Contents

1 Preamble

The type `Seq.t` represents sequences. Sequences are ***parallel collections***: ordered collections of things, with parallelism-friendly operations on them. Don't think of sequences as being implemented by lists or trees (though you could implement them as such); think of them as a new built-in type with only the operations we're about to describe. The differences between sequences and lists or trees is the cost of the operations, which we specify below. In this document, we describe the cost of array-based sequences.

2 Signature

```
1  signature SEQUENCE =
2  sig
3
4    type 'a t   (* abstract *)
5    type 'a seq = 'a t   (* concrete *)
6
7    exception Range of string
8
9
10   (* Constructing a Sequence *)
11
12   val empty : unit -> 'a seq
13   val singleton : 'a -> 'a seq
14   val tabulate : (int -> 'a) -> int -> 'a seq
15   val fromList : 'a list -> 'a seq
16
17
18   (* Deconstructing a Sequence *)
19
20   val nth : 'a seq -> int -> 'a
21   val null : 'a seq -> bool
22   val length : 'a seq -> int
23   val toList : 'a seq -> 'a list
24   val toString : ('a -> string) -> 'a seq -> string
25   val equal : ('a * 'a -> bool) -> 'a seq * 'a seq -> bool
26
27
28   (* Simple Transformations *)
29
30   val rev : 'a seq -> 'a seq
31   val append : 'a seq * 'a seq -> 'a seq
32   val flatten : 'a seq seq -> 'a seq
33   val cons : 'a -> 'a seq -> 'a seq
34
35
36   (* Combinators and Higher-Order Functions *)
37
38   val filter : ('a -> bool) -> 'a seq -> 'a seq
39   val map : ('a -> 'b) -> 'a seq -> 'b seq
40   val reduce : ('a * 'a -> 'a) -> 'a -> 'a seq -> 'a
41   val reduce1 : ('a * 'a -> 'a) -> 'a seq -> 'a
42   val mapreduce : ('a -> 'b) -> 'b -> ('b * 'b -> 'b) -> 'a seq -> '
      b
43   val zip : ('a seq * 'b seq) -> ('a * 'b) seq
44   val zipWith : ('a * 'b -> 'c) -> 'a seq * 'b seq -> 'c seq
45
46
47   (* Indexing-Related Functions *)
```

```sml
    val enum : 'a seq -> (int * 'a) seq
    val mapIdx : (int * 'a -> 'b) -> 'a seq -> 'b seq
    val update : ('a seq * (int * 'a)) -> 'a seq
    val inject : 'a seq * (int * 'a) seq -> 'a seq

    val subseq : 'a seq -> int * int -> 'a seq
    val take : 'a seq -> int -> 'a seq
    val drop : 'a seq -> int -> 'a seq
    val split : 'a seq -> int -> 'a seq * 'a seq


    (* Sorting and Searching *)

    val sort : ('a * 'a -> order) -> 'a seq -> 'a seq
    val merge : ('a * 'a -> order) -> 'a seq * 'a seq -> 'a seq
    val search : ('a * 'a -> order) -> 'a -> 'a seq -> int option


    (* Views *)

    datatype 'a lview = Nil | Cons of 'a * 'a seq

    val showl : 'a seq -> 'a lview
    val hidel : 'a lview -> 'a seq

    datatype 'a tview = Empty | Leaf of 'a | Node of 'a seq * 'a seq

    val showt : 'a seq -> 'a tview
    val hidet : 'a tview -> 'a seq

end
```

3 Documentation

If unspecified, we assume that all functions that are given as arguments (such as the `g` in `reduce g`) have $O(1)$ work and span. In order to analyze the runtime of sequence functions when this is not the case, we need to analyze the corresponding cost graphs.

Constraint: Whenever you use these sequence functions, please make sure you meet the specified preconditions.

If you do not meet the precondition for a function, it may not behave as expected or meet the cost bounds stated below.

