://csg.csail.mit.edu>


Application: Modeling via RTL prototyping on FPGAs

- ◆ Some RTL structures are inefficient to map directly onto FPGAs
 - For example, a 3-ported register file (RF) consumes lot of area as opposed a 1-ported RF used for 3 cycles
 - However, replacing a 3-ported RF naively by a 1-ported RF in a design may lose “cycle-accuracy”, even if the high-level functionality “turns out” to be correct

June 12, 2009

<http://csg.csail.mit.edu>

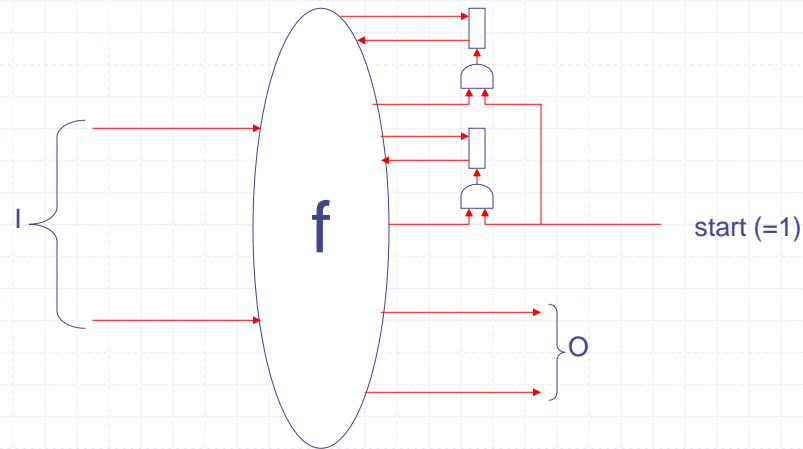
Application: Cycle-accurate modeling



June 12, 2009

<http://csg.csail.mit.edu>

Startable SSMs: SSMs with a “start” signal to update registers



June 12, 2009

<http://csg.csail.mit.edu>

1000 feet view of LI-refinement of an SSM

- ◆ FIFOs are introduced in every input and every output of the SSM
- ◆ Time cycles of the SSM are converted into enqueues into inputs and dequeues from outputs
 - “Cycle-accurate” w.r.t SSM
- ◆ Atomic rules for the operations are defined so that no extraneous dependencies are introduced
 - Ensures deadlock-free operation

June 12, 2009

<http://csg.csail.mit.edu>

Writing the LI-BDN wrapper for an SSM

LI-BDN:

```
rule j (!oj.done)
  oj.done <= True
  oj.enq( fj(ij1.first, ... , ijij.first, s) )

rule finish (o1.done && o2.done && ...)
  o1.done <= False; o2.done <= False; ...
  s <= g(i1.first, i2.first, ... , s)
  i1.deq ; i2.deq ; ...
```

Given the SSM:

$$o_j(t) = f_j(i_{j1}(t), \dots, i_{jij}(t), s(t))$$

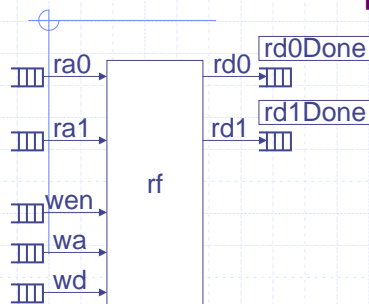
// $i_{j1}, i_{j2}, \dots, i_{jij}$ are in Depends-on(o_j)

$$s(t+1) = g(i_1(t), i_2(t), \dots, s(t))$$

June 12, 2009

<http://csg.csail.mit.edu>

2) Automatically generated LI-BDN for a 3-ported register file



This again uses 3 ports

```
rule RD0 when (~rd0Done)
  rd0.enq(rf_0[ra0.first])
  rd0Done <= True

rule RD1 when (~rd1Done)
  rd1.enq(rf_1[ra1.first])
  rd1Done <= True

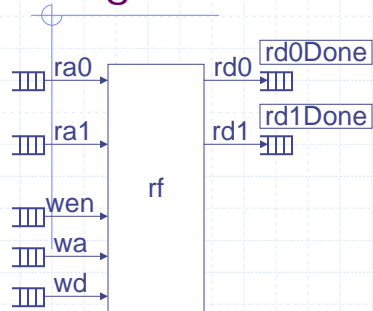
rule finish when (rd0Done ^ rd1Done)
  ra0.deq; ra1.deq
  wen.deq; wa.deq; wd.deq
  rf_2[wa.first] <= wen.first?
  wd.first : rf_2[wa.first]

rd0Done <= False
rd1Done <= False
```

June 12, 2009

<http://csg.csail.mit.edu>

Refinement into a one-ported register file LI-BDN



This uses 1 port

```
rule RD0 when (~rd0Done)
  rd0.enq( rf_0 [ra0.first])
  rd0Done <= True

rule RD1 when (~rd1Done)
  rd1.enq( rf_0 [ra1.first])
  rd1Done <= True

rule finish when (rd0Done ^ rd1Done)
  ra0.deq; ra1.deq
  wen.deq; wa.deq; wd.deq
  rf_0 [wa.first] <= wen.first?
  wd.first : rf_0 [wa.first]

rd0Done <= False
rd1Done <= False
```

June 12, 2009

<http://csg.csail.mit.edu>

Thanks

<http://csg.csail.mit.edu>