

Hierarchical parcel-swapping representation of turbulent mixing. Part 2. Application to channel flow

Alan R. Kerstein[†]

72 Lomitas Road, Danville, CA 94526, USA

(Received 7 January 2014; revised 7 April 2014; accepted 15 May 2014;
first published online 10 June 2014)

A novel concept for simulation of turbulent mixing, termed hierarchical parcel swapping (HiPS), was recently proposed. The method involves either a parameterized representation of the turbulent flow or a more self-contained flow simulation. As a step toward turbulent mixing applications, the latter formulation is used for the first numerical demonstration of model performance. Owing to its suitability for this purpose and its role as a canonical benchmark, channel flow is the target application. Despite its idealized representation of this flow, HiPS is shown to capture salient features of the flow with a notable degree of quantitative accuracy. The implications of this finding with regard to flow physics and with regard to the applicability of HiPS to other problems are discussed.

Key words: turbulence modelling, turbulent boundary layers, turbulent mixing

1. Introduction

Interest in economical yet accurate model representations of turbulence phenomenology is continually growing due the need to predict increasingly challenging turbulent flow and mixing processes. This motivated the recent introduction (Kerstein 2013) of a novel numerical modelling concept termed hierarchical parcel-swapping (HiPS). That paper provides a generic description of the method and its physical interpretation. Here, the first numerical application of HiPS is reported. The target case is fully developed turbulent channel flow, a widely used test case due to its geometrical simplicity and rich phenomenology.

As in many turbulent mixing models, HiPS models turbulent flow as a collection of fluid parcels, with mixing implemented as occurrences of exchange or redistribution of the contents of pairs or groups of parcels that have been identified as adjacent. Within this framework, turbulent advection has two main effects. First, the overall rate of mixing occurrences is based on a controlling advective time scale, as in most if not all parcel-based formulations (Fox 2003). Second, turbulence influences the time-evolution of the partitioning of parcels into adjacent pairs or groups. Prior to the formulation of HiPS, there appeared not to have been any systematic effort to incorporate turbulence phenomenology into the specification of this evolution. Instead,

[†] Email address for correspondence: alan.kerstein@gmail.com

the identification of adjacent parcels was based on the closeness of the set of scalar or species compositions of parcel pairs or groups, such as in Subramaniam & Pope (1998) and Pope (2013), based on metrics defined for that purpose. Although this approach is superior to the earlier and still common practice of ignoring the degree of closeness, and though it appropriately reflects the fact that spatially proximate parcels are close compositionally, it does not draw upon known turbulence phenomenology to further improve model performance.

HiPS is a conceptually and computationally economical method for time-advancing the grouping of parcels in a manner that reflects the stepwise nature of length-scale breakdown in the inertial-range turbulent cascade. It has long been recognized (Meneveau & Sreenivasan 1991; Aurell, Dormy & Frick 1997) that a hierarchical or tree geometry provides an especially suitable representation of this breakdown, with the scale similarity of the breakdown represented by a multiplicative stride between successive tree levels (Sreenivasan & Stolovitzky 1995). The novel element of HiPS is the formulation of system evolution within a hierarchical geometry in a manner that suitably time-advances the grouping of parcels.

In doing so, HiPS offers the additional possibility of associating flow variables (velocity components) with the parcels and redistributing momentum among parcels to represent viscous transport. This can make the model more self-contained, because the flow information can be used to specify time scales governing system evolution that otherwise need to be preassigned. (Such preassignment is effectively a turbulence parameterization.) More generally, HiPS becomes potentially a predictive model of turbulent flow as well as a mixing model.

Another unified framework for modelling turbulent flow and mixing is the transported velocity-composition joint probability distribution function (p.d.f.) method (Pope 1985). However, this is a coarse-grained approach that requires subgrid mixing closure. For that purpose, the type of mixing model described by Fox (2003) and the other cited references is typically used. In this context, unified flow and mixing treatment within HiPS can be viewed as a possible subgrid-scale complement to the coarse-grained unification within the transported joint p.d.f. method for the purpose of improving the subgrid mixing closure. However, consideration of any subgrid-scale use of HiPS must be preceded by demonstration of the performance of HiPS as a self-contained model.

Accordingly, an appropriate first step is to evaluate the turbulence prediction capability of HiPS. For this purpose, a wall-bounded case is chosen due to the importance and difficulty of modelling turbulent boundary layers.

Such an evaluation inevitably confronts the issue of the level of modelling that is required to obtain prediction rather than description. With reference to turbulent boundary layers, the challenge of prediction led Fernholz & Finley (1996) to conclude that limited theoretical understanding necessitates a primarily descriptive focus. It is shown here that HiPS modelling of turbulent channel flow sheds some light on the degree to which a simple model of turbulent near-wall flow can be predictive.

The HiPS approach and its relationship to other approaches, which were explained in detail in Kerstein (2013), are summarized in § 2. Specialization to channel flow is described in § 3. The physical interpretation of this formulation and consequent parameter-setting and data-reduction procedures are discussed in §§ 4–7. Results are presented and interpreted in § 8.

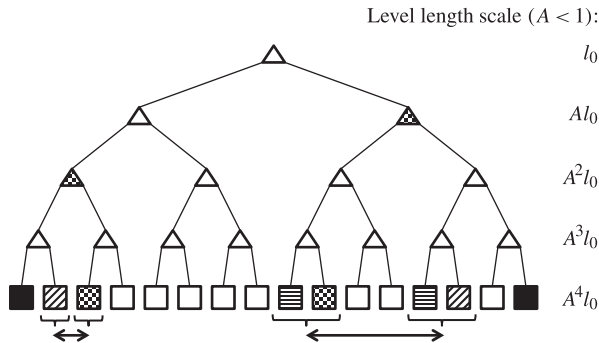


FIGURE 1. A five-level ($N = 4$) binary tree and representative permutations. The squares represent parcels at the base of the tree (level 4, based on top-down level numbering starting from 0); the triangles represent nodes at other levels, and the line segments are node connections. In the model specialization to channel flow, black squares represent wall parcels. Allowed permutations that can move two representative parcels (filled with a checkerboard pattern) closer to a wall parcel based on the proximity definition of parcel separation are marked by brackets and double-arrow lines that connect the sets of parcels, designated to be swapped. Parcels that change places with the checkerboarded parcels are shown as diagonally striped; other displaced parcels are horizontally striped. The node at the apex of each permutation is checkerboarded. The illustrated swaps are the only allowed swaps (i.e. swaps that do not displace a wall parcel) that can bring a checkerboarded parcel into closer proximity to the nearest wall parcel, where proximity is based on the size of the smallest subtree containing both parcels. Additionally, there are two allowed swaps (not shown) that move the leftmost checkerboarded parcel farther from the left wall and one (not shown) that moves the rightmost checkerboarded parcel farther from the right wall, as well as five allowed swaps (not shown) that do not affect the wall proximity of either checkerboarded parcel but change the adjacency pairing of one of them, inducing mixing that can change the state of a checkerboarded parcel.

2. The HiPS representation of turbulent flow

2.1. Geometrical structure

HiPS is briefly summarized here to provide context for the description of its channel-flow specialization in § 3. HiPS consists of three elements: its geometry, the mechanisms of time-advancement and the physical modelling that specifies the rates at which the mechanisms are implemented.

The purpose of the geometrical structure is to specify the proximity of parcel pairs and to provide a mathematical framework in which the time-advancement mechanisms and their rates can be prescribed. In this sense, the geometrical structure is an auxiliary construct. The physical state consists of the parcel properties and their spatial locations as inferred from parcel proximities.

For these purposes, the parcels are assumed to reside at the base of a hierarchical (tree) structure, as shown in figure 1 for the special case of a binary tree. Only binary trees are considered here. For a more general treatment, see Kerstein (2013).

Auxiliary length scales and time variables are associated with nodes of the tree other than the base nodes at which parcels reside. The length scales are fixed in time and are identical within each tree level, varying only from level to level. The time variables can be fixed through parameterization or can be functions of time-varying

parcel properties. Flow prediction requires the latter, which is therefore assumed in what follows.

The tree levels are labelled sequentially from 0 at the apex (top level) to N at the base (bottom level) of the tree. The apex length scale is denoted by l_0 , and the length scale associated with any level- n node is $l_n \equiv A^n l_0$ for $n \in [0, N]$. The role of the scale-reduction factor A will be discussed shortly.

The node length scales serve two purposes. In conjunction with parcel properties, they determine the time variables that control the time-advancement as described in §2.2. Additionally, they specify the parcel size.

The proximity of two parcels, called their ultrametric distance in other contexts (Benzi, Biferale & Trovatore 1997), is defined here as the level index of the smallest subtree that contains both parcels. Adjacent parcels are those with proximity $N - 1$, which corresponds to the smallest length scale above the base of the tree. The length scale l_{N-1} is taken to be the parcel size, S . Additional assumptions discussed in §4 and appendix C are needed to specify fully the spatial arrangement of parcels within the channel.

Here A and N are key structural parameters of the model, as explained by Kerstein (2013) and elaborated here. They control the scale range of the tree; and, additionally, A , which is positive and less than 1, determines the relationship between the number of parcels, which is 2^N for a binary tree, and the parcel size $S = l_{N-1}$. For fixed A , N controls the relationship between parcel size and parcel number. Viewing the decreasing parcel size with increasing N as a refinement of spatial resolution, an effective dimension d of the space covered by parcels is identified by taking the parcel number to be proportional to the parcel size raised to the power $-d$, hence $2^N \sim l_{N-1}^{-d} = l_0^d A^{(1-N)d}$. The value of A that satisfies this proportionality relation is $2^{-1/d}$. Thus, the effective dimension is determined by A . This has diverse implications (Kerstein 2013); notably, small d can be chosen to reduce computational cost, subject to performance considerations. For channel flow, the A -dependence of cost and performance is examined in §8.

2.2. Time-advancement mechanisms

HiPS time-advancement involves a mechanism for successive regroupings of parcels into adjacent pairs, based on the definition of adjacency in §2.1, and a mechanism for mixing the contents of adjacent parcels. In order to focus on the former, which is the novel feature of HiPS, a minimal representation of mixing is adopted here. At the instant that two parcels become adjacent, their contents are fully mixed so that their states become identical. There is no further mixing involving either parcel until one of them becomes adjacent to a different partner. Implications of this and other possible mixing formulations are discussed by Kerstein (2013). Channel-flow implications are discussed in §3.

The tree structure allows a simplified representation of turbulent advection effects over a large range of scales in terms of a relatively small number of possible rearrangements of parcel placements at the base of the tree. Owing to the discreteness of the tree structure (versus the continuity of physical space), these rearrangements are instantaneous events. This is consistent with the ultimate purpose of HiPS, which is to modify repeatedly the grouping of parcels into adjacent pairs.

An advection event involves a permutation of parcel placements within the tree structure, as follows. The event is associated with a chosen tree node, called the apex of the permutation, which can be at any level $n \leq N - 2$. Any such node

connects downward to the apexes of two level- $(n + 1)$ subtrees, and each of those apexes connects downward to the apexes of two level- $(n + 2)$ subtrees. Within each level- $(n + 1)$ subtree, one of the two level- $(n + 2)$ subtrees is chosen. The two chosen subtrees are swapped, meaning that the parcels within each subtree are moved to the other subtree.

In figure 1, two advection events are illustrated. The tree nodes associated with the respective events are the triangles filled with a checkerboard pattern. Each event, denoted by a double-arrow line segment, swaps two sets of parcels. Each set is marked by a bracket, and consists of the parcels contained in a subtree two levels below the apex of the permutation.

The swapping of subtrees two levels below the apex of the permutation, rather than at some other choice of level, is based on the following considerations. A swap of the two subtrees that are one level below the apex does not change the proximity of any pair of parcels, so based on the model representation of spatial structure, such a swap does not change the state of the system and therefore is not a time-advancement mechanism. A swap of two subtrees that are three or more levels below the apex is a valid analogue of advection, but the locality of scale-space interactions in the turbulent cascade, a fundamental principle with strong conceptual and empirical support (Frisch 1995), is best represented by minimizing the difference in scale between the size and separation of the selected subtrees. As defined in HiPS, advection events transparently emulate salient features of cascade phenomenology (Kerstein 2013).

This defines the generic time-advancement mechanisms within HiPS. Supplementary mechanisms representing wall effects and pressure effects specific to the present formulation are described in § 3 and appendix A. To complete the formulation of the generic model, the rules governing the occurrences of advection events need to be specified. The statistical sampling of events used here for HiPS channel-flow simulation is described in §§ 3 and 7 and appendix A. The formulation is analogous in many respects to one-dimensional turbulence (ODT). Appendix A introduces the relevant ODT concepts.

Kerstein (2013) discusses various implications of the HiPS time-advancement mechanisms. Two key points are restated here. One is that only level- $(N - 2)$ permutations, which swap level- N subtrees (i.e. individual parcels), affect the grouping of parcels into adjacent pairs. The other is that larger-scale (smaller- n) permutations, though having no such effect, are nevertheless consequential because they break down the length scales of parcel-state variation, as in the turbulent cascade. This process promotes the eventual mixing-induced fluxes between adjacent parcels, which must have dissimilar property values in order for mixing to induce net property transfers. Event likelihood is dependent on the local instantaneous flow state in a manner that tends to induce more frequent occurrences of smaller-scale permutations, again reflecting turbulent cascade phenomenology. Additional details are provided in appendix A and Kerstein (2013).

HiPS has no physically meaningful laminar-flow limit, because the restrictive definition of parcel adjacency in HiPS prevents molecular transport of momentum and other parcel properties except between the parcels that make up an adjacent pair. The advection events representing turbulent flow induce the ongoing parcel regrouping that enables mixing effects to be communicated throughout the flow domain. Thus, HiPS is generically a model of advectively controlled transport and mixing regimes. Where molecular transport dominates, e.g. in the viscous sublayer of near-wall flow, case-specific modification is needed, such as that introduced for channel flow in § 3.

3. Specialization to channel flow

Due to the present focus on modelling flow rather than scalar mixing, the only parcel properties that are time-advanced are the three velocity components (u, v, w) corresponding to the streamwise, wall-normal and spanwise directions (x, y, z), respectively, in a channel flow. Here $A = 0.5$, corresponding to $d = 1$, is assumed initially, allowing the physical interpretation (see §4) of parcels being aligned along a one-dimensional wall-normal line of sight spanning the channel. (The model formulation and interpretation for other A values are explained in appendix C.)

Channel flow is simulated by introducing two case-specific features: an imposed streamwise pressure gradient dp/dx and a representation of wall effects. Two parcels that have proximity 0 relative to each other are chosen to be wall parcels. Figure 1 shows these parcels located at the endpoints of the linear parcel array. Their placement at array endpoints is strictly a physical interpretation. The model formulation recognizes no geometrical relationships other than proximity as governed by the tree structure.

Wall parcels have zero nominal size. All parcels have the same mass, corresponding to a fixed density ρ in non-wall parcels.

The pressure forcing introduces a third time-advancement mechanism, namely time-stepping of the u variable of each non-wall parcel according to

$$\frac{du}{dt} = -\frac{1}{\rho} \frac{dp}{dx}, \quad (3.1)$$

where dp/dx is constant in space and time and is assumed to be negative so that the mean flow is in the positive x direction. This is the only continuous-in-time (though numerically discretized) advancement process in the channel-flow simulation.

