

Calculating a linear-time solution to the densest-segment problem

SHARON CURTIS

Oxford University, UK
(e-mail: sharon.curtis@ndph.ox.ac.uk)

SHIN-CHENG MU

Academia Sinica, Taiwan
(e-mail: scm@iis.sinica.edu.tw)

Abstract

The problem of finding a densest segment of a list is similar to the well-known maximum segment sum problem, but its solution is surprisingly challenging. We give a general specification of such problems, and formally develop a linear-time online solution, using a sliding-window style algorithm. The development highlights some elegant properties of densities, involving partitions that are decreasing and all right-skew.

1 Introduction

Processing large datasets, such as those containing DNA sequences, requires efficient algorithms. The data analysis problem we examine here involves finding a *segment* (contiguous subsequence) of a list of elements with maximum *density* (to be defined in the next section). While this problem resembles the well-known maximum segment sum problem, a standard textbook example of programme derivation (Kaldewaij, 1990), the internal structure of the segment density problem is more intricate, and finding a linear-time solution is much harder.

Our aim in this paper is the derivation of a linear-time functional programme to solve a generalised segment density problem, in an elegant and clear way. We present two instances of such density problems: MMS (maximum mean segment) and MDS (maximum density segment), which has applications to analysis of DNA sequences. The MMS problem we believe is new; MDS has been solved before (Chung & Lu, 2004; Goldwasser *et al.*, 2005), but in an imperative setting, where it is difficult to see the structure and correctness of the algorithms amongst the details.

Our algorithm derivation uses a traditional programme calculation approach, at least initially. However, as several of the proofs are tricky or sizable, for additional reassurance, the theorems and their proofs are also coded into the dependently-typed language/theorem prover Agda (Norell, 2007) and the Algebra of Programming in Agda (AOPA) library (Mu *et al.*, 2009); the complete proof is around 3,500 lines of Agda code.

In this paper, we outline the main proofs; the rest of the proofs can be found in a supplement containing full details, available online (Curtis & Mu, 2014), along

with a Haskell implementation and Agda code. The paper is structured as follows: in Section 2, we give a description and specification of the densest-segment problem, along with the two instances MMS and MDS. The development of the algorithm then proceeds in two stages: in Section 3, we develop the structure of the algorithm, which processes the input list using a *sliding-window* technique. The second stage is presented in Section 4, which explores properties of the window to develop functions sufficient to allow the algorithm to run in linear time. The resulting datatypes for the window, and the algorithm's implementation, are presented in Section 5 along with an analysis of its performance. We conclude and summarise related work in Section 6.

2 The densest segment problem

Segment problems have been well examined in the literature, especially in the late 1980s and early 1990s, when their study formed part of a wider body of work on constructive algorithmics. Some of this calculational work resulted in imperative-style implementations, e.g. (Rem, 1988; Kaldewaij, 1990; Van Den Eijnde, 1990), while others used a functional style for their development, e.g. (Bird, 1987; Jeuring, 1993; Swierstra & de Moor, 1993).

The segment problems studied were optimisation problems on a list of input elements, where it was desired to find a segment that was optimal with regards to some function. Such a function might evaluate a segment on its length (i.e. wanting a shortest or longest segment, such as in Jeuring & Meertens (1989)), or in other more interesting ways, such as trying to find the largest rectangle under a histogram (Van Den Eijnde, 1990). Often the segments of interest would also have to satisfy some predicate, and depending on the properties of the predicate, it might be possible to find a linear-time algorithm to solve the problem, often in a “sliding-window” style (Zantema, 1992).

The general densest-segment problem we address in this paper describes a particular specialised class of segment problems, using optimisation functions that satisfy a *density property* (see Section 2.1) and a predicate broadly related to segment length. In rest of this section, we give a description of this class of segment problems, before presenting two specific instances: MDS, which features in the existing literature, and MMS, the latter of which we believe is novel.

2.1 Description

The problem's input data is a sequence of elements of type *Elem*. Each segment (contiguous subsequence) of the input elements has a *width*,

$$\text{width} :: [\text{Elem}] \rightarrow \mathbb{R}^{\geq 0},$$

with the property that an empty list has zero width, and adding more elements at either end increases width. Formally,

$$\text{width} [] = 0, \tag{1}$$

$$\text{width } x < \text{width } (x \# y) > \text{width } y, \tag{2}$$

where x and y are any non-empty lists of elements. For example, the function *length* is such a width function, but there are other more interesting possibilities (see the MDS problem in Section 2.2 below).

Not all segments are considered as candidate solutions to the problem, and there is a constraint on segment width: a lower bound given by $L : \mathbb{R}^{>0}$. (Without such a constraint, the solution is often trivial.) This predicate checks whether a segment is wide enough:

$$\begin{aligned} \textit{wide} &:: [Elem] \rightarrow Bool \\ \textit{wide } x &= \textit{width } x \geq L. \end{aligned}$$

Evaluation of segments occurs by means of a *density* function

$$d :: [Elem] \rightarrow \mathbb{R}.$$

Informally, the density of a segment can be thought of as comparing the “weight” of a segment’s elements to its width, in some way. (For specific examples, please see the functions d_{MMS} and d_{MDS} in Sections 2.3 and 2.2 below.) Comparison between segments is carried out using the relation

$$\begin{aligned} (\leq_d) &:: [Elem] \rightarrow [Elem] \rightarrow Bool \\ x \leq_d y &\Leftrightarrow d\ x \leq d\ y. \end{aligned}$$

As a shorthand, we will write $x <_d y$ for $x \leq_d y \wedge y \not\leq_d x$, and write \geq_d and $>_d$ for the converses of \leq_d and $<_d$.

However, the function d is not strictly needed in the calculation and implementation of a solution: what matters is that there exists a relation \leq_d for comparing segments that is reflexive, transitive, and total on non-empty lists (anti-symmetry is not required, as two different segments may be equally dense). It is also required that the relation \leq_d satisfies the following *density property*: for any non-empty lists $x, y :: [Elem]$,

$$x \leq_d x ++ y \Leftrightarrow x \leq_d y \Leftrightarrow x ++ y \leq_d y. \tag{3}$$

The above must also hold if \leq_d is replaced with $<_d$, \geq_d or $>_d$.

In summary, the densest-segment problem is the finding of a maximally dense segment with respect to a relation \leq_d that satisfies the density property, subject to a lower bound L on width.

2.2 Maximum density segment (MDS)

In the MDS problem, each element has a (strictly positive) *size*:

$$\textit{size} :: Elem \rightarrow \mathbb{R}^{>0},$$

and the width of a segment is the sum of the sizes of its elements

$$\begin{aligned} \textit{width}_{MDS} &:: [Elem] \rightarrow \mathbb{R}^{\geq 0} \\ \textit{width}_{MDS} &= \textit{sum} \cdot \textit{map } \textit{size}. \end{aligned}$$

Note that \textit{width}_{MDS} satisfies the width properties (1) and (2) stated above.




Fig. 1. A list of elements, with weights in bold and sizes as indicated. The density d_{MDS} of the whole list is $(9 + 6 + 14 + 20 - 10 + 20 - 2 + 27)/(6 + 2 + 7 + 4 + 5 + 8 + 2 + 6) = 84/40 = 2.1$, which is shown by the height of the dashed line.

Also, each element in the MDS problem has a *weight*:

$$weight_{MDS} :: Elem \rightarrow \mathbb{R},$$

and the density of a non-empty segment is calculated by dividing its total weight by its width

$$d_{MDS} :: [Elem] \rightarrow \mathbb{R}$$

$$d_{MDS} \text{ seg} = (\text{sum} (\text{map } weight_{MDS} \text{ seg})) / (\text{width}_{MDS} \text{ seg}).$$

A straightforward arithmetical calculation shows that $\leq_{d_{MDS}}$ does indeed satisfy the density property (3).

For the MDS problem, densities can be visualised accurately: when each element is depicted as a rectangle as wide as its size, with area corresponding to its weight, the density of a segment is simply its average height, e.g. see Figure 1. Thus, the MDS problem is to find a maximally dense segment with respect to d_{MDS} , subject to a lower bound L on width. For example, the densest segment of the list of elements in Fig. 1, with width at least 9, is the segment $[(14, 7), (20, 4)]$, which has density $34/11 \approx 3.1$.

The MDS problem has applications in bio-informatics, to the analysis of genetic material to find fragments with high densities of certain characteristics. For example, sections of DNA or RNA that are *GC-rich* – with a high density of bases that are guanine (G) or cytosine (C) – are of interest because they are likely to indicate portions containing genes (Han & Zhao, 2009). Another example is that sections of DNA that have a high density of mutations are of interest because they can provide clues pointing to the three-dimensional molecular structure.

The solution of the MDS problem has a varied history: in Huang (1994), a simpler version of the MDS problem was considered, where all elements have unit size. Huang noticed that with a lower bound L on segment width, there is a maximally dense segment no wider than $2L - 1$, which leads to a simple $O(nL)$ algorithm, for n input elements. Later, the complexity of a MDS solution was improved to $O(n \log L)$ (Lin *et al.*, 2002). Goldwasser *et al.* (2002) studied a variation of MDS that also used an upper bound U on segment width, and they presented an $O(n)$ algorithm for the simpler case of MDS when all elements have unit size (also an instance of MMS), and an $O(n \log(U - L + 1))$ algorithm for elements with variable sizes. Another published algorithm claimed to be linear, but wasn't; details




Fig. 2. A sequence of blood glucose readings, where the dotted area indicates the target range. Usually, several months' worth of data would be available.

are recorded in Chung & Lu (2004). Then, both Chung & Lu (2003; 2004) and Goldwasser *et al.* (2005) produced similar algorithms to take linear time to solve MDS with lower and upper bounds.

2.3 Maximum mean segment (MMS)

In the MMS problem, the function to measure the width of a segment is not specified further than it having to obey the width properties (1) and (2).

$$\text{width}_{\text{MMS}} :: [\text{Elem}] \rightarrow \mathbb{R}^{\geq 0}.$$

The density of a non-empty segment, for the MMS problem, is the mean weight of its elements:

$$\begin{aligned} d_{\text{MMS}} &:: [\text{Elem}] \rightarrow \mathbb{R} \\ d_{\text{MMS}} \text{ seg} &= (\text{sum} (\text{map weight seg})) / (\text{length seg}). \end{aligned}$$

A straightforward arithmetical calculation shows that $\leq_{d_{\text{MMS}}}$ satisfies the density property (3). Thus, the MMS problem is to find a maximally dense segment with respect to d_{MMS} , subject to a lower bound L on segment width.