Definition 3.1 (Associative). Fix some type `t`. We say a function `g : t * t -> t` is *associative* if for all `a`, `b`, and `c` of type `t`:

$$g \ (g \ (a,b),c) \cong g \ (a,g \ (b,c))$$

Definition 3.2 (Identity). Fix some type `t`. Given a function `g : t * t -> t`, we say `z` is the *identity* for `g` if for all `x : t`:

$$g \ (x,z) = g \ (z,x) = x$$

3.1 Constructing a Sequence

```
empty : unit -> 'a seq
```

ENSURES: `empty ()` evaluates to the empty sequence (the sequence of length zero).

Work: $O(1)$, Span: $O(1)$.

```
singleton : 'a -> 'a seq
```

ENSURES: `singleton x` evaluates to a sequence of length 1 whose only element is `x`.

Work: $O(1)$, Span: $O(1)$.

```
tabulate : (int -> 'a) -> int -> 'a seq
```

REQUIRES: $n \geq 0$, and for all $0 \leq i < n$, `f i` is valuable.

ENSURES: `tabulate f n` evaluates to a sequence `S` of length `n`, where the i^{th} element of `S` is equal to `f i`.

Note that indices are zero-indexed.

Work $\sum_{i=0}^{n-1} W_{\texttt{f}}(\texttt{i})$, Span $\max_{0 \leq i < n} S_{\texttt{f}}(\texttt{i})$.

Work $O(\texttt{n})$, Span $O(1)$, with constant-time `f`.

```
fromList : 'a list -> 'a seq
```

ENSURES: `fromList L` returns a sequence consisting of the elements of `L`, preserving order. This function is intended primarily for debugging purposes.

Work: $O(|\texttt{L}|)$, Span: $O(|\texttt{L}|)$.

3.2 Deconstructing a Sequence

```
nth : 'a seq -> int -> 'a
```
REQUIRES: $0 \leq \texttt{i} < \texttt{length S}$

ENSURES: `nth S i` evaluates to the $\texttt{i}^{\text{th}}$ element (zero-indexed) of `S`.

Work: $O(1)$, Span: $O(1)$.

```
null : 'a seq -> bool
```
ENSURES: `null S` evaluates to `true` if `S` is an empty sequence, and `false` otherwise.

Work: $O(1)$, Span: $O(1)$.

```
length : 'a seq -> int
```
ENSURES: `length S` (often written as $|\texttt{S}|$) evaluates to the number of items in `S`.

Work: $O(1)$, Span: $O(1)$.

```
toList : 'a seq -> 'a list
```
ENSURES: `toList S` returns a list consisting of the elements of `S`, preserving order. This function is intended primarily for debugging purposes.

Work: $O(|\texttt{S}|)$, Span: $O(|\texttt{S}|)$.

```
toString : ('a -> string) -> 'a seq -> string
```
ENSURES: `toString ts S` evaluates to a string representation of `S`, using the function `ts` to convert each element of `S` into a string.

Work: $O(|\texttt{S}|)$, Span: $O(\log|\texttt{S}|)$.

```
equal : ('a * 'a -> bool) -> 'a seq * 'a seq -> bool
```
ENSURES: `equal eq (S1,S2)` returns whether or not the two sequences are equal according to the equality function `eq`.

Work: $O(\min(|\texttt{S1}|, |\texttt{S2}|))$, Span: $O(\log(\min(|\texttt{S1}|, |\texttt{S2}|)))$.

3.3 Simple Transformations

```
rev : 'a seq -> 'a seq
```

ENSURES: `rev S` returns the sequence containing the elements of `S` in reverse order.

Work: $O\left(|S|\right)$, Span: $O\left(1\right)$.

```
append : 'a seq * 'a seq -> 'a seq
```

ENSURES: `append (S1,S2)` evaluates to a sequence of length $|S1| + |S2|$ whose first $|S1|$ elements are the sequence `S1` and whose last $|S2|$ elements are the sequence `S2`.

