

Structured Wide-Area Programming: Orc Programming Examples

Jayadev Misra

Department of Computer Science
University of Texas at Austin

<http://orc.csres.utexas.edu>

Some Algorithms

- Enumeration and Backtracking
- Using Closures
- List Fold, Map-reduce
- Parsing using Recursive Descent
- Exception Handling
- Process Network
- Quicksort
- Graph Algorithms: Depth-first search, Shortest Path

Enumeration

Given: integer n , list of integers xs

Return all subsequences of xs that sum to n .

$\text{sum}(5, [1, 2, 1, 2]) = [1, 2, 2], [2, 1, 2]$

$\text{sum}(5, [1, 2, 1])$ is silent

```
def sum(0, []) = []
```

```
def sum(_, []) = stop
```

```
def sum(n, x : xs) =  
    sum(n - x, xs) >ys> x : ys  
    | sum(n, xs)
```

Backtracking: Use of Otherwise

Given: integer n , list of integers xs

Return the “first” subsequence of xs that sums to n .

`sum(5, [1, 2, 1, 2]) = [1, 2, 2]`

`sum(5, [1, 2, 1])` is silent

```
def sum(0, _) = []
```

```
def sum(_, []) = stop
```

```
def sum( $n$ ,  $x : xs$ ) =  $x : sum(n - x, xs)$  ; sum( $n$ ,  $xs$ )
```

Backtracking: Eight queens

Place 8 queens on a chessboard so that no queen captures another.

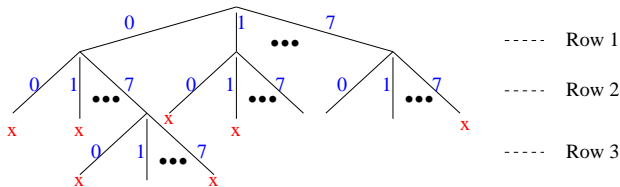


Figure: Backtrack Search for Eight queens

Eight queens; contd.

- *xs*: partial placement of queens (list of values from 0..7)
- *extend(xs)* publishes **all** solutions that are extensions of *xs*.
- *open(xs)* publishes the columns that are **open** in the next row.
- Solve the original problem by calling *extend*([]).

```
def extend(xs) =  
  if (length(xs) = 8) then xs  
  else (open(xs) >j> extend(j : xs))
```

Using Closure

A UNITY Program

$x, y = 0, 0$

$x < y \rightarrow x := x + 1$
| $y := y + 1$

- Program has: variable declarations
a set of functions
- Variables are initialized as given.
- Program is run by: choosing a function arbitrarily,
choosing functions fairly.

Corresponding Orc program

```
val (x, y) = (Ref(0), Ref(0))
```

```
def f1() = Ift(x? <: y?) >> x := x? + 1
```

```
def f2() = y := y? + 1
```

Run the program by:

- choosing a function arbitrarily,
- choosing functions fairly.

Scheduling the UNITY Program

```
def unity(fs) =  
  val arlen = length(fs)  
  val fnarray = Array(arlen)  
  
  { - populate() transfers from list fs to array fnarray - }  
  def populate(_, []) = signal  
  def populate(i, g : gs) = fnarray(i) := g >> populate(i + 1, gs)  
  
  { - Execute a random statement and loop.  
    Randomness guarantees fairness. - }  
  def exec() = random(arlen) >j> fnarray(j)?() >> exec()  
  
  { - Initiate the work - }  
  populate(0, fs) >> exec()
```

Running the example program

val $(x, y) = (\text{Ref}(0), \text{Ref}(0))$

def $f1() = \text{If}(x? <: y?) \gg x := x? + 1$

def $f2() = y := y? + 1$

unity $([f1, f2])$

Associative Fold

- Define $afold(f, xs)$ where f is an associative binary function and xs is a non-empty list.
- Goal is to combine elements in parallel.
- Each iteration reduces adjacent pairs of items to single values.
- Iterations continue until there is a single value.

Associative Fold; contd.

def *afold*(*f*, [*x*]) = *x*

def *afold*(*f*, *xs*) =

def *step*([]) = []

def *step*([*x*]) = [*x*]

def *step*(*x* : *y* : *xs*) = *f*(*x*, *y*) : *step*(*xs*)

afold(*f*, *step*(*xs*))

- $f(x, y) : step(xs)$ is an implicit fork-join.
- $f(x, y)$ executes concurrently with $step(xs)$.
- All calls to f execute concurrently within each iteration of $afold$.