One application of the MMS problem is the analysis of chronological data. For example, consider a blood glucose measurement history from an individual with diabetes (a disorder of blood sugar control): an answer to the question “When were blood sugar levels worst over the past year?” might prove useful in identifying factors impairing control. Typical data consists of a sequence of timestamped blood glucose readings in chronological order, such as that illustrated in Figure 2.

This situation is an instance of the MMS problem: each element is a blood glucose reading, and its weight is how far it deviates from the target blood sugar range. Thus, the density of a segment is how far its readings are outside the target range, on average, over the timespan of that segment. Here, it makes sense to put a lower bound on that timespan (for example, a week) as spurious extreme readings can happen occasionally, so it is not the single worst reading that is sought, but an average over a suitably long period.

We expect that MMS problems can be identified in many other application areas, including analysis of other kinds of medical history data.

3 Development

3.1 Specification & refinement

In this section, we continue on from the datatypes and definitions that have already been given in Section 2.1, to arrive at a specification for the densest-segment problem.

One standard way to produce all the segments of a list is to take all prefixes of suffixes:

$$\begin{aligned}
 \text{prefixes, suffixes, segments} &:: [a] \rightarrow [[a]], \\
 \text{prefixes} &= \text{inits} \\
 &\quad \textbf{where } \text{inits } [] = [[]] \\
 &\quad \quad \text{inits } (x : xs) = [[]] \uplus \text{map } (x :) (\text{inits } xs), \\
 \text{suffixes} &= \text{tails} \\
 &\quad \textbf{where } \text{tails } [] = [[]] \\
 &\quad \quad \text{tails } (x : xs) = (x : xs) : \text{tails } xs, \\
 \text{segments} &= \text{concat} \cdot \text{map } \text{prefixes} \cdot \text{suffixes}.
 \end{aligned}$$

Later, we will write $x \sqsubseteq y$ to denote that x is a prefix of or equal to y , and use \sqsubset for the strict version.

To select maximal segments, we will use a binary operator \uparrow_{\sqsubseteq} , which selects the maximum of two arguments with respect to a total ordering \sqsubseteq :

$$\begin{aligned}
 x \uparrow_{\sqsubseteq} y &= x, \quad \text{if } y \sqsubseteq x; \\
 &= y, \quad \text{otherwise.}
 \end{aligned}$$

In the case of a tie between the two arguments, \uparrow_{\sqsubseteq} selects the left-hand argument. Later, we will also need an operator \downarrow_{\sqsubseteq} , which is defined similarly to break ties in favour of the right-hand argument. We will abbreviate \uparrow_{\leq_d} and \downarrow_{\leq_d} by \uparrow_d and \downarrow_d respectively, to avoid cumbersome subscripts.

Given a binary operator \oplus (such as \uparrow_{\sqsubseteq} or \downarrow_{\sqsubseteq}) that selects the maximum of two values, the following function selects the maximum of a list:

$$\begin{aligned}
 \text{max}_{\oplus} [x] &= x \\
 \text{max}_{\oplus} (x : xs) &= x \oplus (\text{max}_{\oplus} xs).
 \end{aligned}$$

We can now specify the problem of finding a densest segment as

$$\text{max}_{\uparrow_d} \cdot \text{filter } \text{wide} \cdot \text{segments}. \tag{4}$$

At this stage, \downarrow_d could just as easily have been used, and the choice of \uparrow_d will be justified later in Section 3.4 (just after the proof of Lemma 3.5).

The above specification (4) can be a starting point for a calculation to derive a solution. However, during the subsequent development, a number of issues occur, which we find easier to manage if the initial specification is refined first.

The first issue is that the above specification does not describe a total function, because it may be that no segments of the input are wide enough. Unfortunately, this lack of totality forms a major obstacle to deriving an algorithm to solve this problem, as will be discussed later in Section 3.4.

Attempting a derivation using relational algebra rather than functional programming calculus, for example in the style described in Bird & de Moor (1997), does

not help, as the requirement for totality remains. To ensure totality, there needs to be unified treatment of wide segments and those that are not wide enough. Using a *Maybe* type is one possibility, but this leads to a very fiddly derivation.

Instead, we use another way to model treating segments differently depending on whether they are wide enough, by generalising \leq_d to a relation \leq with the following properties:

$$\leq \text{ is reflexive, transitive, and total;} \tag{5}$$

$$\neg \text{wide } x \wedge \text{wide } y \Rightarrow x < y, \tag{6}$$

$$\text{wide } x \wedge \text{wide } y \Rightarrow (x \leq y \Leftrightarrow x \leq_d y), \tag{7}$$

$$\neg \text{wide } x \wedge \neg \text{wide } y \Rightarrow (x \leq y \Leftrightarrow \text{width } x \leq \text{width } y), \tag{8}$$

where $x < y$ is a shorthand for $x \leq y \wedge y \not\leq x$. Informally, the requirement (5) ensures that taking a maximum from a non-empty list of segments is always possible, Equation (6) says that wide segments are always strictly better than those not wide enough, Equation (7) says that amongst wide lists, a denser segment is better, and Equation (8) says that amongst not-wide-enough segments, wider is better (the reason for which will appear as a consequence of the algorithmic development, at the end of Section 3.4). Note also that the definition of \leq ensures that density comparisons are only made when the segments are both wide enough and therefore non-empty.

In this way, selecting a densest segment can be carried out by \max_{\leq} , which we will abbreviate to \max_{\leq} . Our new version of the specification is then

$$\text{mds} = \max_{\leq} \cdot \text{segments}. \tag{9}$$

This is a total specification, unlike Equation (4). It produces a slightly different output, which can be used to solve the original problem: if the resulting segment is not at least L wide, then there were no wide segments. Otherwise, the output is a densest segment with respect to \leq_d , which will be the leftmost such segment, as \leq is used.

3.2 Initial calculation steps

Starting from the *mds* specification, we take a standard formal development route to calculate for the cases $[]$ and $(a : x)$. After obtaining $\text{mds } [] = []$, we then proceed as follows:

$$\begin{aligned} & \text{mds } (a : x) \\ = & \quad \{\text{definition of } \text{mds} \text{ (9)}\} \\ & (\max_{\leq} \cdot \text{segments}) (a : x) \\ = & \quad \{\text{definitions of } \text{segments} \text{ and } \text{suffixes}\} \\ & \max_{\leq} (\text{concat } (\text{map } \text{prefixes } ((a : x) : \text{suffixes } x))) \\ = & \quad \{\text{definitions of } \text{map} \text{ and } \text{concat}\} \\ & \max_{\leq} (\text{prefixes } (a : x) \uplus \text{concat } (\text{map } \text{prefixes } (\text{suffixes } x))) \end{aligned}$$

$$\begin{aligned}
&= \{max \text{ distributes over } \#, \text{ since } \uparrow_{\leq} \text{ is associative}\} \\
&\quad \max_{\leq} (\text{prefixes } (a : x)) \uparrow_{\leq} \max_{\leq} (\text{concat } (\text{map } \text{prefixes } (\text{suffixes } x))) \\
&= \{\text{definitions of } \textit{segments} \text{ and } \textit{mds}\} \\
&\quad \max_{\leq} (\text{prefixes } (a : x)) \uparrow_{\leq} \textit{mds } x \\
&= \{\text{definition, see below}\} \\
&\quad \textit{mdp } (a : x) \uparrow_{\leq} \textit{mds } x.
\end{aligned}$$

Above, we used this definition of a maximally dense prefix (*mdp*) with respect to \leq :

$$\textit{mdp} = \max_{\leq} \cdot \textit{prefixes}. \quad (10)$$

Maximally dense prefixes will feature heavily in the coming calculations, so we now pause the development in order to collect some basic properties of *mdp* that we will need later.

3.3 Properties of *mdp*

First, as can be deduced from its definition, *mdp* returns a prefix of its input for any list of elements x :

$$\textit{mdp } x \sqsubseteq x. \quad (11)$$

Second, this lemma describes the relationship between *mdp* and segment width:

Lemma 3.1

For any list of elements x ,

$$\textit{width } x \leq L \Rightarrow x = \textit{mdp } x, \quad (12)$$

$$\textit{wide } x \Leftrightarrow \textit{wide } (\textit{mdp } x). \quad (13)$$

The implication (12) expresses that taking the *mdp* of a narrow segment returns the segment itself, and Equation (13) expresses that taking the *mdp* of a segment does not affect whether it meets the width constraint or not.

Proof

The implication (12) is proved using a straightforward induction on x . A short calculation yields $\textit{wide } x \Rightarrow \textit{wide } (\textit{mdp } x)$, which together with Equation (12) is sufficient to show Equation (13). (For details, see the supplement.) \square

The function *mdp* has two reasonably straightforward-to-prove properties, which show that extending a segment can only increase the length and density of the densest prefix.

Lemma 3.2 (Monotonicity of *mdp*)

For any lists of elements x, y such that $x \sqsubseteq y$, we have that

$$\textit{mdp } x \sqsubseteq \textit{mdp } y, \quad (14)$$

$$\textit{mdp } x \leq \textit{mdp } y. \quad (15)$$

The next lemma says that a prefix shorter than the mdp of a segment is not maximally dense. This relies on mdp selecting the leftmost amongst maximally dense prefixes.

Lemma 3.3

For any lists of elements x, y such that $x \sqsubset y$, we have that

$$x \sqsubset mdp\ y \Rightarrow x < mdp\ y.$$

Finally, we have the following lemma that says if x is “sandwiched” in between $mdp\ y$ and y , then the maximally dense prefixes of x and y are the same.