Work: $O\left(|S1| + |S2|\right)$, Span: $O\left(1\right)$.

```
flatten : 'a seq seq -> 'a seq
```

ENSURES: `flatten S` flattens a sequence of sequences down to a single sequence (similar to `List.concat` for lists).

Work: $O\left(|S| + \sum_{s \in S} |s|\right)$, Span: $O\left(\log |S|\right)$.

```
cons : 'a -> 'a seq -> 'a seq
```

ENSURES: If the length of `S` is `n`, `cons x S` evaluates to a sequence of length `n+1` whose first item is `x` and whose remaining `n` items are exactly the sequence `S`.

Constraint: Beware of using `cons`. When overused, it often leads to a sequential coding style with little opportunity for parallelism, defeating the purpose of using sequences. See if there's a way to use other sequence functions to achieve your objective.

Work: $O\left(|S|\right)$, Span: $O\left(1\right)$.

3.4 Combinators and Higher-Order Functions

```
filter : ('a -> bool) -> 'a seq -> 'a seq
```

ENSURES: `filter p S` evaluates to a sequence containing all of the elements `x` of `S` such that `p x` $\implies$ `true`, preserving element order.

Work $\sum_{x \in S} W_{\texttt{p}}(x)$, Span $O\left(\log |\texttt{S}|\right) + \max_{x \in S} S_{\texttt{p}}(x)$.

Work $O\left(|\texttt{S}|\right)$, Span $O\left(\log |\texttt{S}|\right)$, with constant-time `p`.

```
map : ('a -> 'b) -> 'a seq -> 'b seq
```

ENSURES: `map f S` $\implies$ `S'` such that $|\texttt{S}| = |\texttt{S'}|$ and for all $0 \le \texttt{i} < |\texttt{S'}|$, `nth S' i` $\cong$ `f (nth S i)`.

Work $\sum_{x \in S} W_{\texttt{f}}(\texttt{x})$, Span $\max_{x \in S} S_{\texttt{f}}(\texttt{x})$.

Work $O\left(|\texttt{S}|\right)$, Span $O\left(1\right)$, with constant-time `f`.

```
reduce : ('a * 'a -> 'a) -> 'a -> 'a seq -> 'a
```

REQUIRES:

- `g` is associative
- `z` is an identity for `g`

ENSURES: `reduce g z S` uses the function `g` to combine the elements of `S` using `z` as a base case (analogous to `foldr g z L` for lists, but with a less-general type).

Work: $O\left(|\texttt{S}|\right)$, Span: $O\left(\log |\texttt{S}|\right)$, with constant-time `g`.

```
reduce1 : ('a * 'a -> 'a) -> 'a seq -> 'a
```

REQUIRES: `g` is associative, and `S` is nonempty.

ENSURES: `reduce1 g S` uses the function `g` to combine the elements of `S`. If `S` is a singleton sequence, the sequence element is returned.

Work: $O\left(|\texttt{S}|\right)$, Span: $O\left(\log |\texttt{S}|\right)$, with constant-time `g`.

```
mapreduce : ('a -> 'b) -> 'b -> ('b * 'b -> 'b) -> 'a seq -> 'b
```

REQUIRES: `z` and `g` meet the preconditions for `reduce`.

ENSURES: `mapreduce f z g S` $\cong$ `reduce g z (map f S)`

Work: $O(|\mathtt{S}|)$, Span: $O(\log|\mathtt{S}|)$, with constant-time `g` and `f`.

```
zip : 'a seq * 'b seq -> ('a * 'b) seq
```

ENSURES: `zip (S1,S2)` evaluates to a sequence of length $\min(\mathtt{length\ S1}, \mathtt{length\ S2})$ whose $\mathtt{i}^{\text{th}}$ element is the pair of the $\mathtt{i}^{\text{th}}$ element of `S1` and the $\mathtt{i}^{\text{th}}$ element of `S2`.