Associative and Commutative Fold

- Transfer list items to a channel (arbitrary order of items).
- Fold any two channel items and put the result in the channel.

```
def acfold(f, xs) =  
  val c = Channel()
```

```
def xfer([]) = stop  
def xfer(x : xs) = (c.put(x), xfer(xs))
```

```
def combine(1) = c.get()
```

```
def combine(m) =  
  c.get() >x> c.get() >y>  
  (c.put(f(x, y)) >> stop | combine(m - 1))
```

```
xfer(xs) | combine(length(xs))
```

map-reduce

- Given is a list of tasks.
- A processor from a processor pool is assigned to process a task. Each task may be processed independently, yielding a result.
- If a processor does not respond within time T , a new processor is assigned to the task.
- After all the results have been computed, the results are reduced by calling *reduce*.

Implementation

- *processlist* processes a list of tasks concurrently.
process(t) processes a single task *t*.
process(t) publishes a result; *processlist* a list of results.
- Function *process* first acquires a processor.
It assigns the task to the processor.
If the processor responds within time *T*, it publishes the result.
Else, it repeats these steps.
- *process(t)* may never complete if the processors keep failing.
- The list of published results are reduced by function *reduce*.

map-reduce

```
def processlist([]) = []
```

```
def processlist(t : ts) = process(t) : processlist(ts)
```

```
def process(t) =
```

```
  val processor = Processorpool()
```

```
  val (result, b) = (processor(t), true) | (Rwait(T), false)
```

```
  if b then result else process(t)
```

```
processlist(tasks) >x> reduce(x)
```


Parsing using Recursive Descent

Consider the grammar:

$$\textit{expr} ::= \textit{term} \mid \textit{term} + \textit{expr}$$
$$\textit{term} ::= \textit{factor} \mid \textit{factor} * \textit{term}$$
$$\textit{factor} ::= \textit{literal} \mid (\textit{expr})$$
$$\textit{literal} ::= 3 \mid 5$$

Parsing strategy

For each non-terminal, say *expr*, define *expr(xs)*:
publish all suffixes of *xs* such that the prefix is a *expr*.

```
def isexpr(xs) = expr(xs) > [] > true ; false
```

To avoid multiple publications (in ambiguous grammars),

```
def isexpr(xs) =  
  val res = expr(xs) > [] > true ; false  
  res
```

----- Test

```
isexpr  
(["(", "(", "3", " * ", "3", ")"), " + ", "(", "3", " + ", "3", ")"])  
  — ((3*3))+(3+3)
```

```
:: true
```

Function for each non-terminal

Given: $expr ::= term \mid term + expr$

Rewrite: $expr ::= term (\epsilon \mid + expr)$

def $expr(xs)$ = $term(xs) \succ ys \succ (ys \mid ys \succ "+" : zs \succ expr(zs))$

def $term(xs)$ = $factor(xs) \succ ys \succ (ys \mid ys \succ "*" : zs \succ term(zs))$

def $factor(xs)$ = $literal(xs)$
| $xs \succ "(" : ys \succ expr(ys) \succ ")" : zs \succ zs$

def $literal(n : xs)$ = $n \succ "3" \succ xs \mid n \succ "5" \succ xs$

def $literal([])$ = $stop$

Exception Handling

Client calls `site server` to request service.

The server “may” request authentication information.

```
def request(x) =  
  val exc = Channel() -- returns a channel site  
  
  server(x, exc)  
  | exc.get() >r> exc.put(auth(r)) >> stop
```

Process Networks

- A process network consists of: processes and channels.
- The processes run autonomously, and communicate via the channels.
- A network is a process; thus hierarchical structure. A network may be defined recursively.
- A channel may have intricate communication protocol.
- Network structure may be dynamic, by adding/deleting processes/channels during its execution.

Channels

- For channel c , treat $c.put$ and $c.get$ as site calls.
- In our examples, $c.get$ is blocking and $c.put$ is non-blocking.
- We consider only FIFO channels.
Other kinds of channels can be programmed as sites.
We show rendezvous-based communication later.