Lemma 3.4 (Sandwich Lemma)

For any lists of elements x and y ,

$$mdp\ y \sqsubseteq x \sqsubseteq y \Rightarrow mdp\ x = mdp\ y. \quad (16)$$

3.4 A sliding-window algorithm

Returning to the algorithmic development, in Section 3.2, we calculated the following alternative definition for mds , which requires finding a densest prefix for all suffixes of the input:

$$\begin{aligned} mds\ [] &= [] \\ mds\ (a : x) &= mdp\ (a : x) \upharpoonright_{\leq} mds\ x. \end{aligned}$$

Despite being tail-recursive, this function is inefficient, as computing mdp for every suffix of the input would take too long. If, however, mdp itself can be expressed as a *foldr*, then the problem can be solved by performing the paired computation $\langle mds, mdp \rangle^1$, which can be carried out as a *foldr*: this would compute an optimal segment ($mds\ x$) and an optimal prefix ($mdp\ x$) for each suffix x of the input. Then, the value $mdp\ (a : x)$ would be readily available at each step, when calculating $mds\ (a : x)$ from $mds\ x$.

[A technical aside: aficionados of the morphism zoo may like to note that the $\langle mds, mdp \rangle$ computation illustrates the *zygomorphism* computational pattern (Malcolm, 1990), where the value of one expression is computed from a fold that pairs it with a second expression that is directly computable from a fold (“zygo” means “yoked”). The zygomorphic nature of our algorithm is why we are using total functions (as discussed in Section 3.1): zygomorphisms exist in the category of sets and total functions (**Fun**), but are not guaranteed to exist in the categories of partial functions (**PFun**) and relations (**Rel**).

Unfortunately, mdp is not a *foldr*. If it were, then $mdp\ (a : x)$ could be calculated from a and $mdp\ x$ alone, but consider the list $[(1, 5), (9, 2), (8, 2)]$ of *(weight, size)* elements for the MDS problem, with a segment width lower bound $L = 1.8$. Here, $mdp\ [(9, 2), (8, 2)] = [(9, 2)]$, but $mdp\ [(1, 5), (9, 2), (8, 2)] = [(1, 5), (9, 2), (8, 2)]$, which is not computable solely from $(1, 5)$ and $[(9, 2)]$.

¹ The “split” operator is defined by $\langle f, g \rangle a = (f\ a, g\ a)$.

One of the surprising features of the *mdp* problem, however, is that *mdp* not being a *foldr* does not seem to matter. In the above example, $mdp [(9, 2), (8, 2)]$ is denser than $mdp [(1, 5), (9, 2), (8, 2)]$, and thus the latter cannot be a solution to the whole problem. This suggests that perhaps $mdp (a : x)$ is not needed when x has a denser shorter prefix. This is the case, and we will prove a slightly generalised version: if $mdp (u \# x)$ is longer than $u \# mdp x$, then $mdp (u \# x)$ is less dense than $mdp x$. The following lemma formalises this property:

Lemma 3.5 (Overlap Lemma)

For all lists of elements u and x , we have

$$u \# mdp x \sqsubset mdp (u \# x) \Rightarrow mdp (u \# x) < mdp x.$$

Proof

We do a case analysis on the width of x .

Case $\neg \text{wide } x$.

By Equation (12), we have $x = mdp x$. The assumption thus simplifies to $u \# x \sqsubset mdp (u \# x)$ which is false, as $mdp (u \# x)$ is a prefix of $u \# x$, from property (11).

Case $\text{wide } x$.

As from Equation (2), prepending only increases width, then by Equation (13), we know that $\text{wide } (mdp x)$ and $\text{wide } (mdp (u \# x))$, and thus from the definition of $<$ and \leq Equation (7), the goal simplifies to proving that $mdp (u \# x) <_d mdp x$.

Note that since $u \# mdp x \sqsubset mdp (u \# x)$, we have $mdp (u \# x) = u \# mdp x \# z$, for some non-empty z . We first prove that $u \# mdp x \# z <_d z$:

$$\begin{aligned} & u \# mdp x \sqsubset mdp (u \# x) \\ \Rightarrow & \{ \text{lemma 3.3} \} \\ & u \# mdp x < mdp (u \# x) \\ \Leftrightarrow & \{ \text{decomposition of } mdp (u \# x), \text{ as above} \} \\ & u \# mdp x < u \# mdp x \# z \\ \Rightarrow & \{ \text{definition of } \leq (7), \text{ as } \text{wide } (mdp x) \text{ and width increases with } \# (2) \} \\ & u \# mdp x <_d u \# mdp x \# z \\ \Rightarrow & \{ \text{density property (3)} \} \\ & u \# mdp x \# z <_d z. \end{aligned}$$

Then, we prove that $z \leq_d mdp x$:

$$\begin{aligned} & u \# mdp x \# z = mdp (u \# x) \\ \Rightarrow & \{ \text{prefix property of } mdp (11) \} \\ & u \# mdp x \# z \sqsubseteq u \# x \\ \Leftrightarrow & \{ \text{property of prefixes} \} \\ & mdp x \# z \sqsubseteq x \\ \Rightarrow & \{ mdp x \text{ maximum wrt } \leq \} \end{aligned}$$

$$\begin{aligned}
 & mdp\ x \# z \leq mdp\ x \\
 \Rightarrow & \{ \text{definition of } \leq \text{ (7), as } wide(mdp\ x) \text{ and width increases with } \# \text{ (2)} \} \\
 & mdp\ x \# z \leq_d mdp\ x \\
 \Rightarrow & \{ \text{density property (3), with } \geq_d \} \\
 & z \leq_d mdp\ x.
 \end{aligned}$$

Thus, by transitivity we conclude that $u \# mdp\ x \# z <_d mdp\ x$, which is the same as $mdp\ (u \# x) <_d mdp\ x$. □

The Overlap Lemma (3.5) suggests an alternative to computing $mdp\ x$ for each suffix x of the input list of elements: each new element a is added to the front of the prefix produced from the previous stage of the computation, and then mdp of this segment is computed.

[An aside: this strategy is what suggests the use of \uparrow_{\leq} in the definition (10) of mdp : in the case that there is more than one densest prefix of a segment, $max_{\uparrow_{\leq}}$ chooses the shortest, thus leaving fewer elements to be processed later on than if the \uparrow_{\leq} operator were to be used. In turn, this justifies the use of \uparrow_{\leq} in the definition (9) of mds as well.]

Formally, this computation can be expressed as the following function wp :

$$\begin{aligned}
 wp\ [] &= [] & (17) \\
 wp\ (a : x) &= mdp\ (a : wp\ x).
 \end{aligned}$$

(The name wp is intended to abbreviate “window processing”.) Note that unlike mdp , the function wp is clearly a *foldr*. We now hope to use the value of wp instead of mdp for each suffix, in order to compute mds . That is, given the following function ms :

$$\begin{aligned}
 ms\ [] &= [] & (18) \\
 ms\ (a : x) &= wp\ (a : x) \uparrow_{\leq} ms\ x,
 \end{aligned}$$

we need to show that $mds = ms$, which occurs in Section 3.6.

The benefit of using ms may not be obvious. Recall that in Section 3.4, we were unable to use the paired computation $\langle mds, mdp \rangle$ to compute mds efficiently because mdp is not a *foldr*. Now, however, as wp is a *foldr*, we can use a paired computation $\langle ms, wp \rangle$, expressible as a *foldr* that makes one call to mdp in each step. Letting $m\ swp = \langle ms, wp \rangle$, a standard *tupling* transformation gives us

$$\begin{aligned}
 m\ swp\ [] &= ([], []) \\
 m\ swp\ (a : x) &= (w' \uparrow_{\leq} m, w') \\
 &\quad \textbf{where } (m, w) = m\ swp\ x \\
 &\quad \quad w' = mdp\ (a : w).
 \end{aligned}$$

This follows the paradigm known as a *sliding-window* algorithm (Zantema, 1992): such a computation maintains a segment of the input sequence (the *window*), that “slides” along the sequence at each step of the algorithm, always in the same direction. Above, w is the window storing the value $wp\ x$ from the previous step of the computation, and the application of $mdp \cdot (a :)$ produces an updated window w'

that is a prefix of $a : w$, thus sliding the window to the left. (The name wp stands for “window processing”.)

[Another aside: this is the point where it is clear how \leq compares segments that are not wide enough, as expressed earlier in Equation (8). When the algorithm has barely started, and is still examining suffixes of the input that are not wide enough, it is not clear what wp should return, but for the first suffix x that is wide enough, $mds\ x$ will be x itself, and therefore wp will need to retain all the elements of x . This means that on segments that are not wide enough, \leq should favour wider segments.]

Given suitable list structures to enable constant time computation of \uparrow_{\leq} , the above algorithm can be implemented in linear time, provided that the computation of $mdp \cdot (a :)$ can be done in amortised constant time. This is our goal in Section 4: choosing suitably efficient data structures for the window.

3.5 Properties of wp and ms

As mdp returns a prefix of its input in Equation (11), a short inductive proof starting from the definition of wp in Equation (17) shows that $wp\ x$ also returns prefixes

$$wp\ x \sqsubseteq x, \quad (19)$$

for any sequence of elements x .

The following lemma shows that wp and ms have similar relationships to the width of segments that mdp does (see Lemma 3.1).

Lemma 3.6

For any list of elements x ,

$$width\ x \leq L \Rightarrow x = wp\ x = mds\ x, \quad (20)$$

$$wide\ x \Leftrightarrow wide\ (wp\ x) \Leftrightarrow wide\ (mds\ x). \quad (21)$$

3.6 Correctness of the sliding-window algorithm

Now, we carry out our task to prove that $mds\ x = ms\ x$.

Proof

We use an induction on x , and the $[]$ case is immediate from the definitions of mds and ms . For the $a : x$ case, the proof obligation (after standard use of the definition of wp and the induction hypothesis that $mds\ x = ms\ x$) is

$$mdp\ (a : x) \uparrow_{\leq} mds\ x = mdp\ (a : wp\ x) \uparrow_{\leq} mds\ x. \quad (22)$$

Noting how \leq is defined, we start from a case analysis on the width of x .

Case $\neg wide\ x$. By Equation (20), $x = wp\ x$, and thus Equation (22) holds.