Wall parcels are not deemed to be part of the flow, so advection events that would displace a wall parcel are forbidden. A wall parcel is time-advanced by mixing with its adjacent neighbour immediately after advection places a new parcel at the adjacent location, and by an additional mechanism that is introduced to enforce the no-slip boundary condition, i.e. that immediately after mixing changes the state of a wall parcel, its state is reset to $u = v = w = 0$.

As explained in §2.2, mixing of adjacent parcels is complete, i.e. the resulting parcel states are identical. Because each wall parcel has the same nominal mass as other parcels, mixing with the adjacent parcel halves the velocity values of that parcel.

The enforcement of the boundary condition after mixing causes the state of a wall parcel to be dissimilar from the state of its adjacent neighbour. This is the only situation in which dissimilarity of adjacent parcels is allowed to persist.

Operationally, a HiPS simulation consists of time-stepped advancement of (3.1), punctuated by occurrences of advection events. If an event modifies parcel pairings, then the newly adjacent pairs are mixed. If mixing changes the state of a wall parcel, then boundary-condition enforcement is implemented immediately afterwards as described. The event occurrence times and the subtrees swapped during an event are randomly sampled in a manner that embodies the HiPS representation of turbulence phenomenology.

When an advection event occurs, not only are the parcels in the selected subtrees swapped as described but, additionally, the u, v and w values of these parcels are modified in a manner that models the fluctuating part of the pressure-gradient term of the momentum equation. Unlike swapping and mixing, which affect all parcel properties, this operation is applied only to velocity components. (In the present

case there are no other time-advanced properties.) It does not induce dissimilarity of adjacent parcels because its effect on parcels within each swapped subtree is uniform. This and other aspects of advection in HiPS are explained further in appendix A.

Without this pressure-fluctuation treatment, there would be no mechanism generating non-zero values of v and w . These components could be omitted, thereby limiting the model to the prediction of u statistics; but they have been included here in order to broaden both the present investigation and the scope of possible future applications.

4. Physical-space interpretation

In § 2.1 it is assumed that the parcel size S is $l_{N-1} = A^{N-1}l_0$ (except that $S = 0$ for wall parcels). Specializing to $A = 0.5$, the parcels are assumed to form a line in the wall-normal direction that spans the channel without gaps or overlaps. On this basis, S times the number $2^N - 2$ of non-wall parcels is equal to the channel height $2h$, where h is the half-height. This gives $h = (2^{N-1} - 1)S = (1 - 2^{1-N})l_0$. This relationship between h and l_0 is used in § 6.2 to determine the Reynolds numbers corresponding to a given HiPS configuration.

Apart from the small effect of the treatment of wall parcels, this physical-space interpretation implies that l_0 corresponds to the separation of the midpoints of the level-1 subtrees on either side of the midpoint of the wall-normal line. Analogously, the parcel size l_{N-1} is the separation of the midpoints of adjacent parcels, subject to the noted small correction. This indicates almost exact consistency of the physical meaning of l_n across scales. This consistency is important for physically correct rendering of the turbulence phenomenology used in appendix A to formulate the event-sampling scheme.

These assumptions and their consequences might seem straightforward, but they are valid only for $A = 0.5$. Generalization to other A values, explained in § 8.2 and appendix C, is less transparent.

In the HiPS channel-flow formulation, the distinct treatment of wall parcels with regard to both advection and boundary-condition enforcement is the sole basis for flow inhomogeneity, analogous to the origin of spatial inhomogeneity in physical channel flow. In the absence of a distinct treatment of particular parcels, proximity defines spatial structure in HiPS only in a relative sense. The distinct treatment of wall parcels enables a degree of representation of channel-flow physical-space structure.

Figure 2 illustrates HiPS channel-flow structure, which is evident also in instantaneous flow snapshots but is more readily discernible in this mean-flow illustration. Because the wall treatment is the sole mechanism for deviation from homogeneity, statistical properties of a parcel depend only on its proximity to each of the walls. In other words, parcels whose wall proximities are the same have indistinguishable statistics.

Thus, the successive intervals of uniform (apart from slight fluctuations due to the finite run-time; see § 8.3) parcel mean-velocity values reflect the absence of any basis for systematic parcel-to-parcel behavioural differences other than their proximities to the wall parcels. In this sense, each interval contains a set of parcels that are equivalent samples of the flow state anywhere within the interval. Therefore such an interval may be viewed as a coarse-grained region containing parcels that constitute a representative sample of fluid states within the region. Reflecting the stated physical-space interpretation, the parcels are located at uniform y increments in physical space in figure 2, but within homogeneous intervals the parcel placements are immaterial. The figure shows half of the y -coordinate range $[0, 2h]$.




FIGURE 2. For case *N8* of table 1, parcel mean streamwise velocities in the y interval $[0, h]$, with uniform parcel spacing as described in the text.

In § 8, results are presented by combining the parcel statistics within each homogeneous interval and locating the result at the y value corresponding to the midpoint of the interval. Determination of the midpoints is based on the assumed uniform spatial distribution of non-wall parcels filling the size- $2h$ domain. In particular, the physical-space picture implies that the midpoint of a parcel adjacent to a wall is at a distance $S/2$ from the wall. In §§ 5 and 8, wall shear is evaluated as the velocity of that parcel divided by $S/2$. In addition to identifying the locations of interval midpoints, the physical-space interpretation of the HiPS channel-flow formulation has other roles in data gathering and interpretation that are discussed in appendices B and C.

This data-reduction approach is consistent with the interpretation of each homogeneous interval as a single control volume (using conventional terminology, though control interval would be literally more accurate here) within a non-uniformly coarse-grained finite-volume scheme. From this perspective, the artificiality of the piecewise-constant structure of the mean flow profile becomes immaterial. As noted above, the parcels within a homogeneous interval are now statistical samples of the state within the interval, and there is no need to interpret them in terms of spatial structure within the interval. Importantly, the hierarchical structure within the interval remains operative during the time-advancement so that the HiPS turbulent-cascade representation governs the flow evolution. (The intervals are homogeneous statistically but not instantaneously.)

Viewed as an effectively coarse-grained approach, the HiPS channel-flow formulation is less unconventional than figure 2 suggests. Additionally, it has some beneficial features from this viewpoint.

For example, the geometrical increase of interval size corresponds to uniform increase of the logarithm of size. As suggested by figure 2 and further illustrated in § 8, this corresponds to roughly uniform increments of \bar{u}^+ in the log-law subrange. Thus, near-optimal spatial partitioning is obtained automatically.

Irrespective of coarse-grained data reduction, there is an effectively constant number density of parcels along the y coordinate, so the coarse-graining has no direct cost benefit. However, this conclusion is specific to $A = 0.5$. For $A < 0.5$, figure 10 shows that the corresponding effective dimensionality $d < 1$ results in a cost-saving decrease in the parcel number density with increasing y . Conceptual and performance concerns limit the degree to which this benefit can be exploited, but results in § 8.2 suggest some scope for reducing A without excessive sacrifice in performance. An additional consideration is that mixing applications of the HiPS channel-flow formulation might have minimum parcel-number-density requirements for adequate statistical sampling of mixture states.

Thus it is seen that the ostensibly unphysical behaviour seen in figure 2 is both reasonable and beneficial from a coarse-graining perspective. This perspective likewise indicates a benefit of the unphysically restrictive definition of parcel adjacency in HiPS, namely that it prevents mixing of parcel pairs which are adjacent in the spatial rendering in figure 2 but have highly dissimilar states because they are in different homogeneous intervals. This prevents spurious mixing of the sort that degrades the performance of conventional coarse-grained models. In those models, the coarse-graining is inherent in the formulation rather than in a data-aggregation strategy, so they typically require subgrid modelling in order to avoid an unacceptable degree of spurious mixing. As noted in § 1, HiPS is designed to improve the performance of subgrid models as well as to provide novel predictive capabilities as a self-contained model.

With regard to the distinction between approaches to coarse-graining, it is shown in § 8.1 that it is advantageous to set the parcel size S to a larger value than the near-wall resolution scale; so, as presently implemented, HiPS is itself inherently coarse-grained to some extent, irrespective of any coarse-graining introduced through the aggregate data-reduction treatment. The detailed evaluations in § 8 show that the model formulation in many respects performs as would be expected from a more costly fully resolved treatment.

5. Channel-flow phenomenology in HiPS

The mean momentum balance for fully developed turbulent channel flow implies that

$$u_\tau^2 \equiv \nu \left. \frac{d\bar{u}}{dy} \right|_w = -\frac{h}{\rho} \frac{dp}{dx}, \quad (5.1)$$

where u_τ is the friction velocity, \bar{u} is the mean velocity, ν is the kinematic viscosity and w denotes wall value. As indicated by (5.1), $(d\bar{u}/dy)|_w$ adjusts so that momentum is fluxed to the walls at the rate that it is supplied by the work done by the imposed pressure gradient in the statistically steady state. The HiPS analogue of Reynolds shear stress and the associated form of the mean momentum balance in HiPS are discussed in § 8 and appendices B and C.

The HiPS mixing mechanism and the wall treatment described in § 3 do not involve gradient-based momentum fluxing governed by the kinematic viscosity, but the viscosity implied by this formulation can be inferred from (5.1) by evaluating $(d\bar{u}/dy)|_w$ from the computed mean velocity profile. In the HiPS channel-flow formulation, this quantity is the \bar{u} value of a parcel adjacent to a wall divided by the distance $S/2$ of its midpoint from the wall. Then the implied viscosity is

$$\nu = -\frac{h}{\rho} \frac{dp/dx}{(d\bar{u}/dy)|_w}. \quad (5.2)$$

This is used to evaluate the friction and bulk Reynolds numbers in § 6.

A self-consistent estimate of the flow structure resulting from the wall treatment follows from the mixing-length picture of the mean structure of statistically steady near-wall flow. The eddy-viscosity representation of the Reynolds shear stress can be expressed as $-v_e(y)(d\bar{u}/dy)$. Substitution of the mixing-length ansatz $v_e(y) \sim y^2(d\bar{u}/dy)$ for the eddy viscosity gives proportionality to $y^2(d\bar{u}/dy)^2$. In the sublayer where the Reynolds shear stress is roughly constant because advective transport dominates viscous-transport and mean-pressure effects, y -independence of the mixing-length result implies $\bar{u}(y) \sim \log y$, showing how the log law arises from elementary considerations.

The physical basis of the ansatz is the notion that turbulent transport at a distance y from the wall is governed by size- y eddies. The mean-square fluid displacement by such an eddy in the y direction scales as y^2 , explaining the appearance of this term in the $v_e(y)$ ansatz. Invoking the random-walk interpretation of turbulent diffusion, $v_e(y)$ is estimated by multiplying this term by the frequency of such eddies, governed by the inverse time scale $d\bar{u}(y)/dy$.

The turbulent transport by size- y eddies generically dominates the turbulent transport by smaller eddies, justifying neglect of the latter, but the justification of the omission of contributions by larger eddies is configuration-specific. A location y contained in an eddy of size $s \gg y$ must be near the edge of that eddy, because the edge can extend at most to the wall, so location y is at most a distance $y \ll s$ from the edge. Putting aside spatial inhomogeneity effects, it is much less likely for a given point to be near the edge of an eddy than to be at an arbitrary location within the eddy. In the $v_e(y)$ scaling estimate, this indicates that a relatively small frequency factor tends to counteract the large mean-square-displacement factor associated with eddies of size $s \gg y$. Thus it is plausible that the contribution of eddy sizes $s \gg y$ to $v_e(y)$ is negligible. More detailed reasoning, as well as empirical evidence, supports the validity of this hypothesis (Schlichting & Gersten 2000).

In HiPS, if displacement of wall parcels were permitted, then there would be no wall-proximity-dependence of the frequency of advection events that contain a given parcel, and hence no basis for log-law mean-velocity scaling near the wall. The prohibition of wall-parcel displacement introduces such a dependence, resulting in logarithmic mean-velocity scaling near the wall, as illustrated in § 8. To quantify the dependence, consider a parcel whose proximity to a wall parcel is n . If an advection event whose apex is at level $m < n - 1$ displaces this parcel, then it displaces all parcels in the level- n subtree containing the parcel; thus it also displaces the wall parcel, and is therefore prohibited. This shows that prohibition of wall-parcel displacement disallows displacements of any given parcel by advection events that affect parcels much farther from the wall than the given parcel. Therefore the largest and thus dominant advection events affecting such a parcel are of the order of the length scale associated with the proximity of the parcel to the wall parcel. On this basis, the prohibition of wall-parcel displacement satisfies one of the prerequisites for logarithmic mean-flow behaviour in HiPS channel-flow simulation. Logarithmic scaling also requires the inverse time scale governing advective transport at a given y to be of order $d\bar{u}(y)/dy$. The approach used to determine the frequencies of advection events in HiPS, described in appendix A, satisfies this requirement.

The immobility of the wall parcels is the HiPS analogue of wall blocking. HiPS transparently illustrates that wall blocking of turbulent motions, which is introduced in the present context by enforcing a simple constraint, is the principal mechanism underlying fundamental features of channel flow such as those discussed in § 8. Figure 1 illustrates the limitation of advective parcel displacements resulting from the wall-blocking constraint.

6. Inputs and parameter assignments

6.1. Parameter dependences of HiPS simulations

The dimensional quantities specifying a particular channel flow are h , $(1/\rho)(dp/dx)$ and ν . Normalized flow properties depend only on the friction Reynolds number $Re_\tau = hu_\tau/\nu$. Evaluation of u_τ and ν using (5.1) and (5.2), respectively, gives $Re_\tau = (d\bar{u}/dy)|_w[-(1/(h\rho))(dp/dx)]^{-1/2}$, showing that the needed physical input is a time scale $T \equiv [-(1/(h\rho))(dp/dx)]^{-1/2}$. Owing to the method by which ν is inferred in HiPS, the model output $(d\bar{u}/dy)|_w$ is also needed to determine Re_τ , so Re_τ itself is a model output rather than an input, contrary to the physical situation.

The rate distribution (A 2) that governs the frequencies of the possible advection events involves an overall multiplicative factor C that scales event frequencies and hence the rates of viscous as well as advective processes, due to the association of viscous transport with advection events in HiPS. Therefore C scales the rates of all model processes except the pressure forcing, whose associated rate of momentum input is governed by the quantity $1/T$. From this perspective, CT is the governing parameter group because changes of C and T that leave CT unchanged are equivalent to a rescaling of time without changing the sequence of flow states that occur. The flow states are unaffected because rescaling of time involves the speeding or slowing of time-advancement processes but not a rescaling of parcel velocity values. Velocity rescaling would be a necessary consequence of time rescaling of the physical equations of motion, but in HiPS the parameter C absorbs the rescaling, so the advection rate can change without changing velocities.

In the HiPS treatment of viscous transport, rescaling of time alters the relationship between velocity gradients and mixing-induced (interpreted as viscous) momentum transfer to the wall, so it affects the viscosity inferred from (5.2) through a rescaling that is derived in § 6.2. This and a rescaling of u_τ are then shown to imply a rescaling of Re_τ , indicating that the time rescaling converts a simulation corresponding to a given Re_τ into a simulation corresponding to a different Re_τ .