Typical Iterative Process

Forever: Read x from channel c , compute with x , output result on e :

def $p(c, e) = c.get() >x> \text{Compute}(x) >y> e.put(y) \gg p(c, e)$

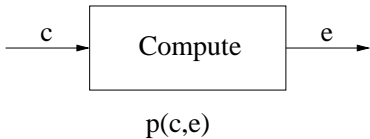


Figure: Iterative Process

Composing Processes into a Network

Process (network) to read from both c and d and write on e :

def $net(c,d,e) = p(c,e) \mid p(d,e)$

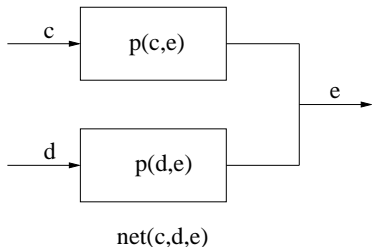


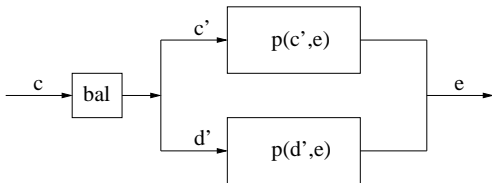
Figure: Network of Iterative Processes

Workload Balancing

Read from c , assign work randomly to one of the processes.

```
def bal(c, c', d') = c.get() >x> random(2) >t>  
  (if t = 0 then c'.put(x) else d'.put(x)) >>  
  bal(c, c', d')
```

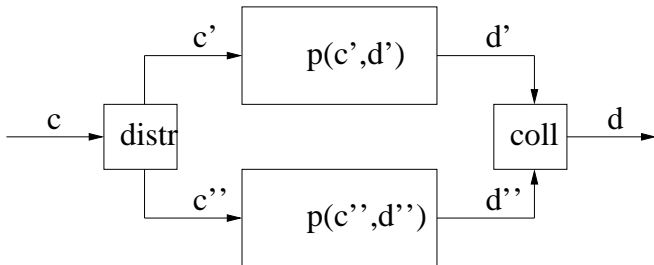
```
def workbal(c, e) = val c' = Channel()  
  val d' = Channel()  
  bal(c, c', d') | net(c', d', e)
```



workBal(c,e)

Deterministic Load Balancing

- Retain input order in the output.
- `distr` alternatively copies input to c' and c'' .
`coll` alternatively copies from d' and d'' to output.



Deterministic Load Balancing

```
def detbal(in, out) =  
  def distributor(c, c', c'') =  
    c.get() >x> c'.put(x) >>  
    c.get() >y> c''.put(y) >>  
    distributor(c, c', c'')
```

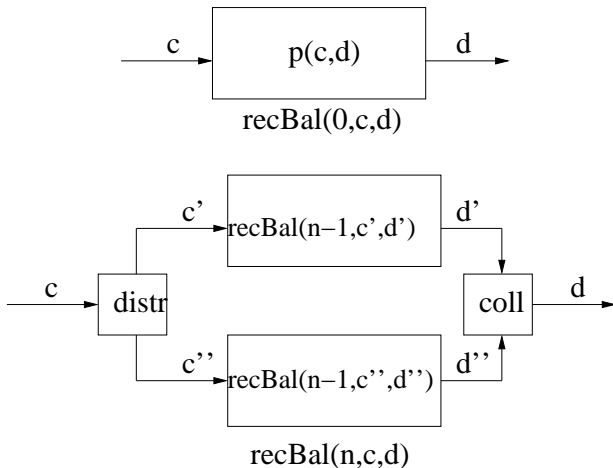
```
def collector(d', d'', d) =  
  d'.get() >x> d.put(x) >>  
  d''.get() >y> d.put(y) >>  
  collector(d', d'', d)
```

```
val (in', in'') = (Channel(), Channel())  
val (out', out'') = (Channel(), Channel())
```

```
  distributor(in, in', in'') | collector(out', out'', out)  
  | p(in', out') | p(in'', out'')
```

Deterministic Load Balancing with 2^n servers

Construct the network recursively.



Recursive Load Balancing Network

def *recbal*(0, *in*, *out*) = *P*(*in*, *out*)

def *recbal*(*n*, *in*, *out*) =
def *distributor*(*c*, *c'*, *c''*) = ...

def *collector*(*d'*, *d''*, *d*) = ...

val (*in'*, *in''*) = (*Channel*(), *Channel*())

val (*out'*, *out''*) = (*Channel*(), *Channel*())

distributor(*in*, *in'*, *in''*) | *collector*(*out'*, *out''*, *out*)
| *recbal*(*n* - 1, *in'*, *out'*) | *recbal*(*n* - 1, *in''*, *out''*)