Otherwise, we split the case for $wide\ x$ into two cases, based on whether the left-hand side of Equation (22) equates to $mdp\ (a : x)$ or $mds\ x$. As the operator \uparrow_{\leq} selects its left-hand argument when comparing two segments that are equal with respect to \leq , the two cases are $mdp\ (a : x) < mds\ x$ and $mdp\ (a : x) \geq mds\ x$. As $wide\ x$

holds, these simplify, using the definition of \leq_d , to the cases $mdp(a : x) <_d mds\ x$ and $mdp(a : x) \geq_d mds\ x$ respectively.

Case *wide* x and $mdp(a : x) <_d mds\ x$.

As the left-hand side of Equation (22) reduces to $mds\ x$, we need to show that the right-hand side also reduces to $mds\ x$. This necessitates proving that $mdp(a : wp\ x) < mds\ x$:

$$\begin{aligned}
 & mdp(a : wp\ x) < mds\ x \\
 \Leftarrow & \quad \{ \text{since } mdp(a : x) <_d mds\ x \} \\
 & mdp(a : wp\ x) \leq mdp(a : x) \\
 \Leftarrow & \quad \{ \text{monotonicity of } mdp \text{ as in (15)} \} \\
 & a : wp\ x \sqsubseteq a : x \\
 \Leftarrow & \quad \{ wp \text{ produces prefixes (19)} \} \\
 & \text{true.}
 \end{aligned}$$

Case *wide* x and $mdp(a : x) \geq_d mds\ x$.

This is the difficult case, and clearly if we can establish that

$$mdp(a : x) \geq_d mds\ x \Rightarrow mdp(a : x) = mdp(a : wp\ x), \tag{23}$$

in the circumstance that *wide* x , then Equation (22) will hold.

To establish Equation (23), we need a separate result, Lemma 3.7 below, which generalises the prepending ($a :$) operation in Equation (23) to the prepending of a list ($z ++$). To use the lemma, we let $z := [a]$ and $y := x$ in Equation (24) of the lemma, and then the assumption required, on the left-hand side of Equation (24), is $mds\ x \leq mdp(a : x)$. This can be obtained by using the facts that $mdp(a : x) \geq_d mds\ x$ and *wide* x , together with Equations (13) and (21). \square

Lemma 3.7

For any finite lists of elements y and z ,

$$mds\ y \leq mdp(z ++ y) \Rightarrow mdp(z ++ y) = mdp(z ++ wp\ y). \tag{24}$$

Proof

The proof proceeds by induction on y . The case for $y := []$ is routine. For the case $a : y$, the assumption is

$$mds(a : y) \leq mdp(z ++ (a : y)),$$

and for the induction hypothesis, we can assume that

$$mds\ y \leq mdp(v ++ y) \Rightarrow mdp(v ++ y) = mdp(v ++ wp\ y),$$

for any finite list of elements v .

We need to show that $mdp(z ++ a : y) = mdp(z ++ wp(a : y))$. One crucial idea is to use the sandwich lemma to turn the proof obligation from an equality to a prefix relation, which then allows us to use the overlap lemma, the key property that

guarantees that we need not consider prefixes that are too long. The proof goes as follows:

$$\begin{aligned}
 & mdp(z \# a : y) = mdp(z \# wp(a : y)) \\
 \Leftarrow & \quad \{ \text{ sandwich lemma (16)} \} \\
 & mdp(z \# a : y) \sqsubseteq z \# wp(a : y) \sqsubseteq z \# a : y \\
 \Leftrightarrow & \quad \{ \text{ by (19), } wp(a : y) \sqsubseteq a : y \} \\
 & mdp(z \# a : y) \sqsubseteq z \# wp(a : y) \\
 \Leftarrow & \quad \{ \text{ induction, with } v = z \# [a], \text{ writing } P \text{ for } mds\ y \leq mdp(z \# a : y) \} \\
 & mdp(z \# a : wp\ y) \sqsubseteq z \# wp(a : y) \wedge P \\
 \Leftrightarrow & \quad \{ \text{ definition of } wp \} \\
 & mdp(z \# a : wp\ y) \sqsubseteq z \# mdp(a : wp\ y) \wedge P \\
 \Leftarrow & \quad \{ \text{ overlap lemma 3.5, contra-positive form (see note below)} \} \\
 & mdp(a : wp\ y) \leq mdp(z \# a : wp\ y) \wedge P \\
 \Leftarrow & \quad \{ \text{ induction} \} \\
 & mdp(a : wp\ y) \leq mdp(z \# a : y) \wedge P \\
 \Leftarrow & \quad \{ \text{ by monotonicity (15), and } a : wp\ y \sqsubseteq a : y \} \\
 & mdp(a : y) \leq mdp(z \# a : y) \wedge P \\
 \Leftrightarrow & \quad \{ \text{ expanding } P \} \\
 & mdp(a : y) \leq mdp(z \# a : y) \wedge mds\ y \leq mdp(z \# a : y) \\
 \Leftrightarrow & \quad \{ \text{ maximum} \} \\
 & mdp(a : y) \upharpoonright_{\leq} mds\ y \leq mdp(z \# a : y) \\
 \Leftrightarrow & \quad \{ \text{ definition of } mds \} \\
 & mds(a : y) \leq mdp(z \# a : y) \\
 \Leftarrow & \quad \{ \text{ assumption} \} \\
 & \text{true.}
 \end{aligned}$$

Note that the step above using the Overlap Lemma relies on the lists $mdp(z \# a : wp\ y)$ and $z \# mdp(a : wp\ y)$ both being prefixes of $z \# a : wp\ y$, in order to simplify the \sqsubseteq in $mdp(z \# a : wp\ y) \sqsubseteq z \# mdp(a : wp\ y)$ to \sqsupseteq . \square

4 Window

Recall the definition of wp in Equation (17), which describes the sliding of the window in the main algorithm:

$$\begin{aligned}
 wp\ [] &= [] \\
 wp\ (a : x) &= mdp(a : wp\ x).
 \end{aligned}$$

Our goal is to choose a suitable data structure for the window so that the computation of $mdp \cdot (a :)$ takes amortised constant time, thus allowing the overall algorithm to be linear.

In this section, we will examine the properties of mdp in order to see how best to represent the window, so that suitable data structures for the window can be chosen in Section 5. In what follows, let the window w at the previous step of the computation be denoted by w , and let $w' = a : w$.

4.1 Window header

Recall that the function mdp takes the best prefix of w' with respect to \uparrow_{\leq} . In the case that $\neg wide\ w'$, then we just have $mdp\ w' = w'$, from Lemma 3.1. In the case when $wide\ w'$, another straightforward step can be taken, given part (6) of the definition of \leq , which says that wide segments are always better than too-short segments. Thus, when $wide\ w'$, it is also the case that $wide\ (mdp\ w')$, as formalised in Lemma 3.1, in Equation (13). This means that when $wide\ w'$, the prefix $mdp\ w'$ must include sufficiently many elements that it is wide enough; these elements can be thought of as a ‘‘header’’ prefix of w' .

Let the function $hsplit$ be such that it splits a segment into its header prefix, and the rest of its elements.

$$hsplit :: [Elem] \to ([Elem], [Elem]).$$

A possible definition of $hsplit$ is given later (see Figure 8 in Section 5.2), but for now, we just need the following property: if $hsplit\ w' = (h, x)$, we have $h \# x = w'$, where h is the header prefix, so that either

- $width\ h < L$ and $x = []$, or
- $width\ (init\ h) < L \leq width\ h$ (i.e. h is the shortest wide prefix), where $init\ (y \# [b]) = y$.

Furthermore, recall that part (7) of the definition of \leq says that the better of two wide segments is determined by \leq_d . Thus, when $wide\ w'$, taking a maximum prefix using the \uparrow_{\leq} operator in $mdp\ w'$ is the same as max_{\uparrow_d} (prefixes w'). This means that for wide w' , we can restrict our attention to prefixes of w' of the form $h \# y$, where $(h, x) = hsplit\ w'$ and $y \sqsubseteq x$, and compare densities with respect to \uparrow_d .

We thus represent a window as a pair, where the first component is the header h ,

$$\mathbf{type}\ Window = ([Elem], \dots),$$

and the second component is some structure yet to be determined, to represent the rest of the window. From this structure, we will need to find the segment that produces the densest prefix of the whole window with respect to \uparrow_d , when prepended by the header h .

4.2 Densest prefixes

Let $(h, x) = hsplit\ w'$; our goal is to find a prefix of w' of the form $h \# y$, that is densest with respect to \uparrow_d . One possibility is that $y = []$ and h itself is the densest such segment; otherwise, there must be a non-empty y such that $h <_d h \# y$, which is equivalent to $h <_d y$, by the density property (3). This leads us to investigate




Fig. 3. A running MDS example: the list of elements [(58,4), (77,11), (36,3), (23,2), (-3,1), (61,3), (12,2), (69,4), (43,4), (88,6), (97,10), (82,5), (77,5), (60,5), (60,7), (45,3), (-50,10), (32,9), (-18,9), (69,9), (66,8)], with weights indicated in bold. Lower bound $L = 16$; the compulsory header is h .

non-empty prefixes of x that are denser than h ; indeed perhaps choosing y to be a maximally dense prefix of x might result in a densest prefix $h \uparrow y$?

Unfortunately, this is not the case: there may be a prefix of x that is slightly less dense than y , but wider, and that produces a denser prefix overall when prepended by the header h . A concrete example is illustrated in Figure 3, where a maximally dense non-empty prefix of x is $y = [(23, 2), (-3, 1), (61, 3), (12, 2), (69, 4)]$, with density 13.5, resulting in a density of 11.1 for $h \uparrow y$. However, extending y on the right by the following elements $z = [(43, 4), (88, 6), (97, 10), (82, 5), (77, 5), (60, 5)]$ yields an improved density of 12 for $h \uparrow y \uparrow z$.

This is disappointing, but maybe a densest prefix can still be of use? Although the example above illustrates that y (a maximally dense prefix of x) can be bettered by a prefix wider (longer) than y , perhaps non-empty prefixes shorter than y can be ruled out? It turns out that this is indeed the case, which will be shown in what follows.