In addition to these parameters, the structural parameters of HiPS have important effects on computed results. In appendix C it is shown that the choice of A has a profound effect on the physical-space interpretation of HiPS. Additionally, the scale $l_{N-1} = A^{N-1}l_0$ of the smallest possible advective displacements (which is the HiPS analogue of the viscous cutoff scale) depends on both A and N .

The HiPS channel formulation provides only a minimal representation of the viscous sublayer near the wall. This is because the restrictive definition of parcel adjacency allows only advection events to bring parcel properties into closer proximity to the wall, except for property mixing between a wall parcel and its adjacent partner. In this sense, the HiPS viscous sublayer extends only to the parcel adjacent to a wall. Accordingly, it would be appropriate to locate that parcel at the value of $y^+ = yu_\tau/\nu$ corresponding to the spatial extent of the viscous sublayer, which terminates at $y^+ \approx 5$. The first-differencing procedure introduced in § 5 to determine $(d\bar{u}/dy)|_w$ and hence to infer ν is quantitatively accurate only if the physical profile of \bar{u} is linear from the wall to the location of the first parcel, a condition that is satisfied only if that interval is entirely contained within the viscous sublayer.

These considerations suggest that N might be chosen so that the parcel adjacent to a wall is located at $y^+ \approx 5$. For reasons explained below, this is not the procedure used to determine N , but for $A = 0.5$ it is shown in § 8.1 that this result is obtained as an outcome rather than as an imposed constraint, indicating a degree of self-consistency of the HiPS channel-flow formulation. For lower A , the extent of the HiPS viscous sublayer increases, but the increase is slight for the cases considered in § 8.2.

In ODT, parameter tuning has been used to obtain the correct extent of the viscous sublayer. This was accomplished through adjustment of the parameter that scales the viscous term in the event-rate expression (Schmidt *et al.* 2003). That term serves to suppress small (especially near-wall) eddies in ODT channel flow and thereby tune the spatial extent of viscous dominance. In ODT, many cells might be used to resolve the viscous sublayer, but in HiPS it cannot be resolved using more than one parcel because there is no viscous interaction between the parcel adjacent to a wall and parcels farther from the wall. From this viewpoint, there is no benefit to introducing near-wall event suppression in HiPS because in the y range in which it would suppress advection, instead of having purely viscous momentum transport there would be no momentum transport at all between homogeneous intervals, resulting in no momentum transport to the wall and thus no convergence of the time-advancement to a statistically steady state.

This difference between the HiPS and ODT resolution criteria illustrates the economy of the HiPS formulation, but this economy is accompanied by a loss of fidelity. Heat or mass transfer at a boundary layer might involve Pr or Sc much different from unity, such that the thermal or scalar wall sublayer might be embedded within the viscous layer (for high Pr or Sc) or *vice versa* (for low Pr or Sc). ODT, but not HiPS (as presently formulated), can capture such near-wall Pr and Sc effects.

6.2. Parameter assignments

Based on the parameter dependences explained in § 6.1, the following procedure for assigning HiPS parameter values is adopted. A particular choice of N is assumed. Again specializing to $A = 0.5$ (see appendix C for generalization to any $A \leq 0.5$), the remaining inputs are C and $T = h/u_\tau$.

In HiPS, event rates R can be non-dimensionalized using any fixed time T_0 , giving $R = (C/T_0)r$, where r is non-dimensional and does not depend on C . Then the non-dimensional quantity CT/T_0 is the sole input governing the steady-state single-time statistics of a HiPS channel-flow simulation. (Due to the associated time rescaling, varying C and T with CT fixed can affect multi-time statistics such as time-lagged correlation functions, which are not considered here. Additionally, the numerical evaluation of fluxes must account for the time rescaling, as explained at the end of appendix B.)

The first step in parameter assignment is to enforce the high- Re_U empirical friction correlation (Zanoun, Nagib & Durst 2009; Schultz & Flack 2013)

$$C_f \equiv 2(u_\tau/U)^2 = 0.0743 Re_U^{-1/4}, \quad (6.1)$$

where U is the bulk velocity and $Re_U \equiv 2hU/\nu$ is the bulk Reynolds number. Cases investigated here are within the Re_U range in which (6.1) is applicable. Schultz & Flack (2013) also proposed a more accurate correlation, but the additional accuracy is not needed for the present purposes and (6.1) simplifies the derivation of relevant relationships.

HiPS single-time statistics for given C and T and the associated Re_τ (determined using output) represent equally well the statistics for some \hat{Re}_τ (evaluated below) for any other values \hat{C} and \hat{T} such that $\hat{C}\hat{T} = CT$. From its definition, $C_f = 2(h/(TU))^2$. Use of (5.2) to evaluate ν in the definition of Re_U gives $Re_U = (2T^2U/h)(d\bar{u}/dy)|_w$. Insertion of these expressions for C_f and Re_U into (6.1) gives $2(h/(\hat{T}U))^2 = 0.0743 (\hat{T}U/h)^{-1/2} [(2h/U)(d\bar{u}/dy)|_w]^{-1/4}$, where \hat{T} now denotes the T value for

which (6.1) is satisfied based on substituting $T = \hat{T}$ into the expressions for C_f and Re_U . This gives

$$\hat{T} = \left(\frac{2}{0.0743} \right)^{2/3} \frac{h}{U} \left(\frac{2h}{U} \frac{d\bar{u}}{dy} \Big|_w \right)^{1/6}. \tag{6.2}$$

This result is used as follows. For any chosen C , $(1/\rho)(dp/dx)$ and h , where the latter two quantities determine $T \equiv [-(1/(h\rho))(dp/dx)]^{-1/2}$, a HiPS run yields values of U and $(d\bar{u}/dy)|_w$. All the quantities needed to evaluate \hat{T} are now available. Based on this evaluation, $\hat{C} = CT/\hat{T}$ completes the determination of the quantities \hat{T} and \hat{C} that would give the same values of U and $(d\bar{u}/dy)|_w$ (and of all other single-time flow statistics) but also obey the friction correlation. Thus, the friction correlation can be enforced by reinterpreting the results of a given run without performing any additional runs. Here, reinterpretation means adjustment of all quantities that are affected by the adjustment $T \rightarrow \hat{T}$, $C \rightarrow \hat{C}$.

To proceed in this way, the adjustment ratio $\hat{T}/T = (2/0.0743)^{2/3}(u_\tau/U) \cdot [(2h/U)(d\bar{u}/dy)|_w]^{1/6}$ is obtained from (6.2) and $T = h/u_\tau$, where u_τ is the original (unadjusted) friction velocity. Here $T = h/u_\tau$ also implies the adjustment $\hat{u}_\tau = (T/\hat{T})u_\tau$. Powers of \hat{T}/T adjust other relevant quantities. The adjustment of Re_τ involves adjustment of ν as well as u_τ . For fixed $(d\bar{u}/dy)|_w$, the mean momentum balance gives $\hat{\nu} \sim \hat{u}_\tau^2 \sim (T/\hat{T})^2 \nu$. This and the u_τ adjustment give $\hat{Re}_\tau = (\hat{T}/T)Re_\tau$. The Re_U adjustment involves only the ν adjustment, giving $\hat{Re}_U = (\hat{T}/T)^2 Re_U$. Finally, $\hat{y}_\tau = \hat{\nu}/\hat{u}_\tau = (T/\hat{T})y_\tau$ is used to evaluate $\hat{y}^+ = y/\hat{y}_\tau$.

It might be supposed that results for various \hat{Re}_τ values are obtained for given N by running cases with different values of the product CT . In fact, all non-zero CT values for given N yield the same wall-scaled results, corresponding to the same \hat{Re}_τ value. The reason is that all event-related velocity-field modifications by advective and viscous processes in HiPS, as well as event rates governed by (A 2), scale linearly with respect to a multiplicative rescaling of the velocity field. Denoting a representative velocity scale by V , this linearity implies that event rates scale as CV . The rate of momentum transfer induced by both advective and viscous processes involves an additional factor of V , giving the scaling CV^2 . Therefore, statistically stationary balance between given pressure forcing and all other processes implies that a change of C changes V so as to leave CV^2 invariant, hence $V \sim C^{-1/2}$. The relevant point is that modification of the input C for fixed T induces a spatially uniform rescaling of all single-time velocity statistics of the fully developed flow.

A particular consequence is that a change of C rescales the mean velocity profile without changing its shape, such that $U \sim (d\bar{u}/dy)|_w \sim C^{-1/2}$. The result needed here is that U and $(d\bar{u}/dy)|_w$ have the same C -dependence.

Suppose then that C is changed for fixed T , which changes U and $(d\bar{u}/dy)|_w$ by the same factor F but leaves u_τ unchanged because $T = h/u_\tau$. So C affects Re_τ only through the dependence $1/\nu \sim (d\bar{u}/dy)|_w$ from (5.2), giving $Re_\tau \sim (d\bar{u}/dy)|_w$, which implies that the change of C changes Re_τ by the factor F . Now the adjustment to enforce (6.1) is applied. Based on the powers of U and $(d\bar{u}/dy)|_w$ in (6.2), the change of C changes \hat{T} and hence the ratio \hat{T}/T by the factor $1/F$. Then, in the adjustment $\hat{Re}_\tau = (\hat{T}/T)Re_\tau$, the change of C introduces the factors $1/F$ and F , which cancel. Thus there is no dependence of \hat{Re}_τ on the input C . Since there is also no CT -dependence, is there likewise no T -dependence.

The upshot is that the adjustment that enforces the empirical friction correlation always yields the same \hat{Re}_τ value. In other words, the adjustment collapses all possible

cases (i.e. choices of C and T) for a given N to the same adjusted case. Henceforth, the hat notation is dropped and all quantities subject to adjustment are assumed to be adjusted.

The reasoning has been specialized to $A = 0.5$. The model extension to $A \neq 0.5$ in appendix C yields the same collapse for any given A , but A -dependence is not eliminated, so the collapse is for given A and N . Results are therefore presented in § 8 for various A as well as N values. There is no simple relationship between Re_τ and the hierarchy scale range A^{-N} , because the adjustment procedure involves outputs that are sensitive to the details of model dynamics.

Generalizing (A2) to incorporate viscous or buoyancy effects (see § 6.1 and appendix A) would introduce nonlinear dependence of event rates on a rescaling of the velocity field, resulting in C -dependence for given A and N , but there would still be no dependence on CT . The absence of CT -dependence is attributable to the constrained role of the mixing operation that represents viscous transport. If this operation were allowed to proceed on the basis of an independent time scale governed by a given viscosity, then there would be more freedom to preselect any value of Re_τ and to tune the model to fit comparison data. These modifications are desirable for some purposes, but here the goal of clearcut assessment of a minimal formulation is best served by avoiding such flexibility.

7. Numerical implementation

A HiPS channel-flow simulation time-advances the solution of a one-dimensional initial-boundary-value problem according to the advancement mechanisms described in §§ 2.2 and 3. Boundary-condition enforcement has been described in § 3.

The time-advancement cycle is governed by the statistical sampling of advection events, which involves sampling of both occurrence times and the particular swaps to be performed. The sampling and associated advancement mechanisms proceed as follows (omitting details explained in appendix A and references cited therein).

- (i) The time t of occurrence of the next prospective event is randomly sampled based on Poisson event-time statistics. The event is prospective because implementation of an event at time t is not guaranteed.
- (ii) Mean-pressure-gradient forcing of the u velocity profile is advanced to time t in one step because (3.1) is an exactly solvable linear differential equation.
- (iii) Having arrived at the event occurrence time t , a particular allowed swap is statistically sampled. The assumed swap probability distribution and the sampling rate used in step (i) together imply a swap sampling rate λ_s that can be different for each of the possible swaps.
- (iv) Based on the physical states of the parcels affected by the selected swap, physical modelling as described in appendix A determines the current value λ of the time-varying physical rate at which the selected swap should occur during the simulation.
- (v) Provided that $\lambda < \lambda_s$ (a requirement that sets a lower bound on the assigned value of λ_s), this determines the probability $P = \lambda/\lambda_s < 1$ that the current prospective event will be implemented (versus rejected).
- (vi) A Bernoulli trial using acceptance probability P determines whether the event is implemented. If not, the remaining steps of the time-advancement cycle are skipped.
- (vii) The pre-swap states at swap-affected parcel locations are saved.

- (viii) The swap is implemented.
- (ix) The post-swap and saved pre-swap states at swap-affected parcel locations are used to update the statistics of the M_i and M_{ii} terms of the turbulent-kinetic-energy (TKE) budget. (In appendix B, where details of the TKE budget are discussed, pre-change and post-change quantities are indicated by carets and tildes, respectively.)
- (x) The post-swap states at swap-affected parcel locations are saved.
- (xi) The mechanism that models pressure-gradient fluctuations is implemented.
- (xii) The saved post-swap states and the current states at swap-affected parcel locations are used to update the statistics of the K_i and K_{ii} terms of the TKE budget.
- (xiii) If the swap modifies any parcel pairings, then the current states at swap-affected parcel locations are saved; otherwise this and the remaining steps are skipped.
- (xiv) The mixing mechanism is implemented.
- (xv) The saved pre-mixing states and the current states at swap-affected parcel locations are used to update the statistics of the V_i and V_{ii} terms of the TKE budget.
- (xvi) If mixing modifies a wall-parcel state, then the velocity components in that parcel are set to zero to enforce the wall boundary conditions. For the affected wall parcel, the statistics of the V_i and V_{ii} terms of the TKE budget are updated.

In addition to the noted gathering of TKE-budget statistics, single-parcel statistics are gathered at event occurrence times, with weighting that accounts for the statistical variability of the time intervals between events. Spatial aggregation of the outputs to obtain the final results is described in § 8.1.

The present focus is on the fully developed, statistically stationary flow, so the initial condition is unimportant except for its effect on the duration of transient relaxation to the fully developed state. A tent-shaped (triangular) initial u profile is used whose average is set equal to a prior estimate of the steady-state bulk velocity. Each of the cases reported in § 8 consists of a single simulated realization in which data is gathered during the statistically stationary advancement after the relaxation of the initial transient flow development.

As explained in § 6.2, the parameter-setting procedure eliminates dependence of computed results on model inputs other than A and N . Therefore any convenient values of C , h , ρ and dp/dx are assigned. Variations of these parameters were run to verify the anticipated insensitivity to their values.

8. Numerical results

8.1. Results for $A = 0.5$

Single-point statistics, such as mean quantities, higher moments and probability density functions, can be gathered straightforwardly subject to the interpretations discussed in § 4, notably that statistics for the parcels within each homogeneous interval can be aggregated. Additionally, data from statistically equivalent homogeneous intervals on either side of $y = h$ are combined. Flux terms in the mean momentum balance and in budgets of quantities such as TKE require special treatment, explained in appendix B.

Figure 2 shows mean u values of individual parcels in order to highlight the intervals of statistical homogeneity, but henceforth parcel statistics are aggregated. Owing to the piecewise-constant character of statistical properties, such as in figure 2,

A	N	Re_τ	κ	t^+	CPU	CPU/ t^+
0.5	9	1967	0.262	249 262	66 min	0.016
0.5	8	1023	0.254	649 500	48 min	0.0044
0.4	7	1664	0.264	402 615	6 min	0.000 89
0.3	6	2336	0.266	474 239	82 s	0.000 17
0.3	5	781	0.252	202 528	8 s	0.000 040

TABLE 1. Simulation cases: cases are denoted by Nm , where m is the N value shown in the second column of the table. Here κ is the von Kármán constant, t^+ is the simulation time in wall units, and CPU is the processing time of the run; CPU/ t^+ is expressed in seconds.

aggregated statistical output is deemed to reside at isolated points, i.e. the midpoints of homogeneous intervals, except for the Reynolds shear stress, which resides at interval endpoints (see below). Thus, the data reduction is based on the treatment of homogeneous intervals as the control volumes of a standard finite-volume scheme. Line segments connecting data points in plots are drawn solely for visual clarity and are not intended to imply spatially continuous profiles.