An Iterative Process: Transducer

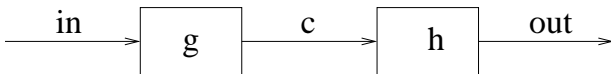
Compute $f(x)$ for each x in channel in and output to out , in order.

```
def transducer(in, out, fn) =  
  in.get() >x> out.put(fn(x)) >> transducer(in, out, fn)
```

Pipeline network

Apply function f to each input: $f(x) = h(g(x))$, for some g and h .

```
def pipe(in, out, g, h) =  
  val c = Channel()  
  transducer(in, c, g) | transducer(c, out, h)
```

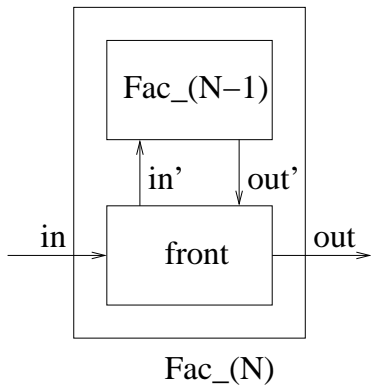


Recursive Pipeline network

Consider computing factorial of each input.

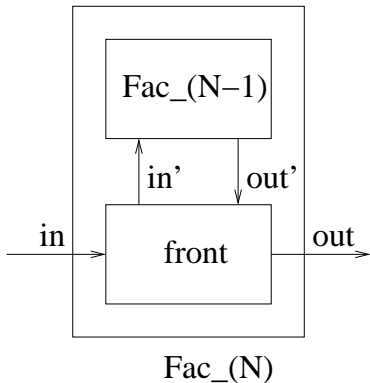
$$fac(x) = \begin{cases} 1 & \text{if } x = 0 \\ x \times fac(x - 1) & \text{if } x > 0 \end{cases}$$

Suppose $x \leq N$, for some given N .



Outline of a program

```
def fac(N, in, out) =  
  val (in', out') = (Channel(), Channel())  
  front(in, out, in', out') | fac(N - 1, in', out')
```



Implementation of Fac_0

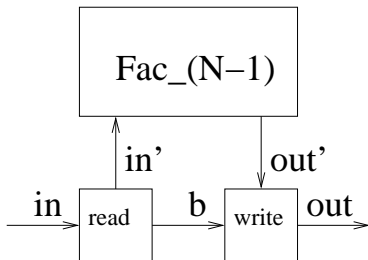
- receive input x , $x = 0$
- output 1
- loop.

```
def fac(0, in, out) =  
  in.get() >> out.put(1) >> fac(0, in, out)
```

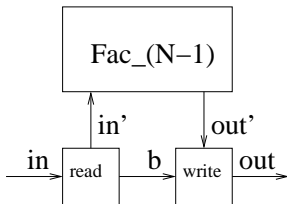
Implementation of *front*

front has two subprocesses, **read** and **write**, doing forever:

- **read** receives input x from in .
 - If $x = 0$, output x on b .
 - If $x > 0$, output x on b , send $x - 1$ on in' .
- **write** receives input x from b :
 - If $x = 0$, output 1.
 - If $x > 0$, receive y from out' , send $x \times y$ on out



Code of *front*

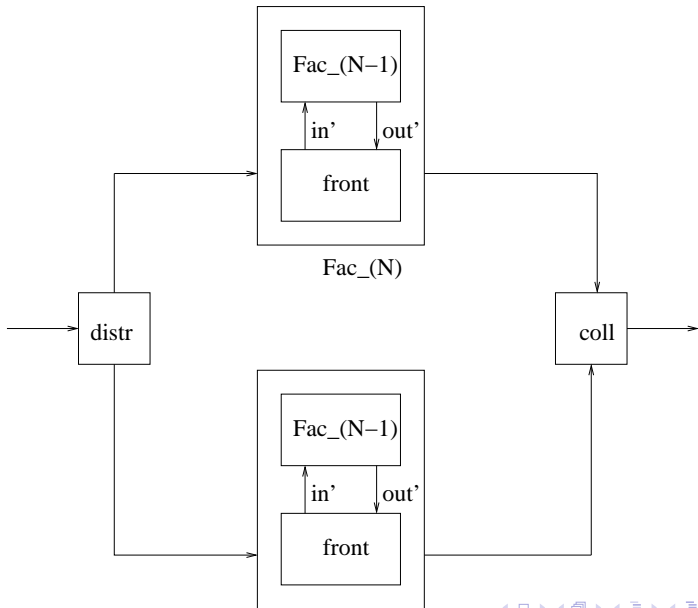


```
def front() =  
  val b = Channel()  
  def read() = in.get() >x> b.put(x) >>  
    if x :> 0 then in'.put(x - 1) else signal >> read()  
  
  def write() = b.get() >x>  
    if x = 0 then out.put(1)  
    else (out'.get() >y> out.put(x * y)) >> write()  
  
read() | write()
```