4.3 Right-skew segments

First, we will need to examine some properties of densest prefixes. Note that if a non-empty list of elements $y = y' \uparrow y''$ is a maximally dense prefix of x , then as y' is also a prefix of x , we have $y' \leq_d y$. Then, from the density property, this is equivalent to $y' \leq_d y''$. This leads us to the following definition:

Definition 4.1 (Right-Skew)

A non-empty list z of elements is called *right-skew* if for all n such that $0 < n < \text{length } z$, we have $\text{take } n \ z \leq_d \text{drop } n \ z$. Let us denote this by the predicate

$$\text{rightskew} (\leq_d) :: \text{Elem} \rightarrow \text{Bool}.$$

The idea of right-skew segments originated with Lin *et al.* (2002); informally, a segment is right-skew when chopping it into two always results in a right-hand side that is at least as dense than the left. However, note that density does not increase in a monotonic fashion: although shorter prefixes of right-skew segments are often less dense than longer prefixes, this is merely a trend, and should not be relied upon. For an example of a right-skew list that illustrates this lack of monotonicity, see Figure 4.

From this reasoning, we have that every densest prefix of a list of elements is right-skew (the converse does not hold):




Fig. 4. The segment $[(23,2), (-3,1), (61,3), (12,2), (69,4)]$ is right-skew.

Lemma 4.2

Let y be a non-empty prefix of a list of elements x , such that y is maximally dense with respect to \leq_d . Then, y is right-skew.

One property of right-skew segments we will need is that the concatenation of two right-skew segments in the correct order is itself right-skew.

Lemma 4.3

If z and z' are both right-skew and $z \leq_d z'$, then $z \uplus z'$ is right-skew.

Returning to the development, recall that given a non-empty list of elements y that is a maximally dense prefix of x , we wish to rule out the use of prefixes shorter than y . The following lemma allows us to do so:

Lemma 4.4

Let z_1, z_2, z_3 be non-empty list of elements such that $z_2 \leq_d z_3$. Then,

$$z_1 \upharpoonright_d (z_1 \uplus z_2) \upharpoonright_d (z_1 \uplus z_2 \uplus z_3) = z_1 \upharpoonright_d (z_1 \uplus z_2 \uplus z_3).$$

In words, this means that the densest among the three lists $z_1, z_1 \uplus z_2$ and $z_1 \uplus z_2 \uplus z_3$ with respect to \upharpoonright_d is either z_1 or $z_1 \uplus z_2 \uplus z_3$. Thus, as a densest prefix y of x is right-skew, and hence any proper division of y into $y = y' \uplus y''$ results in $y' \leq_d y''$, Lemma 4.4 allows us to deduce that of the three lists $h, h \uplus y'$, and $h \uplus y$, the densest with respect to \upharpoonright_d is not $h \uplus y'$. Thus, non-empty prefixes of y can be eliminated from consideration. This can be shown formally, as an application of the following theorem:

Theorem 4.5

Let z and y be non-empty lists of elements with y right-skew. Then,

$$\max_{\upharpoonright_d} (\text{map } (z \uplus) (\text{prefixes } y)) = z \upharpoonright_d (z \uplus y). \tag{25}$$

This helps suggest a strategy for structuring the rest of the window, as follows. Having shown that the interior of any right-skew prefix y of x can be eliminated from consideration, it seems reasonable that choosing the longest possible right-skew prefix of x (let us denote this prefix by y_0) may help reduce the remaining computation within the rest of the window. Furthermore, having established h and $h \uplus y_0$ as possible candidates for the *mdp* of the window, but with no need to consider prefixes in between, it seems reasonable to try the same step with the rest of the window. That is, let y_1 be the longest right-skew prefix of $x \setminus y_0$ (here, the

list subtraction operator is denoted by $\setminus\setminus$, so that $x \setminus\setminus y_0$ contains the elements of x apart from its initial prefix of y_0 . Then applying Theorem 4.5 with $z = h \uplus y_0$ and $y = y_1$ removes prefixes between $h \uplus y_0$ and $h \uplus y_0 \uplus y_1$ from consideration.

Similarly, let y_2 be the longest right-skew prefix of $x \setminus\setminus (y_0 \uplus y_1)$, and so on, resulting in a partition of x into right-skew lists of elements $[y_0, y_1, \dots, y_k]$. Repeated application of Theorem 4.5 leads us to the following theorem:

Theorem 4.6

Given $h : [Elem]$ and $xs :: [[Elem]]$ such that each list in xs is right-skew, we have $max_{\uparrow_d} (map (h \uplus) (prefixes (concat xs))) = max_{\uparrow_d} (map ((h \uplus) \cdot concat) (prefixes xs))$.

This theorem says that to compute $max_{\uparrow_d} (h \uplus y_0 \uplus y_1 \uplus \dots \uplus y_k)$ we only need to consider the ends of each partition: $h, h \uplus y_0, h \uplus y_0 \uplus y_1$, etc.

Later on, we will need an alternative version of the above theorem, as follows:

Corollary 4.7

Given $h : [Elem]$ and $xs :: [[Elem]]$ such that each list in xs is right-skew, we have

$$max_{\uparrow_h} (prefixes (concat xs)) = concat (max_{\uparrow_{h \uplus}} (prefixes xs)),$$

where \uparrow_h and $\uparrow_{h \uplus}$ abbreviate $\uparrow_{\leq_d^h}$ and $\uparrow_{\leq_d^{h \uplus}}$ respectively, where

$$\begin{aligned} x \leq_d^h y &\equiv h \uplus x \leq_d h \uplus y, \\ xs \leq_d^{h \uplus} ys &\equiv h \uplus concat xs \leq_d h \uplus concat ys. \end{aligned}$$

This partitioning of x into right-skew lists will be how the rest of the window is structured, but first we need to formally examine the properties of the partition we have just created, and prove some properties we will need.

4.4 Decreasing right-skew partitions

In this section, we present the concept of *DRSP*, which were proposed by Lin *et al.* (2002). We also investigate the properties of these partitions, and show that they are the same as the partitions proposed above for the window structure.

A decreasing right-skew partition (abbreviated *DRSP*) of a list of elements is a partition of the list into segments, such that each segment is right-skew and the densities of the segments are strictly decreasing from left to right. Formally:

Definition 4.8 (DRSP)

Let xs be a partition of a list of elements x , that is, where $concat xs = x$. The list xs is a *DRSP* of x when

- all (rightskew (\leq_d)) xs , and
- $sdec (\leq_d) xs$,

where $sdec (\leq)$ is the predicate that expresses that a list is strictly decreasing with respect to a linear ordering \leq , i.e. $sdec (\leq) [x_1 \dots x_n]$ holds precisely when $x_1 \triangleright x_2 \triangleright \dots \triangleright x_n$.




Fig. 5. The (unique) DRSP of the list x from Fig. 3 is $[[(23, 2), (-3, 1), (61, 3), (12, 2), (69, 4), [(43, 4), (88, 6)], [(97, 10), (82, 5), (77, 5)], [(60, 5)], [(60, 7), (45, 3)], [(-50, 10), (32, 9), (-18, 9), (69, 9), (66, 8)]]$. The densities of its segments are 13.5, 13.1, 12.8, 12, 10.5, and 2.2 respectively, indicated by the dashed lines representing the average heights of the segments.

Note that the right-skew property ensures that each segment of the partition is non-empty. An example of a DRSP can be seen in Figure 5.

The first important property that we need is that DRSPs are unique:

Theorem 4.9 (DRSP uniqueness)

Let x be a list of elements. There exists precisely one DRSP of x .

This is a surprising property of such partitions, and indeed we did not find the proof of uniqueness in Lin *et al.* (2002) convincing, so have provided our own in the supplement to this paper.

A consequence of DRSP uniqueness that we will need is the necessary existence of a function $drsp_{\leq d} : [Elem] \rightarrow [[Elem]]$ that produces the unique DRSP, given a list of elements; a definition of such a function will be given later (see just after Theorem 4.12).

The following property notes that a DRSP has a “rotational symmetry”: to visualise this, turn Figure 5 upside-down, and it still depicts a DRSP:

Theorem 4.10 (DRSP rotation)

For any non-empty list of elements x ,

$$drsp_{\leq d} x = reverse \cdot map \ reverse \cdot drsp_{\geq d} (reverse \ x).$$

The DRSP rotation property is not needed for this algorithmic development, but it is included here for the record.

The final property of DRSPs concerns how they can be constructed. We will need some preliminaries. First, the longest right-skew prefix of a list of elements is also the longest densest prefix of the list:

Lemma 4.11

Let y be a non-empty list of elements. Then, the longest right-skew prefix of y is also the longest of the maximally dense prefixes of y .

It is this lemma that illustrates how the partition of the right-hand side of the window $x = concat [y_0, y_1, \dots, y_k]$, as suggested in the previous section, results in a DRSP. Recall that the segment y_0 is selected as the longest right-skew prefix of x , and therefore, by Lemma 4.11, it must also be the longest densest prefix of x . This ensures that the next segment y_1 must be of strictly lower density than

y_0 (otherwise $y_0 \# y_1$ would be at least as dense as y_0 , contradicting y_0 being the longest of the maximally dense prefixes of x). Thus $y_0 >_d y_1$, and repeatedly selecting longest right-skew segments results in a partition of right-skew segments of decreasing densities.

Lemma 4.11 describes the first of several possible methods of constructing DRSP:

Theorem 4.12 (DRSP construction)

The DRSP of a list of elements can be constructed by any of the following methods:

- i. repeatedly taking longest maximum-density prefixes
- ii. repeatedly taking longest right-skew prefixes²
- iii. repeatedly taking longest minimum-density suffixes³
- iv. repeatedly taking longest right-skew suffixes.

The following function illustrates one of the above ways (i) to build a DRSP:

$$\begin{aligned} \text{drsp1} &:: [\text{Elem}] \rightarrow [[\text{Elem}]] \\ \text{drsp1} [] &= [] \\ \text{drsp1 } x &= y : \text{drsp1 } (\text{drop } (\text{length } y) x) \\ &\quad \mathbf{where } y = \max \uparrow_d (\text{prefixes } x). \end{aligned}$$

However, for the representation of the window in our algorithm, we are not usually going to be constructing a complete DRSP from a list of elements; the window will already consist of a header and the DRSP of the remaining elements, and we will want to incrementally update the partition as new elements are added.