Work: $O(\min(|\mathtt{S1}|, |\mathtt{S2}|))$, Span: $O(1)$.

```
zipWith : ('a * 'b -> 'c) -> 'a seq * 'b seq -> 'c seq
```

ENSURES: `zipWith f (S1,S2)` $\cong$ `map f (zip (S1,S2))`.

Work $\displaystyle\sum_{i=0}^{min(|\mathtt{S1}|,|\mathtt{S2}|)-1} W_{\mathtt{f}}(\mathtt{S1}[i], \mathtt{S2}[i])$, Span $\displaystyle\max_{i=0}^{min(|\mathtt{S1}|,|\mathtt{S2}|)-1} S_{\mathtt{f}}(\mathtt{S1}[i], \mathtt{S2}[i])$.

Work $O(\min(|\mathtt{S1}|, |\mathtt{S2}|))$, Span $O(1)$, with constant-time `f`.

3.5 Indexing-Related Functions

```
enum : 'a seq -> (int * 'a) seq
```

ENSURES: `enum S` evaluates to a sequence such that for each index $0 \leq$ `i` $<$ `length S`, the `i`[th] index of the result is `(i,nth S i)`.

Work: $O\left(|S|\right)$, Span: $O\left(1\right)$.

```
mapIdx : (int * 'a -> 'b) -> 'a seq -> 'b seq
```

ENSURES: `mapIdx f S` $\cong$ `map f (enum S)`.

Work $\sum_{x \in S} W_f(x)$, Span $\max_{x \in S} S_f(x)$.

Work $O\left(|S|\right)$, Span $O\left(1\right)$, with constant-time `f`.

```
update : 'a seq * (int * 'a) -> 'a seq
```

REQUIRES: $0 \leq$ `i` $<$ `length(S)`

ENSURES: `update (S,(i,x))` returns a sequence identical to `S` but with the `i`[th] element (0-indexed) now `x`.

Work: $O\left(|S|\right)$, Span: $O\left(1\right)$.

```
inject : 'a seq * (int * 'a) seq -> 'a seq
```

REQUIRES: For all `(i,x)` in `U`, $0 \leq$ `i` $<$ `length(S)`.

ENSURES: `inject (S,U)` evaluates to a sequence where for each `(i,x)` in `U`, the `i`[th] element of `S` is replaced with `x`. If there are multiple elements at the same index, one is chosen nondeterministically.

Work: $O\left(|S| + |U|\right)$, Span: $O\left(1\right)$.

```
subseq : 'a seq -> int * int -> 'a seq
```

REQUIRES: $0 \leq$ `i` $<$ `length S`, and `i` $\leq$ `i + l` $<$ `length S`

ENSURES: `subseq S (i,l)` takes the subsequence of `S` with length `l` starting at index `i`.

Work: $O\left(1\right)$, Span: $O\left(1\right)$.

```
take : 'a seq -> int -> 'a seq
```

REQUIRES: $0 \leq$ `i` $\leq$ `length S`

ENSURES: `take S i` evaluates to the sequence containing exactly the first `i` elements of `S`.

Work: $O(1)$, Span: $O(1)$.

```
drop : 'a seq -> int -> 'a seq
```

REQUIRES: $0 \leq$ `i` $\leq$ `length S`

ENSURES: `drop S i` evaluates to the sequence containing all but the first `i` elements of `S`.

Work: $O(1)$, Span: $O(1)$.

```
split : 'a seq -> int -> 'a seq * 'a seq
```

REQUIRES: $0 \leq$ `i` $\leq$ `length S`

ENSURES: `split S i` evaluates to a pair of sequences `(S1,S2)` such that `S1` has length `i` and `append (S1,S2)` $\cong$ `s`.

Work: $O(1)$, Span: $O(1)$.

3.6 Sorting and Searching

```
sort : ('a * 'a -> order) -> 'a seq -> 'a seq
```

REQUIRES: `cmp` is a comparison function (reflexive and transitive)

ENSURES: `sort cmp S` returns a permutation of `S` that is sorted according to `cmp`. The sort is stable: elements that are considered equal by `cmp` remain in the same order they were in `S`.