Program for *fac*

```
def fac(0, in, out) =  
    in.get() >> out.put(1) >> fac(0, in, out)  
  
def fac(N, in, out) =  
    val (in', out') = (Channel(), Channel())  
  
    def front() = ...  
  
front() | fac(N - 1, in', out')
```

Combining Server Farm and Pipeline



Exercise: Combining Server Farm and Pipeline

- A dataset is a list of positive numbers.
The datasets are available on input channel *in*.
Each list length is no more than N , for some given N .
- Required: compute mean and variance of each dataset.
Output the results (as pairs) in order on channel *out*.
- First, divide the processing among about \sqrt{N} servers.
- Next, structure each server as a recursive pipeline.

Recursive Equations for Mean and Variance

- Use the equations:

$$\begin{aligned} \text{sum}([]) &= 0, \\ \text{sum}(x : xs) &= x + \text{sum}(xs) \end{aligned}$$

$$\begin{aligned} \text{length}([]) &= 0, \\ \text{length}(x : xs) &= 1 + \text{length}(xs) \end{aligned}$$

$$\text{mean}(xs) = \text{sum}(xs) / \text{length}(xs)$$

$$\begin{aligned} \text{var}([]) &= 0, \\ \text{var}(xs) &= \text{mean}(\text{map}(\text{square}, xs)) - \text{mean}(xs) ** 2 \end{aligned}$$

- Hint: For each list, compute the sum, sum of squares, and length by a recursive pipeline.
Apply a function to compute mean and variance from these data.

Quicksort

- In situ permutation of an array.
- Array segments are simultaneously sorted.
- Partition of an array segment proceed from left and right simultaneously.
- Combine Concurrency, Recursion, and Mutable Data Structures.

Traditional approaches

- Pure functional programs do not admit in-situ permutation.
- Imperative programs do not highlight concurrency.
- Typical concurrency constructs do not combine well with recursion.

Scan over array a ; swap

- $lr(i)$ returns the smallest index j , $i \leq j \leq t$, where t is given, such that $a(i)? > p$. Returns $t + 1$ if there is no such index.
- $rl(i)$ returns the largest index j , $0 \leq j \leq i$, such that $a(i)? \leq p$. There is guaranteed to be such an index.
- $swap(a, b)$ swaps the contents of two refs, and returns a signal.

def $lr(i) =$ *if* $(i <: t \ \&\& \ a(i)? \leq p)$ *then* $lr(i + 1)$ *else* i

def $rl(i) =$ *if* $(a(i)? > p)$ *then* $rl(i - 1)$ *else* i

def $swap(a, b) = (a?, b?) >(x, y) > (a := y, b := x) \gg$ *signal*

Partition

```
def part(p, s, t) = -- s and t are array boundaries
  def lr(i) = if (i <: t && a(i)? <= p) then lr(i + 1) else i
  def rl(i) = if (a(i)? >: p) then rl(i - 1) else i

  val (s', t') = (lr(s), rl(t))

  ( Ift(s' + 1 <: t') >> swap(a(s'), a(t')) >> part(p, s' + 1, t' - 1)
    | Ift(s' + 1 = t') >> swap(a(s'), a(t')) >> s'
    | Ift(s' + 1 >: t') >> t'
  )
```

Returns m where

$$a(s) \cdots a(m) \leq p,$$
$$a(m+1) \cdots a(t) > p$$

Sorting

```
def sort(s, t) =  
    if s >= t then signal  
    else part(a(s)?, s + 1, t) >m>  
        swap(a(m), a(s)) >>  
        (sort(s, m - 1), sort(m + 1, t)) >>  
    signal  
sort(0, a.length() - 1)
```