The following lemma allows a DRSP to be updated from the left-hand side:

Lemma 4.13

Let z be a right-skew list of elements, and ys a DRSP. Then, $\text{prepend } z \text{ } ys$ is a DRSP, where

$$\begin{aligned} \text{prepend} &:: [\text{Elem}] \rightarrow [[\text{Elem}]] \rightarrow [[\text{Elem}]] \\ \text{prepend } z [] &= [z] \\ \text{prepend } z (y : ys) &= \mathbf{if } z \leq_d y \mathbf{ then } \text{prepend } (z \# y) ys \\ &\quad \mathbf{else } z : y : ys. \end{aligned}$$

Above, the prepend function repeatedly joins the list of elements z with the leftmost segment of ys (justified by Lemma 4.3), until the densities are once more decreasing.

This leads to the following alternative way to build a DRSP, using a function addl that appends one element to the left of an existing DRSP:

$$\begin{aligned} \text{drsp} &:: [\text{Elem}] \rightarrow [[\text{Elem}]] \\ \text{drsp} &= \text{foldr } \text{addl} [], \\ \text{addl} &:: \text{Elem} \rightarrow [[\text{Elem}]] \rightarrow [[\text{Elem}]] \\ \text{addl } a \text{ } xs &= \text{prepend } [a] \text{ } xs. \end{aligned}$$

² A longest right-skew prefixes approach is used in Goldwasser *et al.* (2005) and Lin *et al.* (2002).

³ The algorithm in Chung & Lu (2004) slides the window from left to right, using a mirror-image DRSP structure, the construction of which is based on taking longest minimum-density prefixes.

Note that the right-skew property that *prepend* requires of its first argument applies, as any singleton list $[a]$ is right-skew. The function *addl* will be useful in our final code.

4.5 Use of the DRSP

To summarise where we have reached in the development: given a list of elements $w' = a : w$ for which we want to calculate *mdp* w' , we split the window w' using *hsplit* $w' = (h, x)$, and represent the list x by its DRSP. The data structure for the window is therefore:

type *Window* = (*Elem*), [[*Elem*]],

where the first component is the compulsory header h , and the second is the rest of the window, partitioned into *drsp* x . Furthermore, Theorem 4.6 established that we don't need to consider prefixes that end in the interior of right-skew segments in the partition of x .

Note that the lists in the *Window* datatype should be considered as abstract representations at this stage, as we have not yet finalised whether to use cons-lists, snoc-lists, or some other queue representation, as befits a sliding-window algorithm. We will make these decisions in Section 5.2, to allow the final algorithm to run in linear time.

Having chosen a data structure for the window, we will need some representation-changing functions to convert between a list of elements and its partition into a DRSP:

wbuild :: [*Elem*] → *Window*
wbuild = (*id* × *drsp*) · *hsplit*,
wflatten :: *Window* → [*Elem*]
wflatten (*h*, *xs*) = *h* ++ *concat xs*.

Above, the function *wbuild* constructs a window from a list of elements⁴, while *wflatten* does the opposite, flattening a window back to a list. Thus, we have *wflatten* · *wbuild* = *id*.

To work on a more abstract level, we define a function *wcons*, so that *wcons a* adds an element a to the left side of the window. It can be seen as $(a :)$ lifted to the *Window* datatype, although it performs much more work – repartitioning the header, and updating the DRSP:

wcons :: *Elem* → *Window* → *Window*
wcons a (*h*, *xs*) = (*h'*, *foldr addl xs x*)
where (*h'*, *x*) = *hsplit* (*a* : *h*).

We have made progress: the data structure for the window has been chosen, we can construct and update a DRSP by adding elements incrementally to the front of it. Now, we just need the *mdp* of the window.

⁴ The “product functor” (\times) is defined by $(f \times g) (a, b) = (f a, g b)$.




Fig. 6. In our running example, the options for prepending the header h onto prefixes given by the DRSP results in prefix densities (from left to right) 9.5 11.1, 11.6, 12, 12, 11.8, and 8.2. This bitonic hill shape can be seen outlined in the arrowheads.

4.6 Extracting the densest prefix

To find the densest prefix with respect to \uparrow_d , we need to make use of Theorem 4.6, which guarantees that for a partition of the elements beyond the compulsory header into right-skew segments (such as the DRSP we are using), only the positions of the input corresponding to the ends of these segments need to be considered.

Furthermore, the decreasing densities of the DRSP will enable us to pinpoint exactly where the mdp of the window elements is to be found, as follows. Empirical examination suggests that the densities of the prefixes h , then $h \uparrow y_0$ and so on up to $h \uparrow y_0 \dots \uparrow y_k$ form a simple hill shape, illustrated by our running example in Figure 6. To be precise: the hill for the window prefix densities of a header & DRSP is bitonic, consisting of a strictly increasing ascent on the left, followed by a (non-strictly) decreasing descent. Either side of the hill can be empty, for example it might just be a gentle decreasing slope. For the moment, please assume that our assertion about this hill shape is true, and this will soon be addressed formally in Lemma 4.14.

From this hill shape, it is now easy to identify the mdp : the densest window prefix(es) can be found at the top of the hill, and in particular, the densest prefix wrt \uparrow_d (which is what mdp uses) will be at the top of the hill, at the left-hand end if it turns out that the top of the hill is a plateau (i.e. more than one prefix having maximal density).

Finding the top of the hill can be achieved by a simple traverse from either the left or right side of the hill; however it will be seen later that starting from the right will have a crucial effect on the overall efficiency of the algorithm, enabling each element to be processed once only (more details are given in Section 5.2). Here is the definition of a function that, given a header h and a DRSP, chops off segments of the DRSP from the right-hand side, stopping at the left of the top of the hill:

$$\begin{aligned}
 \text{maxchop} & \quad :: [Elem] \rightarrow [[Elem]] \rightarrow [[Elem]] \\
 \text{maxchop } h \ [] & \quad = [] \\
 \text{maxchop } h (xs \uparrow [x]) & = \text{if } h \uparrow \text{concat } xs <_d x \text{ then } xs \uparrow [x] \\
 & \quad \text{else } \text{maxchop } h \ xs.
 \end{aligned}$$

The maxchop function repeatedly removes a segment x when $h \uparrow \text{concat } xs \not\prec_d x$, which from the density property is equivalent to $h \uparrow \text{concat } xs \geq_d h \uparrow \text{concat } (xs \uparrow [x])$. In other words, maxchop chops from the right whilst going up the right-hand side of the hill, only stopping if it runs out of hill (the $[]$ case) or if it finds a strict

decrease in density, as the given condition $h \# concat \ xs <_d x$ is equivalent to $h \# concat \ xs <_d h \# concat \ (xs \# [x])$, from the density property. To illustrate, *maxchop* applied to our running MDS example in Figure 6 returns the list of segments $[[(23, 2), (-3, 1), (61, 3), (12, 2), (69, 4), [(43, 4), (88, 6)], [(97, 10), (82, 5), (77, 5)]]$.

The following lemma shows that the list of segments that *maxchop* returns is strictly ascending in density when prepended by the header h , and thus shows formally that we really do have a simple bitonic hill shape as described at the beginning of this section.

Lemma 4.14

Let h be a non-empty list of elements, and $xs \# [x]$ be a list of non-empty segments with strictly *decreasing* densities. If $h \# concat \ xs <_d x$, then the following list of segments has strictly *increasing* densities:

$$map \ ((h\#) \cdot concat) \ (prefixes \ (xs \# [x])).$$

We can then use this lemma to show that *maxchop* does indeed compute the *mdp* for us, which is what we require.

Theorem 4.15

Let h a non-empty list of elements, and let xs be a (possibly empty) list of non-empty segments having strictly decreasing densities. Then,

$$h \# concat \ (maxchop \ h \ xs) = max_{1_d} \ (map \ ((h\#) \cdot concat) \ (prefixes \ xs)).$$

Later, in our proofs, we will actually use the following corollary, which is the same property stated differently:

Corollary 4.16

Let h a non-empty list of elements, and let xs be a (possibly empty) list of non-empty segments having strictly decreasing densities. Then,

$$maxchop \ h \ xs = max_{1_{h\#}} \ (prefixes \ xs).$$

($1_{h\#}$ is defined in Corollary 4.7.)

For convenience, we also define a wrapper function for *maxchop* that operates on the whole window:

$$\begin{aligned} wmaxchop &:: Window \rightarrow Window \\ wmaxchop \ (h, xs) &= (h, maxchop \ h \ xs). \end{aligned}$$

5 Putting everything together

Let us remind ourselves of the development so far, the outline of which is summarised in Figure 7.

In Section 2, the original specification was refined to produce the function *mds*, then in Section 3, a sliding-window algorithm *ms* was shown to satisfy that specification, where the value of *ms* is obtained from the paired computation

<p>The refined specification (9): $mds = \max_{\leq} \cdot \text{segments}$</p> <p>We proved that $mds = ms$</p> <p>Computing ms is obtained from $\langle ms, wp \rangle$, which can be expressed as a <i>foldr</i>.</p>	<p>Recap of definitions:</p> $ms [] = []$ $ms (a : x) = wp (a : x) \upharpoonright_{\leq} ms x$ $wp [] = []$ $wp (a : x) = mdp (a : wp x)$ $mdp = \max_{\leq} \cdot \text{prefixes}$
---	--

Fig. 7. Summary of the algorithmic development so far.

$\langle ms, wp \rangle$. Section 4 was devoted to finding a faster way to compute $mdp \cdot (a :)$ in the inductive step for wp , the window-processing function, and as a result we obtained a datatype *Window* to use for representing $wp x$.

In this section, we will put the results of Sections 3 and 4 together to calculate our algorithm to solve the density problem. We will then refine the data structure in Section 5.2 to allow efficient implementation of certain operations. For now, using a standard list data structure allows us to see what is going on in the algorithm more easily.