Work: $O\left(|S|\log|S|\right)$, Span: $O\left(\log^2|S|\right)$, with constant-time `cmp`.

```
merge : ('a * 'a -> order) -> 'a seq * 'a seq -> 'a seq
```

REQUIRES: `cmp` is a comparison function (reflexive and transitive), and `S1` and `S2` are both `cmp`-sorted.

ENSURES: `merge cmp (S1,S2)` returns a sorted permutation of `append (S1,S2)`.

Work: $O\left(|S1|+|S2|\right)$, Span: $O\left(\log(|S1|+|S2|)\right)$, with constant-time `cmp`.

```
search : ('a * 'a -> order) -> 'a -> 'a seq -> int option
```

REQUIRES: `cmp` is a comparison function (reflexive and transitive), and `S` is `cmp`-sorted.

ENSURES: `search cmp x S` $\implies$ `SOME` `i` where `i` is the first index in `S` satisfying `cmp (nth i S, x)` $\cong$ `EQUAL` or `NONE` if no such index exists.

Work: $O\left(\log|S|\right)$, Span: $O\left(\log|S|\right)$, with constant-time `cmp`.

4 Views

4.1 List Views

Recall that list operations have bad parallel complexity, whereas the corresponding sequence operations are much better.

However, sometimes you want to write a ***sequential*** algorithm (e.g., because the inputs aren't very big, or because no good parallel algorithms are known for the problem). Given the sequence interface so far, it is difficult to decompose a sequence as "either empty, or a cons with a head and a tail." You *could* write this jank codemonkey's "length induction instead of structural induction"...

```
case Seq.length s of
  0 =>
| _ => ... uses (Seq.hd s) and (Seq.tl s) ...
```

But nah. We can solve this problem using a ***view***. We'll put an appropriate datatype in the signature, along with corresponding functions that convert sequences to and from this datatype. This allows us to pattern-match on an abstract type, while keeping the actual representation abstract. These definitions enable viewing a sequence like a list:

```
datatype 'a lview = Nil | Cons of 'a * 'a seq

val showl : 'a seq -> 'a lview
val hidel : 'a lview -> 'a seq
```

Note the invariant:

$$\texttt{showl (hidel v)} \Longrightarrow \texttt{v}$$

Because the datatype definition is in the signature, the constructors `Nil` and `Cons` can be used outside the abstraction boundary. The `showl` and `hidel` functions convert between sequences and list views. The following is an example of using this view to perform list-like pattern matching:

```
case Seq.showl s of
  Seq.Nil => ... (* Nil case *)
| Seq.Cons (x,s') => ... uses x and s' ... (* Cons case *)
```

Note that the second argument to `Cons` is another `'a seq`, ***not*** an `lview`. Thus, `showl` lets you do one level of pattern matching at a time: you can write patterns like `Seq.Cons (x,xs)` but not `Seq.Cons (x,Seq.Nil)` (to match a sequence with exactly one element).

We have also provided `hidel`, which converts a view back to a sequence—`Seq.hidel (Seq.Cons (x,xs))` is equivalent to `Seq.cons(x,xs)` and `Seq.hidel Seq.Nil` is equivalent to `Seq.empty()`.

4.2 Tree Views

The analogous `'a tview`, `showt`, and `hidet` are provided in the signature.

5 Thinking About Cost

Let's think about the big picture of parallelism. Parallelism is relevant to situations where many things can be done at once: e.g. using the multiple cores in multi-processor machine, or the many machines in a cluster. Overall, the goal of parallel programming is to describe computation in such a way that it can make use of this ability to do work on multiple processors simultaneously. At the lowest level, this means deciding, at each moment in time, what to do on each processor. This is limited by the data dependencies in a problem or a program. For example, evaluating `(1 + 2) + (3 + 4)` takes three units of work, one for each addition, but you cannot do the outer addition until you have done the inner two. So even with three processors, you cannot perform the calculation in fewer than two timesteps. That is, the expression has work 3 but span 2.