Putting the Pieces together

```
def quicksort(a) =  
  def swap(a, b) = (a?, b?) >(x, y) > (a := y, b := x) >> signal  
  def part(p, s, t) =  
    def lr(i) = if (i <: t && a(i)? <= p) then lr(i + 1) else i  
    def rl(i) = if (a(i)? >: p) then rl(i - 1) else i  
    val (s', t') = (lr(s), rl(t))  
    ( Ift(s' + 1 <: t') >> swap(a(s'), a(t')) >> part(p, s' + 1, t' - 1)  
      | Ift(s' + 1 = t') >> swap(a(s'), a(t')) >> s'  
      | Ift(s' + 1 >: t') >> t'  
    )  
  def sort(s, t) =  
    if s >= t then signal  
    else part(a(s)?, s + 1, t) >m>  
      swap(a(m), a(s)) >>  
      (sort(s, m - 1), sort(m + 1, t)) >>  
      signal  
sort(0, a.length() - 1)
```

Remarks and Proof outline

- Concurrency without locks
- $sort(m, n)$ sorts the segment; does not touch items outside the segment.
- Then, $sort(s, m - 1)$ and $sort(m + 1, t)$ are non-interfering.
- $part(p, s, t)$ does not modify any value outside this segment. May read values.

Depth-first search of undirected graph Recursion over Mutable Structure

N : Number of nodes in the graph.

$conn$: $conn(i)$ the list of neighbors of i

$parent$: Mutable array of length N
 $parent(i) = v$, $v \geq 0$, means v is the parent node of i
 $parent(i) < 0$ means parent of i is yet to be determined

Once i has a parent, it continues to have that parent.

$dfs(i, xs)$: starts a depth-first search from all nodes in xs in order,
 i has a parent (or $i = N$),
 $xs \subseteq conn(i)$,
All nodes in $conn(i) - xs$ have parents already.

Depth-first search

val $N = 6$ -- N is the number of nodes in the graph

val $parent = Table(N, lambda(_) = Ref(-1))$

def $dfs(_, []) = signal$

def $dfs(i, x : xs) =$

if $(parent(x)? \geq 0)$ *then* $dfs(i, xs)$

else $parent(x) := i \gg dfs(x, conn(x)) \gg dfs(i, xs)$

$dfs(N, [0])$ -- depth-first search from node 0

Shortest path problem

- Directed graph; non-negative weights on edges.
- Find shortest path from source to sink.

We calculate just the length of the shortest path.

Shortest Path Algorithm with Lights and Mirrors

- Source node sends rays of light to each neighbor.
- Edge weight is the time for the ray to traverse the edge.
- When a node receives its first ray, sends rays to all neighbors. Ignores subsequent rays.
- Shortest path length = time for sink to receive its first ray.
Shortest path length to node i = time for i to receive its first ray.

Graph structure in function $Succ()$

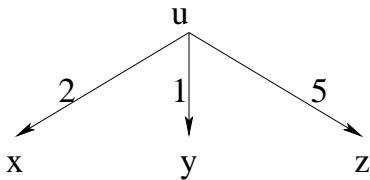


Figure: Graph Structure

$Succ(u)$ publishes $(x, 2)$, $(y, 1)$, $(z, 5)$.

Recording the values

For node u , record its path length in channel u .

u is a bounded channel of length 1.

The first “put” blocks all other puts until the recorded value is read out.

Algorithm

def $eval(u, t) =$ record value t for u \gg
for every successor v with $d = \text{length of } (u, v)$:
wait for d time units \gg
 $eval(v, t + d)$

Goal : $eval(\text{source}, 0)$ |
read the value recorded for the *sink*

Algorithm(contd.)

def $eval(u, t) =$ record value t for u \gg
for every successor v with $d = \text{length of } (u, v)$:
wait for d time units \gg
 $eval(v, t + d)$

Goal : $eval(\text{source}, 0)$ |
read the value recorded for the *sink*

def $eval(u, t) =$ $u.put(t)$ \gg
 $Succ(u) >(v, d)>$
 $Rwait(d)$ \gg
 $eval(v, t + d)$

{ - *Goal :-* } $eval(\text{source}, 0)$ | $sink.get()$

Algorithm(contd.)

```
def eval(u, t) =   u.put(t) >>  
                   Succ(u) >(v, d)>  
                   Rwait(d) >>  
                   eval(v, t + d)
```

```
{- Goal :- }      eval(source, 0) | sink.get()
```

- Any call to *eval(u, t)*: Length of a path from source to *u* is *t*.
- First call to *eval(u, t)*: Length of the shortest path from source to *u* is *t*.
- *eval* does not publish.

Drawbacks of this algorithm

- Running time proportional to shortest path length.
- Executions of *Succ*, *put* and *get* should take no time.