Table 1 lists the cases considered here. As explained in appendix C, $A = 0.5$ is the physically most consistent A value, so results for $A < 0.5$ serve primarily to quantify the tradeoff between the loss of fidelity and the cost savings resulting from reduction of A . As explained in § 6.2, after constraining model parameters to obey the empirical friction correlation (6.1), results depend only on A and the number $N + 1$ of tree levels, because a given tree structure as specified by its A and N values can be consistent with the friction correlation for only one Re_τ value and the corresponding Re_U value prescribed by the friction correlation. For given A , this implies that channel flow can be simulated only for a discrete set of Re_τ values corresponding to the values of the integer N . As noted in § 6.2, HiPS channel flow can be formulated to be more flexible in this regard, but the more restrictive formulation is advantageous for present purposes.

Profiles of Reynolds shear stress for cases $N8$ and $N9$ are shown in figure 3. As explained in appendix B, there are no net property fluxes within a homogeneous interval, so fluxes are evaluated at the boundaries between adjacent intervals. Therefore the symbols in figure 3 are located at interval endpoints, but in plots of other quantities they are located at interval midpoints.

Wall-normalized Reynolds shear stress is not evaluated in HiPS using its definition $\overline{u'v'^+}$. In HiPS, parcels are advected not by their associated velocities but by advection events that are only indirectly controlled by the instantaneous velocity field. The operational definition of Reynolds shear stress that closes the HiPS mean momentum balance is derived in appendix B.

The linearity of the Reynolds shear stress (except at the wall, where it is zero) reflects the fact that momentum transport between a wall parcel and its adjacent partner occurs only by mixing, representing viscous transport, while elsewhere momentum transport between homogeneous intervals occurs only by advection. Through the mean momentum balance, the spatial uniformity of the pressure-gradient forcing requires the sum of mean viscous and advective transport to be linear. Thus, the profiles in figure 3 directly reflect model mechanisms. This behaviour can be compared to the channel-flow near-wall transition region seen in the measured y profile of Reynolds shear stress (Schultz & Flack 2013). The close adherence of the




FIGURE 3. Wall-normalized Reynolds shear stress. Symbols: +, case $N8$; \times , case $N9$; —, exact result in the absence of viscous transport.

HiPS results to the linear y -dependence shows that the data gathering was delayed long enough to allow full relaxation of initial transients and that the data-gathering period was long enough so that finite-sample effects were minor. (Section 8.3 contains further discussion of run-times and statistical precision.)

The solid line corresponding to the exact linear dependence is plotted in figure 3 for comparison with the HiPS results. This line is the spatial continuum result in the absence of viscous transport, corresponding to the infinite-Reynolds-number limit of physical channel flow. In contrast, HiPS is intrinsically a spatially discrete formulation, such that the finite cell size S implies a finite spatial extent of the viscous sublayer, allowing the implied viscosity to be inferred accordingly based on (5.2). The offset of the leftmost non-wall point from $y^+ = 0$ is not easily discerned in the linear coordinates of figure 3, but it is readily apparent in the semi-logarithmic plot examined next. In any case, the HiPS Reynolds-stress profile *per se* has no explicit manifestation of finite-Reynolds-number behaviour.

Like figure 2, figure 4 shows mean velocity profiles but with data aggregation as described, so that plotted data points correspond to the midpoints of homogeneous intervals. The logarithmic scaling that occurs for reasons explained in §5 is represented by line segments. The corresponding values of the von Kármán constant κ shown in table 1 are somewhat smaller than the value $\kappa = 0.4$ obtained from channel-flow measurements (Schultz & Flack 2013). It is unclear whether the slight Re_τ -dependence of κ indicated by table 1 is related to mechanisms that have been proposed to explain such dependence (Buschmann & Gad-el-Hak 2003); this merits further investigation.

In §6.1, it is noted that HiPS near-wall resolution of $y^+ \approx 5$, corresponding to the extent of the viscous sublayer, would be physically appropriate because net viscous momentum transfer occurs only between wall parcels and the parcels adjacent to them. The parameter-setting procedure does not enforce this as a constraint, but nevertheless it is seen that resolution close to this value is obtained. In any case, no




FIGURE 4. Mean velocity profiles. Symbols: +, case *N8*; ×, case *N9*; ○, measurements for $Re_\tau = 2000$ from Schultz & Flack (2013). The κ values in table 1 for the model cases are based on the solid and dashed lines here, which indicate logarithmic scaling.

procedure for setting the near-wall resolution can fully compensate for the fact that, in reality, viscous effects on the mean flow taper off gradually with increasing y^+ . This explains the absence of the experimentally observed shoulder, and further identifies the HiPS near-wall treatment as being largely responsible for both the capabilities and limitations of the HiPS channel-flow formulation.

Overall, the mean velocity profiles are more accurate than the underprediction of κ and the absence of the shoulder might suggest. These two errors tend to cancel because the high slope of the HiPS logarithmic region brings the HiPS profile back into conformance with the measured profile beyond the viscous-to-logarithmic transition region. Irrespective of phenomenological inaccuracies, prediction of the mean velocity profile with less than 30% error at all y^+ with no tuning of the model to fit the profile is a noteworthy outcome relative to previously demonstrated capabilities and has potential practical significance.

In this regard, it is noted in § 6.1 that ODT exhibits good mean-flow behaviour upon adjustment of a parameter that selectively suppresses near-wall eddies. When this parameter is set to a negligibly small value so that the effect is eliminated, the ODT mean profile, curve (a) of figure 3 in Schmidt *et al.* (2003), is remarkably similar to the HiPS profiles in our figure 4. In the HiPS formulation, empirical input is used to eliminate one degree of freedom, but not to tune the mean profile specifically, while the ODT curve (a) result is based on tuning the ODT analogue of the parameter C to optimize agreement with the mean velocity profile. This provides an indication that HiPS performance relative to its low cost (see § 8.3) and simplicity compares favourably to previously demonstrated capability.

HiPS profiles of u and v variance for cases *N8* and *N9* are compared in figure 5 with measurements at Re_τ values close to the model values. HiPS does not capture details of the shapes of the measured profiles, and the v variance is significantly overpredicted. The overall magnitude of the u variance profiles is more reasonable.




FIGURE 5. Velocity variances $\overline{u^2}^+$ (squares) and $\overline{v^2}^+$ (circles) for $Re_\tau = 1000$ (filled symbols) and 2000 (open symbols). The HiPS results for cases $N8$ and $N9$ (see table 1 for exact Re_τ values) are connected by line segments. The other symbols are measurements from Schultz & Flack (2013). Also shown are the $Re_\tau = 590$ DNS results of Moser, Kim & Mansour (1999) for $\overline{v^2}^+$ (dashed curve) and $\overline{w^2}^+$ (dot-dash curve).

ODT results shown in figure 5 of Schmidt *et al.* (2003) substantially underpredict both variance profiles after tuning two parameters to match the mean flow. The HiPS results, not tuned to match flow structure, are more accurate overall.

A key consideration with regard to the v -variance overprediction is that v fluctuations are produced solely by the return-to-isotropy representation of pressure-scrambling effects, as described in appendix A. Without this mechanism, only the u velocity would be non-zero, but in other respects HiPS channel-flow phenomenology would not be significantly affected. The ODT analogue of this mechanism was originally formulated with an additional parameter that can be used to adjust the production of v fluctuations (Kerstein *et al.* 2001). To minimize parameter tuning in the present study, the more restrictive approach of Kerstein & Wunsch (2006), which does not involve this additional tuning, is used here.

One limitation of the pressure-scrambling treatment is that it results in statistical equivalence of the v and w components for the channel-flow configuration. The analogous limitation in ODT is discussed in Schmidt *et al.* (2003). To quantify the effect of this artifact, v and w variance profiles from a direct numerical simulation (DNS) by Moser *et al.* (1999) are shown in figure 5. It is seen that the inability to capture the distinction is somewhat less impactful than the overprediction of either profile. Overall, it is found in HiPS, as in ODT, that velocity components for which the pressure-scrambling treatment is the sole turbulence-production mechanism are not as well predicted as those that are driven by body forces or imposed shear.

ODT near-wall subgrid closure of a channel-flow large-eddy simulation yields more accurate variance profiles (Schmidt *et al.* 2003), indicating that coupling to a coarse-grained 3-D core-flow representation can improve the ODT-resolved near-wall results.




FIGURE 6. Probability density function of wall-normalized wall stress. Symbols: +, case N8; ×, case N9; —○—, DNS for $Re_\tau = 1000$ by Lenaers *et al.* (2012).

Near-wall and general subgrid closure is a major goal of HiPS development (see § 9.1), which can possibly lead to similar accuracy improvement.

Due to the smallness of the v and w fluctuations relative to the u fluctuations, the overprediction of the v and w variances has only a minor impact on the TKE and hence on flow properties more generally. In this regard, it is explained in appendix A that HiPS velocity profiles do not directly advect parcels; rather, they influence advection through the dependence of the rates of occurrence of advection events on velocity fluctuations.

The Re_τ -sensitivity of the wall-scaled variance profiles, though slight, is an indication of deviation from inner scaling, the origin of which has been investigated in detail. Spectral and other analyses of channel-flow DNS have identified structural mechanisms of inner–outer flow coupling that can cause such deviation (del Álamo *et al.* 2004; Hoyas & Jiménez 2006). The occurrence of an analogous deviation in HiPS channel flow suggests that HiPS might embody an idealized representation of the inner–outer coupling.

Alfredsson, Örlü & Schlatter (2011) show that the root-mean-square fluctuation $\tau_w'^+$ of the wall stress is equal to the near-wall limit of u'/\bar{u} . The method by which viscosity is inferred for HiPS channel flow is equivalent to an assumption that the instantaneous value of τ_w^+ at a wall is equal to u_a/\bar{u}_a , where u_a and \bar{u}_a denote the instantaneous and average u values, respectively, of the parcel adjacent to the wall. This relationship is used to gather statistics of τ_w^+ . Probability density functions of τ_w^+ for cases N8 and N9 are compared in figure 6 with results of a channel-flow DNS for $Re_\tau = 1000$ (Lenaers *et al.* 2012).

The HiPS p.d.f.s have negligible Re_τ -dependence and are seen to have long-tailed shapes resembling the DNS p.d.f. but somewhat narrower. DNS p.d.f.s likewise exhibit no noticeable Re_τ -dependence on a linear plot, but in the far tails, discernible only on a logarithmic plot, DNS p.d.f.s exhibit a slight Re_τ -dependence and also a small likelihood of negative τ_w^+ values. The only advancement mechanism in HiPS that can possibly cause such occurrences is the pressure-scrambling operation described in




FIGURE 7. Probability density function of wall-normalized wall stress. Symbols: +, case $N8$; —, fitted lognormal curve.

appendix A, in which a kernel function that has both negative and positive values is added to the velocity values of a set of swapped parcels, such that the negative kernel values could be large enough locally to change the sign of u in some places. Thus, flow reversal can and possibly does occur, but evidently not adjacent to the walls, according to figure 6. Rather, τ_w^+ is seen in figure 6 to have a positive lower bound in the vicinity of 0.2.

ODT produces τ_w^+ p.d.f.s that closely resemble the HiPS p.d.f.s in figure 6. Like the HiPS p.d.f.s, they have shorter tails than the DNS p.d.f.s and have positive lower bounds (F. Meiselbach, personal communication). When ODT is used as a subgrid model in a 3-D flow solver, the τ_w^+ p.d.f.s that are obtained have negative tails, and in fact they quite accurately reproduce the DNS p.d.f.s (C. Glawe, personal communication). An important future goal is to determine whether subgrid implementation of HiPS in a 3-D flow solver will yield comparable performance at lower cost.

Except in the far tails, p.d.f.s of τ_w^+ from DNS are found to be lognormally distributed (Örlü & Schlatter 2011). Figure 7 shows that the HiPS results likewise exhibit lognormality.

Lognormal and other long-tailed p.d.f.s are familiar signatures of turbulence intermittency. Scale-invariant hierarchical stochastic models of turbulence capture various aspects of intermittency (Meneveau & Sreenivasan 1991) through their embodiment of the random-multiplier mechanism (Sreenivasan & Stolovitzky 1995) that idealizes the stepwise multiplicative breakdown of flow-structure scales, notably leading to lognormality under simple assumptions (Gurvich & Yaglom 1967). Although there are fundamental differences between HiPS and other hierarchical models, scale breakdown in HiPS is likewise multiplicative and scale-invariant (Kerstein 2013), so lognormality of property gradients is not a surprising outcome. Lognormality and other facets of turbulence intermittency have been studied in detail using scale-invariant hierarchical stochastic models and will likewise be examined in the future using HiPS. Figure 7, and a related result in § 8.2, are at most preliminary indications.




FIGURE 8. Positive y^+ : terms of the TKE budget, normalized by u_τ^4/ν , for cases $N8$ (crosses) and $N9$ (squares). Negative y^+ : DNS of Moser *et al.* (1999) for $Re_\tau = 590$ (curves) and 395 (crosses). The different line types represent the production term (solid), dissipation term (dashed) and sum of viscous and advective transport (dot-dash).




FIGURE 9. Contributions to the TKE transport term of figure 8: viscous (solid) and advective (dashed); the other symbols are as in figure 8.

In figures 8 and 9, the TKE budgets for cases $N8$ and $N9$ are compared to budgets from the channel-flow DNS of Moser *et al.* (1999). The budget terms, whose evaluation for HiPS channel flow is explained in appendix B, are normalized by dividing them by u_τ^4/ν .

Like the DNS budgets, the HiPS budgets exhibit negligible sensitivity to order-one variation of Re_τ , indicating that it is valid to compare the HiPS results with the lower- Re_τ DNS results. The main quantitative discrepancies in figure 8 are the

overpredictions of the viscous transport and the dissipation at the wall. These are actually a single discrepancy because the two quantities must sum to zero.

This discrepancy might be related to an ambiguity in the attribution of viscous effects to viscous transport and dissipation. Such ambiguities arise because evaluation of spatial derivatives is needed to determine certain budget terms, but these derivatives are not uniquely defined because HiPS statistical quantities have a piecewise-constant character stemming from the grouping of parcels into homogeneous intervals. The deviation from smooth behaviour is greatest near walls, as seen in figure 2. Owing to this spatial structure, there is a non-negligible dependence of several budget terms on the choice of numerical approximant or physical modelling assumption used to resolve each ambiguity. Importantly, this is not a numerical implementation sensitivity of the HiPS formulation *per se*; rather, it is an ambiguity of the after-the-fact interpretation of model results.

Subject to this caveat, the HiPS TKE budget is compared with DNS results. Away from the wall, it is seen that the HiPS results for TKE production, transport and dissipation are quantitatively as well as qualitatively reasonable. The degree of agreement of the HiPS TKE dissipation profile with the DNS result is potentially significant for the intended application of HiPS to turbulent mixing, in which scalar dissipation effects are crucial (Fox 2003).