5.1 Structuring the window

Our algorithm for obtaining a maximally dense segment is $ms = fst \cdot \langle ms, wp \rangle$, from which we can introduce the window structure as follows:

$$\begin{aligned}
 &fst \cdot \langle ms, wp \rangle \\
 = &\{ \text{window identity, see page 21} \} \\
 &fst \cdot \langle ms, wflatten \cdot wbuild \cdot wp \rangle \\
 = &\{ \text{identity, composition} \} \\
 &fst \cdot (id \times wflatten) \cdot \langle ms, wbuild \cdot wp \rangle \\
 = &\{ \text{projections} \} \\
 &fst \cdot \langle ms, wbuild \cdot wp \rangle.
 \end{aligned}$$

Our task is now to push $wbuild$ into the calculation of wp so that all the window processing is done with the *Window* datatype. Thus, we define

$$\begin{aligned}
 mwp &:: [Elem] \rightarrow ([Elem], Window) \\
 mwp &= \langle ms, wbuild \cdot wp \rangle,
 \end{aligned}$$

and we will aim to calculate an inductive definition of mwp .

The base case is straightforward: $mwp [] = ([], ([], []))$. For the inductive case, we calculate

$$\begin{aligned}
 &mwp (a : x) \\
 = &\{ \text{definitions of } ms \text{ and } wp \}
 \end{aligned}$$

$$\begin{aligned}
& (mdp (a : wp \ x) \uparrow_{\leq} ms \ x, wbuild (mdp (a : wp \ x))) \\
= & \quad \{ \text{since } wflatten \cdot wbuild = id \} \\
& (wflatten u \uparrow_{\leq} ms \ x, u) \\
& \textbf{where } u = wbuild (mdp (a : wp \ x)).
\end{aligned}$$

Consider the above binding $u = wbuild (mdp (a : wp \ x))$. If we can somehow manage to push $wbuild$ to the right and show that $u = f (wbuild (wp \ x))$ for some f , we will then have an inductive definition for mwp . To do this, since $wbuild = (id \times drsp) \cdot hsplit$, it will be helpful to know how $hsplit$ interacts with $(a :)$ and mdp . First, the following lemma shows how $hsplit$ and $(a :)$ exchange:

Lemma 5.1

For any element a and list of elements z ,

$$\begin{aligned}
hsplit (a : z) &= (h, z_1 \uparrow z_2) \\
& \textbf{where } (h', z_2) = hsplit \ z \\
& (h, z_1) = hsplit (a : h').
\end{aligned}$$

Second, this is how $hsplit$ interacts with mdp :

Lemma 5.2

For any list of elements x ,

$$\begin{aligned}
hsplit (mdp \ x) &= (h, max_{1_h} (prefixes \ y)) \\
& \textbf{where } (h, y) = hsplit \ x.
\end{aligned}$$

(1_h is defined in Corollary 4.7.)

This lemma merely restates the left-hand side involving the input list of elements x , into a form relating to the window data structure, with its compulsory header h . Recall that mdp chooses an optimal prefix with respect to \leq . If x is not wide enough, then both y and $mdp \ x$ evaluate to $[]$, and $h = x$. Otherwise, the right-hand side says that comparison with respect to \leq is carried out by comparing densities using \leq_d^h (also defined in Corollary 4.7) on prefixes within the window structure.

We are now ready to transform $wbuild (mdp (a : wp \ x))$. Writing z for $wp \ x$, we calculate

$$\begin{aligned}
& wbuild (mdp (a : z)) \\
= & \quad \{ \text{definition of } wbuild \} \\
& (id \times drsp) (hsplit (mdp (a : z))) \\
= & \quad \{ \text{property of } hsplit \text{ in Lemma 5.2, letting } (h, z') = hsplit (a : z) \} \\
& (id \times drsp) (h, max_{1_h} (prefixes \ z')) \\
= & \quad \{ \text{letting } (h', z_2) = hsplit \ z \text{ and } (h, z_1) = hsplit (a : h'), \text{ as in Lemma 5.1} \} \\
& (id \times drsp) (h, max_{1_h} (prefixes (z_1 \uparrow z_2))) \\
= & \quad \{ \text{since } concat \cdot drsp = id \} \\
& (h, drsp (max_{1_h} (prefixes (concat (drsp (z_1 \uparrow z_2))))))
\end{aligned}$$

$$\begin{aligned}
&= \{ \text{Corollary 4.7} \} \\
&\quad (h, drsp (\text{concat} (\max_{1_{h\#}} (\text{prefixes} (drsp (z_1 \# z_2)))))) \\
&= \{ drsp (z_1 \# z_2) = \text{foldr addl} (drsp z_2) z_1, \text{ from the definition of } drsp \text{ in Section 4.4} \} \\
&\quad (h, drsp (\text{concat} (\max_{1_{h\#}} (\text{prefixes} (\text{foldr addl} (drsp z_2) z_1)))) \\
&= \{ \text{definition of } wbuild \text{ (see Section 4.5), letting } (h', zs_2) = wbuild z \} \\
&\quad (h, drsp (\text{concat} (\max_{1_{h\#}} (\text{prefixes} (\text{foldr addl} zs_2 z_1)))) \\
&= \{ \text{definition of } wcons, \text{ letting } (h, zs) = wcons a (wbuild z) \} \\
&\quad (h, drsp (\text{concat} (\max_{1_{h\#}} (\text{prefixes} zs)))) \\
&= \{ \text{Theorem 4.16} \} \\
&\quad (h, drsp (\text{concat} (\text{maxchop } h \text{ } zs))) \\
&= \{ \text{cancellation, as } zs \text{ is a DRSP (see below)} \} \\
&\quad (h, \text{maxchop } h \text{ } zs) \\
&= \{ \text{definition of } wmaxchop \text{ (see end of Section 4.5)} \} \\
&\quad (wmaxchop \cdot wcons a \cdot wbuild) z.
\end{aligned}$$

The cancellation in the penultimate step is valid because of the following: *wbuild* builds a DRSP, and *wcons* and *maxchop* return DRSPs when given a DRSP as input. Then, as DRSPs are unique, applying *drsp* · *concat* to a DRSP has no effect.

In summary, we have shown that

$$wbuild \cdot mdp \cdot (a :) = wmaxchop \cdot wcons a \cdot wbuild. \quad (26)$$

We resume the calculation of the inductive case for *mwp*:

$$\begin{aligned}
&mwp (a : x) \\
&= \{ \text{by previous calculation and (26)} \} \\
&\quad (\text{wflatten } u \uparrow_{\leq} ms \ x, u) \\
&\quad \mathbf{where} \ u = wmaxchop (wcons a (wbuild (wp \ x))) \\
&= \{ \text{adding variables } m \text{ and } w \} \\
&\quad (\text{wflatten } u \uparrow_{\leq} m, u) \\
&\quad \mathbf{where} \ m = ms \ x \\
&\quad \quad w = wbuild (wp \ x) \\
&\quad \quad u = wmaxchop (wcons a \ w) \\
&= \{ \text{definition of } mwp \} \\
&\quad (\text{wflatten } u \uparrow_{\leq} m, u) \\
&\quad \mathbf{where} \ (m, w) = mwp \ x \\
&\quad \quad u = wmaxchop (wcons a \ w).
\end{aligned}$$

We have thus constructed an inductive alternative definition for *mwp*:

$$\begin{aligned}
 mwp &:: [Elem] \rightarrow ([Elem], Window) \\
 mwp [] &= ([], ([], [])) \\
 mwp (a : x) &= (wflatten u \upharpoonright_{\leq m}, u) \\
 &\quad \text{where } (m, w) = mwp\ x \\
 &\quad \quad u = wmaxchop (wcons\ a\ w).
 \end{aligned}$$

As a maximally dense segment is obtained from $fst \cdot mwp$, we have arrived at our algorithm, which is summarised in the following section.

5.2 Data refinement and performance analysis

In this section, we perform some final data structure refinement, to ensure that the final algorithm runs in linear time (and space). For the reader's convenience, Figure 8 summarises the derived programme, which includes an implementation of the function *hsplit*. In this code, the type *Window* is given more abstractly as

$$\text{type } Window = (Header\ Elem, Parts\ (Seg\ Elem)).$$

The types *Header*, *Parts*, and *Seg* are respectively the datatypes with which we implement the header, the partition, and each segment in the partition. These datatypes can be thought of as instances of a *List* class, with polymorphic operators such as $[]$, $(:)$, and $(+)$, and we use these operators in the code for clarity.

Lists of elements of type *Header*, *Parts*, or *Seg*, all need to have their density and width computable in constant time. This can easily be done by pairing the lists with their current values for weight, length, width, etc., and is considered a separate issue from the implementation details that follow.

The data structure used for the *Header* datatype depends on the implementation of the *hsplit* function, as within the algorithm, the header is only altered using *hsplit*. One possible implementation makes use of an auxiliary function *split*, which repeatedly removes the rightmost element of the header until it does not exceed the width limit, while retaining as much width as possible. Since each element is initially added to the header on the left and later removed from the right, at most once, *hsplit* runs in linear time if addition on the left and removal from the right are both amortised constant time operations. This can be done by implementing *Header* as a simple queue using two lists, e.g. see Okasaki (1999).

For the *Parts* and *Seg* datatypes, we need to examine the operations on the DRSP structure of the window. Elements are added one-by-one on the left by *addl*, which just launches *prepend*. The function *prepend* makes an indefinite number of recursive calls, and in each call two segments are joined. However, notice that segments in the DRSP are all non-empty, and are never split once joined. Therefore, given an input of length n , there can be at most $O(n)$ joins in total. The function *prepend* thus runs in linear time, provided that *Parts* allows addition to and removal from the left in amortised constant time, and that *Seg* supports list concatenation in (amortised) constant time. The requirement on *Seg* is easy to fulfill: one may simply use a join list: **data** *Seg* $a = Single\ a \mid Join\ (Seg\ a)\ (Seg\ a)$.

```
ms :: [Elem] → [Elem]
ms = fst · mwp
```

```
type Window = (Header Elem, Parts (Seg Elem))
```

```
mwp      :: [Elem] → ([Elem], Window)
mwp []   = ([], ([], []))
mwp (a : x) = (wflatten u |≤ m, u)
           where (m, w) = mwp x
                 u = wmaxchop (wcons a w)
```

```
hsplit :: Header Elem → (Header Elem, [Elem])
hsplit x = split (x, [])
```

```
split      :: (Header Elem, [Elem]) → (Header Elem, [Elem])
split ([], y) = ([], y)
split (x ++ [a], y) = if width x < L then (x ++ [a], y)
                   else split (x, a : y)
```

```
wcons      :: Elem → Window → Window
wcons a (h, xs) = (h', foldr addl xs x)
           where (h', x) = hsplit (a : h)
```

```
addl      :: Elem → Parts (Seg Elem) → Parts (Seg Elem)
addl a xs = prepend [a] xs
```

```
prepend    :: Seg Elem → Parts (Seg Elem) → Parts (Seg Elem)
prepend x [] = [x]
prepend x (y : xs) = if x ≤d y then prepend (x ++ y) xs
                   else x : y : xs
```

```
wmaxchop   :: Window → Window
wmaxchop (h, xs) = (h, maxchop h xs)
```

```
maxchop    :: Header Elem → Parts (Seg Elem) → Parts (Seg Elem)
maxchop h [] = []
maxchop h (xs ++ [x]) = if (h ++ concat xs) <d x then xs ++ [x]
                       else maxchop h xs
```

Fig. 8. The derived algorithm.