The approach to parallelism that we're advocating in this class is based on raising the level of abstraction at which you can think, by *separating the specification of what work there is to be done from the schedule that maps it onto processors.* As much as possible, you, the programmer, worry about specifying what work there is to do, and the compiler takes care of scheduling it onto processors. Three things are necessary to make this separation of concerns work:

1. The code itself must not bake in a schedule.

2. You must be able to reason about the *behavior* of your code independently of the schedule.

3. You must be able to reason about the *time complexity* of your code independently of the schedule.

Our central tool for avoiding baking in a schedule is *functional programming.* First, we focus on bulk operations on big collections which do not specify a particular order in which the operations on each element are performed. For example, sequences come with an operation `map` that is specified by saying that the value of `map f <x1,x2,...,xn>` is the sequence `<f x1, f x2, ... f xn>`. This specifies the data dependencies (to calculate `map`, you need to calculate `f x1 ...`) without specifying a particular schedule. You can implement the same computation with a loop, saying "do `f x1`, then do `f x2`,...", but this is inlining a particular schedule into the code— which is bad, because it gratuitously throws away opportunities for parallelism. Second, functional programming focuses on pure, mathematical functions, which are evaluated by calculation. This limits the dependence of one chunk of work on another to what is obvious from the data-flow in the program. For example, when you `map` a function `f` across a sequence, evaluating `f` on the first element has no influence on the value of `f` on the second element, etc.—this is not the case for imperative programming, where one call to `f` might influence another via memory updates. It is in general undecidable to take an imperative program and notice, after the fact, that what you really meant by that loop was a bulk operation on a collection, or that this particular piece of code really defines a mathematical function.

So why are we teaching you this style of parallel programming? There are two reasons: First, even if you have to get into more of the gritty details of scheduling to get your code to run fast today, it's good to be able to think about problems at a high level first, and then figure out the details. If you're writing some code for an internship this summer using a low-level parallelism interface, it can be useful to first think about the abstract algorithm—what are the dependencies between tasks? what can possibly be done in parallel?—and then figure out the details. You can use parallel functional programming to design algorithms, and then translate them down to whatever interface you need. Second, it's our thesis that eventually this kind of parallel programming will be practical and common: as language implementations improve, and computers get more and more cores, this kind of programming will become possible and even necessary. You're going to be writing programs

for a long time, and we're trying to teach you tools that will be useful years down the road.

Cost Semantics Reminder

A cost graph is a form of *series-parallel graphs*. A series-parallel graph is a directed graph (we always draw a cost graph so that the edges point down the page) with a designated source node (no edges in) and sink node (no edges out), formed by two operations called sequential and parallel composition. The particular series-parallel graphs we need are of the following form:
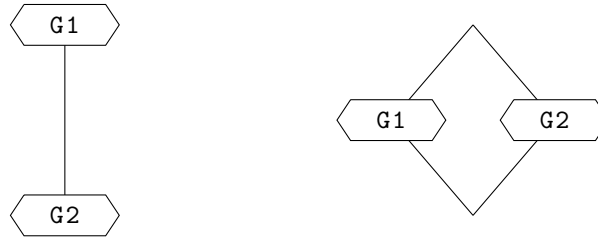
Depicted on the left is *sequential combination*: the graph formed by putting an edge from the sink of `G1` to the source of `G2`. The other, *parallel combination*, is the graph formed by adding a new source and sink, and adding edges from the source to the source of each of `G1` and `G2`, and from the sinks of each of them to the new sink. We also need an *n*-ary parallel combination of graphs `G1 ... Gn`
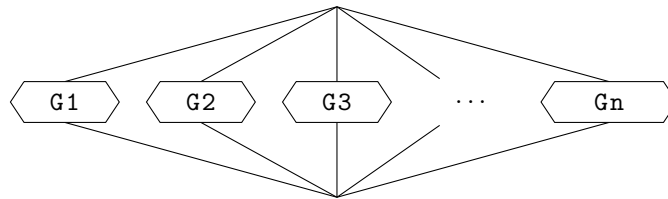
The *work* of a cost graph is the number of nodes. The *span* is the length of the longest path, which we may refer to as the *critical path*, or the *diameter*. We will associate a cost graph with each closed program, and define the work/span of a program to be the work/span of its cost graph.