Figure 9 shows the partitioning of the TKE transport into advective and viscous contributions. The necessarily inexact representation of derivatives is the underlying cause of the evident discrepancies. In particular, the production term, which is subject to numerical error resulting from finite-difference approximation of the derivative of the mean velocity, is an input to the evaluation of the advective transport. As a result, the advective transport contributions do not sum to zero. This does not indicate a violation of conservation laws but, rather, the lack of a fully consistent spatially discrete analogue of the continuum TKE budget. Conventional evaluation of budget terms involves order-of-accuracy errors that can cause deviations from the required summing to zero of advective transport contributions, but for smooth property fields, these errors can be kept small by mesh refinement.

An ODT channel-flow budget of u variance is shown in figure 6 of Schmidt *et al.* (2003). Compared to the HiPS TKE budget in figure 8, the overall accuracy is roughly the same, but the individual transport terms are more accurate than the HiPS results in figure 9. Spatial profiles in ODT are continuous, so budget terms are not subject to the technical difficulties that arise in HiPS data reduction. Considering this and the idealized HiPS representation of turbulent channel flow, the degree to which HiPS performs comparably to ODT is noteworthy.

8.2. Results for $A \neq 0.5$

In § 4 it is shown that the physical-space interpretation of HiPS channel flow for $A = 0.5$, involving one-dimensional alignment of parcels exactly covering the wall-normal line spanning the channel, is consistent (subject to a small deviation related to wall parcels) with the proximity definition of parcel separations. This reflects the effective dimensionality $d = 1$ that is obtained for $A = 0.5$ based on the relation $A = 2^{-1/d}$ (see § 2.1).

Reduction of A results in greater scale-space separation between hierarchy levels. Therefore fewer levels and hence fewer parcels are needed to represent a given scale range, implying a reduction of computational cost (see § 8.3). Results for $A < 0.5$, corresponding to $d < 1$, are examined in order to evaluate the cost–performance tradeoff associated with the reduction of d .




FIGURE 10. For case *N5*, parcel mean streamwise velocities in the y interval $[0, h]$, using uniform parcel spacing within each homogeneous interval.

Non-integer d suggests that the HiPS parcels cover a fractal subset of physical space. For the spatially inhomogeneous flow considered here, it is shown in appendix C that consistent physical-space interpretation requires parcel separation to increase with increasing distance from the wall. Therefore the parcels form a spatially inhomogeneous set; but, as noted below, this does not preclude fractal attributes.

Figure 10, which is the $A = 0.3$ analogue of figure 2, illustrates the low- A physical-space parcel distribution and the associated mean-flow structure. The reasoning that leads to the grouping of parcels into homogeneous intervals is valid for any A , so this behaviour is seen in both figures. The new feature is the reduced number of parcels per unit length in each successive homogeneous interval.

Like the $A = 0.5$ results that have been presented, results shown below for $A < 0.5$ are based on averaging of parcel statistics within each homogeneous interval, with interval averages plotted at the centres of the respective intervals. Therefore the assigned spatial distribution of parcels within each homogeneous interval is immaterial. For clarity, in figure 10 parcels are equally spaced within each homogeneous interval. A spacing that is more consistent with the proximity measure of separation yields a self-affine structure with associated fractal properties (Mandelbrot 1985). Due to the arbitrariness of the choice of spacing within homogeneous intervals, this has no physical consequence in the present context, but for other applications it might be relevant.

The number of parcels within a homogeneous interval increases with increasing distance from a wall. Results averaged over homogeneous intervals containing many parcels have greater statistical precision than averages over intervals with fewer parcels. So, by gathering statistics over sufficient run-time to obtain a desired precision near the walls, the precision elsewhere is better than needed. From this viewpoint, low A is computationally efficient because it concentrates parcels near the wall, where they are most needed. In this sense, it introduces efficient non-uniform meshing, albeit




FIGURE 11. Reynolds shear stress $\overline{u'v'}$. Symbols: \times , case $N5$; $+$, case $N6$; \diamond , case $N7$; —, exact result for $A = 0.5$ in the absence of viscous transport.

by the unphysical mechanism of introducing gaps between parcels, as explained next. This feature does not necessarily extend to cases other than wall-bounded flow.

It is important to note the distinction between fractal scaling of properties of the HiPS flow solution, which is discussed in Kerstein (2013), and the fractal or other low-dimensional spatial structure of the set of parcels that form the support of the flow solution. The parcels whose centres are plotted in figure 10 are all the same size, so their increasing separation implies gaps between the parcels. This introduces the unphysical notion of the region between the channel walls consisting of isolated patches of fluid separated by voids. HiPS allows advective and viscous property transport across such voids because HiPS advection is a non-local process and because parcel adjacency is a topological relationship that does not rely on physical adjacency.

Nevertheless, this artifact reduces the physical realism of HiPS, as reflected by lower accuracy of model results as A is reduced. This is illustrated by the A -dependence of the Reynolds-shear-stress profiles shown in figure 11. The pressure-gradient momentum source is applied only within parcels, so far from walls, where fluid parcels are widely spaced, it is applied in a relatively small fraction of each homogeneous interval. Mean momentum conservation then implies a relatively shallow slope of the Reynolds-shear-stress profile. The profile steepens as the wall is approached. The resulting deviation from linearity increases with decreasing A , reflecting the increasingly unrealistic spatial structure.

Figure 12 and the κ values in table 1 indicate only a slight sensitivity of the mean flow to A . Here κ varies mainly with Re_τ . At the parcels adjacent to the walls, the wall treatment enforces the viscous-layer relation $\bar{u}^+ = y^+$. Except in a y^+ range to the left of any plotted point, the graph of this relation has a slope steeper than the logarithmic-layer slopes, so the vertical offsets of the successive profiles are sensitive mainly to the coarser physical-space resolution as N decreases. The dominant dependence on N




FIGURE 12. Mean velocity profiles. Symbols: +, case *N5*; \diamond , case *N6*; \square , case *N7*; \times , case *N9*. The κ values in table 1 for these cases are based on the line segments here, which indicate logarithmic scaling.

rather than Re_τ is not a physically meaningful model behaviour, but its quantitative impact is not large.

Coarsening resolution causes the model representation of the viscous sublayer to extend beyond its true physical extent $y^+ \approx 5$. As noted in § 6.1, this implies possible inaccuracy in the determination of ν and therefore possible systematic discrepancies in various wall-normalized results. However, the negligible κ -sensitivity suggests that this source of inaccuracy is not significant for the cases considered.

Figure 13 indicates greater A -sensitivity of variance profiles. Again, the sensitivity is not primarily a Re_τ -dependence. The results indicate decreasing accuracy as A is reduced.

Figure 14 indicates moderate sensitivity to A of the τ_w^+ p.d.f. The main trend is a shortening of the tail of the p.d.f. with decreasing N . Lognormality is an asymptotic property resulting from the multiplication of a large number of random variables representing the scale reduction at successive levels, so as the number of hierarchy levels decreases, the long-tailed character of the p.d.f. diminishes.

Figure 15 compares the TKE budgets for the cases with smallest and largest N . The differences are significant, yet case *N5* nevertheless has a rough quantitative correspondence to the DNS results. It is noteworthy that a case that involves 32 parcels and runs in a few seconds (see table 1 and § 8.3) yields such an outcome for a flow metric that challenges any form of reduced modelling. The results for case *N6* (not shown) differ only slightly from the case *N5* results, consistent with the insensitivity to Re_τ for given A that is seen in figure 8.

The overall impression is that reduction of A causes increasing but generally moderate quantitative discrepancies. Some degree of reduction might be beneficial depending on application-specific criteria, including cost considerations such as those discussed in § 8.3.




FIGURE 13. Velocity variances $\overline{u^{2+}}$ (filled symbols) and $\overline{v^{2+}}$ (open symbols): ∇ , case $N5$; \triangle , case $N6$; \diamond , case $N7$; \square , case $N9$; \circ , measurements for $Re_\tau = 2000$ from Schultz & Flack (2013).




FIGURE 14. Probability density function of wall-normalized wall stress for cases $N5$ (dot-dash), $N6$ (dashed), $N7$ (solid) and $N9$ (crosses).

In this regard, the A -dependences indicated here do not necessarily hold for flow configurations or model formulations other than the model formulation and application considered here. The focus here is on simplicity and minimal parameter adjustment to allow model assessment from a fundamental perspective. Additional empiricism can




FIGURE 15. Terms of the TKE budget, normalized by u_τ^4/ν : production (solid), dissipation (dashed) and the sum of viscous and advective transport (dot-dash). Positive y^+ : \diamond , case *N5*; \square , case *N9*. Negative y^+ : DNS results as in figure 8.

be introduced straightforwardly, e.g. by allowing wall parcels to have an adjustable nominal mass rather than the same mass as other parcels. Other possible model variants have been noted. Accordingly, the results presented here should be viewed as a baseline assessment rather than a complete exploration of the range of options.

8.3. Computational cost

In the discussion of figure 10, it is noted that the number of parcels representing a flow for given Re_τ decreases with decreasing A , but this does not necessarily impact statistical precision because successive homogeneous intervals contain multiplicatively increasing numbers of parcels. Therefore it is reasonable to base a cost comparison over a range of A values on CPU time per unit normalized time-advancement using the wall-normalized run-time $t^+ = tu_\tau^2/\nu = t(d\bar{u}/dy)|_w$, where t is the run-time and the final expression follows from (5.1).

The timing results shown in table 1 are based on serial implementation on an i7™ processor. It is found that there is roughly a factor-of-four cost increase per additional hierarchy level, whether comparing equal or unequal A values, indicating that the cost reduction resulting from decreasing A is due mainly to reduction of the number of hierarchy levels needed to reach a given Re_τ range. Each additional hierarchy level doubles the number of parcels and roughly doubles the number of distinct advection events that can occur. The latter effect doubles the number of sampled events per unit time, and the former effect doubles the cost of evaluating a typical sampled event, so the combined effects account for the fourfold cost increase. (Determination of which sampled events are implemented, described in appendix A, is the dominant computational cost.)

The indicated t^+ values are much higher than would be needed for converged statistics in three dimensions, because 3-D simulation allows averaging over the

homogeneous directions. In figure 3, results for case $N9$ deviate slightly but noticeably from linearity, but no such deviation is seen for case $N8$, whose run-time is longer in terms of t^+ , suggesting that roughly 5×10^5 wall times are needed for sufficiently converged statistics so that there is no visually detectable finite-sample variability.

It is possible, however, that a larger multiple of the wall time t^+ is needed for N values larger than considered here. Figure 2 provides a hint of the origin of such N -dependence. The two longest homogeneous intervals exhibit slight jumps of their \bar{u}^+ values across their midpoints. Within each homogeneous interval, the parcel groups on either side of the midpoint have lower proximity relative to each other than parcels within either group. Based on the turbulence phenomenology embedded in the procedure for selecting advection events (see appendix A), this implies less frequent advective transfer between the groups than within the groups. For a given run-time, this implies more sampling variability of the intergroup \bar{u}^+ difference than of the intragroup \bar{u}^+ difference, as seen. As more hierarchy levels are added, longer run-times (as multiples of t^+) might therefore be needed in order to maintain a given variability tolerance. A countervailing consideration is that the longest homogeneous intervals have increasing numbers of parcels with increasing N , so interval averages are increasingly stable. The shorter near-wall intervals might therefore control overall precision. The resulting N -dependence might be different for different output statistics.

The precision target adopted here is more stringent than required for most uses of turbulence simulations, so meaningful results should be obtainable in shorter times than indicated in table 1, especially after application-specific optimization of the numerical algorithm. Cost minimization is a key requirement for subgrid implementation within 3-D flow solvers, which is a long-term goal of HiPS development.

9. Discussion

9.1. Comparison with other approaches

Turbulence modelling based on a hierarchical (tree) geometry has hitherto focused on homogeneous conditions and associated universal properties of the turbulent cascade. HiPS can likewise address these features (Kerstein 2013), but the first quantitative demonstration of model performance, presented here, focuses on the phenomenology of an inhomogeneous turbulent flow. This focus serves to illustrate the novel capabilities of HiPS relative to other hierarchical formulations.

Hierarchical turbulence modelling is one of several contexts that provide perspectives on the present study. Other relevant contexts are turbulence modelling in general and modelling of turbulent mixing.

The degree of physical fidelity required for accurate modelling of turbulent flow and processes within turbulent flow, such as molecular mixing, is sensitive to the impact of small-scale processes, such as mixing, on the flow at larger scales. Viscosity is primarily a small-scale kinetic-energy sink with relatively small effects on large-scale flow structure, but in wall-bounded flow, the wall-blocking mechanism suppresses large-scale influences near walls, resulting in important viscous influences within the thin but crucial interval from the viscous through the logarithmic sublayers. In this respect, modelling of near-wall turbulent flow and modelling of turbulent mixing involve comparable challenges.

To capture the viscous influence on near-wall turbulence, modelling of this flow regime at a deeper level than mixing-length and related dimensional reasoning (an example of which is presented in § 5) requires some degree of resolution of unsteady

viscous effects. Accomplishing this at less cost than DNS requires some combination of sacrificing resolution, simplifying the dynamics and reducing the dimensionality of the flow domain. The predominant dependence of near-wall flow properties on the wall-normal direction suggests that dimension reduction involving flow modelling on a one-dimensional wall-normal domain might be a promising approach. From this viewpoint, near-wall flow is a good test case for a dimension-reduction strategy as well as an intrinsically important target application.

This reasoning motivated the development of ODT, which is the basis of much of the modelling within HiPS as well as a useful comparison case, as noted in § 8. On a one-dimensional wall-normal domain, ODT is fully resolved, allowing physically sound representation of viscosity and other molecular transport processes. One-dimensional unsteady simulation of turbulence requires simplification of the flow dynamics regardless of cost considerations. The baseline ($A = 0.5$) HiPS channel-flow formulation is likewise effectively one-dimensional, with dynamics based largely on the ODT approach. Comparisons shown in § 8 indicate no major performance advantage of either method. In this context, the benefits of HiPS relative to ODT and other alternatives require careful examination.

One important consideration is cost. As implemented here, HiPS is spatially under-resolved, yet its performance is adequate for some purposes and is not matched by any comparably economical approach. (Predictions of some flow properties considered here, such as wall-stress p.d.f.s, are beyond the scope of other comparably economical approaches.) Coarse-grained ODT with subgrid modelling on the one-dimensional domain, though feasible and demonstrated for near-wall flow (Kerstein & Wunsch 2006), holds promise but is not fully established as a viable alternative cost-reduction strategy. Coarse-graining negates some of the salient features of ODT, such as its comprehensive representation of molecular transport, including Pr and Sc dependencies, without attaining conceptual and algorithmic simplicity comparable to HiPS. If adjacent ODT fluid elements were instantly mixed, the flow state would become and remain spatially uniform. The restrictive definition of adjacency in the hierarchical geometry prevents such anomalies, enabling a simple, economical coarse-grained mixing treatment involving instant mixing of newly adjacent parcels.

HiPS has additional scope for cost reduction by decreasing A to values below 0.5, corresponding to a non-integer effective dimensionality below unity. As noted in appendix C, $A < 0.5$ is physically problematic (though not to a degree that precludes meaningful results) for channel-flow simulation, but this concern does not necessarily extend to other flow configurations.