The removal of elements from the window is carried out by the *maxchop* function, which removes segments in *Parts* from the right. Therefore, we need *Parts* to support amortised constant-time removal from both ends. A number of data structures support such operations, for example, Banker's dequeues (Okasaki, 1999), or 2–3 finger trees (Hinze & Paterson, 2006).

We tested our Haskell implementation, compiled using the Glasgow Haskell Compiler version 7.10.1 and run on a Intel Xeon Quad-Core E5620 PC running OS X, at 2.40 GHz. Sample timing measurements are presented in Figure 9.




Fig. 9. Measurements of mean (over 20 lists) running time for the MDS implementation. Sizes were randomly chosen from the range $[1..20]$, weights from the range $[-100, 100]$, and $L = 400$.

Timing measurements were also carried out with different values for the lower bound L , and different ranges and probability distributions for the numerical input. In all cases, the resulting timings were very close to those in Figure 9 (too similar to distinguish, if plotted on the same diagram). Typically, the maximally dense segments produced are only a little wider than the width constraint, and in order to give the window data structure a more thorough test, we also created special cases designed to produce much wider solutions. Again, the timings were very similar to those illustrated above.

Thus, the empirical evidence is consistent throughout with our analysis of a linear-time efficiency for the algorithm.

6 Conclusions

We have derived a linear-time algorithm for solving the generalised segment density problem with a lower bound on segment width. The algorithm scans through the input list using a sliding window, which is split into a header and a partition of right-skew segments with decreasing densities (DRSP), whose properties we exploit to make the linear-time processing possible. While the programme itself barely occupies one page, its proof is anything but simple, involving the discovery of intricate properties, and our complete proof of the algorithm uses approximately 3,500 lines of Agda code.

Two instances of the density segment problem were presented. While the MDS problem has a long history, we believe that the MMS problem is new. These are similar but neither is a generalisation of the other: in MDS, the density function is fundamentally linked to the segment width, but this is not the case for the MMS problem. This shows that, for the densest segment problem, width can be separate from density, allowing the algorithm to analyse data such as blood glucose measurements.

The development of our solution has not been straightforward, and we faced a number of design decisions from the beginning. Should we model the problem functionally, relationally, or use total relations? As there were technical complications involving relations and zygomorphisms (see the technical aside in Section 3.4), a functional approach was strongly indicated. Another decision concerned the maximum operators \uparrow or \downarrow , when deciding which segment to return in case of a tie: which to use, and did it matter? It turns out that the choice does matter – choosing \downarrow would have meant that the algorithm would fail to meet that specification. Furthermore, the algorithm might not yield a result when there is no segment within the width bound, but using a *Maybe* type made the reasoning rather cumbersome. Using an extended ordering \leq allows a much cleaner formulation.

While we had an informal understanding of why this algorithm is correct, formally writing down the properties that make it so turned out to be surprisingly tricky. The correctness of the “outer” algorithm (Section 3.6), treated casually in all previous works, took us a considerable amount of time to formalise. We attempted to come up with a more declarative specification of the prefix returned by *wp* (for example, we guessed that it is the shortest prefix satisfying certain properties), but none of those specifications were correct. Various possibilities were tried before we reached Lemma 3.7 and its supporting definitions. Afterwards, the proof quickly followed. It is often the case that finding the right thing to prove is harder than producing the proof. It was surprising that in Lemma 3.7, we could prove an equality, rather than merely that the two sides yield segments having the same density. In fact, we needed an equality for the inductive proof to work.

All this hard work was not spent in vain. We re-proved the uniqueness of the DRSP, and presented various properties about it, including how it can be constructed. This fills in more properties about the DRSP than given in previous papers.

We initially set out to solve a more general version of the MDS problem, with an additional upper bound on segment width, making the problem even more intricate. As mentioned in Section 1, an algorithm with a wrong time analysis has been published before. Even the algorithms of Chung & Lu (2004) and Goldwasser *et al.* (2005) are not entirely correct: they both fail for a boundary case, specifically when there is no segment in the window whose width is within the bounds. The former algorithm could potentially return an invalid result, for which there is an easy fix (Chung, 2010), while the latter loops and it is harder to see whether it is fixable.

We have developed data structures that are used in the sliding window to produce densest segments with an upper bound on their width, and we have some preliminary results on correctness proofs, which will have to be deferred to another paper.

We believe that the difficulty in developing correct linear-time algorithms is partly due to the complicated nature of the MDS problem, and partly due to the absence of a rigorous approach to programme construction. The imperative algorithms of Chung & Lu and Goldwasser both maintain invariants that are neither explicitly stated nor easy to reconstruct. The invariants rely on states stored in static variables surviving between subroutine calls, which makes reasoning about them extremely hard. In addition, their liberal use of array indices obscures some beautiful structural

properties of segment densities. The MDS problem provides a useful case study of how formal development techniques can help with constructing correct programmes.

Acknowledgments

We would like to thank Richard Bird for stimulating discussions of the MDS problem. Many thanks are also due to the anonymous referees, whose detailed comments helped to improve our presentation.

References

- Bird, R. S. (1987) An introduction to the theory of lists. In *Logic of Programming and Calculi of Discrete Design*, Broy, M. (ed), NATO ASI Series F, vol. 36. Springer-Verlag, pp. 342.
- Bird, R. S. & de Moor, O. (1997) *The Algebra of Programming*. Prentice Hall.
- Chung, K.-M. (2010) Personal communication.
- Chung, K.-M. & Lu, H.-I. (2003) An optimal algorithm for the maximum-density segment problem. In *Annual European Symposium on Algorithms*, Lecture Notes in Computer Science, vol. 2832, pp. 136–147, Springer-Verlag.
- Chung, K.-M. & Lu, H.-I. (2004) An optimal algorithm for the maximum-density segment problem. *SIAM J. Comput.* **34**(2), 373–387.
- Curtis, S. & Mu, S.-C. (2014) *Calculating a linear-time solution to the densest segment problem (paper and supplement materials)*. Available at: <http://www.iis.sinica.edu.tw/~scm/2014/mds/>. Last accessed: November 23rd, 2015.
- Goldwasser, M. H., Kao, M.-Y. & Lu, H.-I. (2002) Fast algorithms for finding maximum-density segments of a sequence with applications to bioinformatics. In *Proceedings of the 2nd Workshop on Algorithms in Bioinformatics (WABI 2002)*, Springer, Berlin, pp. 157–171.
- Goldwasser, M. H., Kao, M.-Y. & Lu, H.-I. (2005) Linear-time algorithms for computing maximum-density sequence segments with bioinformatics applications. *J. Comput. Syst. Sci.* **70**(2), 128–144.
- Han, L. & Zhao, Z. (2009) CpG islands or CpG clusters: How to identify functional gc-rich regions in a genome? *BMC Bioinformatics* **10**(65), doi: 10.1186/1471-2105-10-65.
- Hinze, R. & Paterson, R. (2006) Finger trees: A simple general-purpose data structure. *J. Funct. Program.* **16**(2), 197–217.
- Huang, X. (1994) An algorithm for identifying regions of a DNA sequence that satisfy a content requirement. *Comput. Appl. Biosci.* **10**(3), 219–225.
- Jeurig, J. T. (1993) *Theories for Algorithm Calculation*. Ph.D. thesis, Utrecht University.
- Jeurig, J. T. & Meertens, L. (1989) The least-effort cabinet formation. *Squiggolist* **1**(2), 12–16.
- Kaldewaij, A. (1990) *Programming: The Derivation of Algorithms*. Prentice Hall.
- Lin, Y.-L., Jiang, T. & Chao, K.-M. (2002) Efficient algorithms for locating the length-constrained heaviest segments, with applications to biomolecular sequence analysis. In *Proceedings of the 27th International Symposium on Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science, Springer-Verlag, Berlin, vol. 2420, pp. 459–470.
- Malcolm, G. R. (1990) *Algebraic Data Types and Program Transformation*. Ph.D. thesis, Groningen University, the Netherlands.
- Mu, S.-C., Ko, H.-S. & Jansson, P. (2009) Algebra of programming in Agda: Dependent types for relational program derivation. *J. Funct. Program.* **19**(5), 545–579.

- Norell, U. (2007) *Towards a Practical Programming Language Based on Dependent Type Theory*. Ph.D. thesis, Chalmers University of Technology.
- Okasaki, C. (1999) *Purely Functional Data Structures*. Cambridge University Press.
- Rem, M. (1988) Small programming exercises 20. *Sci. Comput. Program.* **10**(1), 99–105.
- Swierstra, S. D. & de Moor, O. (1993) Virtual data structures. In *IFIP TC2/WG2.1 State-of-the-Art Report on Formal Program Development*, Bernhard M., Helmut P. & Steve S. (eds), Lecture Notes in Computer Science, pp. 355–371, Springer.
- Van Den Eijnde, J. P. H. W. (1990) Left-bottom and right-top segments. *Sci. Comput. Program.* **15**(1), 79–94.
- Zantema, H. (1992) Longest segment problems. *Sci. Comput. Program.* **18**(1), 39–66.