These graphs model *fork-join parallelism*: a computation forks into various subcomputations that are run in parallel, but these come back together at a well-defined joint point. These forks and joins are well-nested, in the sense that the join associated with a later fork precedes the join associated with an earlier fork.

6 Cost Graphs

Intuitively, these sequence operations do the same thing as the operations on lists that you are familiar with. However, they have different time complexity than the list functions: First, sequences admit constant-time access to elements: `nth` takes constant time. Second, sequences have better parallel complexity: many operations, such as `map`, act on each element of the sequence in parallel.

For each function, we (1) describe its behavior abstractly and (2) give a cost graph.
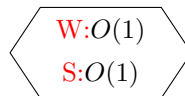
Note: $\langle x_1, ..., x_n \rangle$ is **not** SML syntax. It is mathematical syntax which we will use to represent a sequence value.

Length

Behavior of `length`:

$$\texttt{length } \langle x_1, ..., x_n \rangle \cong \texttt{n}$$

Cost Graph for `length s`:

W:$O(1)$
S:$O(1)$

As a consequence,

Work: **Span:**

$O(1)$ $O(1)$

Nth

Behavior of `nth`:

$$\texttt{nth } \langle x_0, ..., x_{n-1} \rangle \texttt{ i} \cong x_i \text{ if } 0 \leq \texttt{i} < n \text{ or raises } \texttt{Range} \text{ otherwise}$$

Cost Graph for `nth s i`:

W:$O(1)$
S:$O(1)$

As a consequence,

Work: **Span:**
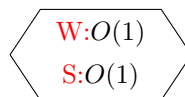
$O(1)$ $O(1)$

Map

Behavior of `map`:

$$\texttt{map f } \langle x_1, ..., x_n \rangle \cong \langle \texttt{f } x_1, ..., \texttt{f } x_n \rangle$$

Cost Graph for `map f s`:
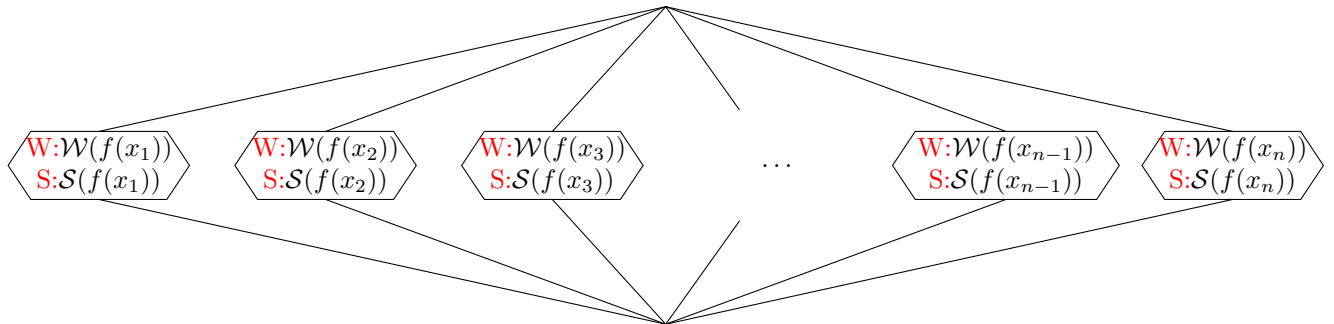
where `s` is $\langle x_1, ..., x_n \rangle$.
As a consequence,