Most relevant in this regard is mixing simulation, which is the main future target application (Kerstein 2013). Although self-contained mixing simulation is of technological as well as scientific interest, the principal use of mixing models is for subgrid mixing closure of 3-D coarse-grained unsteady flow simulations and steady-state solvers. For this purpose, it is generally appropriate to treat the subgrid mixing-field structure within 3-D control volumes as spatially homogeneous. In this case, the problematic channel-flow low- A physical-space interpretation does not apply, though there is still a practical limit to the reduction of A due to the need for enough parcels to provide a sufficient statistical sample of mixture states. Thus, reduction of the effective dimensionality to values below unity is a suitable cost-reduction strategy for mixing simulation.

Parcel-based subgrid mixing simulation as generally implemented (see § 1) is effectively zero-dimensional, so statistical sampling requirements determine the needed

number of parcels. By choosing a HiPS effective dimension that corresponds to the same number of parcels for a given case, the turbulence phenomenology embodied in HiPS can enhance model fidelity with no cost increase other than the minor overhead associated with HiPS time-evolution of parcel grouping. Implementation within existing algorithms is straightforward because it affects only the selection of pairs to be mixed, with no other impact on time-advancement.

An important exception to the appropriateness of a spatially homogeneous subgrid mixing-field structure is subgrid modelling in control volumes adjacent to walls. The only robust multiphysics near-wall subgrid treatment that is currently available is ODT, whose near-wall closure implementation for channel-flow large-eddy simulation has been demonstrated (Schmidt *et al.* 2003). (Here, multiphysics refers to couplings between dynamically active scalars and the flow field.) The results presented here suggest that HiPS has the potential to provide comparable near-wall performance at lower cost, with the additional benefit of compatibility with current modelling frameworks. An important caveat is that HiPS cannot resolve Pr and Sc effects, physical-space flame structure etc., so ODT is ultimately a higher-fidelity approach for the most demanding applications. Concurrent use of HiPS in some control volumes and ODT in others is a cost-effective high-fidelity approach that can be envisioned.

Apart from cost and performance considerations, there is a fundamental scientific motivation for the development of HiPS. Modelling is ultimately a quest for the simplest formulation that can represent complex reality. This motivated the formulation of shell models that simplify Navier–Stokes dynamics by collapsing the wavenumber vector to a discretized scalar, which precludes a direct physical-space interpretation (Kadanoff 1995). Hierarchical (tree) geometry is the minimal extension of this approach towards a physical-space treatment. From this viewpoint, it is natural to investigate the extent to which model development within this framework can lead to formulations that broaden its capabilities and scope of application. The degree to which HiPS can achieve these outcomes has been discussed in detail (Kerstein 2013), and the present study provides the first direct supporting evidence. An important conclusion is that HiPS to some extent bridges the longstanding gap between application-focused modelling and the physics-focused model progression originating with shell models.

9.2. Features and performance of the channel-flow formulation

In a fundamental scientific context, the present study not only serves as a capability demonstration but additionally indicates the extent to which the phenomenology of near-wall turbulent flow can be captured using a formulation that is minimal in the sense explained in § 9.1. To the extent that this is possible, the study indicates that aspects of this phenomenology are not uniquely attributable to the particular vortical structures (streamwise vortices, hairpins etc.) or transient events (sweeping, ejection etc.) known to occur in near-wall turbulent flow. Here, flow inhomogeneity is introduced by imposing two elementary constraints on the parcels selected to represent the channel walls: they cannot be advected, and any non-zero velocity components resulting from momentum transfer to those parcels are immediately set to zero in order to enforce the Dirichlet boundary condition. The advection process is as featureless as possible within a formulation designed to capture the range of scales of turbulent motion. It consists of swaps implemented as rigid translations, such that each swap causes two fluid regions to change places. Viscous transport is minimally implemented as complete mixing whenever two parcels become adjacent due to advection.

In the physical-space interpretation of this formulation, these swaps introduce velocity discontinuities. However, the restrictive definition of parcel adjacency within the hierarchical structure prevents viscous transport across these discontinuities. Therefore the discontinuities do not result in anomalies of the sort that might occur if the model were implemented in physical space.

The salient feature of HiPS flow simulation is the development of a cascade analogous to the inertial-range turbulent cascade, resulting in the breakdown of flow structures from the largest to the smallest scales of the model hierarchy (Kerstein 2013). The wall representation in the channel-flow formulation modifies this phenomenology in a fundamental way. As illustrated in figure 1 and explained in §5, the self-evident requirement that the wall parcels cannot be advected has the consequence that the largest-scale advective motion that can include a given location y is of order y rather than of order h , which is the largest possible scale of motion in the hierarchy. This model property is the basis of key features of HiPS channel flow, such as the logarithmic subrange of the mean velocity profile.

The extent to which the highly idealized turbulence, viscosity and wall-coupling representations in HiPS capture channel-flow phenomenology has been evaluated. As shown in §6.2, the simplicity of the formulation leads to unique model predictions upon invoking the friction correlation as the sole empirical constraint on model parameters. Given this constraint, results depend on the scale-space stride A and the number $N + 1$ of hierarchy levels, but $A = 0.5$ is preferred due to its physical consistency for channel flow (see appendix C). For given A , each N value corresponds to a unique value of Re_τ , so N is a surrogate physical input rather than a tunable parameter. Because N is an integer, the model yields results for only a discrete set of Re_τ values for given A . The minimal degree of empiricism needed to uniquely specify the formulation is an important consideration motivating the choice of channel flow as the first HiPS application.

Various flow statistics quantify the predictive capabilities of the model. It is shown in §5 that the hierarchical geometry in conjunction with the wall treatment imposes a logarithmic range of the mean velocity profile. Apart from the friction correlation, which is an imposed constraint, the most fundamental metric of channel-flow structure is the von Kármán constant κ . The slightly case-sensitive results underpredict κ , but figure 4 shows that the mean velocity profile is more accurate than this underprediction might suggest. Other model predictions are, in general, comparably accurate. Some of the noted exceptions can be attributed to features of the particular HiPS formulation used in this study rather than general limitations of the HiPS approach. Future formulations tailored to other applications will broaden the perspective in this regard. Additionally, it has been noted that HiPS channel flow can be formulated in a more empirical but more flexible way, allowing simulations for any Re_τ rather than just for isolated values.

The model exhibits collapse of various flow statistics in wall units and the known deviation of wall-scaled velocity variance profiles from collapse. In general, the sensitivity or insensitivity of HiPS results to Re_τ is in accordance with experimental and DNS results, suggesting that the model might embody some representation of the inner–outer coupling that causes such sensitivity.

These assessments are based on results for $A = 0.5$. For lower A values, computational cost is considerably reduced, motivating examination of results for $A = 0.4$ and 0.3 despite the increasingly problematic physical-space interpretation of HiPS as A is reduced. The results suggest that modest reduction of A below 0.5 can offer advantageous cost–performance outcomes, depending on application-specific criteria.

Owing to the effective dimensionality $d < 1$ of the spatial support of the flow domain for $A < 0.5$, this algebraically straightforward parameter variation introduces the novel concept of turbulence simulation on a domain of non-integer effective spatial dimensionality, whose implications merit further examination.

9.3. Relationship to turbulent mixing

Near-wall turbulent flow and mixing are related, at least in a restricted sense, by the Reynolds analogy (Eckert & Drake 1972). In HiPS, the relationship is broader because velocity components are more closely analogous to advected scalars than in Navier–Stokes dynamics.

The reason is that HiPS velocity components do not directly advect fluid. Rather, the dynamical role of velocity components (complementary to their role as model observables) is to serve as carriers of kinetic energy, whose spatial distribution influences the stochastic sampling of advection events as described in appendix A. As a result, the usual advective term that distinguishes velocity from other flow properties is absent, and velocity components evolve mathematically as individual advected diffusive scalar fields subject to various sources and sinks, such as those associated with the return-to-isotropy mechanism.

In appendix A it is noted that this mechanism can be, and originally was, omitted from ODT, in which case only the u component is non-zero. Then, for channel flow, u is subject only to advection, viscosity and pressure forcing. Here u is then mathematically equivalent to a passive scalar, say temperature with $Pr = 1$, subject to a heat source that is uniform in space and time, with isothermal boundary conditions at the channel walls. In this case, an ODT channel-flow simulation is equivalent to a scalar mixing simulation.

The same considerations apply to HiPS. For maximum scope of application, the return-to-isotropy mechanism is included in HiPS, so the present study does not unequivocally address scalar mixing. However, related ODT results strongly suggest that, as a practical matter, the present HiPS implementation is roughly equivalent to a mixing simulation. Specifically, the results of one-component ODT simulations of turbulent boundary layers (Kerstein 1999) and free shear flows (Kerstein & Dreeben 2000) closely resemble those of their three-component counterparts (Kerstein *et al.* 2001; Schmidt *et al.* 2003). Thus, at least inferentially, the present results can be viewed as valid indicators of the performance of HiPS as a turbulent mixing model.

Nevertheless, this is at most a preliminary step in the demonstration and application of HiPS mixing capabilities. An important consideration with regard to future steps is that mixing properties relevant to reacting flows are potentially sensitive to all the details of the ensemble of time histories of parcel-pair encounters. A theoretical development that is pertinent in this regard is the temporal multiscaling concept and its associated mathematically and computationally demonstrated properties (L'vov, Podivilov & Procaccia 1997; Biferale, Calzavarini & Toschi 2011). Future investigation of temporal multiscaling using HiPS will complement previous hierarchical shell-model studies of temporal multiscaling (Pandit, Ray & Mitra 2008), as well as indicate HiPS performance in a context relevant to mixing.

Turbulent mixing simulation has different performance requirements than turbulent flow simulation because of molecular-transport and chemical-kinetic uncertainties that limit the attainable fidelity as well as sensitivities to flow details that might be unimportant in constant-property flows. Although formulated from the perspective of mixing simulation, it has been shown here that HiPS has noteworthy turbulence modelling capabilities that offer a novel perspective on the mechanistic origins of some well-known channel-flow behaviours.

Acknowledgements

The author thanks P. Schlatter and M. P. Schultz for providing numerical and experimental results. This work was performed using an adaptation of a computer code written by D. O. Lignell. The approach on which this research is based was formulated by the author during a workshop conducted by the Institute for Pure and Applied Mathematics (IPAM) at UCLA. IPAM is supported by the US National Science Foundation.

Appendix A. Advection in HiPS flow simulation

Advection in HiPS can be either parameterized or implemented as a more self-contained and predictive flow simulation. The focus here is on the latter. A minimal formulation of advection in HiPS flow simulation is outlined in Kerstein (2013). In the present study, an advection treatment is used which is more robust in ways that will be explained. It has not been described previously, so details are presented here.

Advection in HiPS flow simulation is closely analogous to advection in ODT. In particular, the previously outlined minimal HiPS flow formulation is analogous to the ODT formulation introduced by Kerstein (1999), but the approach used here is analogous to the constant-density specialization of the version of ODT used in Kerstein & Wunsch (2006), Kerstein (2009) and Gonzalez-Juez, Kerstein & Lignell (2011).

The notation and conventions of the latter references are adopted, except that here the one-dimensional spatial domain is denoted by y instead of z and component indexing rather than (u, v, w) notation is used where convenient. The analogy between HiPS and ODT is explained so that applicable analysis and interpretation in previous ODT studies can be carried over directly.

On a uniformly meshed one-dimensional spatial domain, advection in ODT is represented by permutations of the spatial ordering of the mesh cells. (A different approach to numerical implementation of permutations is also possible in ODT (Lignell *et al.* 2013), but it is not applicable to HiPS and therefore not discussed here.)

In ODT, each permutation is a ‘triplet map’ applied to a contiguous interval whose size is denoted by l . A HiPS permutation swaps two sets of parcels, such that the sum of the sizes of the swapped intervals is the HiPS analogue of the ODT map size l . The swapped intervals are equal in size and are therefore each of size $l/2$.

In ODT, length scales such as the map size l are defined in the spatial continuum, and spatial discretization is undertaken for purposes of numerical implementation. HiPS is formulated in a discrete space that can be interpreted as continuous space. It is valid and convenient to carry over the continuous-space terminology and formalism of ODT to HiPS provided that a consistent interpretation is used, as in §4 and appendix C.

In ODT, $f(y)$ denotes the inverse of the map, meaning that $f(y)$ is the location that is mapped to location y . To evaluate this in HiPS, a y coordinate is used such that the collection of swapped parcels together form a contiguous size- l interval whose centre is nominally at $y = 0$. Then the associated inverse map is $f(y) = y - (l/2) \operatorname{sgn}(y)$ for $|y| \leq l/2$ and $f(y) = y$ for $|y| > l/2$, where sgn denotes the sign function. Another key construct is the kernel $K(y) \equiv y - f(y)$, which is the map-induced displacement of a location mapped to y .

The length scale l is central to the analysis that follows. It can be defined in several possible ways, e.g. as two times the length scale l_n associated with the hierarchy

level n of the swapped subtrees. For $A = 0.5$, $2l_n$ is the sum of the sizes of the parcels in the swapped subtrees; but for $A < 0.5$, that sum is less than l_n due to the gaps between parcels. The gaps are deemed to be void, with no dynamical role (see § 8.2 and appendix C), so they are excluded from spatial integrals of flow properties. Because l is the normalizing length scale that converts certain of these integrals into averages, l is defined as the sum of the sizes of the swapped parcels rather than as $2l_n$.

Based on these considerations, the y coordinate used in the event-sampling analysis corresponds to the coverage of spatial intervals by parcels rather than the full extent of the intervals, and hence differs from the y coordinate used in appendix C. This is one example of the complications that arise for $A < 0.5$. The analysis that follows is most easily understood by focusing on the simpler case of $A = 0.5$.

In terms of the generic inverse map f and kernel K , the formulation and analysis of ODT event-sampling and event-implementation procedures carry over directly to HiPS, subject to the sole distinction that the mean-square displacement $\langle K^2 \rangle = (1/l) \int K^2(y) dy$ within a triplet-mapped region is $(4/27)l^2$, but the HiPS analogue is $(1/4)l^2$, which is the nominal squared displacement of every swapped parcel. Accordingly, we follow closely the ODT analysis and notation, omitting details that can be found in the cited references.

In spatial-continuum notation, random sampling of event occurrences requires the specification of an event-rate distribution that is denoted by $\lambda(l, y_0; t)$ in ODT, where y_0 is the left boundary of the mapped interval and the argument t reflects the dependence of λ on the instantaneous system state within that interval. In HiPS, the sampling of subtrees to be swapped requires the event-rate distribution to have different arguments. Assuming any convenient indexing of subtrees within the HiPS hierarchy, e.g. as in Kerstein (2013), the indices j and k of the sampled subtrees complete the specification of a swap. Then $\lambda(j, k; t)$ denotes the HiPS event-rate distribution. Because j and k are integers, λ has units of inverse time, unlike λ in the continuum mathematical formulation of ODT. (The spatially discrete numerical implementation of ODT is closer in formulation to HiPS.)