Work:

$$\sum_{x \in s} \mathcal{W}(f(x))$$

Span:

$$\max_{x \in s} \mathcal{S}(f(x))$$

Reduce
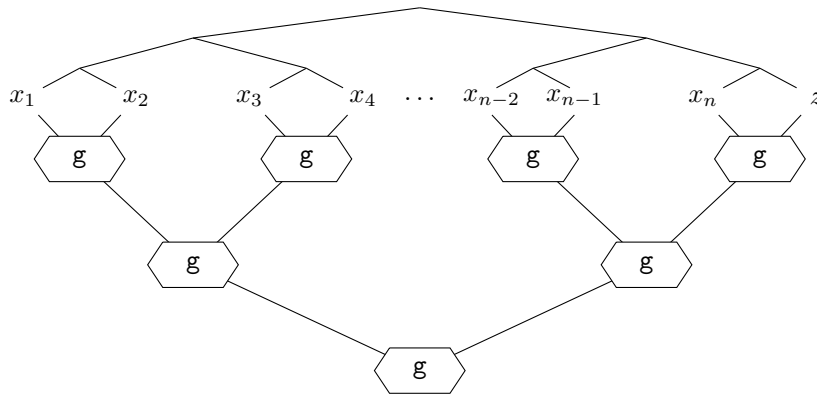
The behavior of `reduce` is

$$\texttt{reduce g z S} \cong \texttt{List.foldr g z (toList S)}$$

That is, `reduce` applies its argument function `g` between every pair of elements in the sequence, using `z` as a rightmost base case.

Cost Graph for `reduce g z` $\langle x_1, \ldots x_n \rangle$:

Observe from the above graph that, – if we assume that `g` is $O(1)$ – then `reduce` has the following time bounds.

Work: **Span:**

$$O(n) \qquad\qquad\qquad\qquad\qquad O(\log n)$$

where n is the length of the sequence. If `g` does not have constant work and span, then refer directly to the cost graph above.
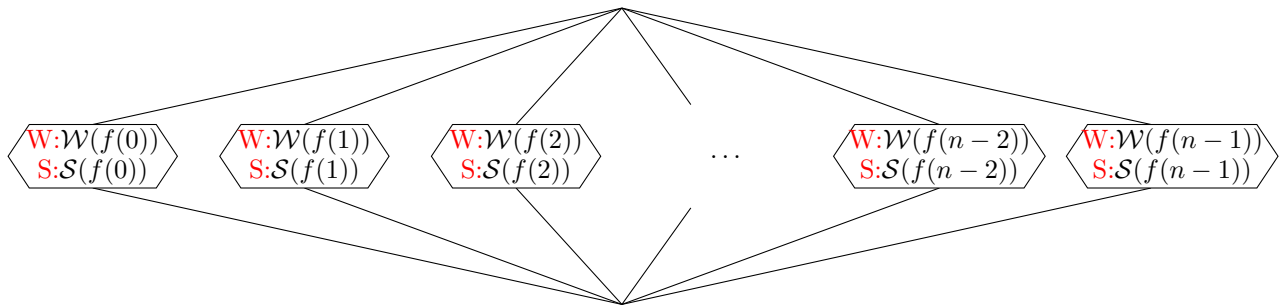
Tabulate

The way of introducing a sequence is *tabulate*, which constructs a sequence from a function that gives you the element at each position, from 0 up to a specified bound.

Behavior of `tabulate f n`:

$$\texttt{tabulate f n} \cong \langle v_0, ..., v_{n-1} \rangle$$

$$\texttt{where} \qquad \texttt{f 0} \cong v_0$$
$$\texttt{f 1} \cong v_1$$
$$\vdots$$
$$\texttt{f (n-1)} \cong v_{n-1}$$

Cost Graph for `tabulate f n`:

As a consequence,

Work:

$$\sum_{i=0}^{n-1} \mathcal{W}(f(i))$$

Span:

$$\max_{i=0}^{n-1} \mathcal{S}(f(i))$$