As in ODT, the available-energy concept is the basis for physical modelling of λ in HiPS. The kinetic energy of velocity component i can be modified by adding $c_i K(y)$ to the velocity profile, where $K(y)$ is the kernel associated with a particular swap. Using the HiPS value of $\int dy K^2(y)$, the resulting kinetic-energy change (in one-dimensional units) is

$$\Delta E_i = \frac{\rho}{2} \int dy [(v_i + c_i K)^2 - v_i^2] = \frac{\rho}{2} \int dy [c_i^2 K^2 + 2c_i v_i K] = \frac{\rho l^3}{8} c_i^2 + \rho l^2 c_i v_{i,K} \quad (\text{A } 1)$$

where $v_{i,K}$ denotes $(1/l^2) \int dy v_i K$. The available energy Q_i is defined to be the maximum of $-\Delta E_i$ with respect to c_i , which occurs at $c_i = -(4/l)v_{i,K}$, corresponding to $Q_i = 2\rho l v_{i,K}^2$. As the terminology suggests, Q_i is the maximum amount of kinetic energy that can be extracted from component i using this procedure.

To model λ , powers of ρ , l and the total available energy $Q = Q_1 + Q_2 + Q_3$ are combined to obtain an inverse time, giving

$$\lambda = C \left(\frac{2\langle K^2 \rangle Q}{\rho l^5} \right)^{1/2}. \quad (\text{A } 2)$$

Here, C is a parameter that can be tuned so that the swap-induced advection is quantitatively consistent with turbulent transport in the corresponding physical flow

(but in the present formulation, the parameter-setting procedure in § 6.2 absorbs this degree of freedom). The quantity $2\langle K^2 \rangle$ is included in order to follow the ODT convention for defining λ .

The ODT expression for λ includes a viscous term with a second tunable parameter Z . This term can be included in HiPS, but is omitted here because it would introduce an unphysical artifact in the present context, as explained in § 6.1.

Every time-advancement operation changes the Q and hence the λ values for some subset of the possible future events. The affected set of rates can be recomputed, but doing so is inefficient. Instead, each type of event, parameterized by j and k , is oversampled. A sampled event is then allowed to occur with probability λ/λ_s , where λ_s is the sampling rate for the chosen event type. To guarantee oversampling, λ_s must always exceed the true rate λ , whose current value for the chosen event type is determined using (A 2). The benefit of this ‘thinning’ procedure (Law & Kelton 2000) is that only one λ value need be computed per event sampling. This more than compensates for the inefficiency of rejecting some of the sampled events, in fact almost all of them because λ_s must be an upper bound on the highly variable quantity λ , so in most cases $\lambda_s \gg \lambda$. Event sampling is done by thinning in both ODT and HiPS.

HiPS time-advancement incorporates an ODT feature that extends model physics without additional tuning, namely its ‘return-to-isotropy’ representation of pressure-scrambling effects. This is done by using the kernel not only to compute Q for λ evaluation but also to modify velocity profiles after map implementation so that the new component available energies are $Q'_1 = Q'_2 = Q'_3 = Q/3$. Enforcing equality of the quantities Q'_i tends to equalize component kinetic eddies, which are correlated with but not identical to their available energies as defined. The c_i values required to enforce this equality are obtained by substituting $\Delta E_i = Q/3 - Q_i$ into (A 1) and solving for c_i . The solution has two branches. The physical branch is the one that satisfies $\Delta E_i \rightarrow 0$ as $c_i \rightarrow 0$.

A broader motivation for introducing this feature is that it is required for a meaningful three-component representation of channel flow. The mean pressure gradient is applied only to the u component. Without the pressure-scrambling representation, the v and w components would never deviate from zero. With it, they are meaningful model observables, though not necessarily accurately predicted (e.g. they are statistically equivalent for the channel-flow configuration, unlike their physical behaviour). The definition $K(y) \equiv y - f(y)$ implies that $\int dy K = 0$ and hence that the pressure-scrambling operation does not change the domain-integrated momentum for constant ρ .

The return-to-isotropy mechanism is one of several extensions of ODT physics that were enabled by introduction of the available-energy concept, superseding the original, kinematically based ODT formulation (Kerstein 1999). Another such extension is the use of the ODT version of (A 1) to enforce kinetic-energy changes equal and opposite to advection-induced changes of gravitational potential energy in buoyant stratified flows (Kerstein & Wunsch 2006). This applies also to HiPS, and is the basis of applications of HiPS to buoyant stratified flow that will be reported in the future.

The return-to-isotropy mechanism conserves total kinetic energy as well as momentum, and additionally conserves the total available energy Q . Since λ is affected by the instantaneous flow state only through its dependence on Q , changes of the component values Q_i that leave Q unaffected do not directly affect the statistics of swap occurrences. However, there is an indirect effect because the u component is subject to pressure forcing whose contribution to kinetic energy depends on the magnitude of u . More generally, any state change potentially affects subsequent energy

dissipation, thereby eventually influencing swap statistics. As a result, the effect of the return-to-isotropy mechanism on the u profile indirectly influences Q values and hence affects swap statistics.

Apart from this pressure-forcing effect, λ is unaffected by the directional orientations of velocity components, and in particular the flow evolution is indifferent to the kinematic distinction between the wall-normal and spanwise components. In three dimensions this distinction is crucial, e.g. because the continuity equation and the no-slip boundary conditions require $dv/dy = 0$ at the wall. In HiPS, 3-D continuity cannot be enforced, and in fact it is immaterial because fluid is advected by swaps rather than swept by velocity fields. By construction, the swaps are globally conservative on the one-dimensional domain, which is sufficient to enforce the HiPS analogue of continuity.

Appendix B. Data reduction

The mathematical formulation and numerical evaluation of the terms of the TKE budget and related quantities follow the overall approach applied to ODT in the appendix of Kerstein *et al.* (2001), the notation of which is adopted here where applicable. Equation numbers that are preceded by 2001 refer to equations in that appendix.

The time-advancement of velocity component v_i in HiPS can be represented formally as

$$\frac{\partial v_i}{\partial t} = V_i + M_i + K_i + C_i, \quad (\text{B } 1)$$

where V_i , M_i and K_i are formal representations of the effects of viscous transport (mixing of parcel-pair states), swaps and kernels, respectively, and C_i is the mean-pressure-gradient contribution, which is constant. The arguments y and t of the terms in (B 1) are not shown. Constant density is assumed, so density is omitted throughout this appendix.

Here C_i is non-zero only in the streamwise direction $i = 1$. It does not appear in the ODT formal time-advancement equation (2001A1) because the analysis there is for flows with no imposed mean pressure gradient.

Each term on the right-hand side of (B 1), except C_i , is a sum of terms of the form $(\tilde{v}_i - \hat{v}_i)\delta(t - t_j)$, where the first factor is the v_i value at y after an advancement operation associated with an event at time t_j minus its value before the advancement operation. (In this appendix, the hat notation has a different meaning than in § 6.2.) What distinguishes V_i , M_i and K_i from one another is that each corresponds to a different type of HiPS advancement operation associated with the event at time t_j .

The HiPS mean momentum equation is obtained by ensemble averaging (B 1) term by term. In (B 1), K_i can be interpreted as some non-unique combination $K_i = T_i + S_i$ of transport and scrambling effects, subject to the requirement $\langle S_i \rangle = 0$ because there are no pressure-fluctuation contributions to the Navier–Stokes mean momentum equation. Then averaging of (B 1) gives

$$\frac{\partial \langle v_i \rangle}{\partial t} = \langle V_i \rangle + \langle M_i \rangle + \langle T_i \rangle + C_i, \quad (\text{B } 2)$$

where the angle brackets denote ensemble averaging. Thus it is recognized that the kernel operation introduces both instantaneous and mean transport of component- i momentum as well as instantaneous, but zero-mean (in the ensemble-average as well

as spatial-average sense), exchange of momentum among components. The derivation and interpretation of (2001A2) are analogous.

The flow structure shown in figure 2 indicates that HiPS provides a (non-uniformly) coarse-grained representation of channel flow. For this reason, statistics are gathered within homogeneous intervals, in effect treating these intervals as coarse-grained finite volumes. The statistically equivalent parcels that form a homogeneous interval can be interpreted as residing at the interval midpoint and are all equally representative of flow statistics within the interval. Only property transfers between homogeneous intervals can generate net fluxes. Fluxes such as the Reynolds shear stress are therefore evaluated at the boundaries between homogeneous intervals.

Viscous transport occurs only between adjacent parcels, so the only net viscous flux in the system is between parcels that are adjacent yet statistically non-equivalent. The only such parcel pairs are the wall parcels and their associated adjacent parcels, which are statistically non-equivalent due to the enforcement of the wall boundary condition as described in § 3. Therefore the term $\langle V_i \rangle$ in (B 2) is dropped. (Its effect at walls will be evaluated shortly.)

Based on these considerations, spatial derivatives are discretized as follows. Using the notation $I_i(y) \equiv \int_y^{y+2h} (\langle M_i \rangle + \langle T_i \rangle)$, (B 2) becomes

$$\frac{\partial \langle v_i \rangle}{\partial t} = -\frac{\partial I_i}{\partial y} + C_i. \quad (\text{B } 3)$$

Comparison to the mean momentum equation for Navier–Stokes channel flow identifies I_1 as the model analogue of the Reynolds shear stress $\langle v_1'v_2' \rangle$.

Integration of (B 3) over a homogeneous interval of width W gives

$$\frac{\partial W \langle v_i \rangle}{\partial t} = I_i(y_-) - I_i(y_+) + WC_i, \quad (\text{B } 4)$$

where spatial homogeneity implies that $\langle v_i \rangle$ is the interval as well as the ensemble average of v_i , and $[y_-, y_+]$ is the spatial extent of the homogeneous interval. Dividing (B 4) by W then identifies $(I_i(y_+) - I_i(y_-))/W$ as the appropriate discrete implementation of the y -derivative of I_i .

The integrand of the I_i definition contains ensemble averages of quantities that are sums of terms of the form $(\tilde{v}_i - \hat{v}_i)\delta(t - t_j)$. Such averages for fixed t are zero for a finite ensemble of numerical simulations and in the formal limit of a countably infinite ensemble. Therefore they are evaluated within finite time windows Δt that should be small for transient evolution but can be as large as desired during statistically steady time-advancement. Division of these averages by the time increment Δt yields an approximate (but sufficiently accurate for small Δt) evaluation of the fixed- t averages. ODT implementation based on (2001A5) is analogous.

The $i = 1$ steady-state specialization of (B 4) is used to evaluate the model analogue I_1 of the Reynolds shear stress. As shown in figure 3, behaviour consistent with the channel-flow mean momentum balance is obtained, except for near-wall behaviour reflecting the simplified treatment of viscous transport. That behaviour is associated with the term $\langle V_i \rangle$ that was dropped in deriving (B 4). Mean viscous transport is non-zero only at the wall and is the only transport mechanism at the wall, so the Reynolds-shear-stress profile in figure 3 is affected only at $y = 0$, where it is zero.

An evolution equation for v_i^2 is obtained by multiplying (B 1) by $2v_i$. As in (B 1), all terms except the last require a formal treatment, now expressed as

$$\frac{\partial v_i^2}{\partial t} = V_{ii} + M_{ii} + K_{ii} + 2v_i C_i, \quad (\text{B } 5)$$

where each term on the right-hand side except the last is a sum of terms of the form $(\tilde{v}_i^2 - \hat{v}_i^2)\delta(t - t_j)$, representing the change of v_i^2 resulting from the viscous, swapping and kernel operations, respectively, associated with the j th advection event. Decomposing the kernel term into transport and scrambling contributions as in the derivation of (B 2), we obtain

$$\frac{\partial \langle v_i^2 \rangle}{\partial t} = \langle V_{ii} \rangle + \langle M_{ii} \rangle + \langle T_{ii} \rangle + \langle S_{ii} \rangle + 2\langle v_i \rangle C_i, \tag{B 6}$$

where the scrambling term $\langle S_{ii} \rangle$ can now be non-zero. The ODT procedure for modelling the non-unique quantity S_{ii} can be applied, but this is immaterial here because it cancels out of the sum over i in TKE budget equation of interest.

Unlike v_i , v_i^2 is not conserved by the mixing mechanism representing viscous momentum transport in HiPS. (This non-conservation reflects the physical mechanism of viscous dissipation of kinetic energy.) One consequence is that $\langle V_{ii} \rangle$ includes viscous dissipation as well as viscous transport of component- i TKE (see below) and hence cannot be omitted.

Subtraction of $2\langle v_i \rangle$ times (B 2) from (B 6) gives

$$\frac{\partial \langle v_i^2 \rangle}{\partial t} = \langle V_{ii} + M_{ii} + T_{ii} + S_{ii} \rangle - 2\langle V_i + M_i + T_i \rangle \langle v_i \rangle, \tag{B 7}$$

where the term involving C_i no longer appears due to cancellation. Addition and subtraction of $-2I_i(\partial \langle v_i \rangle / \partial y)$ on the right-hand side gives

$$\frac{\partial \langle v_i^2 \rangle}{\partial t} = \langle V_{ii} \rangle - 2\langle V_i \rangle \langle v_i \rangle - 2I_i \frac{\partial}{\partial y} \langle v_i \rangle - \frac{\partial}{\partial y} (I_{ii} - 2\langle v_i \rangle I_i) + \langle S_{ii} \rangle, \tag{B 8}$$

where $I_{ii}(y) \equiv \int_y^{2h} (\langle M_{ii} \rangle + \langle T_{ii} \rangle)$.

In (B 3), introduction of the y -derivative of a y -integrated quantity facilitates conservative spatial discretization in (B 4). There is no analogous procedure for unambiguous spatially discrete evaluation of the spatial derivatives involving $\langle v_i \rangle$. These derivatives were introduced into (B 8) by adding and subtracting $-2I_i(\partial \langle v_i \rangle / \partial y)$ in order to enable identification of the HiPS analogues of the turbulent production and turbulent transport terms of the TKE budget. The numerical error inherent in evaluating the spatial derivatives involving $\langle v_i \rangle$ has the effect of modifying the relative contributions of these two terms, but it does not compromise the numerical closure of the budget provided that the numerical method preserves the algebraic cancellation. In numerical implementation, the term that is expressed as a y -derivative to identify it as turbulent transport is therefore decomposed by applying the derivative term by term so that the term that cancels the turbulent production term (which is immediately to the left of the turbulent transport term) is isolated. In order to enforce the cancellation required for numerical closure of the budget, the algebraically cancelling terms are identically evaluated by applying first-differencing within the spatially non-uniform finite-volume representation of the grouping of parcels into homogeneous intervals. However, enforcing this cancellation to obtain numerical closure is incompatible with numerically evaluating the turbulent transport term as an exact differential. Therefore that term does not numerically integrate to zero. Alternatively, numerical integration of that term to zero could be enforced, but then the budget would not close.

This tradeoff arises in any numerical evaluation of the TKE budget. Because the physical model in this case imposes piecewise-constant statistical properties in finite

intervals, the usual remedy of driving the integral of the turbulent transport toward zero through mesh refinement is not possible. This explains the noticeable non-zero value of this integral in figure 9. A different choice of numerical approximation of the y -derivative would modify the profile but would not eliminate the anomaly.

It has been noted that advection (i.e. swaps) within homogeneous intervals is immaterial and therefore can be omitted from the TKE budget evaluations. The kernel-induced spatial redistribution of the energy of a given component, represented by the T portions of the quantities I_i and I_{ii} , is a fluxing operation and therefore likewise immaterial within homogeneous intervals. However, kernels also redistribute energy among velocity components, as represented by the scrambling term $\langle S_{ii} \rangle$ in the budget (B 8) of the variance of an individual velocity component. Therefore this term is included when forming the variance budget of one velocity component. However, all kernel-induced effects cancel in the summation over components that yields the budget of TKE.

Due to this cancellation, the TKE budget is insensitive to any modelling error associated with the kernel representation of pressure-fluctuation effects. TKE budget results are more accurate overall than the results for the variances of individual velocity components. This provides indirect evidence that the kernel mechanism introduces significant modelling error.

For the same reasons that $\langle V_i \rangle$ vanishes away from walls, there is no mean viscous transport of the variance of v_i away from walls. Therefore the term $\langle V_{ii} \rangle$ represents solely the viscous dissipation, except for wall effects that are treated as follows.

A parcel adjacent to a wall parcel is subject to mixing with the wall parcel followed by a change of the wall-parcel state through enforcement of the Dirichlet boundary condition. The latter step imposes a difference between the mean states of the two parcels. The mixing-induced flux represents viscous transport, so in (B 2) it introduces a non-zero value of $\langle V_1 \rangle$ that exactly balances the mean advective flux to the parcel adjacent to the wall. Likewise, $\langle V_i \rangle$ in (B 8) represents viscous transport of mean momentum and is non-zero only for $i = 1$.

In contrast, the term $\langle V_{ii} \rangle$ of the v_i variance budget has a non-zero viscous-transport contribution for all i at the wall. As in the evaluation of $\partial \langle v_i \rangle / \partial y$, there is a degree of numerical ambiguity in the allocation of $\langle V_{ii} \rangle$ to viscous transport and viscous dissipation. To enable this allocation, it is assumed that mixing-induced changes of wall-parcel states contribute solely to viscous transport.

The boundary-condition enforcement has an equal-and-opposite effect that is attributed to viscous dissipation, which is required in order to balance viscous transport at the wall. The parcel adjacent to the wall is assigned viscous transport equal and opposite to viscous transport at the wall so that viscous transport integrates to zero. Thus, the ambiguity in the allocation of $\langle V_{ii} \rangle$ provides a degree of freedom that is used to enforce the integral constraint on the viscous transport. No such degree of freedom is available for the turbulent-transport term.

Finally, summation of (B 8) over i and multiplication by $1/2$ gives the TKE budget. In the statistically steady state, the time-derivative term is zero.

Based on these considerations, the budgets shown in figures 8 and 9 are obtained. Budget terms are normalized by dividing them by u_τ^4/ν . As in all reported results, adjusted values are used. As explained above, fluxes are evaluated by summing property changes during a finite time window and dividing by the time increment. Because the adjustment described in § 6.2 is a time rescaling, the adjustment is applied to this time increment.

Appendix C. Extension to $A \neq 0.5$

The mathematical formulation of HiPS is indifferent to the value of the ratio A of the length scales associated with successive tree levels. However, the choice of A has a significant impact on the physical interpretation of HiPS, as shown previously in a general context (Kerstein 2013). Here, implications specific to HiPS channel-flow simulation are explained.

A conceptually simple physical interpretation for general A would be to assign the size of non-wall parcels so that their sum is $2h$ and to treat them as a contiguous set covering the interval $2h$. As shown below, for $A < 0.5$ this would imply physical lengths of homogeneous intervals that are inconsistent with the length scales l_n corresponding to the wall proximities of the successive homogeneous intervals. This would not preserve the scale similarity that is the basis for the physical modelling of event rates and associated scale breakdown in accordance with inertial-range phenomenology.

In § 2.1, the parcel size is assigned to be the length scale l_{N-1} corresponding to the wall proximity of the parcel adjacent to the wall. (If any multiple of this length scale were chosen, the prefactor would be absorbed into the relationship between l_0 and h because parcel sizes are required in § 4 to sum to $2h$ for $A = 0.5$.) Likewise, each homogeneous interval is assigned its wall-proximity length scale. This is consistent with the physical-space interpretation that the sum of the sizes of these intervals is $2h$. Thus, the homogeneous intervals exactly cover the one-dimensional physical domain, but for $A < 0.5$ the parcels only partially cover it.

To see this, note first that the sizes of the homogeneous intervals (largest to smallest) on one side of the channel midpoint are $Al_0, A^2l_0, A^3l_0, \dots, A^{N-1}l_0$. Their sum is $((A - A^N)/(1 - A))l_0$. Setting this equal to h gives $l_0 = ((1 - A)/(A - A^N))h$. In practice it is convenient to choose a value of the model parameter l_0 and then determine the corresponding value of h using $h = ((A - A^N)/(1 - A))l_0$. For $A = 0.5$, this gives $h = (1 - 2^{1-N})l_0$.

The coverage of this size- h region by parcels is the product of the parcel size and the number of parcels in this region, giving $[(2A)^{N-1} - A^{N-1}]l_0$. For $A = 0.5$, this gives $(1 - 2^{1-N})l_0$, which is h for this A value. This reflects the exact coverage that has been enforced for $A = 0.5$. However, the sum of parcel sizes is less than the sum of homogeneous-interval sizes for $A < 0.5$.

Although the parcel size can be chosen to be some multiple of l_{N-1} , the smallest homogeneous interval, which is adjacent to a wall, cannot be smaller than one parcel; nor can it be larger because its midpoint must coincide with the midpoint of the parcel it contains, and one of its endpoints is bounded by the wall. This locks the relationship between parcel and interval sizes, reflecting the absorption of any multiplicative factor into the relation determining l_0 . The conclusion concerning the relative coverage of $[0, h]$ by parcels versus homogeneous intervals is thus general.

The incomplete coverage by parcels reflects the effective dimensionality $d < 1$ for $A < 0.5$ (see § 2.1). In physical terms, the size- $2h$ domain is deemed to consist of an active region (the region covered by parcels, which is now a disconnected set) and a void region containing no fluid. This picture is grossly unphysical in the limit of vanishing A ; however, for A values less than 0.5 but of order unity, it is possible in principle for this picture to yield reasonable results. The performance of HiPS channel-flow simulation for $A < 0.5$ is examined in § 8.2.

The fractional coverage of the y domain by parcels is $[(2A)^{N-1} - A^{N-1}]/(2A - A) \cdot [(1 - A)/(1 - A^{N-1})]$. Therefore only this fraction of the domain is subject to pressure-gradient forcing, modifying the momentum balance accordingly. In particular, (5.1) and (5.2) are modified by multiplying h by this fraction.

Viscous transport occurs only between wall parcels and the parcels adjacent to them. Elsewhere, momentum production is balanced solely by advective transport. As noted in § 8.1, the mean momentum balance then implies the relationship $\overline{u'v'^+} = 1 - y/h$ away from the wall, consistent with the behaviour in the region of physical channel flow where advective transport is dominant. This reasoning is specific to $A = 0.5$. The deviation from linear dependence of $\overline{u'v'^+}$ on y for $A < 0.5$ is shown and discussed in § 8.2.

For $A \leq 0.5$, parcel proximity provides a basis for associating parcels with locations along the y coordinate. On this basis, the covered region within a homogeneous interval is self-affine, as noted in § 8.2. This has no physical relevance for present purposes because parcel statistics within a homogeneous interval are aggregated for data-reduction purposes and are associated with the y location at the midpoint of the interval. Based on the results above, the interval midpoints are

$$y_j = \frac{h}{2} \frac{A^{1-j} + A^{-j} - 2}{A^{1-N} - 1} = \frac{l_0}{2} \frac{A^N}{1 - A} (A^{-j} + A^{1-j} - 2),$$

where the rightmost expression is more convenient because l_0 is a model input whereas h is an inferred quantity.

For the largest homogeneous interval, corresponding to $j = N - 2$, y_j approaches $h/2$ in the limit of vanishing A . This reflects the diverging multiplicative increase of interval size with increasing j , such that the last interval fills the entire range $[0, h]$ in this (unphysical) limit, with midpoint $h/2$.

One might suppose that $A > 0.5$, corresponding to higher effective dimensionality of the set of parcels, is physically more realistic than $A \leq 0.5$. Here $A = 2^{-1/2}$, corresponding to dimension $d = 2$, might seem especially promising in this regard. However, this is not the case because of the manner in which wall effects are represented. For any N , the number of wall parcels remains the same and hence constitutes a zero-dimensional set, i.e. a set of isolated points. Only the effectively one-dimensional case $A = 0.5$ is fully consistent with this zero-dimensional wall representation. Owing to the lack of conceptual advantages of $A > 0.5$ and the increase in computational cost associated with increasing A , results are reported in § 8 only for $A \leq 0.5$.

REFERENCES

- DEL ÁLAMO, J. C., JIMÉNEZ, J., ZANDONADE, P. & MOSER, R. D. 2004 Scaling of the energy spectra of turbulent channels. *J. Fluid Mech.* **500**, 135–144.
- ALFREDSSON, P. H., ÖRLÜ, R. & SCHLATTER, P. 2011 The viscous sublayer revisited – exploiting self-similarity to determine the wall position and friction velocity. *Exp. Fluids* **51**, 271–280.
- AURELL, E., DORMY, E. & FRICK, P. 1997 Binary tree models of high-Reynolds-number turbulence. *Phys. Rev. E* **56**, 1692–1698.
- BENZI, R., BIFERALE, L. & TROVATORE, E. 1997 Ultrametric structure of multiscale energy correlations in turbulent models. *Phys. Rev. Lett.* **79**, 1670–1673.
- BIFERALE, L., CALZAVARINI, E. & TOSCHI, F. 2011 Multi-time multi-scale correlation functions in hydrodynamic turbulence. *Phys. Fluids* **23**, 085107.
- BUSCHMANN, M. H. & GAD-EL-HAK, M. 2003 Generalized logarithmic law and its consequences. *AIAA J.* **41**, 40–48.
- ECKERT, E. R. G. & DRAKE, R. M. 1972 *Analysis of Heat and Mass Transfer*. McGraw-Hill.
- FERNHOLZ, H. H. & FINLEY, P. J. 1996 The incompressible zero-pressure-gradient turbulent boundary layer: an assessment of the data. *Prog. Aerosp. Sci.* **32**, 245–311.

- FOX, R. O. 2003 *Computational Models for Turbulent Reacting Flows*. Cambridge University Press.
- FRISCH, U. 1995 *Turbulence: The Legacy of A. N. Kolmogorov*. Cambridge University Press.
- GONZALEZ-JUEZ, E., KERSTEIN, A. R. & LIGNELL, D. O. 2011 Fluxes across double-diffusive interfaces: a one-dimensional-turbulence study. *J. Fluid Mech.* **677**, 218–254.
- GURVICH, A. S. & YAGLOM, A. M. 1967 Breakdown of eddies and probability distributions for small-scale turbulence. *Phys. Fluids Suppl.* **10**, S59–S65.
- HOYAS, S. & JIMÉNEZ, J. 2006 Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Phys. Fluids* **18**, 011702.
- KADANOFF, L. P. 1995 A model of turbulence. *Phys. Today* **48**, 11–13.
- KERSTEIN, A. R. 1999 One-dimensional turbulence: model formulation and application to homogeneous turbulence, shear flows, and buoyant stratified flows. *J. Fluid Mech.* **392**, 277–334.
- KERSTEIN, A. R. 2009 One-dimensional turbulence stochastic simulation of multi-scale dynamics. In *Interdisciplinary Aspects of Turbulence*, Lecture Notes in Physics, vol. 756, pp. 291–333. Springer.
- KERSTEIN, A. R. 2013 Hierarchical parcel-swapping representation of turbulent mixing. Part 1. Formulation and scaling properties. *J. Stat. Mech.* **153**, 142–161.
- KERSTEIN, A. R., ASHURST, W. T., WUNSCH, S. & NILSEN, V. 2001 One-dimensional turbulence: vector formulation and application to free shear flows. *J. Fluid Mech.* **447**, 85–109.
- KERSTEIN, A. R. & DREEBEN, T. D. 2000 Prediction of turbulent free shear flow statistics using a simple stochastic model. *Phys. Fluids* **12**, 418–424.
- KERSTEIN, A. R. & WUNSCH, S. 2006 Simulation of a stably stratified atmospheric boundary layer using one-dimensional turbulence. *Boundary-Layer Meteorol.* **118**, 325–356.
- LAW, A. M. & KELTON, W. D. 2000 *Simulation Modeling and Analysis*. McGraw-Hill.
- LENAERS, P., LI, Q., BRETHOUWER, G., SCHLATTER, P. & ÖRLÜ, R. 2012 Rare backflow and extreme wall-normal velocity fluctuations in near-wall turbulence. *Phys. Fluids* **24**, 035110.
- LIGNELL, D. O., KERSTEIN, A. R., SUN, G. & MONSON, E. I. 2013 Mesh adaption for efficient multiscale implementation of one-dimensional turbulence. *Theor. Comput. Fluid Dyn.* **27**, 273–295.
- L'VOV, V. S., PODIVILOV, E. & PROCACCIA, I. 1997 Temporal multiscaling in hydrodynamic turbulence. *Phys. Rev. E* **55**, 7030–7035.
- MANDELBROT, B. B. 1985 Self-affine fractals and fractal dimension. *Phys. Scr.* **32**, 257–260.
- MENEVEAU, C. & SREENIVASAN, K. R. 1991 The multifractal nature of turbulent energy dissipation. *J. Fluid Mech.* **224**, 429–484.
- MOSER, R. D., KIM, J. & MANSOUR, N. N. 1999 DNS of turbulent channel flow up to $Re_\tau = 590$. *Phys. Fluids* **11**, 943–945.
- ÖRLÜ, R. & SCHLATTER, P. 2011 On the fluctuating wall shear stress in zero pressure-gradient turbulent boundary layer flows. *Phys. Fluids* **23**, 021704.
- PANDIT, R., RAY, S. S. & MITRA, D. 2008 Dynamic multiscaling in turbulence. *Eur. Phys. J. B* **64**, 463–469.
- POPE, S. B. 1985 Pdf methods for turbulent reactive flows. *Prog. Energy Combust. Sci.* **11**, 119–192.
- POPE, S. B. 2013 A model for turbulent mixing based on shadow-position conditioning. *Phys. Fluids* **25**, 110803.
- SCHLICHTING, H. & GERSTEN, K. 2000 *Boundary Layer Theory*. Springer.
- SCHMIDT, R. C., KERSTEIN, A. R., WUNSCH, S. & NILSEN, V. 2003 Near-wall LES closure based on one-dimensional turbulence modeling. *J. Comput. Phys.* **186**, 317–355.
- SCHULTZ, M. P. & FLACK, K. A. 2013 Reynolds-number scaling of turbulent channel flow. *Phys. Fluids* **25**, 025104.
- SREENIVASAN, K. R. & STOLOVITZKY, G. 1995 Turbulent cascades. *J. Stat. Phys.* **78**, 311–333.
- SUBRAMANIAM, S. & POPE, S. B. 1998 A mixing model for turbulent reactive flows based on Euclidean minimum spanning trees. *Combust. Flame* **115**, 487–514.
- ZANOUN, E.-S., NAGIB, H. & DURST, F. 2009 Refined c_f relation for turbulent channels and consequences for high- Re experiments. *Fluid Dyn. Res.* **41**, 021405.